

SEQUENTIALLY OPTIMIZED PROJECTIONS IN X-RAY IMAGING

M. BURGER[†], A. HAUPTMANN[‡], T. HELIN[§], N. HYVÖNEN[¶], AND J.-P. PUSKA[¶]

Abstract. This work applies Bayesian experimental design to selecting optimal projection geometries in (discretized) parallel beam X-ray tomography assuming the prior and the additive noise are Gaussian. The introduced greedy exhaustive optimization algorithm proceeds sequentially, with the posterior distribution corresponding to the previous projections serving as the prior for determining the design parameters, i.e. the imaging angle and the lateral position of the source-receiver pair, for the next one. The algorithm allows redefining the region of interest after each projection as well as adapting parameters in the (original) prior to the measured data. Both A and D-optimality are considered, with emphasis on efficient evaluation of the corresponding objective functions. Two-dimensional numerical experiments demonstrate the functionality of the approach.

Key words. X-ray tomography, parallel beam tomography, optimal projections, Bayesian experimental design, A-optimality, D-optimality, sequential optimization

AMS subject classifications. 62K05, 65F22

1. Introduction. Tomographic image reconstruction is a classical inverse problem with its history spanning over 100 years since the seminal work of Radon [27]. It gives rise to a wealth of applications in modern imaging for which research remains active to date. The contemporary mathematical questions are often driven by practical challenges such as limited data — consider, e.g., having a limited number of imaging angles or a limited field-of-view. There exists a number of proposed methods for image reconstruction given such incomplete data, however, much less is known about how to optimally design the imaging configuration if there are constraints on, e.g., the number of angles, if individual projections do not cover all of the imaged object, or if there is a particular *region of interest* (ROI) that may vary as a function of time. Such design is desirable in any application where the exposure to radiation must be minimized or the acquisition of data is otherwise expensive.

Optimal experimental design (OED) is also an extensive field of research with a massive body of literature; for a general introduction we refer to [5]. In particular, OED in inverse problems has taken leaps during the last decade due to the increase in computational resources. In the present paper, we adopt the Bayesian paradigm to inverse problems, that is, the observed data is used to update the prior knowledge about the unknown into a posterior probability distribution [22, 31]. OED in Bayesian inference is formulated as minimization of a Bayesian cost with respect to the design parameters [10]; different cost functions correspond to different optimality criteria. Here we focus on the two most common approaches: *A-optimality*, where a typical quadratic loss is considered, and *D-optimality*, which maximizes the Kullback–Leibler divergence of the posterior distribution with respect to the prior.

[†]Friedrich-Alexander Universität Erlangen-Nürnberg, Department Mathematik, Cauerstrasse 11, 91058 Erlangen, Germany (martin.burger@fau.de). The work of MB was supported by the Bundesministerium für Bildung und Forschung under the project id 05M16PMB (MED4D).

[‡]University of Oulu, Research Unit of Mathematical Sciences, Oulu, Finland; University College London, Department of Computer Science, London, United Kingdom (andreas.hauptmann@oulu.fi). The work of AH was supported by the Academy of Finland (decision 312123).

[§]LUT University, School of Engineering Science, P.O. Box 20, FI-53851 Lappeenranta, Finland. The work of TH was supported by the the Academy of Finland (decisions 320082 and 326961).

[¶]Aalto University, Department of Mathematics and Systems Analysis, P.O. Box 11100, FI-00076 Aalto, Finland (nuutti.hyvonen@aalto.fi, juha-pekka.puska@aalto.fi). The work of NH and JP was supported by the Academy of Finland (decision 312124).

For general high-dimensional inference problems, the minimization of the Bayes cost requires a huge computational effort since just evaluating the cost corresponds to (numerically) integrating the considered cost function against the joint probability distribution of the unknown and the data. Considerable effort has been put into formulating approximative schemes to solve the corresponding optimization task; in the framework of inverse problems this work is discussed below. However, for linear Bayesian inference problems with Gaussian prior and likelihood, it is well-known that the integral defining the Bayes cost for A or D-optimality can be evaluated analytically as the trace or the determinant of the posterior covariance, respectively, leading to more attractive optimization problems [10]. This is the general setting considered in this work.

1.1. Our contribution. Our paper focuses on discussing benefits of OED in linear X-ray imaging. Our contributions to this subject are as follows:

- We propose a greedy exhaustive sequential method for optimizing the imaging configuration in X-ray tomography in the spirit of [11] and discuss its efficient numerical implementation. More precisely, our work proposes how to optimize the next imaging angle and the lateral position of the source-receiver pair if the geometric specifications of the previous projections are known. At each step, the posterior covariance corresponding to the previous projections is fed as the prior covariance for the optimization of the next imaging angle and the lateral position. We discuss both A and D-optimal designs. As the dimension of the new measurement at each step is considerably lower than that of the unknown, the computational cost of evaluating the necessary posterior traces and determinants can be significantly lowered by resorting to the Woodbury matrix identity and the matrix determinant lemma without employing Monte Carlo estimators (cf. [3, 4, 7]).
- If the covariance matrices for the (original) Gaussian prior and the additive Gaussian noise process are known, the sequentially optimal projections can be determined before performing any actual measurements. However, we also introduce an optimization algorithm based on a varying ROI. Indeed, the ROI can be adapted to the reconstruction after a new projection image becomes available either algorithmically or based on expert knowledge, and subsequently the next optimal projection can be determined accounting for the change in the ROI. A related idea for the A-optimality criterion is suggested in [11].
- We develop a data-driven variant of our optimization method that simultaneously adapts the (original) prior distribution to the data observed at each step. As a numerical example, we demonstrate that the method is able to identify the true correlation length used in generating an imaged target.

It should be emphasized that the forms of the optimization target functions corresponding to both A and D-optimality criteria are well-known [10]. However, linear X-ray imaging provides an interesting example for further analysis of the setting, not only due to its high-dimensionality, but also since the design space (i.e. the projection geometry) is continuous by nature, which leads to challenging computational considerations. We adopt a sequential approach to the optimization due to its computational tractability and in order to include feedback from the observed data in the design process. Notice that sequential optimization is in general suboptimal if the total number of projections is fixed *a priori* and no feedback from the measured data is included in the process [19].

1.2. Literature review. In recent years, OED for large-scale inverse problems has gained substantial attention; see, e.g., [2, 3, 4, 6, 12, 13, 14, 16, 20, 23, 24, 25]. In particular, we acknowledge that the study of A and D-optimal designs has been extended to infinite-dimensional Hilbert spaces in [1]. The computational feasibility of experimental design in large-scale problems is often based on linearization and the use of Monte Carlo trace estimators such as the ones developed in [7]. However, there is also an interesting avenue of research presented in [17, 18, 19], where the authors develop design methods based on dynamical programming.

Optimization of projection angles in X-ray imaging has previously been considered in [28]. The authors explore two interesting approaches for empirical A-optimal design in constrained problems based on training data: On the one hand, they solve a sparse design out of a predetermined set of angles, which is an approach aligned with previous literature. On the other hand, they propose a gradient-based optimization scheme. However, there seems to be no guarantee of being able to find a global optimum and the presented numerical experiments are limited to two angles. The present paper is closely related to [28] as our approach is also based on a prior. As novelties compared to [28], we consider the lateral position of the source-receiver pair as a second design parameter, our sequential design enables incorporating feedback from the observed data in the process, and the exhaustive optimization algorithm facilitates circumventing difficulties arising from the nonconvexity of the objective functions. In fact, our numerical examples clearly demonstrate that the optimization of the projection geometry in X-ray imaging suffers from the existence of local optima.

In terms of adaptivity, our work draws inspiration from [11], where A-optimal sequential design for dynamical inverse problems is considered. Our setup can be derived from [11] by assuming static noise-free dynamics. However, in [11] the adaptation of design is obtained by introducing a so-called monitor function that measures the difference in the last two reconstructions. In our framework, the approach of [11] would roughly correspond to choosing the ROI to be an area where such a difference is high.

From an application's point of view, our proposed method fits in the larger field of sparse tomography [15, 30], but moreover to applications with strict dose limitations and obstructions to the imaging region [9]. Such limitations might be due to safety reasons, experimental environment, or even expensive beamtime at synchrotron facilities [8]. For instance, in image-guided radiotherapy angles are chosen to radiate only a cancerous region of interest, without damaging essential surrounding tissue, using the knowledge obtained during the measurements, cf. [21]. Here, the proposed OED framework could provide insights into treatment planning and creation of automated algorithms to exclude critical regions from the beam direction.

This paper is organized as follows. In Section 2, we introduce our discretized parallel beam framework for two-dimensional X-ray imaging. Section 3 reviews the basics of finite-dimensional Bayesian OED and considers efficient numerical evaluation of the objective functions corresponding to A and D-optimality. The greedy exhaustive optimization algorithm and its adaptive variants are presented in Section 4. Finally, the numerical experiments are documented in Section 5, and some conclusions are drawn in Section 6. For completeness and to facilitate reading the paper, an appendix presents brief derivations of the target functions for A and D-optimal designs.

2. Measurement model and its discretization. The measurements are modeled as parallel beam tomography, where multiple parallel rays are directed into the object $D \subset \mathbb{R}^2$, and the resulting intensities of the rays are measured at detectors [26].

The attenuation is described by the equation

$$I = I_0 \exp\left(-\int_L u \, ds\right), \quad (2.1)$$

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 $u : D \rightarrow \mathbb{R}_+$ is the absorption. Obviously, (2.1) can equivalently be given as

$$\log(I_0) - \log(I) = \int_L u(x) \, ds.$$

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.

After discretizing the imaged domain into pixels, the forward operator, mapping the discretized absorption to a *single* set of log-intensity measurements at the detectors can be approximated by a matrix $R \in \mathbb{R}^{m \times n}$, where m and n are the number of detectors and the number of pixels, respectively (see, e.g., [30]); typically the dimension of the unknown is higher than the number of pixels in a *single* projection image, i.e. $m \ll n$. Each ray is parametrized by its (signed) distance s to the origin and its angle ϕ , say, relative to the positive horizontal axis. For a single projection image with m detectors, one thus needs to specify $m + 1$ parameters (s_1, \dots, s_m, ϕ) . We assume the rays/detectors to be equally spaced, but in our setting the complete parallel beam source-receiver pair does not typically cover the whole domain, that is, a single measurement with a given angle only produces a projection image of a part of D .

Take note that R is very sparse, which makes corresponding matrix multiplications relatively inexpensive computationally. Moreover, at least if $m \lesssim \sqrt{n}$, it is reasonable to expect that R has a full (effective) rank, i.e., $\text{rank}(R) = m$, with all its singular values being of the same order of magnitude. In particular, no projection matrix considered in our numerical experiments has a Euclidean norm based condition number that is higher than 2.

We assume the domain D consists of three distinct regions:

- *ROI*: the area about which we want to recover information,
- *Obstruction*: a nuisance region that obstructs the propagation of X-rays,
- *Background*: the rest of the object.

By resorting to this kind of a division, one can model settings where only a certain part of the target object is of interest. Such situations could arise, e.g., in a medical application if one has already determined the location of an organ, tumor or other abnormality that is to be further monitored. Moreover, in an industrial application one may only be interested in a certain component or section of some structure. In addition, the introduced division enables modeling obstructions that are difficult to image through, such as thick bone structures or other highly attenuating materials.

3. Bayesian inversion and experimental design. In Bayesian inversion all parameters carrying uncertainty are treated as random variables [22, 31]. The prior probability distributions for these parameters reflect the available information before the measurements are carried out. A measurement is modeled as a realization of a random variable depending on both the noise process and the random parameters in the forward model, as well as on the so-called design parameters that are deterministic

and define the measurement setup. The Bayes formula is then employed to form the posterior probability density that updates the prior based on the information in the measurement. In our setting, the design variables are the projection angle and the lateral position of the parallel beam source-receiver pair, and our ultimate aim is to (sequentially) choose these parameters according to either the A or D-optimality criterion of Bayesian experimental design [1, 10].

In this section, we only consider one step of the optimization process, that is, we assume to have a prior covariance (or its inverse) at hand and we consider optimally choosing the design parameters that define a single parallel beam projection. The computational efficiency is in the focus of our attention. The actual sequential optimization algorithm is presented in Section 4.

3.1. Bayes' formula. Let $y \in \mathbb{R}^m$ represent the (noisy) projection image and $p \in \mathbb{R}^2$ be the design variable, and suppose our prior information on the (discretized) absorption distribution is encoded in a probability density $\pi_{\text{pr}} : \mathbb{R}^n \rightarrow \mathbb{R}_+$. By Bayes' formula, the posterior density for the (randomized) absorption X reads

$$\pi(x | y; p) = \frac{\pi(y | x; p)\pi_{\text{pr}}(x)}{\pi(y; p)}, \quad x \in \mathbb{R}^n, \quad (3.1)$$

where $\pi(y | \cdot; p) : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is the so-called likelihood function and the normalizing term in the denominator is the marginal density of the measurement Y evaluated at the available data y .

In this work we assume the prior is Gaussian, i.e. $X \sim \mathcal{N}(x_{\text{pr}}, \Gamma_{\text{pr}})$, and the measurement can be modeled as a realization of the random variable

$$Y = R(p)X + N, \quad (3.2)$$

where $N \sim \mathcal{N}(0, \Gamma_{\text{noise}})$ is independent of X . Here, $\Gamma_{\text{pr}} \in \mathbb{R}^{n \times n}$ and $\Gamma_{\text{noise}} \in \mathbb{R}^{m \times m}$ are symmetric positive definite covariance matrices, $x_{\text{pr}} \in \mathbb{R}^n$ is the prior mean for X , and we have explicitly indicated the nonlinear dependence of the discrete measurement matrix $R(p)$ on the positioning of the parallel beam source-receiver pair. Under these simplifying assumptions, the posterior in (3.1) is also Gaussian with the covariance matrix and mean [22]

$$\Gamma_{\text{post}}(p) = (\Gamma_{\text{pr}}^{-1} + R(p)^T \Gamma_{\text{noise}}^{-1} R(p))^{-1}, \quad (3.3a)$$

$$\hat{x}(p) = \Gamma_{\text{post}}(p)(\Gamma_{\text{pr}}^{-1} x_{\text{pr}} + R(p)^T \Gamma_{\text{noise}}^{-1} y), \quad (3.3b)$$

respectively, as can be deduced by a straightforward completion of squares in (3.1). Using the Woodbury matrix identity, these equations can alternatively be written as

$$\Gamma_{\text{post}}(p) = \Gamma_{\text{pr}} - \Gamma_{\text{pr}} R(p)^T (R(p) \Gamma_{\text{pr}} R(p)^T + \Gamma_{\text{noise}})^{-1} R(p) \Gamma_{\text{pr}}, \quad (3.4a)$$

$$\hat{x}(p) = x_{\text{pr}} + \Gamma_{\text{pr}} R(p)^T (R(p) \Gamma_{\text{pr}} R(p)^T + \Gamma_{\text{noise}})^{-1} (y - R(p) x_{\text{pr}}), \quad (3.4b)$$

which are computationally far more attractive than (3.3) if $m \ll n$, as is the case in our setting. However, if one is interested in the inverse covariance $\Gamma_{\text{post}}(p)^{-1}$, then (3.3) should be used.

3.2. A and D-optimality. In Bayesian optimal experimental design, one often considers minimizing the expected squared distance of the unknown in a given (semi)norm around some chosen point estimate, which corresponds to the so-called

A-optimal design; in all our considerations, the point estimate of interest is the posterior mean. The other commonly used alternative is to look for the D-optimal design, that is, to maximize the information gain when the prior is replaced by the posterior. See, e.g., [1, 10] for more details.

Assuming the employed seminorm is induced by the positive semidefinite matrix $A^T A$ for a given $A \in \mathbb{R}^{l \times n}$, in our simple, i.e. Gaussian, linear and finite-dimensional, setting, A-optimality corresponds to choosing a design parameter $p_A \in \mathbb{R}^2$ satisfying

$$p_A = \arg \min_p \operatorname{tr}(A \Gamma_{\text{post}}(p) A^T). \quad (3.5)$$

On the other hand, D-optimality is equivalent to finding

$$p_D = \arg \min_p \log(\det \Gamma_{\text{post}}(p)), \quad (3.6)$$

where the logarithm follows from the definition of D-optimality but it also makes the numerical treatment of (3.6) considerably more stable as discussed in the following subsection. For completeness, the deductions of (3.5) and (3.6) are presented in Appendix A.

Since all pixels are of the same size in our discretized model, choosing $A = I \in \mathbb{R}^{n \times n}$ to be the identity matrix leads to using (a scaled) $L^2(D)$ -norm as the distance measure in (3.5). If one is only interested in the expected L^2 -accuracy of the posterior mean over some specific ROI within D , one simply only needs to replace the diagonal elements of A corresponding to the pixels in the complement of the ROI by zeros. On the other hand, being only interested in the information gain in the ROI for D-optimality corresponds to replacing $\Gamma_{\text{post}}(p)$ in (3.6) by the matrix obtained by deleting the rows and columns of $\Gamma_{\text{post}}(p)$ corresponding to the uninteresting pixels. We refer to Appendix A for the proof of this statement as well as for the exact form of the information gain when the prior is replaced by the posterior, as it equals $-\log \det(\Gamma_{\text{post}}(p))$ only up to an affine transformation.

3.3. Evaluation of the target functionals. Since the target functionals in (3.5) and (3.6) typically have multiple local minima, we choose a straightforward approach and perform an exhaustive search over a two-dimensional grid. In consequence, the ability to efficiently evaluate the optimization targets becomes the top priority. To this end, observe that the (effective) rank of the perturbation $R(p)^T \Gamma_{\text{noise}}^{-1} R(p)$ in (3.3) is presumably m since the measurement matrix $R(p)$ corresponding to a single parallel beam projection is expected to be of full rank with all its m singular values being of the same order of magnitude (cf. Section 2). Hence, our plan is to employ the matrix determinant lemma and the alternative formula for the posterior covariance in (3.3) in order to only need to compute traces of inverses and log-determinants for matrices of size $m \times m$. In particular, there is no apparent reason for resorting to Monte Carlo techniques, such as those in [3, 4, 7], because the low (effective) rank of the perturbation $R(p)^T \Gamma_{\text{noise}}^{-1} R(p)$ in (3.3) can be exploited explicitly.

With a suitable numbering of the pixels, the posterior covariance matrix can be written in a block form as

$$\Gamma_{\text{post}} = \begin{bmatrix} \Gamma_{\text{ROI}} & \Gamma_{\text{mix}} \\ \Gamma_{\text{mix}}^T & \Gamma_{\text{rest}} \end{bmatrix} = \begin{bmatrix} (\Gamma_{\text{post}})_{11} & (\Gamma_{\text{post}})_{12} \\ (\Gamma_{\text{post}})_{21} & (\Gamma_{\text{post}})_{22} \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad (3.7)$$

where $\Gamma_{\text{ROI}} = (\Gamma_{\text{post}})_{11}$ and $\Gamma_{\text{rest}} = (\Gamma_{\text{post}})_{22}$ are the (marginal) covariance matrices for the pixels in the ROI and the rest of the image, respectively. We use a similar

indexing for analogous block decompositions of other $n \times n$ matrices as well. Moreover, as in (3.7), we often simplify the notation by not explicitly marking the dependence on p .

REMARK 3.1. *Because our final product is a sequential optimization algorithm, the prior covariance for an optimization step is usually the optimized posterior from the previous one. Hence, having an explicit representation for $\Gamma_{\text{prior}}^{-1}$ at hand is arguably the most natural assumption for a single step of the algorithm; see (3.3). However, assuming m is of moderate size, Γ_{prior} can also be formed explicitly via (3.3) without consuming too much time, and thus we assume to explicitly know Γ_{prior} in the following. On the other hand, if the noise covariance is diagonal, as it is in all our numerical experiments, knowing Γ_{noise} is essentially the same as knowing $\Gamma_{\text{noise}}^{-1}$.*

3.3.1. D-optimality. As discussed in Section 3.2, the aim of D-optimal design with a preassigned ROI is to minimize

$$\Phi_{\text{D}}(p) := \log(\det \Gamma_{\text{ROI}}(p)) \quad (3.8)$$

over p in a certain subset of \mathbb{R}^2 . Recall that our aim is to efficiently evaluate $\Phi_{\text{D}}(p)$ at numerous $p \in \mathbb{R}^2$.

To begin with, observe that

$$\begin{aligned} \det(\Gamma_{\text{post}}) &= \det \begin{bmatrix} \Gamma_{\text{ROI}} & \Gamma_{\text{mix}} \\ 0 & \Gamma_{\text{rest}} - \Gamma_{\text{mix}}^T \Gamma_{\text{ROI}}^{-1} \Gamma_{\text{mix}} \end{bmatrix} \\ &= \det(\Gamma_{\text{ROI}}) \det(\Gamma_{\text{rest}} - \Gamma_{\text{mix}}^T \Gamma_{\text{ROI}}^{-1} \Gamma_{\text{mix}}), \end{aligned}$$

where Γ_{ROI} is invertible as a nonempty diagonal block of a positive definite matrix. In particular, $\Gamma_{\text{rest}} - \Gamma_{\text{mix}}^T \Gamma_{\text{ROI}}^{-1} \Gamma_{\text{mix}}$ is the Schur complement of Γ_{ROI} , meaning that

$$\Gamma_{\text{post}}^{-1} =: \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} = \begin{bmatrix} * & * \\ * & (\Gamma_{\text{rest}} - \Gamma_{\text{mix}}^T \Gamma_{\text{ROI}}^{-1} \Gamma_{\text{mix}})^{-1} \end{bmatrix}.$$

Altogether we thus have

$$\det(\Gamma_{\text{ROI}}) = \frac{\det(\Gamma_{\text{post}})}{\det(\Gamma_{\text{rest}} - \Gamma_{\text{mix}}^T \Gamma_{\text{ROI}}^{-1} \Gamma_{\text{mix}})} = \frac{\det(\Sigma_{22})}{\det(\Sigma)},$$

and so one needs to consider evaluating the logarithms of $\det(\Sigma)$ and $\det(\Sigma_{22})$, both of which depend on the design parameter p .

In the rest of this section, we explicitly mark which matrices depend on p . According to (3.3),

$$\begin{aligned} \Sigma(p) &= \Gamma_{\text{pr}}^{-1} + R(p)^T \Gamma_{\text{noise}}^{-1} R(p), \\ \Sigma_{22}(p) &= (\Gamma_{\text{pr}}^{-1})_{22} + R_2(p)^T \Gamma_{\text{noise}}^{-1} R_2(p), \end{aligned}$$

where $R(p) = [R_1(p), R_2(p)] \in \mathbb{R}^{m \times n}$ is a columnwise partitioning, with $R_1(p)$ corresponding to the ROI and $R_2(p)$ to the rest of the image. By virtue of the matrix determinant lemma,

$$\begin{aligned} \det(\Sigma(p)) &= \det(R(p) \Gamma_{\text{pr}} R(p)^T + \Gamma_{\text{noise}}) \det(\Gamma_{\text{noise}}^{-1}) \det(\Gamma_{\text{pr}}^{-1}), \\ \det(\Sigma_{22}(p)) &= \det(R_2(p) ((\Gamma_{\text{pr}}^{-1})_{22})^{-1} R_2(p)^T + \Gamma_{\text{noise}}) \det(\Gamma_{\text{noise}}^{-1}) \det((\Gamma_{\text{pr}}^{-1})_{22}), \end{aligned}$$

where only the first determinants, respectively, depend on p . Moreover, the matrices associated to these p -dependent determinants are only of size $m \times m$, i.e. small.

To evaluate $\log \det(\Gamma_{\text{ROI}}(p))$, one thus needs to precompute Cholesky decomposition for the (large) positive definite matrix $(\Gamma_{\text{pr}}^{-1})_{22}$ that is independent of p . Since m , i.e. the number of columns in $R(p)^T$, is assumed to be relatively low, this further enables building the (small) Cholesky decompositions

$$C(p)C(p)^T = R(p)\Gamma_{\text{pr}}R(p)^T + \Gamma_{\text{noise}}, \quad (3.9a)$$

$$\tilde{C}(p)\tilde{C}(p)^T = R_2(p)((\Gamma_{\text{pr}}^{-1})_{22})^{-1}R_2(p)^T + \Gamma_{\text{noise}} \quad (3.9b)$$

for all p on the employed grid. Finally, a straightforward algebraic manipulation gives

$$\log(\det(\Gamma_{\text{ROI}}(p))) = 2 \sum_{j=1}^m (\log(\tilde{c}_{jj}(p)) - \log(c_{jj}(p))) + c, \quad (3.10)$$

where $c_{jj}(p)$ and $\tilde{c}_{jj}(p)$, $j = 1, \dots, m$, are the diagonal elements of $C(p)$ and $\tilde{C}(p)$, respectively, and $c \in \mathbb{R}$ is independent of p .

Observe that it is far more stable to numerically evaluate the sum of logarithms in (3.10) than the products of the diagonal elements of the Cholesky factors needed for computing $\det(\Gamma_{\text{ROI}}(p))$ itself. Moreover, the constant c in (3.10) can be evaluated by considering Cholesky decompositions for $(\Gamma_{\text{pr}}^{-1})_{22}$ and either Γ_{pr} or Γ_{pr}^{-1} ; the former was already employed when building (3.9), and the latter neither poses any difficulties if n is not huge. In fact, one can even evaluate the actual information gain when the prior is replaced by the posterior for any p without too severely compromising the computational efficiency; see Appendix A for further details. This makes it possible to compare the information gains between different rounds of the sequential optimization algorithm introduced in Section 4 below.

3.3.2. A-optimality. As discussed in Section 3.2, when aiming at an A-optimal design, one needs to minimize

$$\Phi_A(p) := \text{tr}(A\Gamma_{\text{post}}(p)A^T)$$

over p in a certain subset of \mathbb{R}^2 for a given $A \in \mathbb{R}^{l \times n}$. In our numerical tests presented in Section 5, $A = I_{\text{ROI}} \in \mathbb{R}^{n \times n}$ is an identity matrix with the diagonal elements corresponding to the complement of the ROI replaced by zeros. Nevertheless, we perform the following calculations for an unspecified A to maintain generality.

As in the case of D-optimality, we introduce a Cholesky decomposition

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

and then form an auxiliary matrix

$$B(p) = C(p)^{-1}R(p)\Gamma_{\text{pr}}A^T \in \mathbb{R}^{m \times l}.$$

According to (3.4),

$$\Phi_A(p) = \text{tr}(A\Gamma_{\text{pr}}A^T - B(p)^TB(p)) = \text{tr}(A\Gamma_{\text{pr}}A^T) - \text{tr}(B(p)^TB(p)).$$

Hence, fundamental properties of the matrix trace allow the representation

$$\Phi_A(p) = c' - \sum_{j=1}^m \sum_{k=1}^l B_{jk}(p)^2.$$

Since $c' = \text{tr}(A\Gamma_{\text{pr}}A^T)$ does not depend on p , it does not affect the minimization of $\Phi_A(p)$. However, evaluating c' does not considerably slow down the minimization process as a whole because it needs only be done once; knowing c' allows comparing the minimal values of Φ_A between different iterations of the sequential minimization algorithm introduced in the following section.

4. Sequential optimization of measurements. In this section we present the algorithm for sequentially optimizing the parallel beam projections for X-ray tomography. We start with the basic algorithm that aims at finding the optimal projections prior to having any measurements at hand. Subsequently, we consider the modifications needed if one wants to estimate a hyperparameter in the prior based on the observed data or to change the ROI for the $(k + 1)$ th optimization round based on the reconstruction, i.e. the *conditional mean* (CM) estimate, computed from the first k projections.

4.1. Basic algorithm. The algorithm is initialized by determining the discretization of D into n pixels, choosing the spacing of the m detectors/rays for a single parallel beam projection, and specifying the covariance matrices Γ_0 and Γ_{noise} for the Gaussian prior of the absorption coefficient and for the Gaussian noise model, respectively.¹ One also needs to choose the ROI, select the number of optimization rounds, and determine how the design variable $p \in \mathbb{R}^2$ parametrizes a parallel beam projection. In our numerical experiments, the first component of p corresponds to the projection angle and the second one to the distance of the bisection of the parallel beam source-receiver pair from the origin.

The optimization is performed sequentially, i.e., so that the posterior probability density after the k th projection is used as the prior for choosing the $(k + 1)$ th projection. At each step the target functional for A or D-optimality is evaluated for all p on a two-dimensional optimization grid that also needs to be predetermined. The particular p yielding the minimum value for the considered minimization target is chosen as the optimal parameter, and subsequently the posterior covariance, i.e. the prior covariance for the next optimization step, is formed according to (3.3) or (3.4).

The algorithm proceeds altogether as follows:

ALGORITHM 1.

Choose the covariances Γ_0 and Γ_{noise} for the initial Gaussian prior $\mathcal{N}(x_0, \Gamma_0)$ and the noise model $\mathcal{N}(0, \Gamma_{\text{noise}})$. Select the ROI, the grid for p and the number of optimization rounds K . Set $\Gamma_{\text{pr}} = \Gamma_0$ and $P = []$.

for $k = 1, \dots, K$ **do**

for all p on the optimization grid **do**

Evaluate $\Phi_D(p)$ or $\Phi_A(p)$ as outlined in Section 3.3.

end for

Find the minimizer p_k of the considered optimization target.

Append $P = [P, p_k]$.

Form the posterior covariance $\Gamma_{\text{post}}(p_k)$.

Set $\Gamma_{\text{pr}} = \Gamma_{\text{post}}(p_k)$.

end for

return P

The columns of the output matrix $P \in \mathbb{R}^{2 \times K}$ define the K sequentially optimized parallel beam projections. To be more precise, the $(k + 1)$ th column of P defines

¹The mean of the prior is not needed for the sequential optimization algorithm, unless it affects the noise model.

the A or D-optimal projection given the previous k projections, that is, each of the optimized projections is (only) *locally* optimal. In particular, there is no reason to expect that the found K parallel beam projections would be *globally optimal* [19]: It is a simpler and computationally less demanding task to optimize the projections one by one than to simultaneously find K parallel beam projections that are jointly A or D-optimal. However, there *is* anyway reason to expect that the projections defined by the columns of P are more optimal than, e.g., a random choice.

Due to the assumptions that the measurement model is linear with additive noise and the prior and noise process are independent Gaussians, Algorithm 1 can be run prior to performing any measurements. Hence, one may expect to have lots of computational time and resources for performing the sequential optimization of Algorithm 1 in an ‘offline mode’. However, it is also possible to change the ROI for the next optimization step or to estimate some free parameters in the prior based on the previous data or reconstruction as explained in the following two subsections.

4.2. Adaptive region of interest. If the sequentially optimal angles are not determined before the measurements but as a part of the online imaging procedure, the ROI may be altered between the optimization rounds based on the observations of the expert running the algorithm and, e.g., interesting or alarming features in the previous reconstruction. In this case, one also needs to give the initial prior mean as an input and incorporate the computation of the CM estimate in each step of the algorithm; see (3.3) and (3.4). Moreover, the natural output is no longer the sequence of optimal projections but the final reconstruction, and the outer iterations should be stopped by the operator of the algorithm.

ALGORITHM 2 (adaptive ROI).

Choose the covariances Γ_0 and Γ_{noise} as well as the mean x_0 for the initial Gaussian prior $\mathcal{N}(x_0, \Gamma_0)$ and the noise model $\mathcal{N}(0, \Gamma_{\text{noise}})$. Select the ROI and the optimization grid for p . Initialize $\Gamma_{\text{pr}} = \Gamma_0$, $x_{\text{pr}} = x_0$ and $k = 0$.

while *satisfactory reconstruction has not been reached* **do**

Set $k \leftarrow k + 1$.

for *all p on the optimization grid* **do**

Evaluate $\text{tr}(I_{\text{ROI}}\Gamma_{\text{post}}(p)I_{\text{ROI}})$ or $\log(\det(\Gamma_{\text{ROI}}(p)))$ as outlined in Section 3.3.

end for

Find the minimizer p_k of the considered optimization target.

Observe new data $y = R(p_k)x + n_k$, where n_k is a realization of the noise.

Form the posterior covariance $\Gamma_{\text{post}}(p_k)$ and the CM estimate $\hat{x}(p_k)$.

Redefine (heuristically) the ROI based on $\hat{x}(p_k)$.

Set $\Gamma_{\text{pr}} = \Gamma_{\text{post}}(p_k)$ and $x_{\text{pr}} = \hat{x}(p_k)$.

end while

Set $x_{\text{rec}} = x_{\text{pr}}$ and $\Gamma_{\text{rec}} = \Gamma_{\text{pr}}$.

return *x_{rec} and Γ_{rec}*

It follows from basic Bayesian analysis that the sequentially updated reconstruction $x_{\text{rec}} \in \mathbb{R}^n$ and the spread estimator Γ_{rec} produced by Algorithm 2 are the same one would obtain by first collecting the data for all optimized projection angles and only then computing the CM estimate and the posterior covariance in a single step.

4.3. Data-driven hyperparameter estimation. Suppose the *original* prior covariance $\Gamma_0 = \Gamma_0(\rho) \in \mathbb{R}^{n \times n}$ depends on a hierarchical parameter $\rho \in \mathbb{R}^l$, the true value of which is unknown. Our aim is to introduce an algorithm that alternates between determining the optimal design parameter and the maximum likelihood es-

timate for the hyperparameter based on the measured data. The latter step can be considered as empirically adapting the prior at the same time as the sequential design process is performed.

Denote by

$$\Gamma_{\text{ms}}(\rho; p) = \Gamma_{\text{noise}} + R(p)\Gamma_0(\rho)R(p)^T \in \mathbb{R}^{m \times m}$$

the covariance matrix of the measurement given the prior covariance $\Gamma_0(\rho)$ and the projection matrix $R(p)$. The marginal density of the measurement y conditioned on ρ is hence given by

$$\pi(y | \rho; p) \propto \exp\left(-\frac{1}{2}\left(\log(\det \Gamma_{\text{ms}}(\rho; p)) - (y - R(p)x_0)^T \Gamma_{\text{ms}}(\rho; p)^{-1} (y - R(p)x_0)\right)\right),$$

where $x_0 \in \mathbb{R}^n$ is the (original) prior mean. In consequence, having $k \in \mathbb{N}$ independent measurements y_1, \dots, y_k corresponding to the design parameters p_1, \dots, p_k at hand leads to

$$\begin{aligned} \pi(y_1, \dots, y_k | \rho; p_1, \dots, p_k) &= \prod_{i=1}^k \pi(y_i | \rho; p_i) \\ &\propto \exp\left(-\frac{1}{2} \sum_{i=1}^k \left(\log(\det \Gamma_{\text{ms}}(\rho; p_i)) - (y_i - R(p_i)x_0)^T \Gamma_{\text{ms}}(\rho; p_i)^{-1} (y_i - R(p_i)x_0)\right)\right). \end{aligned} \quad (4.1)$$

After obtaining the latest measurement y_k corresponding to the design parameter p_k , one can thus determine the latest *maximum likelihood* (ML) estimate ρ_k for the hyperparameter by maximizing (4.1) with respect to ρ .

In addition to the computational cost corresponding to maximizing (4.1), an extra price one needs to pay for the empirical estimation of a hyperparameter is the need to appropriately update the posterior for X to make it compatible with the current ML estimate ρ_k . In other words, one needs to recompute the posterior from a scratch: $\Gamma_{\text{post}}(\rho_k; p_1, \dots, p_k)$ is formed, e.g., via (3.3) or (3.4) with $\Gamma_{\text{pr}} = \Gamma_0(\rho_k)$, Γ_{noise} built from k diagonal blocks, and $R(p)$ replaced by the ‘total projection matrix’ $[R(p_1)^T, \dots, R(p_k)^T]^T$.

ALGORITHM 3 (data-driven hyperparameter estimation).

Choose the covariance Γ_{noise} for the noise model $\mathcal{N}(0, \Gamma_{\text{noise}})$ and the mean x_0 for the Gaussian prior. Initialize $\rho_0 \in \mathbb{R}^l$, $\Gamma_{\text{pr}} = \Gamma_0(\rho_0)$, and set $k = 0$. Select the ROI and the grid for p .

while *e.g., a satisfactory reconstruction has not been reached* **do**

Set $k \leftarrow k + 1$.

for *all p on the optimization grid* **do**

Evaluate $\Phi_{\text{D}}(p)$ or $\Phi_{\text{A}}(p)$ as outlined in Section 3.3.

end for

Find the minimizer p_k of the considered optimization target.

Observe the data $y_k = R(p_k)x + n_k$, where n_k is a realization of the noise.

Solve for the ML estimate $\rho_k = \arg \max_{\rho \in \mathbb{R}^l} \pi(y_1, \dots, y_k | \rho, p_1, \dots, p_k)$.

Form the posterior covariance $\Gamma_{\text{post}}(\rho_k; p_1, \dots, p_k)$ ‘‘from a scratch’’.

Set $\Gamma_{\text{pr}} = \Gamma_{\text{post}}(\rho_k; p_1, \dots, p_k)$.

end while

Computing the ML estimate ρ_k is relatively straightforward: When $k = 1$, one needs to pay some attention to make sure the employed minimization technique for

finding ρ_1 is not very slow and does not diverge in case there is not much information on the whereabouts of the true parameter value to begin with. However, when ρ_k for $k \in \mathbb{N}$ is available, it can be used as the initial guess for finding ρ_{k+1} by the Newton's method that exhibits fast convergence in our numerical experiments, where ρ is the correlation length for a certain parametrized covariance structure. If the elementwise derivatives of $\Gamma_0(\rho)$ are known, the required first two derivatives for the term in the exponent of $\pi(y_1, \dots, y_k \mid \rho, p_1, \dots, p_k)$ can be straightforwardly evaluated by resorting to differentiation formulas for the determinant and the matrix inverse. Altogether, finding the needed ML estimates poses no severe computational difficulties if m is relatively small and $l = 1$, as is the case in our numerical studies.

5. Numerical experiments. In all our tests, $D = [0, 1]^2$ is the unit square divided into a uniform mesh of $n = N^2$ pixels. The maximal distance of X-rays passing through D from its center point $(0.5, 0.5)$ is set to 0.5. If D contains an obstruction that blocks X-rays, the rows corresponding to the lines passing through the obstruction are deleted from the considered projection matrices, as the corresponding measurements do not carry any information on the unknown absorption and can thus be ignored. Hence, the number of active detectors may vary as a function of the position of the parallel beam source-receiver pair relative to the possible obstruction, which is taken into account when optimizing the optimal projections and computing related reconstructions. (It should be intuitively clear that optimal parallel beam projections try to avoid imaging through an obstruction whenever possible.)

The (initial) prior covariance between the remaining pixels is assumed to be of the form

$$(\Gamma_0)_{i,j} = \gamma^2 \exp\left(-\frac{\|x_i - x_j\|^2}{2\ell^2}\right), \quad (5.1)$$

where $\|\cdot\|$ denotes the Euclidean norm, $\ell > 0$ is the so-called correlation length and $\gamma > 0$ is the pixelwise standard deviation. Here x_i and x_j are the coordinates of the pixels (or more specifically the pixel centers) with indices i and j . Loosely speaking, the larger ℓ is, the less oscillating are random draws from the (initial) prior $\mathcal{N}(x_0, \Gamma_0)$. The components of the additive noise in (3.2) are assumed to be independent with a common standard deviation $\sigma > 0$.

5.1. Test 1: Basic algorithm. Our first numerical tests consider the basic version of our optimization routine, i.e. Algorithm 1. The aim is to demonstrate that the algorithm generates sequentially A-optimal parallel beam projections that are intuitively acceptable and clearly outperform random selections. Moreover, according to our tests, the optimization of projections can be performed using a considerably coarser discretization compared to the one used for computing the actual reconstructions, without significantly compromising the overall performance of the approach. Regarding D-optimality, it is numerically demonstrated that the sequential optimization scheme may sometimes lead to globally suboptimal sets of projections, although in many cases the sets of sequentially A and D-optimal projections are actually qualitatively very similar. For completeness it should also be admitted that the possibility that sequentially A-optimal projections may also sometimes suffer from apparent global nonoptimality cannot be excluded based solely on the presented results.

Test 1.1: Sanity check. Let us start by applying Algorithm 1 to the simplest possible setting: the ROI is D , the parallel beam source-receiver pair is of the maximal width 1, and there are no obstructions inside D . The number of pixels per edge of D is




Fig. 5.1: TEST 1.1. Left: Random draw from the zero-mean Gaussian prior with the covariance structure (5.1) for $\gamma = 1$ and $\ell = 0.05$. Right: Mean $L^2(D)$ reconstruction errors over a sample of 1000 absorption targets as functions of the number of (noisy) projections. The diamonds correspond to the A-optimal projections and the circles to the D-optimal ones. The errorbars give the averages and standard deviations of the corresponding mean $L^2(D)$ reconstructions errors for 1000 sequences of projections with random angles.

chosen to be $N = 100$ corresponding to altogether $n = 10^4$ pixels, and the number of individual detectors is $m = 45$ per a single projection image. The pointwise standard deviation of the prior covariance in (5.1) is assumed to be $\gamma = 1$ and the correlation length is chosen as $\ell = 0.05$. A random draw from the prior probability distribution of the absorption is visualized in the left-hand image of Figure 5.1. Here and in what follows, we assume the prior has zero mean, which is physically unrealistic as such, but mathematically this assumption simply corresponds to translating the coordinate system in \mathbb{R}^n by some given physically sensible, component-wise positive prior mean x_0 . The standard deviation of the additive white noise is $\sigma = 0.05$.

Figure 5.2 presents the A and D-optimality target functionals as functions of the projection angle over the first ten iterations of Algorithm 1. To be quite precise, a modified A-optimality target $\frac{1}{N}\sqrt{\Phi_A}$ is considered as it corresponds to the expected $L^2(D)$ reconstruction error, and in case of D-optimality, the depicted quantity is the actual information gain when the prior is replaced by the posterior. (The maximization of the information gain is equivalent to minimizing Φ_D , as explained in Appendix A.) For both optimality criteria, the optimal angles are distributed rather uniformly over the interval $[-90, 90]^\circ$ in such a way that the most recent projection angle is always located around the midpoint of the widest angular interval with no previous projections, which is intuitively what one would expect. It is noteworthy that after two projections the target functionals have multiple local optima; there is no reason to expect this would not be the case for more complicated imaging configurations as well.

The right-hand image of Figure 5.1 shows the mean $L^2(D)$ reconstruction errors after each of the first ten A and D-optimal projections over a sample of 1000 target absorptions drawn from the assumed prior. To be more precise, the noisy measurements are first simulated for each of the target absorptions and all optimal angles. Then, the reconstructions, i.e. posterior means, are formed and compared to the corresponding targets when including the simulated noisy projections in the reconstruction process one by one in the same order as they were introduced by Algorithm 1. The A and D-




Fig. 5.2: TEST 1.1. Left: Modified A-optimality target $\frac{1}{N}\sqrt{\Phi_A}$, corresponding to the expected $L^2(D)$ reconstruction error, as a function of the projection angle for the first ten projections. The dots denote the locally A-optimal angles and the algorithm proceeds from top to bottom (cf. Figure 5.1). Right: Information gain $-\Phi_D$ (cf. (A.4)) as a function of the projection angle for the first ten projections. The dots denote the optimal angles and the algorithm proceeds from bottom to top.

optimal angles seem to perform equally well on average, even though the mean $L^2(D)$ reconstruction error is precisely the quantity the A-optimal angles were designed to minimize.

For comparison, random projection angles are also considered: the mean $L^2(D)$ reconstruction errors over the considered sample of absorptions are computed for 1000 sequences of random projection angles, the components of which are picked independently from the uniform distribution over $[-90, 90]^\circ$. The right-hand image of Figure 5.1 illustrates the averages of the resulting mean $L^2(D)$ reconstruction errors together with their standard deviations over the considered random sequences of projection angles. According to Figure 5.1, the A and D-optimal projections clearly outperform the random selections — at least in this simple geometric setting.

Test 1.2: Discoidal ROI. The second test is essentially a repetition of the first one, but with a parallel beam source-receiver pair of narrower width 0.5 and the ROI specified to be a disk of radius 0.25 centered at $(0.6, 0.6)$; see the left-hand image of Figure 5.3. Observe, in particular, that the positioning of the source-receiver pair can now also be adjusted laterally and its width is the same as the diameter of the ROI. The standard deviation of the noise is set to $\sigma = 0.02$, and the number of individual detectors in the narrower parallel beam sensor is $m = 23$. The other parameters are as in Test 1.1, that is, $N = 100$, $\gamma = 1$ and $\ell = 0.05$.

Figures 5.4 and 5.5 visualize the first six sequentially A and D-optimal projections, respectively, with the corresponding pixelwise (posterior) standard deviations shown in the background. The A-optimal projections behave as intuitively as in Test 1.1: all projections approximately cover the ROI and they are also distributed relatively uniformly over all angles. Although the D-optimal projections also cover the ROI, they correspond to subsequent rotations of about 10° relative to one of the previous projections. Even though such projections are indeed locally D-optimal, they definitely do not form a globally D-optimal set of projections: for example, the six sequentially A-optimal projections in Figure 5.4 jointly produce a lower value for the D-optimality target Φ_D than the sequentially D-optimal ones in Figure 5.5. This demonstrates that our sequential optimization procedure does not always produce parallel beam

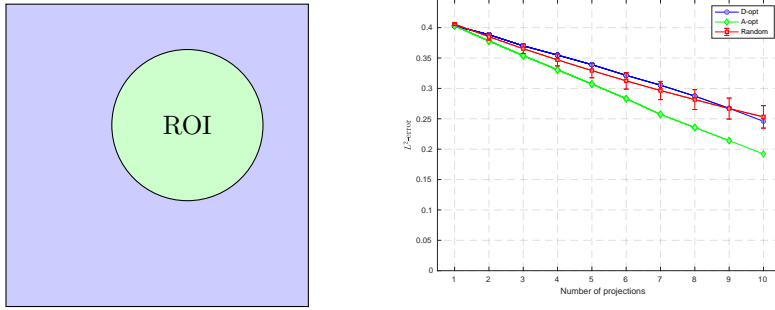


Fig. 5.3: TEST 1.2. Left: ROI. Right: Mean $L^2(D)$ reconstruction errors over a sample of 1000 absorption targets as functions of the number of (noisy) projections. The diamonds correspond to the A-optimal projections and the circles to the D-optimal ones. The errorbars give the averages and standard deviations of the corresponding mean $L^2(D)$ reconstruction errors for 1000 sequences of random projection that cover the ROI.

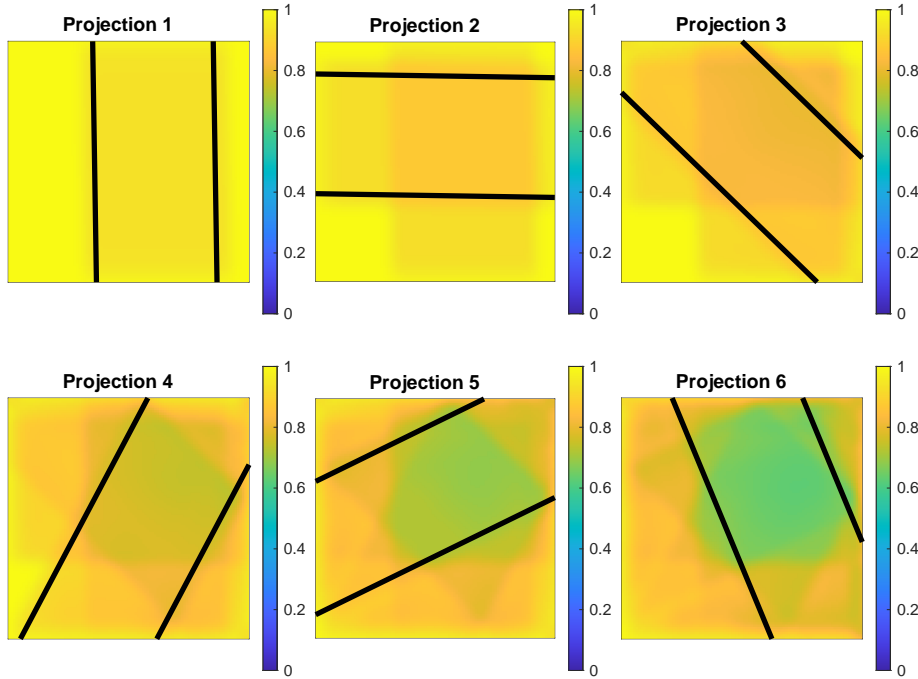


Fig. 5.4: TEST 1.2. The first six sequentially A-optimal parallel beam projections with the corresponding pixelwise (posterior) standard deviations in the background.

projections that are close to the global optimum.

The right-hand image of Figure 5.3 illustrates the mean $L^2(D)$ reconstruction errors after each of the first ten (noisy) A and D-optimal projections over a sample of 1000 target absorptions drawn from the assumed prior. As in Figure 5.1 of Test 1.1, these mean errors are compared with the averages and standard deviations of the corresponding mean $L^2(D)$ reconstruction errors for 1000 (semi)randomly se-

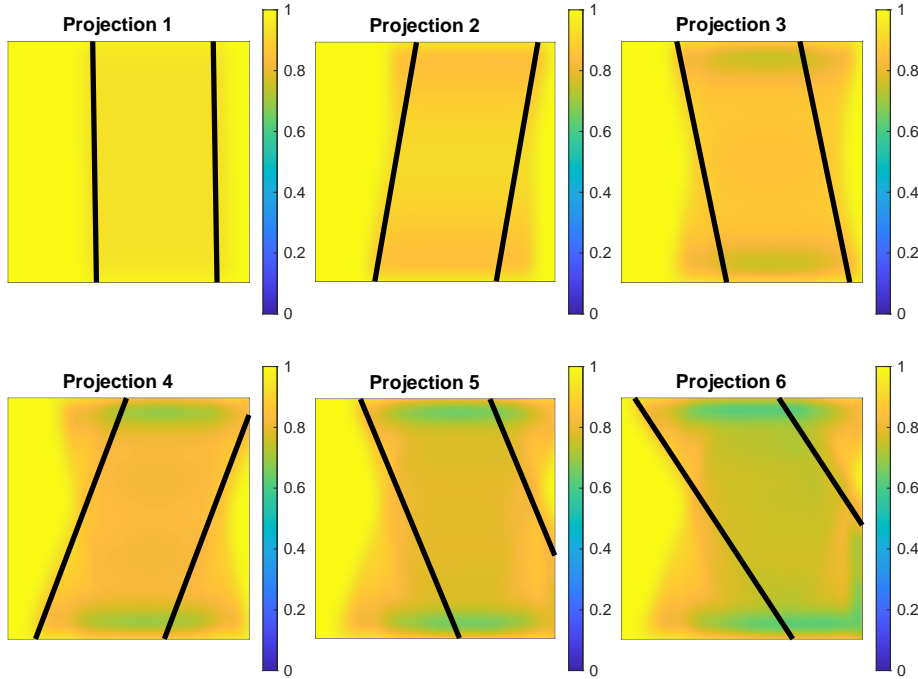


Fig. 5.5: TEST 1.2. The first six sequentially D-optimal parallel beam projections with the corresponding pixelwise (posterior) standard deviations in the background.

lected projection sequences. To be more precise, the angles of the projections in these sequences are picked randomly from the uniform distribution over $[-90, 90]^\circ$, but the lateral position of the parallel beam source-receiver pair is subsequently adjusted so that all projections cover the ROI in order to allow a fair comparison with the optimized projections. The sequentially A-optimal projections clearly produce the lowest mean $L^2(D)$ reconstruction errors, but this time around the random selections slightly outperform the sequentially D-optimal projections in the considered performance metric.

Test 1.3: Obstruction and effect of coarse discretization. Next we test Algorithm 1 with an obstruction inside D and also examine the effect that a coarse discretization has on the optimal angles. Only A-optimality is considered — the results for D-optimality would be qualitatively the same in this case. The obstruction blocking the X-rays is defined via $D_{\text{obst}} = \{x \in D \mid x_1 < 0.5, 0.45 < x_2 < 0.55\}$, with $D \setminus D_{\text{obst}}$ being the ROI; see the left-hand image in Figure 5.6. The parallel beam source-receiver pair has width 0.5, and the parameters defining the prior and the additive measurement noise are also the same as in the previous test: $\sigma = 0.02$, $\gamma = 1$ and $\ell = 0.05$. We employ two different levels of discretization: moderately dense, with $N = 100$ and $m = 23$, and very coarse, with $N = 25$ and $m = 6$.

Figure 5.7 presents the first six sequentially A-optimal projections for the denser discretization, with the corresponding pixelwise (posterior) standard deviations shown in the background. As expected, the algorithm tries to avoid imaging through the obstruction and emphasizes to begin with reducing the uncertainty in the right half




Fig. 5.6: TEST 1.3. Left: Obstruction. Right: Mean $L^2(D)$ reconstruction errors over a sample of 1000 absorption targets as functions of the number of (noisy) projections. The diamonds correspond to the A-optimal projections deduced using the dense discretization and the squares to those deduced using the coarse discretization. The reconstruction and corresponding errors were computed using the dense discretization.

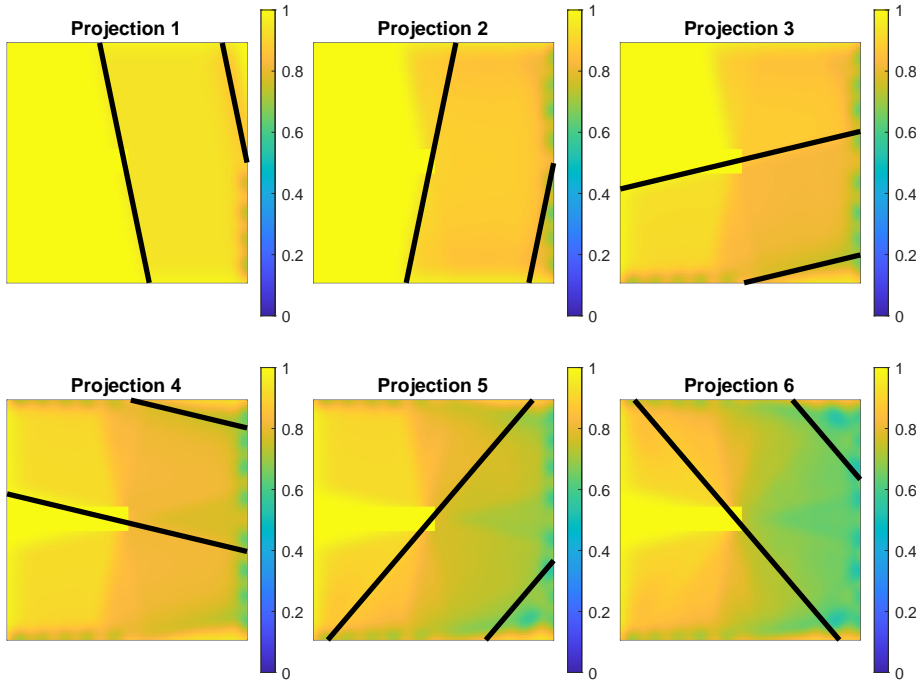


Fig. 5.7: TEST 1.3. The first six sequentially A-optimal parallel beam projections for the dense discretization with the corresponding pixelwise (posterior) standard deviations in the background.

of D , where imaging from almost all directions is possible. The projections produced by Algorithm 1 for the coarser discretization, not illustrated here, are qualitatively similar, but not exactly the same as the ones shown in Figure 5.7.

The right-hand image of Figure 5.6 depicts the mean $L^2(D)$ reconstruction errors after each of the first ten (noisy) A-optimal projections over a sample of 1000 target

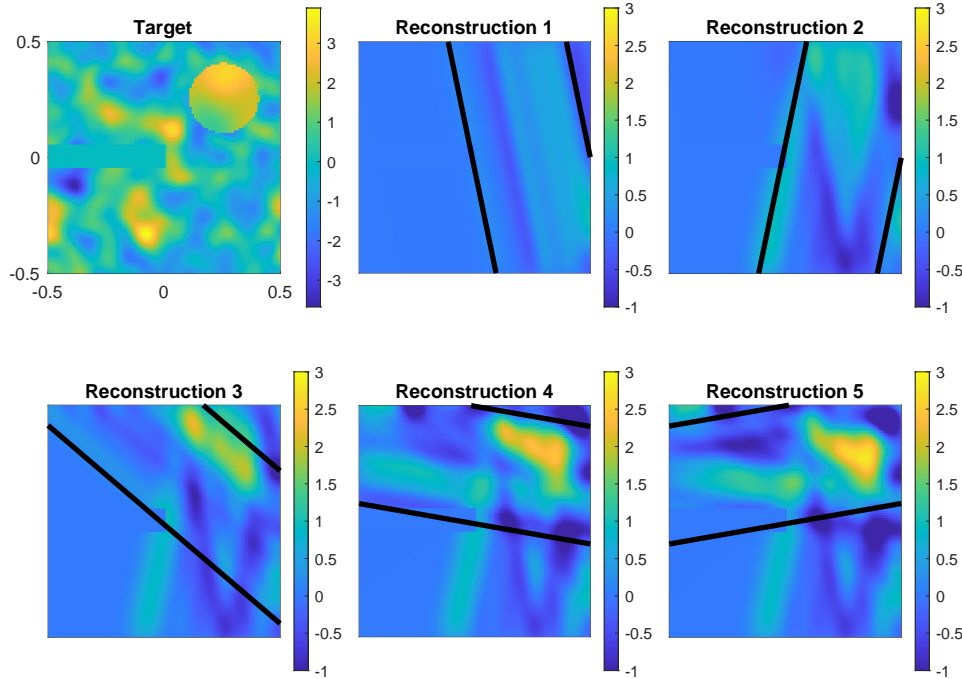


Fig. 5.8: TEST 2. Top left: True absorption and the obstruction D_{obst} . Rest of the top row: The first two sequentially A-optimal projections and the resulting reconstructions from noisy measurements. Bottom row: The next three sequentially A-optimal projection after redefining the ROI to be the top right quadrant of D . The resulting reconstructions from noisy measurements are shown in the background.

absorptions drawn from the assumed prior. Both sets of optimal projections, i.e. those deduced using the dense discretization ($n = 100$ and $m = 23$) and the ones deduced using the coarse discretization ($n = 25$ and $m = 6$), are considered. However, all reconstructions and the corresponding $L^2(D)$ errors are computed using the dense discretization. Based on the right-hand image of Figure 5.6, it seems quite obvious that the level of discretization used for finding the sequentially A-optimal projections does not play a major role in the performance of the overall approach — at least for the studied parameter values and the simple geometry.

5.2. Test 2: Adaptive region of interest. Let us then study a setting where the information on the location of the ROI is updated during the imaging process, that is, we test Algorithm 2. We adopt the geometry in the left-hand image of Figure 5.6 as well as the parameter values in (the denser discretization of) Test 1.3: $\gamma = 1$, $\ell = 0.05$, $\sigma = 0.02$, $N = 100$ and $m = 23$. The width of the parallel beam source-receiver pair is 0.5, and $D \setminus D_{\text{obst}}$ is the ROI to begin with.

The top left image in Figure 5.8 shows the true absorption inside D : in the top right corner, there is an inhomogeneity that is characterized by a higher level of absorption compared to the rest of the domain. The other two images in the top row of Figure 5.8 present the first two A-optimal projections and the corresponding reconstructions, i.e. the posterior means, computed from simulated noisy measure-

ments. Notice that these first projections are the same as the corresponding ones in Figure 5.6. Since the second reconstruction hints there may be something interesting going on in the top right quadrant of D , the ROI is redefined accordingly. The subsequent three A-optimal projections are illustrated in the bottom row of Figure 5.8 together with the resulting reconstructions. Comparing these results with those in Figure 5.6, it seems that the redefinition of the ROI had the desired effect as all of the last three projections clearly provide information on the top right quadrant of D .

5.3. Test 3: Simultaneous estimation of the prior correlation length.

Our final numerical experiment applies Algorithm 3 to deducing the value of an unknown correlation length $\rho = \ell$ in the prior covariance (5.1). The other free parameters are chosen to be $\gamma = 1$, $\sigma = 0.05$, $N = 75$ and $m = 39$. The parallel beam source-receiver pair has width 1, i.e. its position cannot be adjusted laterally, there are no obstructions inside D , and only A-optimality is considered. The presented numerical results would be qualitatively similar if the unknown parameter in the prior were instead the pointwise standard deviation or/and if D-optimality were used as the criterion for choosing the projection angles.

The following simple test is repeated 1000 times: A random value for the true correlation length ℓ is picked from the uniform distribution over $[0.04, 0.06]$ and a corresponding target absorption is subsequently drawn from the zero-mean Gaussian density with the prior covariance (5.1). Then, Algorithm 3 is run with the conservative initial guess $\ell_0 = 0.15$ for the correlation length. The estimate for the correlation length after the first projection, i.e. ℓ_1 , is determined by performing ten steps of the golden section line search over the interval $[0.01, 0.2]$ and subsequently applying the Newton's method to fine-tune the location of the optimum. When determining ℓ_k for $k \geq 2$, mere Newton's method is used with ℓ_{k-1} as the initial guess. In all cases, the Newton's method is stopped when its step size is less than 10^{-4} , and as a safety measure, Newton steps with absolute value larger than 0.01 are scaled down to 0.01 (retaining the sign of the step). No obvious problems with the convergence of the implemented minimization routine for finding ℓ_k were observed during computations.

The left-hand imaged in Figure 5.9 shows the mean signed errors in the estimate for ℓ during the first ten iterations of Algorithm 3 over the aforementioned sample of 1000 absorption targets. The corresponding standard deviations are visualized as errorbars. The mean signed error is almost zero independently of the number of (sequentially A-optimal) projections, whereas the standard deviation of the estimate over the sample of absorption targets decreases from approximately 0.008 after the first projection to about 0.002 after the tenth one. Altogether, the deduction of the prior correlation length based on noisy projection data seems to function well.

Certain mean $L^2(D)$ errors in the reconstructions computed based on noisy parallel beam projections of the studied sample of target absorptions are shown as functions of the number of projections in the right-hand image of Figure 5.9. The reconstructions and the corresponding $L^2(D)$ errors are computed in two different ways: (i) by accounting for the change in the estimate for the prior correlation length predicted by Algorithm 3 after each new sequentially A-optimal projection and (ii) by sticking with the conservative initial guess of 0.15 for the correlation length throughout the optimization of the projection angles and in the computation of the corresponding reconstructions. It is obvious that including the estimation of the correlation length as a part of the reconstruction algorithm leads to superior reconstruction accuracy, although this also makes the algorithm significantly more expensive computationally. If the initial guess ℓ_0 for the prior correlation length were (significantly) more ac-




Fig. 5.9: TEST 3. Left: Average signed error in the estimate for the prior correlation length over a sample of 1000 target absorptions as a function of the number of sequentially A-optimal projection angles. The errorbars indicate the intervals of two standard deviations computed over the studied sample. Right: Mean $L^2(D)$ reconstruction errors when the prior correlation length is simultaneously estimated by Algorithm 3 (squares) and the corresponding mean $L^2(D)$ errors if the conservative initial guess $l_0 = 0.15$ is used both for optimizing the projection angles and for forming the reconstructions (diamonds).

curate, it is debatable whether fine-tuning the estimate for ℓ would be worthwhile computationally.

6. Concluding remarks. We have introduced a greedy, exhaustive optimization algorithm that applies Bayesian OED to optimizing parallel beam projections in X-ray imaging. Although the sequential algorithm does not guarantee global optimality for the set of projections it predicts, our numerical experiments demonstrate that it usually exhibits satisfactory global behavior as well. More importantly, it fits to the medical imaging paradigm where the patient is exposed to the minimal amount of radiation without knowing *a priori* the necessary number of imaging angles. To this end, we have discussed modifications of the algorithm that allow tuning the ROI or some parameters in the (original) prior based on the measured data.

The main obstacle for adopting our algorithm in practical use is arguably its computational complexity: as the optimization target typically has many local minima (cf. Figure 5.2), we seek the optimal parameters via evaluating the target on a dense enough grid in order to guarantee finding the global optimum for the next projection. In a three-dimensional setting, such a simple approach becomes prohibitively expensive as the number of pixels in a single projection image can be in the order of $m = 10^5$, and thus a more sophisticated adaptive optimization routine needs to be developed for more realistic imaging geometries. On the positive side, the overall performance of our approach does not seem to depend heavily on the level of discretization, which may permit a procedure where the next projection geometry is optimized using a coarse discretization, but the actual reconstruction corresponds to a much denser one.

In addition, only a very simple Gaussian prior for the absorption distribution was employed in our numerical studies. The introduced algorithm should thus be tested with more realistic priors and, even more importantly, its functionality should be verified if no Gaussian prior is well aligned with the imaged target.

Appendix A. Target functionals for Bayesian experimental design. The purpose of this appendix is to deduce the minimization targets (3.5) and (3.6) and

to indicate how the latter is actually related to the information gain when a prior is replaced by the posterior. It is also explained why Γ_{post} should be replaced by $(\Gamma_{\text{post}})_{11} = \Gamma_{\text{ROI}}$ in (3.6) if one is only interested in the information gain over the ROI, cf. (3.7).

A.1. D-optimality. The Kullback–Leibler divergence of the distribution of X from that of $X | y$, i.e.,

$$\begin{aligned} D_{\text{KL}}(X | y \| X; p) &= \int_{\mathbb{R}^n} \log \left(\frac{\pi(x | y; p)}{\pi_{\text{pr}}(x)} \right) \pi(x | y; p) dx \\ &= \int_{\mathbb{R}^n} (\log(\pi(x | y; p)) - \log(\pi_{\text{pr}}(x))) \pi(x | y; p) dx, \end{aligned} \quad (\text{A.1})$$

has an interpretation as the increase in the level of information when the prior distribution of X is replaced by the posterior induced by the data y . The aim of D-optimal design is to maximize the expectation of $D_{\text{KL}}(X | y \| X; p)$ over all measurements, which means the precise form of the *minimization* target is

$$\tilde{\Phi}_D(p) = \mathbb{E}^y(-D_{\text{KL}}(X | y \| X; p)) = - \int_{\mathbb{R}^m} D_{\text{KL}}(X | y \| X; p) \pi(y; p) dy;$$

see (3.8) for comparison.

The expression

$$h(X | y) = - \int_{\mathbb{R}^n} \log(\pi(x | y; p)) \pi(x | y; p) dx \quad (\text{A.2})$$

appearing in (A.1) is called the *differential entropy* of $X | y$. Since $X | y$ is Gaussian, $h(X | y)$ allows an explicit representation [29]

$$h(X | y) = \frac{n}{2} + \frac{n}{2} \log(2\pi) + \frac{1}{2} \log(\det(\Gamma_{\text{post}}(p))). \quad (\text{A.3})$$

On the other hand, the expected value of the second term on the right-hand side of (A.1) satisfies

$$\begin{aligned} \mathbb{E}^y \left(\int_{\mathbb{R}^n} \log(\pi_{\text{pr}}(x)) \pi(x | y; p) dx \right) &= \int_{\mathbb{R}^m} \int_{\mathbb{R}^n} \log(\pi_{\text{pr}}(x)) \pi(x, y; p) dx dy \\ &= \int_{\mathbb{R}^n} \int_{\mathbb{R}^m} \pi(y | x; p) dy \log(\pi_{\text{pr}}(x)) \pi_{\text{pr}}(x) dx \\ &= \int_{\mathbb{R}^n} \log(\pi_{\text{pr}}(x)) \pi_{\text{pr}}(x) dx \\ &= -\frac{n}{2} - \frac{n}{2} \log(2\pi) - \frac{1}{2} \log(\det \Gamma_{\text{pr}}), \end{aligned}$$

where the last expression is simply $-h(X)$ for the Gaussian random variable X . In particular, take note that this expression does not depend on the design variable p .

Because $h(X | y)$ does not depend on y , we finally arrive at the simple expression

$$\begin{aligned} \tilde{\Phi}_D(p) &= \mathbb{E}^y(h(X | y)) - h(X) \\ &= \frac{1}{2} (\log(\det \Gamma_{\text{post}}(p)) - \log(\det \Gamma_{\text{pr}})), \end{aligned} \quad (\text{A.4})$$

which explains why the minimizers of $\tilde{\Phi}_D(p)$ and $\Phi_D(y)$ in (3.8) are the same if ROI is the whole D . Note that $-\tilde{\Phi}_D(p)$ corresponds to the actual information gain that is to be maximized.

Finally, observe that the above calculation can be repeated as such if $\pi_{\text{pr}}(x)$ and $\pi(x|y;p)$ are replaced by the corresponding marginal densities over the ROI. Hence, the precise form of the minimization target when the information gain is to be maximized over the ROI is obtained by replacing $\Gamma_{\text{post}}(p)$ by $(\Gamma_{\text{post}}(p))_{11}$ and Γ_{pr} by $(\Gamma_{\text{pr}})_{11}$ in (A.4), cf. (3.8).

A.2. A-optimality. To begin with, let $A \in \mathbb{R}^{l \times n}$ and observe that the CM estimate $\hat{x}(p) = \hat{x}(y;p) \in \mathbb{R}^n$, given in (3.3), naturally depends on the data $y \in \mathbb{R}^m$. The expected squared distance of the unknown X from $\hat{x}(Y;p)$ in the squared seminorm induced by $A^T A$ over the joint distribution of X and Y reads as follows:

$$\begin{aligned} \Phi_A(p) &= \int_{\mathbb{R}^m} \int_{\mathbb{R}^n} (x - \hat{x}(y;p))^T A^T A (x - \hat{x}(y;p)) \pi(x, y; p) \, dx \, dy \\ &= \int_{\mathbb{R}^m} \int_{\mathbb{R}^n} \text{tr}((x - \hat{x}(y;p))^T A^T A (x - \hat{x}(y;p))) \pi(x|y;p) \, dx \, \pi(y;p) \, dy \end{aligned}$$

as the trace of a scalar is the scalar itself. Because the trace is linear and invariant under cyclic permutations, we get

$$\begin{aligned} \Phi_A(p) &= \text{tr} \left(A \int_{\mathbb{R}^m} \int_{\mathbb{R}^n} (x - \hat{x}(y;p))(x - \hat{x}(y;p))^T \pi(x|y;p) \, dx \, \pi(y;p) \, dy \, A^T \right) \\ &= \text{tr} \left(A \int_{\mathbb{R}^m} \Gamma_{\text{post}}(p) \pi(y;p) \, dy \, A^T \right) \\ &= \text{tr}(A \Gamma_{\text{post}}(p) A^T), \end{aligned}$$

where the last step follows from $\Gamma_{\text{post}}(p)$ being independent of y in our simple, Gaussian and linear setting.

REFERENCES

- [1] ALEXANDERIAN, A., GLOOR, P. J., GHATTAS, O., ET AL. On Bayesian A- and D-optimal experimental designs in infinite dimensions. *Bayesian Anal.* 11, 3 (2016), 671–695.
- [2] ALEXANDERIAN, A., PETRA, N., STADLER, G., AND GHATTAS, O. A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized l_0 -sparsification. *SIAM J. Sci. Comput.* 36, 5 (2014), A2122–A2148.
- [3] ALEXANDERIAN, A., PETRA, N., STADLER, G., AND GHATTAS, O. A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems. *SIAM J. Sci. Comput.* 38, 1 (2016), A243–A272.
- [4] ALEXANDERIAN, A., AND SAIBABA, A. K. Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems. *SIAM J. Sci. Comp.* 40, 5 (2018), A2956–A2985.
- [5] ATKINSON, A., DONEV, A., TOBIAS, R., ET AL. *Optimum experimental designs, with SAS*, vol. 34. Oxford University Press, 2007.
- [6] ATTIA, A., ALEXANDERIAN, A., AND SAIBABA, A. K. Goal-oriented optimal design of experiments for large-scale Bayesian linear inverse problems. *Inverse Problems* 34, 9 (2018), 095009.
- [7] AVRON, H., AND TOLEDO, S. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *J. ACM* 58, 2 (2011), 1–34.
- [8] BATENBURG, J. K., DE CARLO, F., MANCINI, L., AND SIJBERS, J. Advanced x-ray tomography: experiment, modeling, and algorithms. *Meas. Sci. Technol.* 29, 8 (2018).

- [9] BORG, L., JØRGENSEN, J. S., FRIKEL, J., AND SPORRING, J. Reduction of variable-truncation artifacts from beam occlusion during in situ x-ray tomography. *Meas. Sci. Technol.* *28*, 12 (2017), 124004.
- [10] CHALONER, K., AND VERDINELLI, I. Bayesian experimental design: A review. *Stat. Sci.* (1995), 273–304.
- [11] FOHRING, J., AND HABER, E. Adaptive A-optimal experimental design for linear dynamical systems. *SIAM/ASA J. Uncertainty Quantification* *4*, 1 (2016), 1138–1159.
- [12] HABER, E., HORESH, L., AND TENORIO, L. Numerical methods for experimental design of large-scale linear ill-posed inverse problems. *Inverse Problems* *24*, 5 (2008), 055012.
- [13] HABER, E., HORESH, L., AND TENORIO, L. Numerical methods for the design of large-scale nonlinear discrete ill-posed inverse problems. *Inverse Problems* *26*, 2 (2009), 025002.
- [14] HABER, E., MAGNANT, Z., LUCERO, C., AND TENORIO, L. Numerical methods for A-optimal designs with a sparsity constraint for ill-posed inverse problems. *Comput. Opt. Appl.* *52*, 1 (2012), 293–314.
- [15] HÄMÄLÄINEN, K., KALLONEN, A., KOLEHMAINEN, V., LASSAS, M., NIINIMÄKI, K., AND SILTANEN, S. Sparse tomography. *SIAM J. Sc. Comput.* *35*, 3 (2013), B644–B665.
- [16] HANNUKAINEN, A., HYVÖNEN, N., AND PERKKIÖ, L. Inverse heat source problem and experimental design for determining iron loss distribution. *arXiv preprint arXiv:2003.10395* (2020).
- [17] HUAN, X., AND MARZOUK, Y. Gradient-based stochastic optimization methods in Bayesian experimental design. *Int. J. Uncertain. Quan.* *4*, 6 (2014).
- [18] HUAN, X., AND MARZOUK, Y. M. Simulation-based optimal Bayesian experimental design for nonlinear systems. *J. Comput. Phys.* *232*, 1 (2013), 288–317.
- [19] HUAN, X., AND MARZOUK, Y. M. Sequential Bayesian optimal experimental design via approximate dynamic programming. *arXiv preprint arXiv:1604.08320* (2016).
- [20] HYVÖNEN, N., SEPPÄNEN, A., AND STABOULIS, S. Optimizing electrode positions in electrical impedance tomography. *SIAM J. Appl. Math.* *74* (2014), 1831–1851.
- [21] JAFFRAY, D. A. Image-guided radiotherapy: from current concept to future perspectives. *Nat. Rev. Clin. Oncol.* *9*, 12 (2012), 688.
- [22] KAIPIO, J., AND SOMERSALO, E. *Statistical and computational inverse problems*, vol. 160. Springer Science & Business Media, 2006.
- [23] KHODJA, M., PRANGE, M., AND DJIKPESE, H. Guided Bayesian optimal experimental design. *Inverse Problems* *26*, 5 (2010), 055008.
- [24] LONG, Q., MOTAMED, M., AND TEMPONE, R. Fast Bayesian optimal experimental design for seismic source inversion. *Comput. Methods Appl. Mech. Eng.* *291* (2015), 123–145.
- [25] LONG, Q., SCAVINO, M., TEMPONE, R., AND WANG, S. Fast estimation of expected information gains for Bayesian experimental designs based on laplace approximations. *Comput. Methods Appl. Mech. Eng.* *259* (2013), 24–39.
- [26] NATTERER, F., AND WÜBBELING, F. *Mathematical methods in image reconstruction*. SIAM Monographs on Mathematical Modeling and Computation. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.
- [27] RADON, J. Über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten. *Ber. Verh. Schsischen Akad. Wiss. Leipzig. Math.-Phys. Kl.* *69* (1917), 262–277.
- [28] RUTHOTTO, L., CHUNG, J., AND CHUNG, M. Optimal experimental design for inverse problems with state constraints. *SIAM J. Sci. Comput.* *40*, 4 (2018), B1080–B1100.
- [29] SHANNON, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* *27* (1948), 379–423.
- [30] SILTANEN, S., KOLEHMAINEN, V., JÄRVENPÄÄ, S., KAIPIO, J. P., KOISTINEN, P., LASSAS, M., PIRTILÄ, J., AND SOMERSALO, E. Statistical inversion for medical x-ray tomography with few radiographs: I. general theory. *Phys. Med. Biol.* *48* (may 2003), 1437–1463.
- [31] STUART, A. M. Inverse problems: A Bayesian perspective. *Acta Numer.* *19* (2010), 451559.