
On Unifying Randomized Methods for Inverse Problems

A PREPRINT

Jonathan Wittmer
Oden Institute
University of Texas at Austin
Austin, TX 78712

C G Krishnanunni
Department of Aerospace Engineering and Engineering Mechanics
University of Texas at Austin
Austin, TX 78712

Hai V. Nguyen
Department of Aerospace Engineering and Engineering Mechanics
University of Texas at Austin
Austin, TX 78712

Tan Bui-Thanh
Oden Institute
Department of Aerospace Engineering and Engineering Mechanics
University of Texas at Austin
Austin, TX 78712

January 3, 2023

ABSTRACT

This work unifies the analysis of various randomized methods for solving linear and nonlinear inverse problems by framing the problem in a stochastic optimization setting. By doing so, we show that many randomized methods are variants of a sample average approximation. More importantly, we are able to prove a single theoretical result that guarantees the asymptotic convergence for a variety of randomized methods. Additionally, viewing randomized methods as a sample average approximation enables us to prove, for the first time, a single non-asymptotic error result that holds for randomized methods under consideration. Another important consequence of our unified framework is that it allows us to discover new randomization methods. We present various numerical results for linear, nonlinear, algebraic, and PDE-constrained inverse problems that verify the theoretical convergence results and provide a discussion on the apparently different convergence rates and the behavior for various randomized methods.

Keywords Randomization, Bayesian Inversion, Ensemble Kalman Filter, randomized maximum a posteriori

1 Introduction

Solving large-scale ill-posed inverse problems that are governed by partial differential equations (PDEs), though tremendously challenging, is of great practical importance in science and engineering. Classical deterministic inverse methodologies, which provide point estimates of the solution, are not capable of accounting for the uncertainty in the inverse solution in a principled way. The Bayesian formulation provides a systematic quantification of uncertainty by posing the inverse problem as one of statistical inference. The Bayesian framework for inverse problems proceeds as follows: given observational data $\mathbf{d} \in \mathbb{R}^k$ and their uncertainty, the governing forward problem and its uncertainty, and a prior probability density function describing uncertainty in the parameters $\mathbf{u} \in \mathbb{R}^n$, the solution of the inverse problems is the posterior probability distribution $\pi(\mathbf{u}|\mathbf{d})$ over the parameters. Bayes' Theorem explicitly

gives the posterior density as

$$\pi(\mathbf{u}|\mathbf{d}) \propto \pi_{\text{like}}(\mathbf{d}|\mathbf{u}) \times \pi_{\text{prior}}(\mathbf{u})$$

which updates the prior knowledge $\pi_{\text{prior}}(\mathbf{u})$ using the likelihood $\pi_{\text{like}}(\mathbf{d}|\mathbf{u})$. The prior encodes any knowledge or assumptions about the parameter space that we may wish to impose before any data are observed, while the likelihood explicitly represents the probability that a given set of parameters \mathbf{u} might give rise to the observed data \mathbf{d} .

Even when the prior and noise probability distributions are Gaussian, the posterior need not be Gaussian, due to possible nonlinearity embedded in the likelihood. For large-scale inverse problems, exploring non-Gaussian posteriors in high dimensions to compute statistics is a grand challenge since evaluating the posterior at each point in that parameter space requires a solution of the parameter-to-observable map, including a potentially expensive forward model solve. Using numerical quadrature to compute the mean and covariance matrix, for example, is generally infeasible in high dimensions. Usually the method of choice for computing statistics is Markov chain Monte Carlo (MCMC), which judiciously samples the posterior distribution, so that sample statistics can be used to approximate the exact ones.

The Metropolis-Hastings (MH) algorithm, first developed by Metropolis *et al.* [1] and then generalized by Hastings [2], is perhaps the most popular MCMC method. Its popularity and attractiveness come from the ease of implementation and minimal requirements on the target density and the proposal density [3, 4]. The problem, however, is that standard MCMC methods often require millions of samples for convergence; since each sample requires an evaluation of the parameter-to-observable map, this could entail millions of expensive forward PDE simulations—a prohibitive proposition. On one hand, with the rapid development of parallel computing, parallel MCMC methods [5, 6, 7, 8, 9] are studied to accelerate the computation. While parallelization allows MCMC algorithms to produce more samples in a shorter time with multiple processors, such accelerations typically do not improve the mixing and convergence of MCMC algorithms. More sophisticated MCMC methods that exploit the gradient and higher derivatives of the log posterior (and hence the parameter-to-observable map) [10, 11, 12, 13, 14, 15, 16, 17, 18, 19] can, on the other hand, improve the mixing, acceptance rate, and convergence of MCMC. Another sample-based family of approaches that provide uncertainty quantification and are well-suited for parallelization on large clusters is the various forms of Stein variational gradient descent [20, 21, 22, 23]. Of related interest are particle filter methods such as those found in [24, 25, 26, 27] that evolve particles through a dynamical system over time, updating both an estimate of the state and uncertainty.

One approach to addressing the computational challenge in high-dimensional statistical inverse problems pose is to use randomization, either to reduce the dimension of the optimization problem used in estimating the maximum a posteriori (MAP) point [28], or to aid in sampling from the posterior distribution [29]. Several methods have been proposed which utilize randomization to accelerate the solution of inverse problems [28, 29, 30, 31, 32, 33]. *As the main contribution of this paper*, we derive unified results of randomized inverse approaches that apply to a broad class of linear and nonlinear inverse problems not only in the asymptotic regime, but also for the non-asymptotic setting. The asymptotic convergence and a non-asymptotic error bound of various existing methods follows immediately as special cases of the general result.

2 A unified analysis of randomized inverse problems through a sample average approximation lens

For the remainder of this paper, we will use lower case letters for scalar quantities (α), boldface lower case letters for vectors (\mathbf{u}) and boldface upper case letters for matrices (\mathbf{A}). Further, we will use superscript lower case letters to denote sample index and subscript uppercase letters to denote the total number of samples, i.e. λ^i is the i^{th} sample and \mathbf{u}_N is a quantity depending on N samples. Lastly, descriptions or method names will be in uppercase superscripts, such as \mathbf{u}^{MAP} which is \mathbf{u} at the MAP point, for example. This should be clear from the context.

Therefore, let $\mathbf{u}, \mathbf{u}_0 \in \mathbb{R}^n$. The posterior measure ν in this case has the density $\pi(\mathbf{u}|\mathbf{d})$ with respect to the Lebesgue measure:

$$\pi(\mathbf{u}|\mathbf{d}) \propto \pi_{\text{like}}(\mathbf{d}|\mathbf{u}) \times \pi_{\text{prior}}(\mathbf{u}), \quad (1)$$

where the likelihood is given by $\pi_{\text{like}}(\mathbf{d}|\mathbf{u}) \propto \exp(-\Phi(\mathbf{u}, \mathbf{d})) = \exp\left(-\frac{1}{2}\|\mathbf{d} - \mathcal{F}(\mathbf{u})\|_{\Sigma^{-1}}^2\right)$ and the prior by $\pi_{\text{prior}} \propto \exp\left(-\frac{1}{2}\|\mathbf{u} - \mathbf{u}_0\|_{\Gamma^{-1}}^2\right)$. Here, $\mathcal{F}(\mathbf{u})$ is known as the parameter-to-observable (PtO) map, an evaluation of which typically requires a solution of the forward model (e.g. partial differential equations) governing the underlying physics. The *maximum a posteriori* (MAP) problem reads

$$\mathbf{u}^{\text{MAP}} := \arg \min_{\mathbf{u}} \mathcal{J}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}) := \frac{1}{2} \|\mathbf{d} - \mathcal{F}(\mathbf{u})\|_{\Sigma^{-1}}^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|_{\Gamma^{-1}}^2, \quad (2)$$

where $\Gamma \in \mathbb{R}^{n \times n}$ is the prior covariance matrix and $\Sigma \in \mathbb{R}^{k \times k}$ is the noise covariance matrix.

To the end of the paper, we denote by \mathbb{E} the expectation with subscript as the random variable with respect to which the expectation is taken. When the random variable is clear from the context we simply omit the subscript for brevity. Let $\sigma, \varepsilon, \delta,$ and λ be finite dimensional independent random variables with bounded second moments such that:

$$\mathbb{E}[\sigma] = 0, \quad \mathbb{E}[\varepsilon\varepsilon^T] = \Sigma^{-1}, \quad \mathbb{E}[\delta] = 0, \quad \mathbb{E}[\lambda\lambda^T] = \Gamma^{-1}. \quad (3)$$

Let us define $\xi = [\sigma, \varepsilon, \delta, \lambda]^T \in \Xi$ with joint probability distribution $\pi = \pi_\sigma \times \pi_\varepsilon \times \pi_\delta \times \pi_\lambda$. Consider the following stochastic cost function:

$$\begin{aligned} \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \xi) &:= \frac{1}{2} \|\varepsilon^T (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\lambda^T (\mathbf{u} - \mathbf{u}_0 - \delta)\|_2^2 \\ &= \frac{1}{2} (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))^T \varepsilon \varepsilon^T (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u})) \\ &\quad + \frac{1}{2} (\mathbf{u} - \mathbf{u}_0 - \delta)^T \lambda \lambda^T (\mathbf{u} - \mathbf{u}_0 - \delta). \end{aligned} \quad (4)$$

Define

$$\mathcal{J}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}) := \mathbb{E}_\pi [\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \xi)],$$

and the sample average approximation (SAA) of \mathcal{J} to be

$$\mathcal{J}_N := \frac{1}{N} \sum_{j=1}^N \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \xi^j) \quad (5)$$

where ξ^j are i.i.d. samples from π . Assume that both \mathcal{J} and \mathcal{J}_N have a minimum and let us define

$$\mathbf{u}^{\text{MAP}} := \arg \min_{\mathbf{u}} \mathcal{J}, \quad \text{and} \quad \hat{\mathbf{u}}_N^{\text{MAP}} := \arg \min_{\mathbf{u}} \mathcal{J}_N. \quad (6)$$

Below we study asymptotic and non-asymptotic convergence of $\hat{\mathbf{u}}_N^{\text{MAP}}$ to \mathbf{u}^{MAP} .

2.1 Asymptotic convergence analysis for inverse problems

Theorem 1 (Asymptotic convergence of randomized nonlinear inverse problems). *Assume that $\mathcal{F}(\mathbf{u})$ is such that $\tilde{\mathcal{J}}$ is a convex, twice continuously differentiable function in \mathbf{u} for almost every ξ , and measurable¹. Then the following hold true:*

- i) Minimizing \mathcal{J} is equivalent to minimizing \mathcal{J} in the sense: $\arg \min_{\mathbf{u}} \mathcal{J} = \arg \min_{\mathbf{u}} \mathcal{J}$.
- ii) $\hat{\mathbf{u}}_N^{\text{MAP}} \xrightarrow[N \rightarrow \infty]{a.s.} \mathbf{u}^{\text{MAP}}$.

Proof. For the first assertion, consider only the first term of \mathcal{J} as the second term follows analogously. We have

$$\begin{aligned} &\frac{1}{2} \mathbb{E}_\pi \left[(\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))^T \varepsilon \varepsilon^T (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u})) \right] \\ &= \frac{1}{2} \mathbb{E}_{\pi_\sigma \times \pi_\varepsilon} \left[(\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))^T \varepsilon \varepsilon^T (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u})) \right] \\ &= \frac{1}{2} \mathbb{E}_{\pi_\sigma} \left[(\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))^T \mathbb{E}_{\pi_\varepsilon} [\varepsilon \varepsilon^T] (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u})) \right] \\ &= \frac{1}{2} \mathbb{E}_{\pi_\sigma} \left[(\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))^T \Sigma^{-1} (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u})) \right] \\ &= \frac{1}{2} (\mathbf{d} - \mathcal{F}(\mathbf{u}))^T \Sigma^{-1} (\mathbf{d} - \mathcal{F}(\mathbf{u})) + \mathbb{E}_{\pi_\sigma} [\sigma^T \Sigma^{-1} \sigma]. \end{aligned} \quad (7)$$

The final term in (7) is constant with respect to \mathbf{u} and can be ignored, leaving only the first term of \mathcal{J} . Applying the same procedure to the second term of \mathcal{J} shows that minimizing \mathcal{J} is equivalent to minimizing \mathcal{J} .

We invoke [34, Theorem 5.4] to prove the second assertion. It is sufficient to verify the following conditions:

¹Here, measurable is with respect to the the σ -algebra given by the product σ -algebras of the deterministic $(\mathbf{u}, \mathbf{u}_0, \mathbf{d})$ and random variables ξ .

- (i) $\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \boldsymbol{\xi})$ is random lower semicontinuous,
- (ii) for almost every $\boldsymbol{\xi} \in \Xi$, $\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \boldsymbol{\xi})$ is convex in \mathbf{u} ,
- (iii) $\mathcal{J}(\mathbf{u}; \mathbf{u}_0, \mathbf{d})$ is lower semicontinuous in \mathbf{u} and there exists a point $\bar{\mathbf{u}} \in \mathbb{R}^n$ such that $\mathcal{J}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}) < \infty$ for all \mathbf{u} in a neighborhood of $\bar{\mathbf{u}}$;
- (iv) the set of optimal solutions of the true problem is nonempty and bounded; and
- (v) the law of large numbers (LLN) [35, 36] holds pointwise for \mathcal{J}_N .

Clearly, $\tilde{\mathcal{J}}$ is a continuous function for every $\boldsymbol{\xi}$, thus random lower semicontinuous as well. By assumption, $\tilde{\mathcal{J}}$ is also convex for almost every $\boldsymbol{\xi}$. Due to the boundedness assumptions (3) and the fact that \mathcal{J} is a continuous and convex function, \mathcal{J} is also a continuous and convex function. Furthermore, taking, for example, $\bar{\mathbf{u}} = \mathbf{u}_1$ in (17) it is straightforward to see that $\mathcal{J}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}) < \infty$ for any ball with finite radius centered at $\bar{\mathbf{u}}$. The last two conditions are clear. \square

An important special case of this theorem occurs when we consider an inverse problem with a linear parameter-to-observable map. When the forward map $\mathcal{F}(\mathbf{u})$ is linear, the convexity and continuous differentiability assumptions are satisfied. While requiring convexity is a strong assumption in general, this is not an insurmountable issue for regularized inverse problems. Note that the Hessian of \mathcal{J} is given by

$$\nabla_{\mathbf{u}}^2 \mathcal{J} = \nabla_{\mathbf{u}}^2 \mathcal{F}(\mathbf{u}) \boldsymbol{\Sigma}^{-1} (\mathcal{F}(\mathbf{u}) - \mathbf{d}) + \nabla_{\mathbf{u}} \mathcal{F}(\mathbf{u})^T \boldsymbol{\Sigma}^{-1} \nabla_{\mathbf{u}} \mathcal{F}(\mathbf{u}) + \boldsymbol{\Gamma}^{-1}.$$

Thus the prior covariance matrix can be chosen such that $\nabla_{\mathbf{u}}^2 \mathcal{J}$ is semi-positive definite. Indeed, this is the major role that the prior covariance plays in regularizing the ill-posed inverse problem.

Lastly, note that instead of treating all of the random variables $\boldsymbol{\sigma}$, $\boldsymbol{\varepsilon}$, $\boldsymbol{\lambda}$, and $\boldsymbol{\delta}$ as a single random variable, sampling them together as in (5), we could have chosen to approximate each random variable separately, applying [34, Theorem 5.4] four times to complete the asymptotic proof.

$$\tilde{\mathbf{u}}_N^{\text{MAP}} := \arg \min_{\mathbf{u}} \frac{1}{N_1 N_2 N_3 N_4} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \sum_{k=1}^{N_3} \sum_{l=1}^{N_4} \tilde{\mathcal{J}}_N(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, [\boldsymbol{\sigma}^i; \boldsymbol{\varepsilon}^j; \boldsymbol{\lambda}^k; \boldsymbol{\delta}^l]) \quad (8)$$

This flexibility in deciding how samples will be drawn will aid in both non-asymptotic convergence analysis and will provide great freedom in designing a variety of randomized methods to solve inverse problems in section 5.

To alleviate some notational burden, let us define a few new quantities.

$$\mathbf{S}_N := \frac{1}{N} \sum_{i=1}^N \boldsymbol{\varepsilon}^i (\boldsymbol{\varepsilon}^i)^T, \quad \mathbf{L}_N := \frac{1}{N} \sum_{i=1}^N \boldsymbol{\lambda}^i (\boldsymbol{\lambda}^i)^T, \quad (9a)$$

$$\bar{\boldsymbol{\sigma}}_N := \frac{1}{N} \sum_{i=1}^N \boldsymbol{\sigma}^i, \quad \bar{\boldsymbol{\delta}}_N := \frac{1}{N} \sum_{i=1}^N \boldsymbol{\delta}^i. \quad (9b)$$

Written in terms of norms, equation (8) becomes

$$\tilde{\mathbf{u}}_N^{\text{MAP}} = \arg \min_{\mathbf{u}} \frac{1}{2} \|\mathbf{d} + \bar{\boldsymbol{\sigma}}_N - \mathcal{F}(\mathbf{u})\|_{\mathbf{S}_N}^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0 - \bar{\boldsymbol{\delta}}_N\|_{\mathbf{L}_N}^2, \quad (10)$$

Note that (10) is equivalent to (5) when at most one of $\boldsymbol{\varepsilon}$ and $\boldsymbol{\sigma}$ are randomized and at most one of $\boldsymbol{\lambda}$ and $\boldsymbol{\delta}$ are randomized. This is because the only difference between the two cost functions is how samples of $\boldsymbol{\varepsilon}$ interact with samples of $\boldsymbol{\sigma}$ and likewise, how samples of $\boldsymbol{\lambda}$ interact with samples of $\boldsymbol{\delta}$.

2.2 Non-asymptotic error analysis for nonlinear inverse problems

In addition to proving a general asymptotic convergence of randomized inverse problems, it is also possible to derive a general non-asymptotic error bound with the slightly stronger assumption that the random variables are subgaussian [37]. The general form of this non-asymptotic bound is useful in that it is easy to identify the key components that go into forming the bound — giving insight into the performance of various methods by enabling easy simplification in the case that certain quantities are not randomized. The derived bound gives a probabilistic worst-case for finite sample

size N when all of the above randomizations are implemented at the same time. By fixing some of the quantities, the given bound can be simplified in a straightforward manner — yielding a more insightful bound.

Additionally, we follow the standard vector norm convention of $\|\mathbf{u}\|_\infty := \max(|u_1|, \dots, |u_n|)$ and $\|\mathbf{u}\|_1 := \sum_{i=1}^n |u_i|$ for a vector $\mathbf{u} \in \mathbb{R}^n$. Matrix norms are understood to be induced norms [38]. Due to the equivalence of norms in finite dimensional spaces, all the results are also valid for other norms, albeit with different constants that possibly depend on the dimension. We use *w.p.* as the abbreviation for *with probability*.

Proposition 2 (Convergence of mean-zero subgaussian random vector). *Let δ^i , for $i = 1, \dots, N$, be independent subgaussian random vectors in \mathbb{R}^n such that $\mathbb{E}[\delta^i] = \mathbf{0}$, $\mathbb{E}[\delta^i(\delta^i)^T] = \Gamma$, and $\mathbb{E}[(\delta^i)^T \delta^i] < \infty$. Denote the empirical mean $\bar{\delta} := \frac{1}{N} \sum_{i=1}^N \delta^i$. Further, let $\zeta(N, \beta) := \exp(-c\beta^2 N)$ for some $\beta > 0$ and c is a constant possibly depending on the dimension n but not on N . Then*

$$\|\bar{\delta}\|_\infty \leq \beta \left\| \Gamma^{\frac{1}{2}} \right\|_\infty \quad \text{w.p. at least } 1 - \zeta(N, \beta). \quad (11)$$

Proof. Define $\delta^i = \Gamma^{\frac{1}{2}} \tau^i$, where $\tau^i \sim \mathcal{N}(0, \mathcal{I})$. Thus, $\bar{\tau} = \frac{1}{N} \sum_{i=1}^N \tau^i \sim \mathcal{N}(0, \frac{\mathcal{I}}{N})$. First from² [39, Theorem 1] we have

$$\mathbb{P}[\|\bar{\tau}\|_\infty > \beta_1] \leq \mathbb{P}[\|\bar{\tau}\|_1 > \beta_1] \leq \exp\left(-\frac{\beta_1^2}{4cn}\right),$$

where c is an absolute positive constant. Therefore, abusing notation to consolidate the constant terms, we have

$$\mathbb{P}\left[\|\bar{\tau}\|_\infty > \frac{\beta_1}{\sqrt{N}}\right] \leq \exp(-c\beta_1^2)$$

since $\sqrt{N} \cdot \bar{\tau}$ has identity covariance. Next, set $\beta = \frac{\beta_1}{\sqrt{N}}$ and note that

$$\mathbb{P}\left[\|\bar{\delta}\|_\infty \leq \beta \left\| \Gamma^{\frac{1}{2}} \right\|_\infty\right] \geq \mathbb{P}[\|\bar{\tau}\|_\infty \leq \beta] \geq 1 - \exp(-c\beta^2 N),$$

and this concludes the proof. \square

Proposition 3 (Convergence of subgaussian random vector outer product). *Let ω^i be a subgaussian random vector in \mathbb{R}^n such that $\mathbb{E}[\omega^i(\omega^i)^T] = \Gamma$ for $i = 1, \dots, N$. Define Ω_N to be the random matrix formed by stacking ω^i in the columns and scaling by $1/\sqrt{N}$. It follows that*

$$\|\Omega_N \Omega_N^T - \Gamma\|_\infty \leq \beta \|\Gamma\|_\infty \quad \text{w.p. at least } 1 - 2\zeta(N, \beta), \quad (12)$$

where c is a constant depending on n but not on N .

Proof. This result follows from straightforward algebraic manipulation of [37, Theorem 4.6.1]. \square

An additional fact needed to prove a non-asymptotic bound is that the product of three subgaussian random variables is α -subexponential with $\alpha = 2/3$. The following discussion on α -subexponential random variables is based on [40]. For a more complete treatment of the topic, [40] can be consulted.

It has been established that the product of two subgaussian random variables is subexponential [37, Lemma 2.7.7]. A centered random variable X is said to be α -subexponential (or sub-Weibull [41, 42]) if it satisfies

$$\mathbb{P}[|X| \geq \beta] \leq 2 \exp(-c\beta^\alpha),$$

for any $\beta > 0$, $\alpha \in (0, 2]$, and some positive constant c . To show that a random variable satisfies this condition, it is sufficient to show that the following Orlicz (quasi-) norm [40, 37] is finite:

$$\|X\|_{\psi_\alpha} := \inf \{\beta > 0 : \mathbb{E} \exp((|X|/\beta)^\alpha) \leq 2\} < \infty. \quad (13)$$

When $\alpha < 1$, this is a quasi-norm since it does not satisfy the triangle inequality.

Proposition 4 (α -subexponential from product of three subgaussian variables). *Let X_1, X_2, X_3 be subgaussian random variables. Then $Y = X_1 X_2 X_3$ is an α -subexponential random variable with $\alpha = 2/3$.*

²While [39, Theorem 1] is derived for Gaussian random matrices, it also applies to subgaussian random matrices because subgaussian random variables have the same bound for the expected value of their moment generating function (see [37, Proposition 2.5.2] for the details).

Proof. It suffices to show that $\mathbb{E} \exp\left(\left(|Y|\right)^{2/3}\right) \leq 2$. Without loss of generality, assume $\|X_i\|_{\psi_2} = 1$. Then $\mathbb{E} \exp\left(X_i^2\right) \leq 2$ and we have

$$\begin{aligned} \mathbb{E} \exp\left(\left(|Y|\right)^{2/3}\right) &= \mathbb{E} \exp\left(\left(|X_1|^{2/3} |X_2|^{2/3} |X_3|^{2/3}\right)\right) \\ &\leq \mathbb{E} \exp\left(\frac{X_1^2}{3} + \frac{X_2^2}{3} + \frac{X_3^2}{3}\right) \quad (\text{Young's inequality for 3 variables [43]}) \\ &= \mathbb{E} \exp\left(\frac{X_1^2}{3}\right) \exp\left(\frac{X_2^2}{3}\right) \exp\left(\frac{X_3^2}{3}\right) \\ &\leq \frac{1}{3} [\exp(X_1^2) + \exp(X_2^2) + \exp(X_3^2)] \quad (\text{Young's inequality again}) \\ &\leq 2. \end{aligned}$$

□

Corollary 5. Let $\delta \in \mathbb{R}^n$ be a zero-mean random vector such that $\mathbb{E}[\delta\delta^T] = \Gamma$ with α -subexponential entries. Define $\zeta_\alpha(N, \beta) := \exp(-c\beta^\alpha N^{\alpha/2})$. Then

$$\mathbb{P}\left[\left\|\frac{1}{N} \sum_{i=1}^n \delta^i\right\|_\infty \leq \beta \left\|\Gamma^{\frac{1}{2}}\right\|_\infty\right] \geq 1 - 2n\zeta_\alpha(N, \beta). \quad (14)$$

Note that while inequality (14) has the dimension n in front of the exponential, this is fixed and the probability of committing an error greater than some tolerance still decreases exponentially in the number of samples N . Though the decaying is at the slower rate $\propto \exp(-N^{\alpha/2})$ compared to the subgaussian error rate $\propto \exp(-N)$, it is not surprising because subgaussian is a special case of α -subexponential when $\alpha = 2$ [40]. These results can be combined to derive a probabilistic non-asymptotic error bound for nonlinear inverse problems.

Lemma 6 (Non-asymptotic error analysis for randomized nonlinear inverse problems). Let $\text{vec}(\Sigma^{-1})$ denote a vectorization of a matrix Σ^{-1} . Define

$$\mathbf{P} := [\text{vec}(\Sigma^{-1}); \text{vec}(\Gamma^{-1}); \mathbf{e}; \mathbf{z}]$$

as a vector concatenating all four vectors $\text{vec}(\Sigma^{-1})$, $\text{vec}(\Gamma^{-1})$, \mathbf{e} and \mathbf{z} , where $\Sigma \in \mathbb{R}^{k \times k}$, $\Gamma \in \mathbb{R}^{n \times n}$, $\mathbf{e} \in \mathbb{R}^k$, and $\mathbf{z} \in \mathbb{R}^n$. Define the function

$$g(\mathbf{P}; \mathbf{u}) := \nabla_{\mathbf{u}} \mathcal{F}(\mathbf{u}) [\Sigma^{-1}(\mathcal{F}(\mathbf{u}) - \mathbf{d}) - \mathbf{e}] + \Gamma^{-1}(\mathbf{u} - \mathbf{u}_0) - \mathbf{z}.$$

Assume that the problem $g(\mathbf{P}; \mathbf{u}) = 0$, with \mathbf{P} as parameters and \mathbf{u} as solution, is Lipschitz well-posed [44] with Lipschitz constant \mathcal{L} , and we define $\mathcal{G}(\mathbf{P})$ as the solution \mathbf{u} . Let

$$\begin{aligned} \mathbf{P}^{\text{MAP}} &:= [\text{vec}(\Sigma^{-1}); \text{vec}(\Gamma^{-1}); 0; 0], \\ \tilde{\mathbf{P}}_N &:= [\text{vec}(\mathbf{S}_N); \text{vec}(\mathbf{L}_N); \mathbf{S}_N \bar{\boldsymbol{\sigma}}_N; \mathbf{L}_N \bar{\boldsymbol{\delta}}_N], \\ \hat{\mathbf{P}}_N &:= \left[\text{vec}(\mathbf{S}_N); \text{vec}(\mathbf{L}_N); \frac{1}{N} \sum_{i=1}^N \boldsymbol{\epsilon}^i (\boldsymbol{\epsilon}^i)^T \boldsymbol{\sigma}^i; \frac{1}{N} \sum_{i=1}^N \boldsymbol{\lambda}^i (\boldsymbol{\lambda}^i)^T \boldsymbol{\delta}^i \right], \end{aligned}$$

where $\mathbb{E}[\boldsymbol{\sigma}\boldsymbol{\sigma}^T] = \Sigma$ and $\mathbb{E}[\boldsymbol{\delta}\boldsymbol{\delta}^T] = \Gamma$. Then

$$\left\|\mathbf{u}^{\text{MAP}} - \hat{\mathbf{u}}_N^{\text{MAP}}\right\|_\infty \leq \beta \mathcal{L} \left(\left\|\Sigma^{-1}\right\|_\infty + \left\|\Gamma^{-1}\right\|_\infty + (1 + \beta) \left\|\Sigma^{-\frac{1}{2}}\right\|_\infty + (1 + \beta) \left\|\Gamma^{-\frac{1}{2}}\right\|_\infty \right) \quad (15)$$

w.p. at least $1 - 10\zeta(N, \beta)$,

and

$$\left\|\mathbf{u}^{\text{MAP}} - \hat{\mathbf{u}}_N^{\text{MAP}}\right\|_\infty \leq \beta \mathcal{L} \left(\left\|\Sigma^{-1}\right\|_\infty + \left\|\Gamma^{-1}\right\|_\infty + \left\|\Sigma^{-\frac{1}{2}}\right\|_\infty + \left\|\Gamma^{-\frac{1}{2}}\right\|_\infty \right) \quad (16)$$

w.p. at least $1 - 4\zeta(N, \beta) - 2k\zeta_{2/3}(N, \beta) - 2n\zeta_{2/3}(N, \beta)$.

Proof. Noting that $\mathbf{u}^{\text{MAP}} = \mathcal{G}(\mathbf{P}^{\text{MAP}})$ and $\tilde{\mathbf{u}} = \mathcal{G}(\tilde{\mathbf{P}}_N)$, we have by the Lipschitz well-posedness assumption,

$$\begin{aligned} \|\mathbf{u}^{\text{MAP}} - \tilde{\mathbf{u}}_N^{\text{MAP}}\|_\infty &= \left\| \mathcal{G}(\mathbf{P}^{\text{MAP}}) - \mathcal{G}(\tilde{\mathbf{P}}_N) \right\|_\infty \leq \mathcal{L} \left\| \mathbf{P} - \tilde{\mathbf{P}}_N \right\|_\infty \\ &\leq \mathcal{L} \left(\left\| \boldsymbol{\Sigma}^{-1} - \mathbf{S}_N \right\|_\infty + \left\| \boldsymbol{\Gamma}^{-1} - \mathbf{L}_N \right\|_\infty + \left\| \mathbf{S}_N \bar{\boldsymbol{\sigma}}_N \right\|_\infty + \left\| \mathbf{L}_N \bar{\boldsymbol{\delta}}_N \right\|_\infty \right). \end{aligned}$$

We can bound $\mathbf{S}_N \bar{\boldsymbol{\sigma}}_N$ (and similarly for $\mathbf{L}_N \bar{\boldsymbol{\delta}}_N$) as follows

$$\begin{aligned} \left\| \mathbf{S}_N \bar{\boldsymbol{\sigma}}_N \right\|_\infty &= \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \left(\boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{S}_N \boldsymbol{\Sigma}^{\frac{1}{2}} - \mathcal{I} + \mathcal{I} \right) \boldsymbol{\Sigma}^{-\frac{1}{2}} \bar{\boldsymbol{\sigma}}_N \right\|_\infty \\ &\leq \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \left(\boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{S}_N \boldsymbol{\Sigma}^{\frac{1}{2}} - \mathcal{I} \right) \right\|_\infty \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \bar{\boldsymbol{\sigma}}_N \right\|_\infty + \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \right\|_\infty \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \bar{\boldsymbol{\sigma}}_N \right\|_\infty. \end{aligned}$$

Note that $\boldsymbol{\Sigma}^{-\frac{1}{2}} \bar{\boldsymbol{\sigma}}_N$ is the sample average of a mean-zero subgaussian random variable with identity covariance and $\mathbb{E}[\boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{S}_N \boldsymbol{\Sigma}^{\frac{1}{2}}] = \mathcal{I}$. Therefore, applying Proposition 3 to $\left\| \boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{S}_N \boldsymbol{\Sigma}^{\frac{1}{2}} - \mathcal{I} \right\|_\infty$ and Proposition 2 to $\left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \bar{\boldsymbol{\sigma}}_N \right\|_\infty$ along with the union bound, we obtain

$$\left\| \mathbf{S}_N \bar{\boldsymbol{\sigma}}_N \right\|_\infty \leq (\beta^2 + \beta) \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \right\|_\infty \quad \text{w.p. at least } 1 - 3\zeta(N, \beta).$$

The proof of the bound $\left\| \mathbf{L}_N \bar{\boldsymbol{\delta}}_N \right\|_\infty \leq (\beta^2 + \beta) \left\| \boldsymbol{\Gamma}^{-\frac{1}{2}} \right\|_\infty$ w.p. at least $1 - 3\zeta(N, \beta)$ follows analogously. Applying Proposition 3 to the terms $\left\| \mathbf{S}_N - \boldsymbol{\Sigma}^{-1} \right\|_\infty$ and $\left\| \mathbf{L}_N - \boldsymbol{\Gamma}^{-1} \right\|_\infty$ along with the union bound, inequality (15) follows immediately.

To prove the bound in (16) for $\left\| \mathbf{u}^{\text{MAP}} - \hat{\mathbf{u}}_N^{\text{MAP}} \right\|_\infty$, it is sufficient to bound $\|e\|_\infty$ and $\|z\|_\infty$ according to the definition of $\hat{\mathbf{P}}_N$ since the bounds for $\left\| \mathbf{S}_N - \boldsymbol{\Sigma}^{-1} \right\|_\infty$ and $\left\| \mathbf{L}_N - \boldsymbol{\Gamma}^{-1} \right\|_\infty$ have already been proven. By Proposition 4, $\varepsilon^i (\varepsilon^i)^T \boldsymbol{\sigma}^i$ is a zero-mean, $2/3$ -subexponential random variable with covariance $\boldsymbol{\Sigma}^{-1}$. Then by Corollary 5,

$$\mathbb{P} \left[\|e\|_\infty \leq \beta \left\| \boldsymbol{\Sigma}^{-\frac{1}{2}} \right\|_\infty \right] \geq 1 - 2k\zeta_{2/3}(N, \beta).$$

Similarly, applying Corollary 5 to $\|z\|_\infty$ yields

$$\mathbb{P} \left[\|z\|_\infty \leq \beta \left\| \boldsymbol{\Gamma}^{-\frac{1}{2}} \right\|_\infty \right] \geq 1 - 2n\zeta_{2/3}(N, \beta).$$

Using the union bound to combine the bounds on $\|e\|_\infty$ and $\|z\|_\infty$ with the bounds previously $\left\| \mathbf{S}_N - \boldsymbol{\Sigma}^{-1} \right\|_\infty$ and $\left\| \mathbf{L}_N - \boldsymbol{\Gamma}^{-1} \right\|_\infty$, inequality (16) follows. \square

Lemma 6 provides a general strategy for analyzing the convergence of various randomized methods. When a variable is not randomized, we simply drop the corresponding terms in equations (15) and (16), and adjust the probability accordingly. Additionally, as will be discussed later, the Lipschitz constant may be affected by the choice of randomization strategy. By choosing to randomize only some variables, the inverse solution may have lower worst-case sensitivity (Lipschitz constant). Note that since we are mainly concerned with small deviations, it is sufficient to consider well-posed inverse problems where the solution depends continuously on \mathbf{d} everywhere and depends on \mathbf{P} in a locally Lipschitz manner in a neighborhood of $[\text{vec}(\boldsymbol{\Sigma}^{-1}); \text{vec}(\boldsymbol{\Gamma}^{-1}); 0; 0]$. Note that it is unlikely for the solution to depend continuously on \mathbf{P} everywhere since $\boldsymbol{\Gamma}^{-1}$ acts as regularization. Otherwise, the original problem would not be ill-posed. In particular, the regularizing role that $\boldsymbol{\Gamma}^{-1}$ plays may cause the solution of the inverse problem to be especially sensitive to perturbations of $\boldsymbol{\Gamma}^{-1}$, greatly increasing the (local) Lipschitz constant compared to the case where $\boldsymbol{\Gamma}^{-1}$ is not randomized. This is clear from the linear case where the condition number of the problem takes the place of the Lipschitz constant. It is well-known that the choice of $\boldsymbol{\Gamma}^{-1}$ has a significant impact on the condition number [45, 46]. Lastly, note that in the linear case, standard perturbation theory for linear systems can be used to find an explicit bound for the relative error in terms of the condition number of $\mathcal{A}^T \boldsymbol{\Sigma}^{-1} \mathcal{A} + \boldsymbol{\Gamma}^{-1}$.

3 Rediscovery of randomized inverse methods

In this section we derive several different known randomization schemes as special cases of the more general randomization scheme proposed in (4). For the rest of this section, we will explicitly write each of the random variables

(subset of $\xi = [\sigma, \epsilon, \delta, \lambda]^T$) that the stochastic cost function depends on. To keep notation clean, we will use $\tilde{\mathcal{J}}$ to represent the stochastic cost function for all randomization schemes, where the form of the cost function will be clear from the context. If σ does not appear as an argument of $\tilde{\mathcal{J}}$, we replace it with 0 in (4). Likewise, we replace $\epsilon\epsilon^T$ with Σ^{-1} , δ with 0, and $\lambda\lambda^T$ with Γ^{-1} in (4) when the corresponding random variable does not appear as an argument of $\tilde{\mathcal{J}}$. Additionally, we will present each method in the general nonlinear setting, but we will also explicitly write the sample average solution in the linear case as closed form solutions are available, yielding additional insights. To that end, we write two equivalent formulations of the MAP estimate for linear inverse problems.

Proposition 7. *When the PtO map is linear, i.e. $\mathcal{F}(\mathbf{u}) = \mathcal{A}\mathbf{u}$, the solution of the MAP problem (2) is given by either*

$$\mathbf{u}_1 = (\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \Gamma^{-1})^{-1} (\mathcal{A}^T \Sigma^{-1} \mathbf{d} + \Gamma^{-1} \mathbf{u}_0) \quad (17a)$$

or

$$\mathbf{u}_2 = \mathbf{u}_0 + \Gamma \mathcal{A}^T (\Sigma + \mathcal{A} \Gamma \mathcal{A}^T)^{-1} (\mathbf{d} - \mathcal{A} \mathbf{u}_0). \quad (17b)$$

Proof. The first of these identities is derived directly from the optimality condition of (2). Specifically,

$$\nabla \mathcal{J} = (\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \Gamma^{-1}) \mathbf{u} - \mathcal{A}^T \Sigma^{-1} \mathbf{d} - \Gamma^{-1} \mathbf{u}_0 = 0. \quad (18)$$

The second formulation can be derived from \mathbf{u}_1 using the Sherman-Morrison-Woodbury formula [47] under the conditions that Γ^{-1} and Σ^{-1} are invertible. \square

Remark 8. *This last assumption concerning the invertibility of Γ^{-1} , while seemingly trivial in light of the fact that Γ^{-1} is written as the inverse of a matrix, will be important in the following discussion of randomization.*

Additionally, let us define new random variables:

$$\mathbf{d}^i := \mathbf{d} + \sigma^i \quad \text{and} \quad \mathbf{u}_0^i := \mathbf{u}_0 + \delta^i, \quad (19)$$

where σ^i and δ^i are the first and the third components of ξ^i defined in Theorem 1. These quantities will be useful in the following discussion. By the LLN we have

$$\frac{1}{N} \sum_{i=1}^N \mathbf{d}^i \xrightarrow[N \rightarrow \infty]{a.s.} \mathbf{d} \quad \text{and} \quad \frac{1}{N} \sum_{i=1}^N \mathbf{u}_0^i \xrightarrow[N \rightarrow \infty]{a.s.} \mathbf{u}_0.$$

3.1 Randomized MAP approach

Assuming that the order of minimization and expectation can be interchanged³, we can write

$$\arg \min_{\mathbf{u}} \mathbb{E}_{\pi_{\sigma} \times \pi_{\delta}} [\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \sigma, \delta)] = \mathbb{E}_{\pi_{\sigma} \times \pi_{\delta}} \left[\arg \min_{\mathbf{u}} \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \sigma, \delta) \right]. \quad (20)$$

The sample average approximation of the RHS can then be written

$$\mathbf{u}_N^{\text{RMAP}} := \frac{1}{N} \sum_{i=1}^N \arg \min_{\mathbf{u}} \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \sigma^i, \delta^i). \quad (21)$$

This randomization approach coincides with the randomized MAP approach [29] when $\mathbb{E}_{\pi_{\sigma}} [\sigma \sigma^T] = \Sigma$ and $\mathbb{E}_{\pi_{\delta}} [\delta \delta^T] = \Gamma$ (also known as the randomized maximum likelihood [49, 50, 51]). In the linear case, we can write

$$\mathbf{u}_N^{\text{RMAP}} = \frac{1}{N} \sum_{i=1}^N (\mathbf{u}^{\text{RMAP}})^i,$$

where, thanks to Proposition 7,

$$\begin{aligned} (\mathbf{u}^{\text{RMAP}})^i &= (\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \Gamma^{-1})^{-1} [\mathcal{A}^T \Sigma^{-1} (\mathbf{d} + \sigma^i) + \Gamma^{-1} (\mathbf{u}_0 + \delta^i)] \\ &= (\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \Gamma^{-1})^{-1} [\mathcal{A}^T \Sigma^{-1} \mathbf{d}^i + \Gamma^{-1} \mathbf{u}_0^i]. \end{aligned}$$

Since the sample average approximation (21) of the right hand side of (20) converges to its expectation, the analysis from Section 2 applies. We have the following result for nonlinear PtO map.

³The conditions under which the interchange is valid can be consulted in [48, Theorem 14.60].

Corollary 9 (Asymptotic convergence of RMAP). *Suppose that the nonlinear PtO map \mathcal{F} satisfies the assumptions of Theorem 1. Let $\sigma^i \sim \pi_\sigma$ and $\delta^i \sim \pi_\delta$ with bounded covariances where $\mathbb{E}_{\pi_\sigma}[\sigma] = 0$ and $\mathbb{E}_{\pi_\delta}[\delta] = 0$. Then*

$$\mathbf{u}_N^{\text{RMAP}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } N \rightarrow \infty.$$

Note that if *only* the MAP point \mathbf{u}^{MAP} is needed, then the randomized MAP approach is not useful: in fact very expensive while only giving an approximate solution for \mathbf{u}^{MAP} . However, the approach could be appealing for Bayesian settings. Indeed, by choosing $\mathbb{E}_{\pi_\sigma}[\sigma\sigma^T] = \Sigma$ and $\mathbb{E}_{\pi_\delta}[\delta\delta^T] = \Gamma$, each solution $(\mathbf{u}^{\text{RMAP}})^i$ is a bona fide sample of the posterior distribution in the linear case. For nonlinear cases, $(\mathbf{u}^{\text{RMAP}})^i$ are biased samples of the posterior [29], but can be corrected via Metropolization [52, 53]. Note that for linear inverse problems, the RMAP approach is the same as the randomize-then-optimize approach in [51] (see an explanation from [29]). This randomized MAP method is embarrassingly parallel and is well-suited for implementation on distributed computing systems. While we could randomize the data and prior mean without exchanging expectation and optimization and convergence would be maintained, such a method would be of little use because we would obtain an inaccurate approximation of \mathbf{u}^{MAP} while having the same cost.

3.2 Randomized misfit approach (left sketching)

In this section we show that the randomized misfit approach (RMA) [28] is a special case of our randomization in (4). Indeed, if we let $\varepsilon \sim \pi_\varepsilon$ where $\mathbb{E}_{\pi_\varepsilon}[\varepsilon] = 0$ and $\mathbb{E}_{\pi_\varepsilon}[\varepsilon\varepsilon^T] = \Sigma^{-1}$, then

$$\begin{aligned} \mathbf{u}^{\text{RMA}} &:= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_\varepsilon} [\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \varepsilon)] \\ &= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_\varepsilon} \left[\frac{1}{2} \|\varepsilon^T (\mathbf{d} - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|_{\Gamma^{-1}}^2 \right]. \end{aligned}$$

The SAA of \mathbf{u}^{RMA} can be written as

$$\mathbf{u}_N^{\text{RMA}} := \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \varepsilon^i) \quad (24a)$$

$$= \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \|\tilde{\mathbf{d}}^i - \tilde{\mathcal{F}}^i(\mathbf{u})\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|_{\Gamma^{-1}}^2, \quad (24b)$$

where

$$\tilde{\mathcal{F}}^i := \varepsilon^i{}^T \mathcal{F} \quad \text{and} \quad \tilde{\mathbf{d}}^i := \varepsilon^i{}^T \mathbf{d}.$$

That is, the random samples, which can be combined into a random matrix, sketch the PtO map and the data from the left. Random sketching has been used extensively to reduce the cost of solving inverse problems [30, 54, 55]. The following is a direct consequence of Theorem 1.

Corollary 10 (Asymptotic convergence of RMA). *Let $\mathbf{u}_N^{\text{RMA}}$ be as defined above. Then*

$$\mathbf{u}_N^{\text{RMA}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } N \rightarrow \infty.$$

Calculating the optimality condition in the linear case results in

$$\mathbf{u}^{\text{RMA}} = (\mathcal{A}^T \mathbb{E}_{\pi_\varepsilon}[\varepsilon\varepsilon^T] \mathcal{A} + \Gamma^{-1})^{-1} (\mathcal{A}^T \mathbb{E}_{\pi_\varepsilon}[\varepsilon\varepsilon^T] \mathbf{d} + \Gamma^{-1} \mathbf{u}_0). \quad (25)$$

By letting

$$\tilde{\mathcal{A}} := \varepsilon^T \mathcal{A} \quad \text{and} \quad \tilde{\mathbf{d}} := \varepsilon^T \mathbf{d},$$

we can rewrite (25) as

$$\mathbf{u}^{\text{RMA}} = \left(\mathbb{E}_{\pi_\varepsilon} [\tilde{\mathcal{A}}^T \tilde{\mathcal{A}}] + \Gamma^{-1} \right)^{-1} \left(\mathbb{E}_{\pi_\varepsilon} [\tilde{\mathcal{A}}^T \tilde{\mathbf{d}}] + \Gamma^{-1} \mathbf{u}_0 \right)$$

If we combine the RMA and randomized MAP approaches into a single stochastic optimization problem, we discover a new method which we will denote RMA+RMAP. Specifically, consider the problem directly arising from randomization of (4) and define the solution using the RMA+RMAP method to be

$$\begin{aligned} \mathbf{u}^{\text{RMA+RMAP}} &:= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_{\sigma} \times \pi_{\varepsilon} \times \pi_{\delta}} \left[\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \sigma, \varepsilon, \delta) \right] \\ &= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_{\sigma} \times \pi_{\varepsilon} \times \pi_{\delta}} \left[\frac{1}{2} \|\varepsilon^T (\mathbf{d} + \sigma - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0 - \delta\|_{\Gamma^{-1}}^2 \right], \end{aligned}$$

and the corresponding SAA solution

$$\begin{aligned} \mathbf{u}_N^{\text{RMA+RMAP}} &:= \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \sigma^i, \varepsilon^i, \delta^i) \\ &= \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{2} \|\varepsilon^{iT} (\mathbf{d} + \sigma^i - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0 - \delta^i\|_{\Gamma^{-1}}^2 \right]. \end{aligned}$$

The following result is immediate from Theorem 1.

Corollary 11 (Asymptotic convergence of RMA+RMAP). *Let $\mathbf{u}_N^{\text{RMA+RMAP}}$ be as defined above. Then*

$$\mathbf{u}_N^{\text{RMA+RMAP}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } N \rightarrow \infty.$$

Allowing for the interchange of optimization and expectation as in the RMAP approach along with independent sample average approximations of each random variable, the following variant sequences also converge.

$$\mathbf{u}_N^{\text{RMA+RMAP}_1} = \frac{1}{N} \sum_{i=1}^N \arg \min_{\mathbf{u}} \left[\frac{1}{2} \|\varepsilon^{iT} (\mathbf{d} + \sigma^i - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0 - \delta^i\|_{\Gamma^{-1}}^2 \right]. \quad (28)$$

$$\mathbf{u}_{N,M}^{\text{RMA+RMAP}_2} = \frac{1}{M} \sum_{i=1}^M \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{j=1}^N \left[\frac{1}{2} \|\varepsilon^{jT} (\mathbf{d} + \sigma^i - \mathcal{F}(\mathbf{u}))\|_2^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0 - \delta^i\|_{\Gamma^{-1}}^2 \right]. \quad (29)$$

We would like to point out (29) is perhaps the most intuitive way to combine RMA and RMAP. Randomization of the noise covariance matrix acts as a random projection (left sketching) while the randomized prior mean and data aid in sampling from the posterior. Note that (29) arises as a variant of the loss function defined in equation (10) by exchanging the optimization and expectation of only σ and δ . On the other hand, (28) would likely yield inaccurate results as it is the sum of solutions where the PtO map and data have been projected onto a one dimensional subspace: thus the prior dominates each solution. In the linear case, $\mathbf{u}_{N,M}^{\text{RMA+RMAP}_2}$ can be written as

$$\mathbf{u}_{N,M}^{\text{RMA+RMAP}_2} = \frac{1}{M} \sum_{i=1}^M \left(\frac{1}{N} \sum_{j=1}^N (\tilde{\mathcal{A}}^j)^T \tilde{\mathcal{A}}^j + \Gamma^{-1} \right)^{-1} \left(\frac{1}{N} \sum_{j=1}^N (\tilde{\mathcal{A}}^j)^T (\varepsilon^j)^T \mathbf{d}^i + \Gamma^{-1} \mathbf{u}_0^i \right).$$

Clearly we also have convergence to the MAP point of other combinations, such as randomizing only one of the data or prior mean. To avoid a combinatorial explosion in the number of corollaries, we omit all the possibilities here.

3.3 Randomized prior

Here we propose a randomization scheme based on randomizing Γ^{-1} though $\boldsymbol{\lambda} \sim \pi_{\boldsymbol{\lambda}}$ where $\mathbb{E}_{\pi_{\boldsymbol{\lambda}}} [\boldsymbol{\lambda}] = \mathbf{0}$ and $\mathbb{E}_{\pi_{\boldsymbol{\lambda}}} [\boldsymbol{\lambda} \boldsymbol{\lambda}^T] = \Gamma^{-1}$. Let

$$\begin{aligned} \mathbf{u}^{\text{RS.U1}} &:= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_{\boldsymbol{\lambda}}} \left[\tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \boldsymbol{\lambda}) \right] \\ &= \arg \min_{\mathbf{u}} \mathbb{E}_{\pi_{\boldsymbol{\lambda}}} \left[\frac{1}{2} \|\mathbf{d} - \mathcal{F}(\mathbf{u})\|_{\Sigma^{-1}}^2 + \frac{1}{2} \|\boldsymbol{\lambda}^T (\mathbf{u} - \mathbf{u}_0)\|_2^2 \right]. \end{aligned} \quad (30)$$

The reason for designating this method ‘‘RS_U1’’ is due to its relationship with the right sketching approach (see Section 4.1). Then the SAA reads

$$\begin{aligned} \mathbf{u}_N^{\text{RS.U1}} &:= \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{u}; \mathbf{u}_0, \mathbf{d}, \varepsilon^i) \\ &= \arg \min_{\mathbf{u}} \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{2} \|\mathbf{d} - \mathcal{F}(\mathbf{u})\|_{\Sigma^{-1}}^2 + \frac{1}{2} \|(\varepsilon^i)^T (\mathbf{u} - \mathbf{u}_0)\|_2^2 \right]. \end{aligned} \quad (31)$$

Corollary 12 (Asymptotic convergence of randomized prior). *Let $\mathbf{u}_N^{\text{RS,U1}}$ be as defined above. Then*

$$\mathbf{u}_N^{\text{RS,U1}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } N \rightarrow \infty.$$

In the linear case, the optimality condition yields the following solution

$$\mathbf{u}^{\text{RS,U1}} = \left(\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \right)^{-1} \left(\mathcal{A}^T \Sigma^{-1} \mathbf{d} + \mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \mathbf{u}_0 \right).$$

This approach can be thought of as sketching the prior from the left.

4 Optimize, transform, then randomize

An important observation about the methods discussed so far is that the linear settings are solved using \mathbf{u}_1 given in (17a). That is, to show the equivalence of the solution of the randomized cost function to the solution of the corresponding method in the literature, one only needs to consider the optimal solution (17a). Additionally, the sample average approximation of the cost function is exactly the same as replacing the expectations in form \mathbf{u}_1 with their respective sample average approximations. The next methods require form \mathbf{u}_2 in (17b) to be used to see the equivalence of the randomized solution and the corresponding method given in the literature – where the Sherman-Morrison-Woodbury formula is applied to the optimality condition before making sample average approximations. As \mathbf{u}_2 is only equivalent to \mathbf{u}_1 in the linear case, we will restrict the following discussion to linear inverse problems.

4.1 Right sketching and the Ensemble Kalman filter

Since we are now considering schemes derived from randomizing \mathbf{u}_2 , we introduce a new random variable, ω , defined such that $\mathbb{E}_{\pi_\omega} [\omega] = 0$ and $\mathbb{E}_{\pi_\omega} [\omega \omega^T] = \Gamma$. By taking advantage of the asymptotic convergence of the SAA of $\mathbb{E}_{\pi_\omega} [\omega \omega^T]$, we have

$$\frac{1}{N} \sum_{i=1}^N \omega^i (\mathcal{A} \omega^i)^T \xrightarrow[N \rightarrow \infty]{a.s.} \mathbb{E}_\omega [\omega \omega^T \mathcal{A}^T] = \Gamma \mathcal{A}^T, \quad (32a)$$

$$\frac{1}{N} \sum_{i=1}^N (\mathcal{A} \omega^i) (\mathcal{A} \omega^i)^T \xrightarrow[N \rightarrow \infty]{a.s.} \mathbb{E}_\omega [\mathcal{A} \omega \omega^T \mathcal{A}^T] = \mathcal{A} \Gamma \mathcal{A}^T. \quad (32b)$$

Combining (17b) and (32) gives

$$\mathbf{u}_N^{\text{RS}} := \mathbf{u}_0 + \left(\frac{1}{N} \sum_{i=1}^N \omega^i (\mathcal{A} \omega^i)^T \right) \left(\Sigma + \frac{1}{N} \sum_{i=1}^N (\mathcal{A} \omega^i) (\mathcal{A} \omega^i)^T \right)^{-1} (\mathbf{d} - \mathcal{A} \mathbf{u}_0), \quad (33)$$

which is the same as sketching the PtO map \mathcal{A} from the right or sketching the transpose of the PtO map from the left.

Lemma 13 (Asymptotic convergence of right sketching). *Let \mathbf{u}_N^{RS} be defined in (33) and assume that $\mathbb{E}_{\pi_\lambda} [\lambda \lambda^T]$ (where λ is defined in Section 2) is invertible. Then*

$$\mathbf{u}_N^{\text{RS}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } N \rightarrow \infty.$$

Proof. Beginning with equation (2) and randomizing only Γ^{-1} through λ , the optimality condition is

$$\mathbf{u}^* = \left(\mathcal{A} \Sigma^{-1} \mathcal{A} + \mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \right)^{-1} \left(\mathcal{A}^T \Sigma^{-1} \mathbf{d} + \mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \mathbf{u}_0 \right).$$

Since $\mathbb{E}_{\pi_\lambda} [\lambda \lambda^T]$ is assumed to be invertible, this can be rewritten using the Sherman-Morrison-Woodbury formula in the form \mathbf{u}_2 as

$$\mathbf{u}^* = \mathbf{u}_0 + \left(\mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \right)^{-1} \mathcal{A}^T \left(\Sigma + \mathcal{A} \left(\mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \right)^{-1} \mathcal{A}^T \right)^{-1} (\mathbf{d} - \mathcal{A} \mathbf{u}_0).$$

Before making a sample average approximation, note that

$$\left(\mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \right)^{-1} = (\Gamma^{-1})^{-1} = \Gamma = \mathbb{E}_{\pi_\omega} [\omega \omega^T].$$

Then,

$$\mathbf{u}^* = \mathbf{u}^{\text{RS}} = \mathbf{u}_0 + \mathbb{E}_{\pi_\omega} [\omega \omega^T] \mathcal{A}^T (\Sigma + \mathcal{A} \mathbb{E}_{\pi_\omega} [\omega \omega^T] \mathcal{A}^T)^{-1} (\mathbf{d} - \mathcal{A} \mathbf{u}_0).$$

since matrix multiplication and matrix inversion are continuous functions,

$$\mathbf{u}_N^{\text{RS}} \xrightarrow{a.s.} \mathbf{u}^{\text{RS}} = \mathbf{u}^*.$$

by the continuous mapping theorem [56, Theorem 2.3]. \square

The key step here is recognizing that, asymptotically, sampling from π_λ and solving using form \mathbf{u}_1 (17a) gives the same results as sampling from π_ω and solving using form \mathbf{u}_2 (17b). However, Lemma 13 does not imply that \mathbf{u}_N^{RS} is equivalent to sampling from π_λ and solving using form \mathbf{u}_1 for a finite N . Indeed, when $N < \dim(\mathbf{u})$, \mathbf{u}_N^{RS} cannot be rewritten in the form \mathbf{u}_1 since

$$\text{rank} \left(\frac{1}{N} \sum_{j=1}^N \omega^j (\omega^j)^T \right) \leq N < \dim(\mathbf{u}).$$

This implies that the sample average of $\mathbb{E}_{\pi_\omega} [\omega \omega^T]$ is not invertible, breaking an assumption of Lemma 13 and showing that \mathbf{u}_N^{RS} does not satisfy the optimality condition of (2). Here, it is important that Γ^{-1} is invertible, otherwise $\mathbf{u}_1 \neq \mathbf{u}_2$ and the randomized schemes discussed here have no hope of converging to \mathbf{u}^* .

If in addition to right sketching we use (19) to randomize \mathbf{d} and \mathbf{u}_0 in (33), and define

$$(\mathbf{u}_N^{\text{ENKF}})^i := \mathbf{u}_0^i + \left(\frac{1}{N} \sum_{j=1}^N \omega^j (\mathcal{A} \omega^j)^T \right) \left(\Sigma + \frac{1}{N} \sum_{j=1}^N (\mathcal{A} \omega^j) (\mathcal{A} \omega^j)^T \right)^{-1} (\mathbf{d}^i - \mathcal{A} \mathbf{u}_0^i), \quad (34)$$

we rediscover the well-known ensemble Kalman filter (EnKF) update formula for a single member of the ensemble [57]. Notice here that the sketching of \mathcal{A} from the right is fixed for each random sample \mathbf{d}^i and \mathbf{u}_0^i . In the language of the EnKF, the sample prior covariance matrix is fixed for all members of the ensemble.

Corollary 14 (Asymptotic convergence of EnKF). *Let*

$$\mathbf{u}_{M,N}^{\text{ENKF}} := \frac{1}{M} \sum_{i=1}^M (\mathbf{u}_N^{\text{ENKF}})^i. \quad (35)$$

Then

$$\mathbf{u}_{M,N}^{\text{ENKF}} \xrightarrow{a.s.} \mathbf{u}^{\text{MAP}} \quad \text{as } M, N \rightarrow \infty.$$

Proof. The result follows immediately from Theorem 1, Lemma 13, and the interchange of expectation and integration used in Corollary 9. \square

As with right sketching, this result only holds asymptotically, with special sensitivity to N , since the validity of $(\mathbf{u}_N^{\text{ENKF}})^i$ as an optimal solution of (2) requires N to be large enough to ensure invertibility of all matrices involved. It should be noted that the EnKF is often used as part of an iterative method for solving inverse problems rather than used directly [33]. The reason for this can be understood by investigating what right sketching (and thus the EnKF) is doing to the prior covariance matrix and viewing this through the lens of regularization.

4.1.1 Right sketching from the left as randomized regularization

Consider again the form \mathbf{u}_1 given in (17a):

$$\mathbf{u}_1 = (\mathcal{A}^T \Sigma^{-1} \mathcal{A} + \Gamma^{-1})^{-1} (\mathcal{A}^T \Sigma^{-1} \mathbf{d} + \Gamma^{-1} \mathbf{u}_0).$$

While the randomized prior (31), right sketching (33) and ensemble Kalman filter (34) methods still fall under the asymptotic analysis given in Section 2.1 for linear inverse problems, a practical and theoretical issue arises due to the regularizing role that Γ^{-1} plays. The inverse of the prior covariance, Γ^{-1} , can be considered to be a regularization operator when viewed through the lens of deterministic inverse problems and is indeed equivalent to a Tikhonov regularization strategy [58]. In the deterministic setting, the role of regularization is often to “damp out” highly oscillatory modes caused by the rapidly decaying spectrum of \mathcal{A} — modes that are highly polluted by noise. While asymptotic analysis (see Theorem 1) establishes the convergence of these aforementioned methods, it is incapable of explaining

why these methods could fail for finite sample size N . This is where non-asymptotic analysis shines. Indeed, Lemma 6 shows that the successful (small error) probability requires quite a large number of samples. According to Remark 4.7.2 of [37], the number of samples required to accurately estimate the covariance matrix is proportional to n/β^2 where n is the dimension of the matrix and β is the tolerance. This is not surprising from a regularization point of view as the sample covariance needs to closely approximate the true covariance in order to adequately perform its role as a regularizer.

As a concrete example, consider the simple case where $\Gamma^{-1} = \alpha \mathcal{I}$ with $\alpha > 0$. Letting $\Sigma^{-\frac{1}{2}} \mathcal{A} = \mathbf{U} \mathbf{S} \mathbf{V}^T$ be the SVD of the whitened PtO map, the first term of \mathbf{u}_1 can be written

$$\left(\mathbf{V} \mathbf{S}^2 \mathbf{V}^T + \Gamma^{-1} \right)^{-1} = \left(\mathbf{V} (\mathbf{S}^2 + \alpha \mathcal{I}) \mathbf{V}^T \right)^{-1} = \mathbf{V} \mathbf{D} \mathbf{V}^T,$$

where \mathbf{D} is the diagonal matrix with the i th diagonal element given by $\mathbf{D}_{ii} = \frac{1}{\mathbf{S}_{ii}^2 + \alpha}$. Comparing to the case of no regularization, we can see that the inverse of the prior covariance shifts the spectrum of $\mathcal{A}^T \Sigma^{-1} \mathcal{A}$ upward by the constant α . Furthermore, upon inverting, $\alpha > 0$ ensures that the denominator of $\frac{1}{\mathbf{S}_{ii}^2 + \alpha}$ is not too close to 0, keeping the inverse solution from blowing up as $\mathbf{S}_{ii}^2 \rightarrow 0$. Now, consider the RS_U1 randomization of Γ^{-1} proposed in (31) — the same randomization as right sketching when viewed in the \mathbf{u}_1 form (sketching the prior from the left):

$$\Gamma^{-1} = \mathbb{E}_{\pi_\lambda} [\lambda \lambda^T] \approx \frac{1}{N} \sum_{i=1}^N (\lambda^i) (\lambda^i)^T.$$

As we saw before, this randomization converges as $N \rightarrow \infty$, but the convergence rate $\mathcal{O}(1/\sqrt{N})$ of a SAA is notoriously slow. So how does this slow convergence affect the regularization strategy? Clearly when $N < \dim(\mathbf{u})$, the regularization is not full rank and there may be $\dim(\mathbf{u}) - N$ modes of $\mathcal{A}^T \Sigma^{-1} \mathcal{A}$ left unregularized, assuming the random matrix has linearly independent columns. Even in the case when $N \geq \dim(\mathbf{u})$, slow convergence of the SAA leaves modes underregularized leading to oscillatory solutions as seen in Figures 11d and 11e for the 1D deconvolution problem. This can also be seen explicitly in Figure 1a where the spectrum of the sample average inverse covariance is plotted against the spectrum of the true prior inverse covariance for various N .

In Figure 1a, we consider the case where $\Gamma = \mathcal{I}$ and $\dim(\mathbf{u}) = 1000$. Because randomizing the inverse of the prior




Figure 1: Convergence of spectrum of the sample average approximation of the inverse prior covariance for the case where $\Gamma = \mathcal{I} \in \mathbb{R}^{1000 \times 1000}$ (a) and Γ is the BiLaplacian (b). When $N = 100$, there are fewer samples than the dimension of the parameter and some modes are left completely unregularized. Even when there are more samples than the dimension of the parameter, this does not guarantee acceptable convergence for SAA. This shows that the sample average of the inverse prior covariance converges slowly to the true inverse prior covariance. However, when the spectrum of the inverse prior covariance decays, the sample average approximation more closely matches the true inverse prior covariance with fewer samples.

covariance results in a poor performing regularizer, solutions using right sketching or a single step of the ensemble

Kalman filter exhibit highly oscillatory behavior when choosing N to be of reasonable size, at least for the identity prior. In problems where a decaying prior spectrum is desirable, randomization of the prior has a less pronounced effect on the quality of the inverse solution. For example, the advection-diffusion PDE constrained inverse problem detailed in Section 5.3 with the BiLaplacian prior shows good results with right sketching. The similarity of the sample average spectrum to the spectrum of the true prior inverse covariance can be seen in Figure 1b for the BiLaplacian prior. Additionally, problems where the PtO map has a slowly decaying spectrum as in the X-ray tomography problem (Section 5.2) may also be less sensitive to inaccurate approximations to Γ^{-1} .

5 Numerical results

In this section we show numerical results for a variety of inverse problems demonstrating the asymptotic convergence of various methods. As the possible number of randomized variants would be unnecessarily burdensome to enumerate, we will focus on a few key methods: randomized misfit approach (24), randomized MAP (21), the combination of RMA and RMAP (29), right sketching (33), the ensemble Kalman filter (35), and randomizing everything (10) (listed as ALL). It is important to keep in mind that we are not advocating for or against the use of any particular method — this section is to serve as numerical validation of the asymptotic convergence of each method. Additionally we discuss the differing convergence behavior of each method for different problems. In particular, we find empirically that methods randomizing Γ^{-1} such as right sketching, the EnKF, and ALL generally have very poor performance and require many more samples than the dimension of the problem in order to provide suitable results for several problems. The reason for this has been discussed at length in Section 4.1.1. These methods do however exhibit asymptotic convergence to the MAP solution as predicted by our theoretical results.

To explore the performance and convergence of the various methods, we consider a variety of prototype problems with different characteristics. The 1D deconvolution problem with scaled identity prior covariance is a relatively simple inverse problem that provides easily digestible visualizations of the convergence for each method. X-ray tomography is a mildly ill-posed two dimensional imaging problem with fewer observations than parameters. The fact that it is only mildly ill-posed exposes interesting effects in the context of randomization. We also show the convergence of each method for a linear time dependent PDE-constrained inverse problem with PDE-based prior covariance on a domain with a hole. Finally, we conclude with an example demonstrating convergence on a non-linear elliptic PDE-constrained inverse problem.

In problems with more than one randomization, such as EnKF and RMA+RMAP, each expectation can be approximated by a separate sample average. However, exploring the effect of choosing a different number of samples for each random variable is outside the scope of this paper and serves only to obscure the asymptotic convergence property that we aim to show in this section. Therefore, all methods assume that the number of random samples is the same for all random variables. To be more concrete, we set $\mathbf{u}_{M,N}^{\text{ENKF}} = \mathbf{u}_{N,N}^{\text{ENKF}}$ in (35). In addition, the relative errors presented are with respect to \mathbf{u}^{MAP} , not the true solution, emphasizing the errors induced by randomization rather than errors due to other effects. This is due to the fact that the theory presented shows convergence to \mathbf{u}^{MAP} .

5.1 1D Deconvolution problem

Deconvolution, the inverse problem associated with the convolution process, finds enormous application in the signal and image processing domains [59, 60, 61]. For demonstration, we consider the 1D deconvolution problem with a 1-periodic function given by:

$$f(x) = \sin(2\pi x) + \cos(2\pi x) \quad x \in [0, 1].$$

The domain is divided into $n = 1000$ sub-intervals. The kernel is constructed as [62]:

$$\Psi(x) = C_a (x + a)^2 (x - a)^2,$$

where $a = 0.235$ and the constant C_a is chosen to enforce the normalization condition [62]. Synthetic observations are generated with 5% additive Gaussian noise. We choose to randomize using the Achlioptas distribution [28, 63], an example of an l -percent sparse random variable with $l = 2/3$ and entries in $\{-1, 0, 1\}$ with equal probability. The reconstructed functions obtained by different randomization approaches are shown in Figure 11 and the relative errors are given in Table 1. It can be seen that the right sketching and EnKF methods give the least accurate results as evidenced in Table 1. This is because randomizing the inverse of the prior covariance results in poor performance as a regularizer, providing numerical confirmation of the discussion in Section 4.1.1. Other methods perform reasonably well. While not all methods perform equally well, all methods converge as more samples are taken and this is consistent with our asymptotic convergence results.

Method	Relative error (%)				
	N = 10	N = 100	N = 1000	N = 10000	N = 100000
RMAP	7.74	2.39	0.80	0.25	0.07
RMA	35.5	4.2	1.8	0.42	0.18
RMA+RMAP	41.7	9.0	1.7	0.80	0.15
RS	917	295	106	31.4	9.8
ENKF	896	352	100	31.9	9.7
ALL	158	33875	844	33.3	9.3

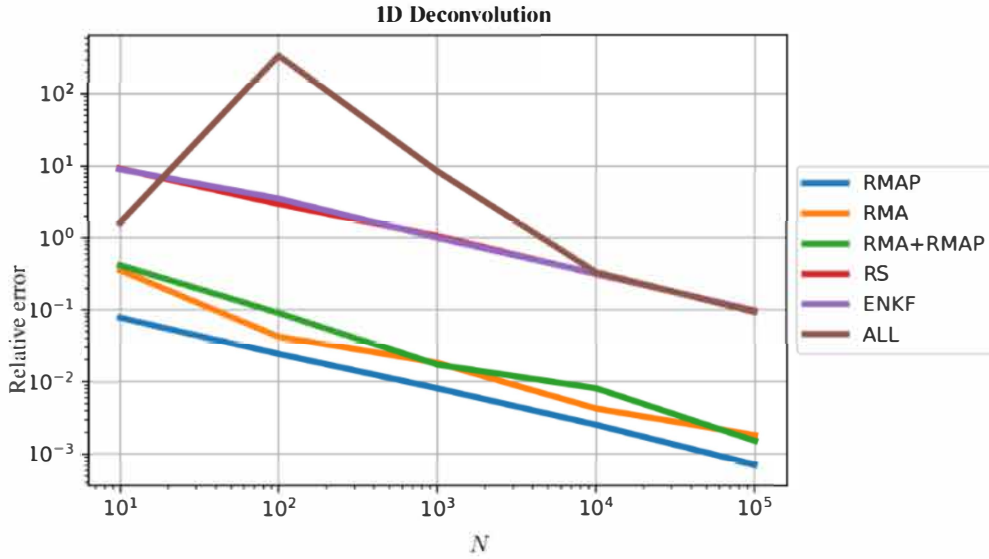
Table 1: Relative error for various randomized methods compared to the u^{MAP} solution for 1D deconvolution.

Figure 2: Relative error plot for 1D Deconvolution problem with Achlioptas random variable.

5.2 X-ray tomography

In x-ray tomographic imaging, X-ray projections of an object are captured at multiple angles and the inverse problem is to recover the internal structure of the object from the projection data [62]. We consider the canonical *phantom* image of size 64×64 pixels with 45 measurement angles uniformly divided over the range $[0, \pi]$. With this number of measurement angles, the PiO map has shape 64×45 by 64^2 resulting in fewer observations than parameters (pixels). A scaled identity prior covariance is once again considered. Measurements are corrupted with 1% additive Gaussian noise.

Method	Relative error (%)				
	N = 10	N = 100	N = 1000	N = 10000	N = 50000
RMAP	6.44	6.44	6.44	6.44	0.04
RMA	94.17	74.95	39.42	31.07	4.02
RMA+RMAP	96.64	77.94	40.35	30.89	4.02
RS	191.87	373.07	176.02	51.87	22.43
ENKF	324.27	352.58	178.58	52.25	21.97
ALL	95.12	71.30	80.50	60.36	23.60

Table 2: Relative error for various randomized methods compared to the u^{MAP} solution for the X-ray tomography problem.




Figure 3: Relative error plot for X-ray tomography problem.

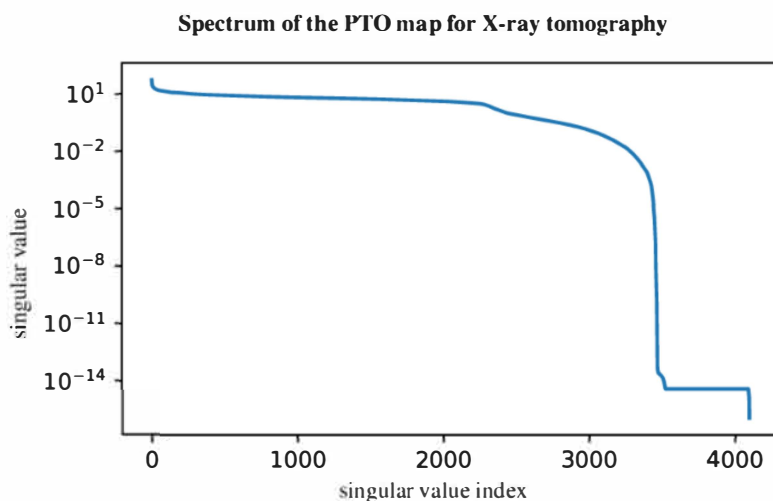


Figure 4: The singular values of the parameter-to-observable map for an X-ray tomography problem decay rapidly at first and then slowly until the last few singular vectors. This shows that the effective rank of the PtO map is close to the minimum dimension.

The results are shown in Figure 12 and Table 2 shows the relative error for different methods. Two observations are in order. First, results show asymptotic convergence of all methods, though convergence is noticeably slower for RMA and RMA+RMAP than in previous problems. This occurs because X-ray tomography is only a mildly ill-posed inverse problem with the spectrum of the PtO map decaying slowly after an initial fast decay (Figure 4). This means that the effective rank of the PtO map is close to the dimension of the data in the case presented. While mildly ill-posed problems are usually easier to work with, this can present a challenge for randomized methods, particularly methods such as RMA that randomize the misfit term. Recall that for any two matrices X and Y , $\text{rank}(XY) \leq \text{rank}(X)$. By projecting the misfit term onto a lower dimensional subspace, important information is lost in the case where \mathcal{A} is has effective rank close to the dimension of the data. This indicates that such a method is better suited for problems that are severely ill-posed.

The second observation is that the error for randomized prior methods initially increases then decreases as the number of samples increases. In previous problems, we have used a direct linear solver to find the solution to the stochastic optimization problem. In this problem, we use an iterative conjugate gradient (CG) solver to showcase how solver choice interacts with the randomized approaches. The main difference that can be seen here between a direct solver and an iterative solver such as CG is that a direct solver will invert possibly tiny eigenvalues of an ill-conditioned matrix while an iterative solver will often stop early, depending on the convergence parameters set, acting as an iterative regularizer [64, 65, 66]. This effect is particularly pronounced on the randomized prior methods such as RS and EnKF where the low rank randomized prior causes the iterative solver to stop earlier with fewer samples. This causes the error of RS and EnKF to increase initially until the regularization has sufficient rank, then the methods converge asymptotically to u^{MAP} . In the case of X-ray tomography, full-rank regularization is not needed due to the mildly ill-posed nature of the problem.

5.3 Initial condition inversion in an advection-diffusion problem

We now consider a linear inverse problem governed by a parabolic PDE based on the method used in [67]. The parameter to observable map (advection-diffusion equation) maps an initial condition $u \in L^2(\Omega)$ to pointwise spatio-temporal observations of the concentration field $y(x, t)$. The advection-diffusion equation is given by:

$$\begin{aligned} y_t - \kappa \Delta y + \vec{v} \cdot \nabla y &= 0 & \text{in } \Omega \times (0, T), \\ y(\cdot, 0) &= u & \text{in } \Omega, \\ \kappa \nabla y \cdot \mathbf{n} &= 0 & \text{on } \partial\Omega \times (0, T), \end{aligned} \quad (36)$$

where, $\Omega \subset \mathbb{R}^2$ is a bounded domain, $\kappa > 0$ is the diffusion coefficient, $T > 0$ is the final time. The velocity field \vec{v} is computed by solving the following steady-state Navier-Stokes equation with the side walls driving the flow:

$$\begin{aligned} -\frac{1}{\text{Re}} \Delta \vec{v} + \nabla q + \vec{v} \cdot \nabla \vec{v} &= 0 & \text{in } \Omega, \\ \nabla \cdot \vec{v} &= 0 & \text{in } \Omega, \\ \vec{v} &= \vec{g} & \text{on } \partial\Omega. \end{aligned} \quad (37)$$

where q is the pressure, and Re is the Reynolds number. The Dirichlet boundary condition $\vec{g} \in \mathbb{R}^2$ is prescribed as $\vec{g} = [0, 1]$ on the left side of the domain, and $\vec{g} = [0, 0]$ elsewhere. Velocity boundary conditions are not prescribed on the right side of the boundary. The values of the forward solution y on a set of locations $\{x_1, x_2, \dots, x_m\}$ at the final time T are extracted and used as the observation vector $\mathbf{d} \in \mathbb{R}^k$ for solving the initial condition inverse problem. Synthetic observations are generated by corrupting this observation vector with 1% additive Gaussian noise. The observation data and the velocity profile used in the study are shown in Figure 5. Upon discretization, the operator \mathcal{A} maps the initial condition $u \in \mathbb{R}^n$ to the observation $\mathbf{d} \in \mathbb{R}^k$.

In addition, we define the prior covariance matrix to be the PDE-based BiLaplacian prior defined as:

$$\mathbf{\Gamma} = (\delta I + \gamma \nabla \cdot (\theta \nabla))^{-2}, \quad (38)$$

where, δ governs the variance of the samples, while the ratio $\frac{\gamma}{\delta}$ governs the correlation length. θ is a symmetric positive definite tensor to introduce anisotropy in the correlation length.

Following [67], a mixed formulation employing $P2$ Lagrange elements for approximating the velocity field and $P1$ elements for pressure is adopted for solving (37) to obtain the velocity field. The computed velocity field is then used to solve the advection-diffusion equation, (36). $P1$ Lagrange elements are used for the variational formulation of the advection-diffusion equation.

The observation vector \mathbf{d} is computed at time $t = 3s$ with $m = 200$ observation points. For this problem, there are $n = 2868$ degrees of freedom. The diffusion coefficient is $\kappa = 0.001$ and the parameters of the BiLaplacian prior (38) are $\delta = 8$, $\gamma = 1$, and $\Theta = \mathcal{I}$.

The MAP solution u^{MAP} is shown in Figure 7. The condition number of $\mathbf{\Gamma}^{-1}$ is of the order of 10^6 . The results of different randomization schemes are shown in Figure 13 in the appendix. Table 3 gives the relative error with respect to u^{MAP} . As expected, randomized MAP gives the most accurate results followed by left sketching and the randomized misfit approach. In contrast to the previous examples considered, the right sketching and EnKF approaches gives reasonably good results as evident from Table 3 and Figure 13 in the appendix.

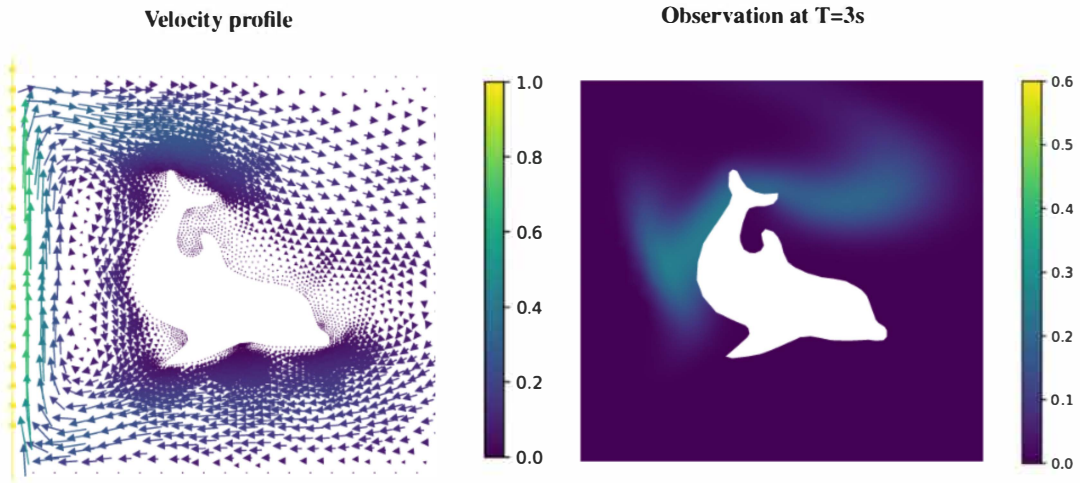


Figure 5: The velocity profile and observation data used for inversion

Method	Relative error (%)			
	N = 10	N = 100	N = 1000	N = 10000
RMAP	5.02	1.96	0.49	0.15
RMA	53.38	15.16	5.30	1.15
RMA+RMAP	80.14	14.48	7.05	1.30
RS	59.58	26.78	7.33	5.06
ENKF	74.09	18.76	7.10	5.43
ALL	91.58	197.66	193.25	9.66

Table 3: Relative error for various randomized methods compared to the u^{MAP} solution for 2D linear advection-diffusion initial condition inverse problem.




Figure 6: Relative error plot for 2D linear advection-diffusion initial condition inversion.

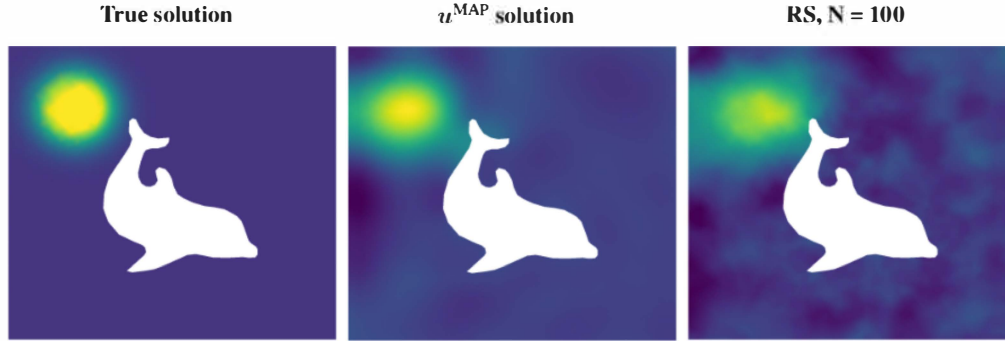


Figure 7: From left to right are true solution, u^{MAP} solution and right sketching solution for linear advection-diffusion initial condition inverse problem. With only 100 random samples, right sketching can obtain a reasonably good initial condition reconstruction (26% relative error). This behavior indicates the fast convergence of the randomly sampled prior inverse covariance to the true prior inverse covariance.

This is due to the faster convergence of the randomized prior covariance to the true prior covariance for the BiLaplacian prior. This also points to the fact that care should be exercised when choosing a randomized method for a particular problem. For inverse problems where a prior with a decaying spectrum is desirable, the EnKF or RS approach may perform well.

5.4 Nonlinear parameter inversion in a steady-state heat equation

To show the convergence of various methods for nonlinear inverse problems, we consider a nonlinear PDE constrained parameter inversion problem. In the previous section, we considered an initial condition problem where the data depended linearly on the parameter, even though the statement of the problem itself was rather involved. Now we consider a simple to state but extremely ill-posed nonlinear inverse problem. Given a steady-state temperature distribution $T(x, y)$ and boundary conditions, invert for the conductivity everywhere in the domain. The governing equations are given by

$$\begin{aligned} \nabla \cdot (e^{\kappa} \nabla T) &= f \quad \text{in } \Omega, \\ T(x, 0) &= 2(1 - x), \\ T(x, 1) &= 2x, \\ \nabla T \cdot \mathbf{n} &= 0 \quad \text{on } \partial\Omega \setminus \{y = 0, y = 1\}. \end{aligned}$$

While this equation is linear in the temperature distribution T , the (log-) thermal conductivity that we are inverting for, κ , appears non-linearly. That is, the parameter-to-observable map is nonlinear. The heat equation makes for an excellent test problem for inverse solvers since the dependence of the steady-state temperature distribution on the conductivity is rather weak.

We again follow a mixed formulation where the temperature distribution is modeled using $P2$ Lagrange elements and the parameter is modeled with $P1$ Lagrange elements. With a mesh size of 64×64 elements, this results in discrete variables $\mathbf{y} \in \mathbb{R}^{16,641}$ and $\kappa \in \mathbb{R}^{4,225}$. The BiLaplacian prior defined in (38) is also used here with $\delta = 0.5$, $\gamma = 0.1$, and the anisotropic diffusion tensor

$$\theta = \begin{bmatrix} \theta_1 \sin^2 \alpha & (\theta_1 - \theta_2) \sin \alpha \cos \alpha \\ (\theta_1 - \theta_2) \sin \alpha \cos \alpha & \theta_2 \cos^2 \alpha \end{bmatrix},$$

where $\theta_1 = 2.0$, $\theta_2 = 0.5$ and $\alpha = \pi/4$. Lastly, we consider an inhomogeneous case where

$$f = 50 \sin^2(\pi x) \cos^2(\pi y).$$

A few remarks are in order to understand the rather unimpressive results in Table 4. First, recall that the results shown here are for an extremely difficult problem. The inverse of the diffusion equation is notoriously ill-posed, as it amounts to the inverse of a compact operator. That is, small perturbations in the data can lead to drastically different inversion results. Indeed, we show an even more difficult problem where the task is to infer 4,225 parameters from only 100 measurements which are recorded in only half of the domain as shown in Figure 8. With so few measurements, adding noise to the data as in RMAP, RMA+RMAP, ENKF, and ALL may not lead to desirable results.




Figure 8: The true log-conductivity (κ) and 100 sparse observations. Observing the temperature distribution only in the lower half of the domain makes inverting for the log-conductivity in the entire domain a more difficult task.

Method	Relative error (%)			
	N = 10	N = 100	N = 1000	N = 5000
RMAP	19.50	14.66	16.39	16.17
RMA	24.97	8.12	4.19	3.08
RMA+RMAP	17.51	17.82	16.27	16.49
RS	47.06	110.88	107.46	14.44
ENKF	20.82	18.26	17.54	19.55
ALL	26.13	15.63	19.45	19.62

Table 4: Relative error of MAP solution for various randomization schemes compared to the u^{MAP} solution.




Figure 9: Relative error plot for nonlinear elliptic parameter inverse problem.

Secondly, while there is still $\sim 20\%$ error for these methods, the MAP estimate for each of these is still reasonably good visibly as seen in Figure 10. The high relative error is in part due to the small norm of the \mathbf{u}^{MAP} solution. Simply shifting the parameter up by a constant changes the norm of the denominator in the relative error formula

$$\text{relative error} := \left\| \mathbf{u}_N^{(\text{method})} - \mathbf{u}^{\text{MAP}} \right\| / \left\| \mathbf{u}^{\text{MAP}} \right\|.$$

In addition to the numerical relative error, it is important to consider the “eyeball norm”. Figure 10 shows that the estimated solutions are still quite close to the \mathbf{u}^{MAP} solution, especially given the extreme ill-posedness. Despite this shortcoming of data-randomization methods, recall that the main advantage of additively randomizing the data and prior mean is to *aid in sampling* from the posterior. In other words, we are most interested in accelerating uncertainty quantification, not getting the SAA MAP estimate error down to machine precision. These randomization schemes may still find use in such applications.




Figure 10: Solutions to the nonlinear diffusion inverse problem for two different methods. Visually, these methods give nearly identical results compared to the \mathbf{u}^{MAP} solution even though numerically, they have relative error $\sim 20\%$.

6 Conclusions

By viewing the randomized solution of inverse problems through the lens of stochastic programming and the sample average approximation, we developed a unified framework through which we can analyze the asymptotic convergence of randomized solutions of linear inverse problems to the solution obtained with its deterministic counterpart. This framework allowed us to prove the asymptotic and non-asymptotic convergence of the minimizer of a general stochastic cost function to the minimizer of the expected value of the stochastic cost function. Several well-known methods for introducing randomness into linear and nonlinear inverse problems were recovered as special cases of this general framework. Viewing the solution to randomized inverse problems through the lens of the sample average approximation also allowed us to prove a novel non-asymptotic error analysis that applies to all randomized methods discussed. We also show that while all of the methods presented converge asymptotically, the results can be quite poor if an insufficient number of samples are drawn. While this observation is easily understood through our non-asymptotic error analysis, it is not possible from an asymptotic view point. In particular, we showed that randomizing the prior covariance matrix may not be a good idea for certain priors due to the regularizing role that the prior plays in the solution of inverse problems. This is due to the potentially slow convergence of random matrices to their expected value, depending on the spectrum of the expected value matrix. The convergence of all schemes was shown numerically for a variety of linear and nonlinear inverse problems, including 1D and 2D problems governed by algebraic and PDE constraints.

7 Acknowledgements

We would like to thank the Texas Advanced Computing Center (TACC) at the University of Texas at Austin for providing HPC resources that contributed to the results presented in this work. URL: <http://www.tacc.utexas.edu>

A Figures

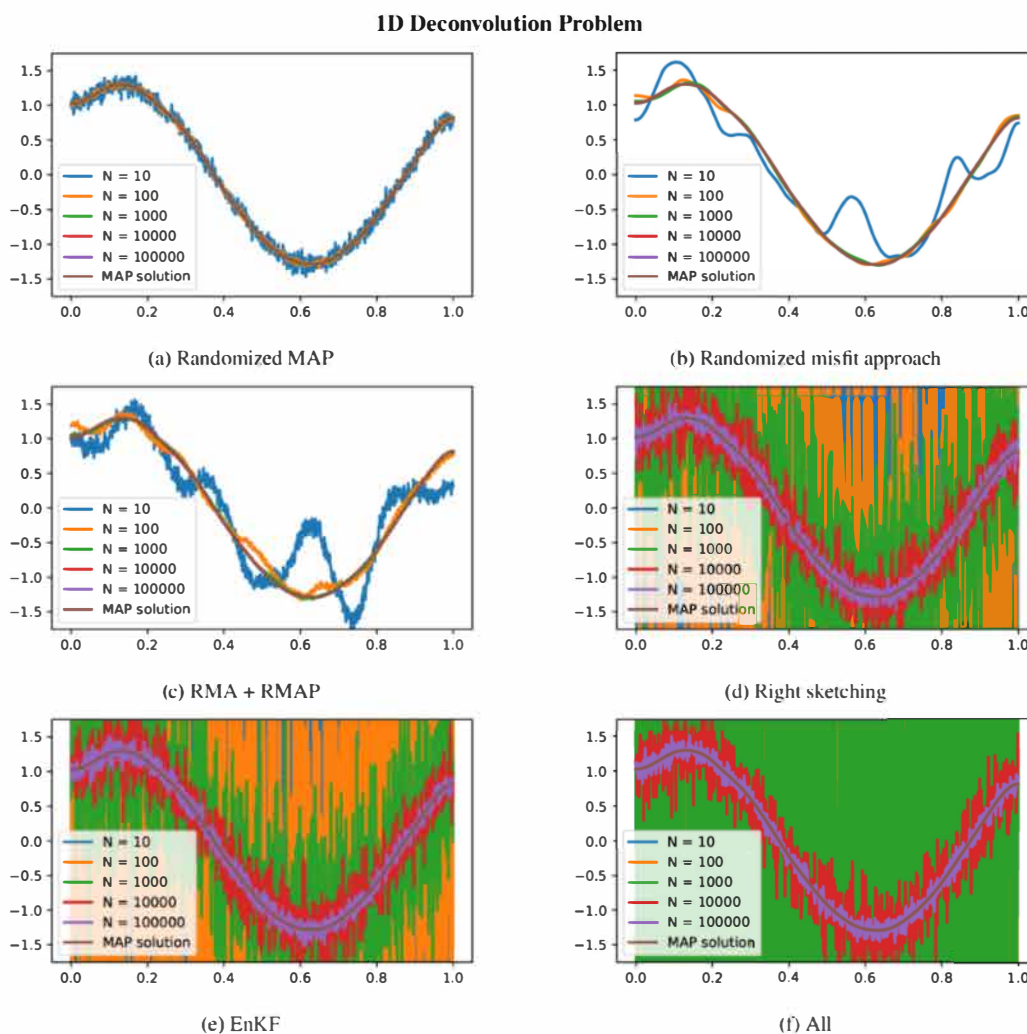


Figure 11: Solutions to 1D deconvolution problem with mesh size $n = 1000$ using various randomization schemes with scaled identity prior. This prior works sufficiently well for those randomization schemes that do not randomize the prior covariance (RMAP, RMA, RMA+RMAP), but performs poorly for RS, EnKF, and ALL which randomize the prior covariance. Random sampling is performed via an Achlioptas ($2/3$ -sparse) random variable




Figure 12: Solutions for various randomization approaches for an x-ray tomography problem with Gaussian random variables.




Figure 13: Solutions for various randomization approaches for a linear advection-diffusion initial condition inverse problem with Gaussian random variables.




Figure 14: Solutions for various randomization approaches for a nonlinear diffusion parameter inversion problem with Gaussian random variables.




Figure 15: Reconstructed state for various randomization approaches for a nonlinear diffusion problem. Even for the right sketching method which did not give good parameter reconstructions until $N = 10000$ samples, the state in the lower half of the domain, where the 100 measurements are taken, look similar to all the other methods.

References

- [1] Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*. 1953;21(6):1087–1092. Available from: <http://dx.doi.org/10.1063/1.1699114>.
- [2] Hastings WK. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*. 1970;57(1):97–109.
- [3] Haario H, Laine M, Miravete A, Saksman E. DRAM: Efficient adaptive MCMC. *Statistics and Computing*. 2006;16:339–354.
- [4] Robert CP, Casella G. *Monte Carlo Statistical Methods* (Springer Texts in Statistics). Secaucus, NJ, USA: Springer-Verlag New York, Inc.; 2005.
- [5] Wang K. *Parallel Markov Chain Monte Carlo Methods for Large Scale Statistical Inverse Problems* [Ph.D. thesis]. Texas A&M University; 2014.
- [6] Byrd J. *Parallel Markov Chain Monte Carlo* [Ph.D. thesis]. University of Warwick; 2010.
- [7] Wilkinson DJ. Parallel Bayesian Computation. In: Kontoghiorghes EJ, editor. *Handbook of Parallel Computing and Statistics*. Marcel Dekker/CRC Press; 2005. p. 481–512.
- [8] Brockwell A. Parallel Markov chain Monte Carlo simulation by pre-fetching. *Journal of Computational and Graphical Statistics*. 2006;15(1):246–261.
- [9] Strid I. Efficient parallelisation of Metropolis–Hastings algorithms using a prefetching approach. *Computational Statistics & Data Analysis*. 2010;54(11):2814–2835.
- [10] Duane S, Kennedy AD, Pendleton B, Roweth D. Hybrid Monte Carlo. *Phys Lett B*. 1987;195:216–222.
- [11] Neal RM. MCMC using Hamiltonian dynamics. In: Brooks S, Gelman A, Jones G, Meng XL, editors. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall / CRC Press; 2010. .
- [12] Girolami M, Calderhead B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2011;73(2):123–214. Available from: <http://dx.doi.org/10.1111/j.1467-9868.2010.00765.x>.
- [13] Beskos A, Pinski FJ, Sanz-Serna JM, Stuart AM. Hybrid Monte Carlo on Hilbert spaces. *Stochastic Processes and their Applications*. 2011;121:2201–2230.
- [14] Bui-Thanh T, Girolami MA. Solving Large-Scale PDE-constrained Bayesian Inverse Problems with Riemann Manifold Hamiltonian Monte Carlo. *Inverse Problems*. 2014;Special Issue:114014.
- [15] Cui T, Martin J, Marzouk YM, Solonen A, Spantini A. Likelihood-informed dimension reduction for nonlinear inverse problems. *arXiv preprint arXiv:14034680*. 2014;.
- [16] Cui T, Law KJ, Marzouk YM. Dimension-independent likelihood-informed MCMC. *Journal of Computational Physics*. 2016;304:109–137.
- [17] Martin J, Wilcox LC, Burstedde C, Ghattas O. A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*. 2012;34(3):A1460–A1487.
- [18] Bui-Thanh T, Ghattas O. A Scaled Stochastic Newton Algorithm for Markov Chain Monte Carlo Simulations. Submitted to *SIAM Journal of Uncertainty Quantification*. 2012;.
- [19] Petra N, Martin J, Stadler G, Ghattas O. A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet inverse problems. *SIAM Journal on Scientific Computing* (to appear). 2014;
- [20] Liu Q, Wang D. Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances in neural information processing systems*. 2016;29.
- [21] Han J, Liu Q. Stein variational gradient descent without gradient. In: *International Conference on Machine Learning*. PMLR; 2018. p. 1900–1908.
- [22] Chen P, Ghattas O. Projected Stein variational gradient descent. *Advances in Neural Information Processing Systems*. 2020;33:1947–1958.
- [23] Zhuo J, Liu C, Shi J, Zhu J, Chen N, Zhang B. Message passing Stein variational gradient descent. In: *International Conference on Machine Learning*. PMLR; 2018. p. 6018–6027.

- [24] Carpenter J, Clifford P, Fearnhead P. Improved particle filter for nonlinear problems. *IEE Proceedings-Radar, Sonar and Navigation*. 1999;146(1):2–7.
- [25] Van Der Merwe R, Doucet A, De Freitas N, Wan E. The unscented particle filter. *Advances in neural information processing systems*. 2000;13.
- [26] Yang T, Mehta PG, Meyn SP. Feedback particle filter. *IEEE transactions on Automatic control*. 2013;58(10):2465–2480.
- [27] Soto A. Self adaptive particle filter. In: *IJCAI*. Citeseer; 2005. p. 1398–1406.
- [28] Le EB, Myers A, Bui-Thanh T, Nguyen QP. A data-scalable randomized misfit approach for solving large-scale PDE-constrained inverse problems. *Inverse Problems*. 2017;33(6):065003. <https://arxiv.org/pdf/1603.01562.pdf>.
- [29] Wang K, Bui-Thanh T, Ghattas O. A Randomized Maximum A Posteriori Method for Posterior Sampling of High Dimensional Nonlinear Bayesian Inverse Problems. *SIAM Journal of Scientific Computing*, Accepted. 2018;<https://arxiv.org/pdf/1602.03658.pdf>.
- [30] Chen K, Li Q, Newton K, Wright SJ. Structured random sketching for pde inverse problems. *SIAM Journal of Matrix Analysis and Applications*. 2020;41.
- [31] Avron H, Sindhvani V, Woodruff D. Sketching structured matrices for faster nonlinear regression. *Advances in neural information processing systems*. 2013;26.
- [32] Wang J, Lee J, Mahdavi M, Kolar M, Srebro N. Sketching meets random projection in the dual: A provable recovery algorithm for big and high-dimensional data. In: *Artificial Intelligence and Statistics*. PMLR; 2017. p. 1150–1158.
- [33] Iglesias MA, Law KJ, Stuart AM. Ensemble Kalman methods for inverse problems. *Inverse Problems*. 2013;29(4):045001.
- [34] Shapiro A, Dentcheva D, Ruszczyński A. *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics; 2009.
- [35] Feller W. *An Introduction to Probability Theory and Its Applications*. vol. 2. 2nd ed. John Wiley & Sons; 1971.
- [36] Durrett R. *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press; 2019.
- [37] Vershynin R. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press; 2018.
- [38] Trefethen LN, Bau III D. *Numerical linear algebra*. vol. 50. Siam; 1997.
- [39] Gao X, Zhang M, Luo J. Tail Bounds for Norm of Gaussian Random Matrices with Applications. *Journal of Mathematics*. 2022;2022.
- [40] Sambale H. Some notes on concentration for α -subexponential random variables. *arXiv preprint arXiv:200210761*. 2020;.
- [41] Vladimirova M, Girard S, Nguyen H, Arbel J. Sub-Weibull distributions: Generalizing sub-Gaussian and sub-Exponential properties to heavier tailed distributions. *Stat*. 2020;9(1):e318.
- [42] Zhang H, Wei H. Sharper sub-weibull concentrations. *Mathematics*. 2022;10(13):2252.
- [43] (<https://math.stackexchange.com/users/3776/jj>) JJ. Young’s inequality for three variables;. URL:<https://math.stackexchange.com/q/406922> (version: 2013-05-30). Mathematics Stack Exchange. Available from: <https://math.stackexchange.com/q/406922>.
- [44] Latz J. On the well-posedness of Bayesian inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*. 2020;8(1):451–482.
- [45] Diao HA, Wei Y, Qiao S. Structured condition numbers of structured Tikhonov regularization problem and their estimations. *Journal of Computational and Applied Mathematics*. 2016;308:276–300.
- [46] Chu D, Lin L, Tan RC, Wei Y. Condition numbers and perturbation analysis for the Tikhonov regularization of discrete ill-posed problems. *Numerical Linear Algebra with Applications*. 2011;18(1):87–103.
- [47] Deng CY. A generalization of the Sherman–Morrison–Woodbury formula. *Applied Mathematics Letters*. 2011;24(9):1561–1564.
- [48] Rockafellar RT, Wets RJB. *Variational Analysis*. Berlin, Heidelberg, New York: Springer Verlag; 1998.
- [49] Kitanidis PK. Quasi-linear geostatistical theory for inversing. *Water Resour Res*. 1995;31(10):2411–2419.

- [50] Oliver DS, Reynolds AC, Liu N. Inverse theory for petroleum reservoir characterization and history matching. Cambridge University Press; 2008.
- [51] Bardsley J, Solonen A, Haario H, Laine M. Randomize-then-optimize: A method for sampling from posterior distributions in nonlinear inverse problems. submitted. 2013;.
- [52] Bui-Thanh T, Nguyen QP. FEM-based discretization-invariant MCMC methods for PDE-constrained Bayesian inverse problems. *Inverse Problems & Imaging*. 2016;10(4):943.
- [53] Chen Y, Dwivedi R, Wainwright MJ, Yu B. Fast mixing of Metropolized Hamiltonian Monte Carlo: Benefits of multi-step gradients. *J Mach Learn Res*. 2020;21:92–1.
- [54] Liu M, Kumar R, Haber E, Aravkin AY. Simultaneous shot inversion for nonuniform geometries using fast data interpolation. *arXiv: Optimization and Control*. 2018;.
- [55] Clarkson KL, Woodruff DP. Low-Rank Approximation and Regression in Input Sparsity Time. *J ACM*. 2017 jan;63(6). Available from: <https://doi.org/10.1145/3019134>.
- [56] Van der Vaart AW. Asymptotic statistics. vol. 3. Cambridge university press; 2000.
- [57] Evensen G. The Ensemble Kalman Filter: Theoretical formulation and practical implementation. *Ocean Dynamics*. 2003;53:343–367.
- [58] Engl HW, Kunisch K, Neubauer A. Convergence rates for Tikhonov regularisation of non-linear ill-posed problems. *Inverse Problems*. 1989;5(4):523. Available from: <http://stacks.iop.org/0266-5611/5/i=4/a=007>.
- [59] Kundur D, Hatzinakos D. Blind image deconvolution. *IEEE signal processing magazine*. 1996;13(3):43–64.
- [60] Swedlow JR. Quantitative fluorescence microscopy and image deconvolution. In: *Methods in cell biology*. vol. 114. Elsevier; 2013. p. 407–426.
- [61] Ryan O, Debbah M. Free deconvolution for signal processing applications. In: *2007 IEEE International Symposium on Information Theory*. IEEE; 2007. p. 1846–1850.
- [62] Mueller JL, Siltanen S. Linear and nonlinear inverse problems with practical applications. SIAM; 2012.
- [63] Achlioptas D. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of Computer and System Sciences*. 2003;66(4):671–687. Special Issue on PODS 2001. Available from: <https://www.sciencedirect.com/science/article/pii/S0022000003000254>.
- [64] Piccolomini EL, Zama F. The conjugate gradient regularization method in Computed Tomography problems. *Applied Mathematics and Computation*. 1999;102(1):87–99. Available from: <https://www.sciencedirect.com/science/article/pii/S0096300398100073>.
- [65] Landi G, Loli Piccolomini E, Tomba I. A stopping criterion for iterative regularization methods. *Applied Numerical Mathematics*. 2016;106:53–68. Available from: <https://www.sciencedirect.com/science/article/pii/S0168927416300368>.
- [66] Hanke M, Nagy JG. Restoration of atmospherically blurred images by symmetric indefinite conjugate gradient techniques. *Inverse Problems*. 1996 apr;12(2):157–173. Available from: <https://doi.org/10.1088/0266-5611/12/2/004>.
- [67] UT Austin, UC Merced. hIPPYlib: Inverse Problem PYthon library;. Available from: <https://hippylib.github.io/>.