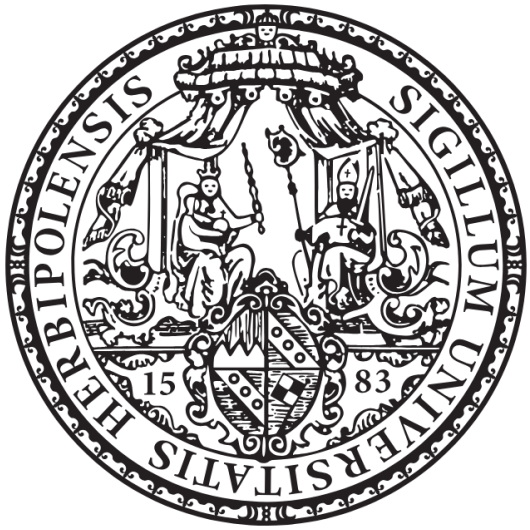
Experimental design for inverse problems preserving sensitivity through randomized down-sampling

– Bachelor's Thesis –

Submitted by

Anna Kircher



Under the supervision of

Prof. Dr. Christian Klingenberg

Mathematical Fluid Dynamics
Faculty of Mathematics and Computer Science
Julius-Maximilians-Universität Würzburg

Würzburg, November 11, 2026

Zusammenfassung

Bei einem inversen Problem werden unbekannte Parameter aus indirekten, verrauschten Messdaten rekonstruiert. Die Güte der Parameterrekonstruktion hängt entscheidend vom experimentellen Design ab, also von der Auswahl der Messkonfigurationen, unter denen die Daten erhoben werden. Da der zugrundeliegende Designraum oft sehr groß ist, wird die Auswertung aller möglichen Designpunkte schnell unpraktikabel, sodass eine Reduktion auf eine kleinere Teilmenge erforderlich wird. Im Gegensatz zu klassischen Methoden des optimalen experimentellen Designs, welche eine bezüglich der spektralen Eigenschaften der Fisher-Informationsmatrix (FIM) optimale Teilmenge suchen, verfolgt diese Thesis das Ziel, eine hinreichend informative Teilmenge zu identifizieren, welche die Sensitivitätseigenschaften des vollständigen Designs erhält, ohne dabei entscheidende Informationen für die Parameterrekonstruktion zu verlieren. Die FIM dient dabei als zentrales Werkzeug zur Quantifizierung dieser Sensitivität. Um eine geeignete Teilmenge zu konstruieren, werden Methoden der Randomized Numerical Linear Algebra, insbesondere des Matrix Sketchings, genutzt, welche die Informationseigenschaften der vollständigen FIM durch gezielte Stichprobenziehung approximativ erhalten. Um die Approximationsgüte der skizzierten FIM theoretisch abzusichern, werden probabilistische Fehlerschranken hergeleitet sowie Garantien für die Wohlgestelltheit des reduzierten Problems bewiesen. Da die hierfür optimale Sampling-Verteilung in der Praxis oft nicht direkt zugänglich ist, wird abschließend ein gradientenfreier Ensemble Kalman Sampler eingesetzt, der eine Approximation dieser Verteilung ermöglicht.

Abstract

In an inverse problem, unknown parameters are reconstructed from indirect, noisy measurement data. The quality of the parameter reconstruction depends crucially on the experimental design, that is, on the selection of the measurement configurations under which the data are collected. Since the underlying design space is often very large, evaluating all possible design points quickly becomes impractical, making it necessary to reduce the space to a smaller subset. In contrast to classical methods of optimal experimental design, which seek a subset that is optimal with respect to the spectral properties of the Fisher Information Matrix (FIM), this thesis aims to identify a sufficiently informative subset that preserves the sensitivity properties of the full design without losing critical information for parameter reconstruction. The FIM serves as the central tool for quantifying this sensitivity. To construct a suitable subset, methods of randomized numerical linear algebra, in particular matrix sketching, are used, which approximately preserve the information properties of the full FIM through targeted sampling. To theoretically verify the approximation quality of the sketched FIM, probabilistic error bounds are derived, and guarantees for the well-posedness of the reduced problem are proven. Since the optimal sampling distribution for this purpose is often not directly accessible in practice, a gradient-free Ensemble Kalman Sampler is finally employed, which enables an approximation of this distribution.

Acknowledgements

I would like to thank my supervisor, Professor Dr. Christian Klingenberg, for suggesting this interesting thesis topic, as well as for his guidance and support throughout the writing process. His calm demeanor and experience created a pleasant work environment, and I appreciate the trust and encouragement he showed me. I am also grateful for the opportunity to attend the weekly meetings with his research group, which introduced me to new mathematical topics and provided an opportunity for exchange with fellow students and researchers.

My gratitude also extends to Dr. Kathrin Hellmuth for her valuable explanations and feedback. Despite the distance and time difference between Germany and California, she always responded very quickly and offered helpful advice whenever I reached out to her.

Finally, I would like to thank my family for their constant support throughout my studies. Special thanks go to my brother Felix for his reassuring insights along the way.

Contents

1 Introduction **1**

1.1 Inverse Problem Terminology 1
1.2 The Fisher Information Matrix 2
1.3 The Down-Sampling Approach 4

2 Stochastic Preliminaries **6**

2.1 Stochastic Processes and Brownian Motion 6
2.2 Doob Martingales 8
2.3 The Azuma-Hoeffding Inequality 9
2.4 Stochastic Differential Equations 10

3 Matrix Sketching in Randomized Numerical Linear Algebra **13**

3.1 Matrix Multiplication 14
3.2 Matrix Sketching 15
3.3 Theoretical Guarantees 17

4 Translation to Experimental Design **26**

4.1 Sketching Setup 26
4.2 Weighted Matrix Sketching 28
4.3 Weighted Sketching Setup 31
4.4 Sensitivity Guarantees under Sketching 33

5 Sampling **38**

5.1 The Bayesian Inverse Problem 38
5.2 Ensemble Kalman Sampling 40
5.3 Mean-Field Limit 43
5.4 Experimental Design through Sampling 45
5.5 Drawbacks of the Sampling Approach 50

6 Conclusion and Outlook **54**

Bibliography **56**

1 Introduction

Inverse problems arise across a wide range of scientific disciplines, from medical imaging to geophysics, astrophysics, and industrial applications, whenever unknown quantities must be inferred from indirect and noisy measurements [FSU19]. In practice, the space of possible measurement configurations often is very large, so that processing the full set of measurements becomes impractical. Consequently, a central challenge in this setting is how to select measurement configurations that are informative enough to allow for a reliable reconstruction. This task is known as experimental design.

This chapter introduces the terminology and foundational concepts required throughout this thesis. Section 1.1 formalizes the notion of an inverse problem. Section 1.2 introduces the Fisher Information Matrix as the central tool for quantifying the sensitivity of measurements to the unknown parameter. Section 1.3 then motivates the down-sampling approach pursued in this thesis.

1.1 Inverse Problem Terminology

An inverse problem differs fundamentally from a direct problem. While the forward problem describes the causal direction, meaning inferring the resulting effects from known causes, the inverse problem reverses this reasoning. The goal of an inverse problem is to infer the underlying, uncertain causes from the observed effects.

We consider a classic statistical inverse problem formulated as an explicit forward model with additive measurement noise as in [HKL26]:

$$y(\xi) = \mathcal{F}(p, \xi) + \eta(\xi). \tag{1.1}$$

Here, the forward operator $\mathcal{F} : \Xi \times \mathbb{R}^K \to \mathbb{R}$ maps the design variable $\xi \in \Xi$ and the parameter $p \in \mathbb{R}^K$ to the data. The parameter vector p represents the unknown quantities of interest that cannot be measured directly. Conversely, the design variable ξ belongs to the design set Ξ, which includes all permissible experimental configurations that can be controlled. Thus, the forward operator $\mathcal{F}$ models the underlying

deterministic physics.

However, in physical applications, observations are always affected by experimental uncertainties. This is accounted for by the additive noise term $\eta(\xi) \sim \mathcal{N}(0, \Gamma)$ with a positive definite covariance matrix $\Gamma \in \mathbb{R}^{N\times N}$, $\Gamma \succ 0$. The noise covariance matrix Γ models both the measurement variances and potential correlations between different measurements.

Ultimately, evaluating the forward model (1.1) yields the observation vector $y \in \mathbb{R}^N$.

The associated inverse problem aims to reverse the process. Given the observed data y under a design ξ, we seek to infer the underlying parameter p. Under this formulation, due to the presence of the noise η, the exact parameter vector p cannot be determined with absolute certainty. Instead, we compute a statistical estimator $\hat{p}$ via a dedicated reconstruction strategy $\mathcal{R}$:

$$\hat{p} = \mathcal{R}(y, \Xi, \Gamma), \tag{1.2}$$

which maps the measurement data y, the underlying design space Ξ, and the noise covariance matrix Γ into the parameter space.

1.2 The Fisher Information Matrix

Based on [HKL26] we introduce the Fisher Information Matrix (FIM). The FIM is a foundational metric in experimental design that characterizes the local sensitivity of the data with respect to the parameter vector p within the neighborhood of a presumed ground-truth p_*. For a finite design space $\Xi = \{\xi_1, \dots, \xi_N\}$, it is defined as:

$$\mathcal{I}(\Xi) = G^\top \Gamma^{-1} G \ \in \mathbb{R}^{K\times K}, \tag{1.3}$$

where $G \in \mathbb{R}^{N\times K}$ denotes the sensitivity matrix containing the compiled gradients of the forward operator with respect to the parameter vector p, evaluated at p_* for each of the design points $\xi_1, \dots, \xi_N$:

$$G = \begin{pmatrix} \nabla_p \mathcal{F}(p_*, \xi_1) \\ \vdots \\ \nabla_p \mathcal{F}(p_*, \xi_N) \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathcal{F}(p_*,\xi_1)}{\partial p_1} & \frac{\partial \mathcal{F}(p_*,\xi_1)}{\partial p_2} & \dots & \frac{\partial \mathcal{F}(p_*,\xi_1)}{\partial p_K} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathcal{F}(p_*,\xi_N)}{\partial p_1} & \frac{\partial \mathcal{F}(p_*,\xi_N)}{\partial p_2} & \dots & \frac{\partial \mathcal{F}(p_*,\xi_N)}{\partial p_K} \end{pmatrix}.$$

By construction, the FIM $\mathcal{I}(\Xi)$ is symmetric and positive semidefinite, since Γ^{-1} is positive definite. This can be shown as follows. For any $x \in \mathbb{R}^K$

$$x^\top \mathcal{I}(\Xi) x = x^\top G^\top \Gamma^{-1} G x = \left\| \Gamma^{-\frac{1}{2}} G x \right\|_2^2 \geq 0,$$

so $I(\Xi)$ is positive semidefinite. Equality holds exactly when $Gx = 0$. Consequently, $I(\Xi)$ is positive definite exactly if $\ker(G) = \{0\}$, meaning G has full column rank.

Intuitively a row i of G measures how strong the data at design point ξ_i reacts to a change of the parameter p around p_*. If the values in the row are large in absolute terms, even a small perturbation in p results in a significant change in the data. If the values in the row are close to zero, a perturbation in p at ξ_i is nearly invisible in the data.

The FIM has a direct optimization-theoretic interpretation. For the nonlinear least-squares functional $\mathcal{L}_\Gamma(p_*) = \frac{1}{2} \|y - \mathcal{F}(p_*, \cdot)\|_\Gamma^2$, one computes

$$\nabla_p^2 \mathcal{L}_\Gamma(p_*) = \left(\nabla_p \mathcal{F}(p_*, \cdot)\right)^\top \Gamma^{-1} \left(\nabla_p \mathcal{F}(p_*, \cdot)\right) - \sum_{i=1}^{N} \left(\Gamma^{-1}\big(y - \mathcal{F}(p_*, \cdot)\big)\right)_i \nabla_p^2 \mathcal{F}(p_*, \xi_i).$$

Since $\mathbb{E}\big[\big(\Gamma^{-1}(y - \mathcal{F}(p_*, \cdot))\big)_i\big] = \big(\Gamma^{-1}\mathbb{E}[\eta]\big)_i = 0$, the expected value is

$$\nabla_p^2 \bar{\mathcal{L}}_\Gamma(p_*) = \mathbb{E}[\nabla_p^2 \mathcal{L}_\Gamma(p_*)] = G^\top \Gamma^{-1} G = \mathcal{I}(\Xi).$$

This means that the FIM coincides with the expected value of the Hessian of the reconstruction loss at the ground truth. Recall that $\bar{\mathcal{L}}_\Gamma$ is called locally strongly convex with parameter $m > 0$ around p_* if $\nabla_p^2 \bar{\mathcal{L}}_\Gamma(p) \succeq m \,\mathrm{Id}$ for all p in a neighborhood of p^*, where Id is the identity. That is, if the minimum eigenvalue is $\lambda_{\min}(\nabla_p^2 \bar{\mathcal{L}}_\Gamma(p_*)) \geq m > 0$. Since $\nabla_p^2 \bar{\mathcal{L}}_\Gamma(p_*) = \mathcal{I}(\Xi)$, this condition is equivalent to $\lambda_{\min}(\mathcal{I}(\Xi)) \geq m$. Therefore, a large minimum eigenvalue of the FIM directly implies strong local convexity of the expected loss landscape around p_* in all parameter directions. Consequently, it ensures high parameter sensitivity on average, meaning that deviations from p_* lead, in expectation, to a rapid increase in loss. Conversely, a small or vanishing $\lambda_{\min}(\mathcal{I}(\Xi))$ indicates a flat direction along which $\bar{\mathcal{L}}_\Gamma$ increases only weakly. Additionally, good conditioning ensures that the expected loss landscape is equally curved in all directions, which prevents numerical issues during reconstruction. Good conditioning is characterized by a condition number $\kappa(\mathcal{I}(\Xi)) = \frac{\lambda_{\max}(\mathcal{I}(\Xi))}{\lambda_{\min}(\mathcal{I}(\Xi))}$ close to one, where $\lambda_{\max}(\mathcal{I}(\Xi))$ denotes the largest eigenvalue of the FIM.

The Cramér-Rao inequality. The spectral properties of the FIM are directly linked to the achievable accuracy of parameter reconstruction. According to the Cramér-Rao inequality, the inverse of the FIM provides a lower bound for the covariance matrix of any unbiased estimator:

$$\mathrm{Cov}(\hat{p}) \succeq \mathcal{I}(\Xi)^{-1}.$$

For the simple case in which p is a single parameter, $\mathcal{I}(\Xi)$ is a number i and the inequality translates to

$$\mathrm{Var}(\hat{p}) \geq \frac{1}{i}.$$

The larger i is, the smaller the minimum variance, which indicates a potentially accurate estimate. If p has $K > 1$ components, then $\mathcal{I}(\Xi)$ is a $K \times K$ matrix, and the same idea that applies to a single number applies to each direction in the parameter space. Decomposing the FIM $\mathcal{I}(\Xi)$ into its eigenvalues λ_j and corresponding eigenvectors v_j for $j = 1, \ldots, K$, it holds

$$v_j^\top \mathrm{Cov}(\hat{p})\, v_j \;\geq\; \frac{1}{\lambda_j},$$

where the left side represents the variance of the estimate along v_j. A large eigenvalue λ_j allows a precise estimation along v_j, while a small eigenvalue results in an unavoidably high degree of uncertainty in that direction, regardless of how well the rest of the reconstruction is performed. Thus, an FIM with large eigenvalues, especially a large smallest eigenvalue, indicates low estimation variance and a high degree of reliability in parameter reconstruction, and is therefore desirable.

1.3 The Down-Sampling Approach

In practical applications the cardinality of the design set Ξ is often very large or even infinite. For simplicity, we assume that Ξ is finite with a large cardinality $|\Xi| = N \gg K$.

Yet, according to [HKL26] the fundamental challenge lies not in the finiteness of $|\Xi|$, but in the necessity of reducing this broad design space to a compressed subset

$$\Xi_c \subset \Xi \quad \text{with} \quad |\Xi_c| = c \ll N = |\Xi|.$$

Restricting the system to a compressed design space Ξ_c yields the subset observations $y_c = y|_{\Xi_c}$, the subsampled forward operator $\mathcal{F}_c = \mathcal{F}|_{\Xi_c}$, and a corresponding reduced covariance matrix $\Gamma_c \succ 0$. Applying the reconstruction strategy to the reduced data defines the estimator

$$\hat{p}_c = \mathcal{R}(y_c, \Xi_c, \Gamma_c). \tag{1.4}$$

Our goal is to construct Ξ_c such that $\hat{p}_c \approx \hat{p}$, i.e. such that $\hat{p}_c$ is a good approximation of the full estimator $\hat{p}$ from (1.2). This is not automatically the case for an arbitrary subset. A poorly chosen Ξ_c may omit exactly those design points at which the data is most sensitive to p, so that $\hat{p}_c$ can deviate substantially from $\hat{p}$. Therefore, the central question of this thesis is how to select Ξ_c so that the approximation holds.

As Hellmuth et al. [HKL26] explain in their study, efficiently handling inverse problems in a large design space Ξ requires a targeted reduction to a compressed design subset Ξ_c, since the full data is often not needed to characterize the parameter p. Primarily, this reduction minimizes both experimental and computational costs. Furthermore, the process is motivated by the overrepresentation of certain information, resulting in statistical redundancy within the full data vector. Crucially, careful selection of informative data points also counteracts data dilution, through which the contributions of highly sensitive sensors could be statistically overshadowed by a large amount of less relevant data points.

The area of experimental design addresses the question of how to collect data to inform mathematical models and make accurate predictions. The FIM, introduced in the previous Section 1.2, is an important quantity in experimental design. It is the core metric used to mathematically formalize the information content of the data. Optimal experimental design (OED) aims to maximize the parameter information gain while minimizing uncertainty under resource constraints. To that end, the objective of OED is to optimize the FIM according to spectral properties of this matrix. Prominent examples include E-, A-, and D-optimal designs, which maximize the minimum eigenvalue of $\mathcal{I}(\Xi)$, minimize the trace of $\mathcal{I}(\Xi)^{-1}$, and maximize the determinant of $\mathcal{I}(\Xi)$, respectively. However, traditional OED methods are often computationally prohibitive for large-scale applications.

To overcome this limitation, this thesis adopts the framework proposed in [HKL26], which aims to identify a sufficient, rather than optimal, subset Ξ_c that preserves the sensitivity properties and information content of the existing FIM

$$\mathcal{I}(\Xi_c) \approx \mathcal{I}(\Xi).$$

The primary advantage of this approach lies in the use of methods from Randomized Numerical Linear Algebra (RNLA). As outlined in [Mah16], RNLA is an interdisciplinary field that introduces non-trivial randomness into algorithmic design to address fundamental matrix problems, such as matrix multiplication, least-squares regression, or low-rank matrix approximation. By substituting exact computations with bounded random approximations, these techniques accelerate large-scale matrix problems while guaranteeing stable parameter reconstruction with a high degree of probability.

2 Stochastic Preliminaries

This chapter collects the foundational probability theory that is used in this thesis. Basic theory that will not be needed later on has been left out. Section 2.1 introduces probability spaces, stochastic processes, filtrations, and Brownian motion. Section 2.2 introduces the concept of a martingale, and Section 2.3 the Azuma-Hoeffding inequality. Together they underlie the proof of the concentration inequality in Theorem 1 of Chapter 3. Section 2.4 introduces stochastic differential equations, which underlie the construction of the sampling algorithm of Chapter 5. We follow the expositions of [Pav14] for the Sections 2.1 and 2.4, and of [Mah16] for the Sections 2.2 and 2.3.

2.1 Stochastic Processes and Brownian Motion

Probability spaces. All randomness in this thesis, such as the measurement noise η, the randomly drawn row or pair indices $i_1, \dots, i_c$, and later the Brownian increments driving the sampler of Chapter 5, is modeled on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Here Ω is the sample space, which is the set of all possible outcomes of the underlying random process. The event space $\mathcal{F}$ is a σ-algebra on Ω, which is the collection of events, where an event is a subset of outcomes in the sample space. The probability measure $\mathbb{P} : \mathcal{F} \to [0, 1]$ assigns a probability to each event in the event space.

A random variable X is a measurable function $X : \Omega \to \mathbb{R}^d$. Its role is to translate an abstract outcome $\omega \in \Omega$ into a numerical quantity.

The basic definitions and results from probability theory are described in [Pav14, Appendix B].

Stochastic processes. Based on the knowledge about probability spaces, we define a stochastic process using the definition from [Pav14, Definition 1.1].

Definition 1 (Stochastic process)**.** Let T be an ordered index set, $(\Omega, \mathcal{F}, \mathbb{P})$ a probability space, and $(E, \mathcal{G})$ a measurable state space. A *stochastic process* is a collection of random variables $X = \{X_t; t \in T\}$ such that for each fixed $t \in T$, X_t is a random

variable from $(\Omega, \mathcal{F}, \mathbb{P})$ to $(E, \mathcal{G})$. The set Ω is known as the sample space, where E is the state space of the stochastic process X_t.

In simple terms, a stochastic process is a collection of random variables, one for each time point t, which together describe the random behavior of a system. Consider, for example, the temperature at a given location over the course of a day, or, as is relevant to this work, the position $u_j(t)$ of an ensemble particle during the sampling iteration in Chapter 5. Throughout this thesis, $T = \mathbb{R}_+$ is a continuous time parameter, and $E = \mathbb{R}^d$ is equipped with the Borel σ-algebra.

For a fixed sample point $\omega \in \Omega$, the function $X_t(\omega) : T \mapsto E$ is called a sample path (or a realization or trajectory) of the process X [Pav14, Chapter 1]. This is the time series that is observed when the random experiment is carried out once.

Filtrations and adapted processes. Some of the constructions below depend on formalizing the intuitive idea that information accumulates over time. This is done with the concept of a filtration. According to [Pav14, Chapter 2], we define a filtration to be a nondecreasing family $\{\mathcal{F}_t, t \in \mathbb{R}_+\}$ of sub-σ-algebras of $\mathcal{F}$:

$$\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F} \quad \text{for } s \leq t.$$

The set $\mathcal{F}_t$ should be thought of as everything that is known at time t.

The filtration generated by the stochastic process X itself is

$$\mathcal{F}_t^X := \sigma(X_s; s \leq t).$$

A stochastic process $(X_t)_{t\geq 0}$ is called adapted to $(\mathcal{F}_t)_{t\geq 0}$ if X_t is $\mathcal{F}_t$-measurable for every $t \in T$ [Pav14, Section 3.2]. This means that the value of X_t depends only on the information available at time t.

Brownian motion. Furthermore, we recall the definition of Brownian motion, since Chapter 5 formulates the Ensemble Kalman Sampler as a stochastic differential equation driven by Brownian motion. The following definition is based on [Pav14, Definition 1.8].

Definition 2 (Brownian motion)**.** A one-dimensional standard *Brownian motion* $W(t) : \mathbb{R}_+ \to \mathbb{R}$ is a real-valued stochastic process with almost surely continuous paths such that $W(0) = 0$, it has independent increments, and for every $t > s \geq 0$, the increment $W(t) - W(s)$ has a Gaussian distribution with mean 0 and variance $t - s$.

A standard d-dimensional standard Brownian motion $W(t) : \mathbb{R}_+ \to \mathbb{R}^d$ is a vector of d independent one-dimensional Brownian motions:

$$W(t) = (W_1(t), \ldots, W_d(t)),$$

where $W_i(t)$, $i = 1, \ldots, d$, are independent one-dimensional Brownian motions.

In other words, Brownian motion is a continuous random path with no memory and whose fluctuations over a time interval of length $t - s$ grow like $\sqrt{t-s}$ in every coordinate, independently of direction. It is the mathematical model of unbiased, uncorrelated noise accumulating over time, and the object denoted by W in the stochastic differential equation in (5.4). Note that Brownian paths are not differentiable.

2.2 Doob Martingales

We now turn to martingales, the key structural property that will let us control the error of the randomized sketching procedure of Chapter 3 even though the c row indices $i_1, \ldots, i_c$ drawn are not independent. They are in fact independent and identically distributed (i.i.d.) in Algorithm 1 itself, but martingales become essential once dependence between draws is allowed, and they provide a uniform proof technique that also carries over unchanged to the pairwise setting of Algorithm 2 in Chapter 4.

Following [Mah16, Definition 4], we first define a martingale with respect to a sequence of random variables.

Definition 3 (Martingale)**.** A sequence of random variables $Z_0, Z_1, \ldots$ is a martingale with respect to a sequence $X_0, X_1, \ldots$ if, for all $t \geq 0$,

- $Z_t = Z_t(X_0, X_1, \ldots, X_t)$, i.e., it is a function of the X_i
- $\mathbb{E}[|Z_t|] < \infty$
- $\mathbb{E}[Z_{t+1} \mid X_0 \cdots X_t] = Z_t$

A sequence is a martingale if it is a martingale with respect to itself.

The last property is the central feature of a martingale. It states that, given everything known up to step t, the expected value of the process at the next step is exactly its current value. This means there is no systematic drift, upward or downward.

The canonical example of a martingale is a gambler who plays a sequence of fair games. This illustration is taken from [Mah16]. Let X_i denote the amount that the gambler

wins in the i-th game and $Z_i = X_1 + \cdots + X_i$ the total winnings of the gambler after the i-th game. Since the game is fair, it holds $\mathbb{E}[X_i] = 0$. Then $\mathbb{E}[Z_{i+1} \mid X_1, \ldots, X_i] = Z_i + \mathbb{E}[X_{i+1}] = Z_i$. So Z_i is a martingale with respect to X_i.

Relevant for our purposes is a way of building a martingale out of an arbitrary target random variable. It is known as the Doob martingale. The following definition is from [Mah16, Definition 5].

Definition 4 (Doob Martingale)**.** Let $X_0, X_1, \ldots, X_n$ be a sequence of random variables, and let Y be a random variable with $\mathbb{E}[|Y|] < \infty$. Generally, Y is such that $Y = Y(X_1, \cdots, X_n)$. The sequence

$$Z_t = \mathbb{E}[Y \mid X_0 \cdots X_t], \quad t = 0, 1, \ldots, n.$$

is called the *Doob martingale*.

The intuition behind the Doob martingale is that Z_t is the best estimator based on the current state of information of the quantity of interest Y, updated as more of the randomness $X_1, \ldots, X_n$ is revealed. We start with no information at all $Z_0 = \mathbb{E}[Y]$, and end with full information $Z_n = Y$, since Y is determined once all of $X_1, \ldots, X_n$ are known. As shown in [Mah16, Lemma 10], a Doob martingale Z_t constructed in this way is a martingale with respect to X_t.

The Doob martingale is what will be used in the proof of Theorem 1 in Chapter 3. There, the quantity of interest is the sketching error $Y = F(i_1, \ldots, i_c) := \|A^\top A - C^\top C\|_F$, which is a function of the c randomly drawn row indices. Its Doob martingale $V_t = \mathbb{E}[F \mid i_1, \ldots, i_t]$ starts at the unconditional expected error $V_0 = \mathbb{E}[F]$, which is bounded in Lemma 4, and progresses step by step until it reaches the actual error $V_c = F(i_1, \ldots, i_c)$ we wish to control. Since the Doob martingale is a martingale, it is guaranteed that $(V_t)_{t=0}^c$ is a martingale. This is exactly the prerequisite needed to apply the concentration inequality in the next section.

2.3 The Azuma-Hoeffding Inequality

A martingale by itself only guarantees the absence of systematic drift. It says nothing about how far Z_n can typically deviate from Z_0. We aim to turn the martingale structure of the previous Section 2.2 into a quantitative, high-probability error bound of the type appearing in Theorem 1. We achieve this using the Azuma-Hoeffding inequality, adapted from [Mah16, Theorem 3].

Lemma 1 (Azuma-Hoeffding inequality)**.** *Let $X_0, \dots, X_n$ be a martingale such that $|X_k - X_{k-1}| \leq c_k$. Then, for all $t \geq 0, \lambda > 0$,*

$$\mathbb{P}[X_t - X_0 \geq \lambda] \leq \exp\left(\frac{-\lambda^2}{2\sum_{k=1}^{t} c_k^2}\right).$$

Two aspects of the Azuma-Hoeffding inequality make it the right tool for our purposes. First, unlike classical concentration results, such as Hoeffding's inequality, it requires no independence between the individual steps of the process. It needs only the martingale property and the bounded differences condition $|X_k - X_{k-1}| \leq c_k$. Second, the bound decays exponentially in λ^2. Concretely, if we fix a target failure probability δ and solve $\exp\left(\frac{-\lambda^2}{2\sum_{k=1}^{n} c_k^2}\right) = \delta$, we obtain $\lambda = \sqrt{2\sum_{k=1}^{n} c_k^2 \log(\delta^{-1})}$. This is the required deviation λ, and hence the resulting error bound, grows only like $\sqrt{\log(\delta^{-1})}$. This is the reason that will let the Theorems 1, 2, and 3 depend through $\sqrt{\log(\delta^{-1})}$ on the inverse failure probability δ^{-1}.

2.4 Stochastic Differential Equations

This section collects the theory on stochastic differential equations (SDEs) needed to make sense of the construction of the Ensemble Kalman Sampler in Chapter 5, including the Itô-SDE, the associated Fokker–Planck equation, and the Euler–Maruyama method for its numerical discretization. Thereby, we follow the explanations of [Pav14].

Itô stochastic differential equations. Based on [Pav14, Chapter 3], we introduce SDEs driven by Gaussian white noise. An SDE with state space $\mathbb{R}^d$ is an equation of the form

$$\mathrm{d}X_t = b(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}W_t, \qquad X_0 = x, \tag{2.1}$$

where $b : \mathbb{R}^d \to \mathbb{R}^d$ is called the drift and $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ the diffusion coefficient. The operator $(W_t)_{t\geq 0}$ denotes a d-dimensional Brownian motion as in Definition 2 in Section 2.1. Since W_t is almost surely nowhere differentiable, (2.1) is only formal and is correctly interpreted as the stochastic integral equation

$$X_t = x + \int_0^t b(X_s)\,\mathrm{d}s + \int_0^t \sigma(X_s)\,\mathrm{d}W_s. \tag{2.2}$$

The second integral is an Itô stochastic integral. It is constructed as an L^2-limit of Riemann sums of $\sigma(X_s)$ evaluated at the left endpoint of each subinterval of a partition

of $[0, t]$:

$$\int_0^t \sigma(X_s)\,\mathrm{d}W_s := \lim_{L\to\infty} \sum_{l=0}^{L-1} \sigma(X_{t_l})(W_{t_{l+1}} - W_{t_l}) \tag{2.3}$$

In contrast to deterministic Riemann integrals, the evaluation point within each subinterval fundamentally alters the result of a stochastic integral. The choice of evaluation point at the left endpoint is essential, because it ensures that the integrand $\sigma(X_{t_l})$ is $\mathcal{F}_t$-measurable, i.e., adapted to the filtration generated by the Brownian motion up to time t_l. This non-anticipating property is what makes the Itô integral a martingale with respect to $(\mathcal{F}_t)_{t\geq 0}$. Consequently, the stochastic noise term does not introduce uncontrolled systematic drift into the system.

An SDE should be read as an ordinary differential equation (ODE) perturbed by noise. The deterministic drift term $b(X_t)\,\mathrm{d}t$ pushes the state X_t in a prescribed direction, exactly as in an ODE, while the stochastic term $\sigma(X_t)\,\mathrm{d}W_t$ perturbs the sample path with Brownian noise. When $\sigma \equiv 0$, (2.1) reduces to an ODE.

The Fokker-Planck equation. Because X_t solving (2.1) is random, it is natural to ask not for the value of a single trajectory but for the probability density ρ_t describing the distribution of X_t at time t. This density evolves according to a deterministic partial differential equation, the Fokker–Planck equation (also called the forward Kolmogorov equation),

$$\partial_t \rho_t = \mathcal{L}^* \rho_t, \qquad \mathcal{L}^*\rho := -\nabla \cdot (b\rho) + \frac{1}{2}\nabla \cdot \nabla \cdot (\Sigma\rho), \qquad \Sigma := \sigma\sigma^\top. \tag{2.4}$$

Basic properties of the Fokker–Planck Equation can be found in [Pav14, Section 4.1]. What matters for Chapter 5 is the notion of an invariant distribution. As explained in [Pav14, Section 2.4] a probability density ρ_s is called invariant for the process X_t if it solves $\mathcal{L}^*\rho_s = 0$, i.e., if it is a stationary solution of the Fokker-Planck equation (2.4). If the law of X_0 is already ρ_s, then the law of X_t remains ρ_s for all $t \geq 0$.

Euler-Maruyama method. Since the SDE (2.1) generally does not have an exact solution, it must be approximated numerically in practice, as explained in [Pav14, Section 5.2]. The basic method for this is the Euler-Maruyama method, which is the stochastic analog of the explicit Euler method for ordinary differential equations. For a time discretization with step size Δt, the method approximates the solution of (2.1) using the recursive rule

$$X_{t+1} = X_t + b(X_t)\,\Delta t + \sigma(X_t)\,\Delta W_t, \qquad t = 0, \ldots, T-1,$$

where $\Delta W_t = W_{t+1} - W_t$ denotes the t-th Brownian increment. As already introduced in Section 2.1, these increments are independent, Gaussian-distributed random variables with mean zero and variance Δt, so that $\Delta W_t = \eta_t \sqrt{\Delta t}$, where $\eta_t \sim \mathcal{N}(0, 1)$ i.i.d. Thus, the scheme can also be written in the form

$$X_{t+1} = X_t + b(X_t)\,\Delta t + \sigma(X_t)\,\eta_t \sqrt{\Delta t}.$$

Given a known X_t, drawing η_t and evaluating the drift b and diffusion σ yields the approximation X_{t+1}. This scheme forms the basis for the discretization (5.20) of the Ensemble Kalman Sampler in Section 5.4.

3 Matrix Sketching in Randomized Numerical Linear Algebra

To achieve information preservation, as described in Section 1.3, we deploy the theoretical frameworks of Randomized Numerical Linear Algebra (RNLA). The following part is based on [HKL26] and [Mah16].

First, we will clarify the concept of matrix sketching outlined in [Mah16]. Matrix sketching is a key technique in RNLA. In this process, a large data matrix is compressed into a significantly smaller version, called a *sketch*. The main goal is to preserve the essential structural properties of the original matrix while reducing its dimensions. Instead of performing computationally intensive operations on the complete, potentially massive dataset, these operations are performed on the sketch. This provides a good approximation of the result, bypassing the traditional memory and processing bottlenecks.

From an algorithmic perspective, there are two main techniques to construct a sketch in complementary ways: random projection and randomized sampling. Random projections perform a type of random rotation that transforms the data into a new basis in which the information is uniformly distributed. This helps to prevent problems such as axis alignment, where important information is concentrated in just a few coordinates. As a result, it is sufficient to sample uniformly or randomly. This means that random projections are data-agnostic. In other words, they can be constructed without knowing the actual data. Random sampling, on the other hand, generally depends on the data. The algorithm must analyze the structure of the data in order to create an importance distribution for the sampling. This distribution is used to carefully extract a small, highly informative subset of the original rows, columns, or elements in order to preserve the structural properties of the matrix.

In this work, we focus exclusively on the randomized sampling approach, as it directly aligns with the selection of design points in experimental design.

We provide a brief overview of this chapter. Section 3.1 recasts the matrix-matrix

product as a sum of rank-one matrices, laying the groundwork for a Monte Carlo interpretation that sketching exploits. Next, Section 3.2 introduces the randomized row-sampling scheme and the resulting sketching algorithm. Section 3.3 then establishes the theoretical guarantees on the approximation error, leading to a high-probability error bound for the sketched matrix product.

3.1 Matrix Multiplication

We consider the matrix-matrix product $B := A^\top A \in \mathbb{R}^{K\times K}$ where $A \in \mathbb{R}^{N\times K}$.

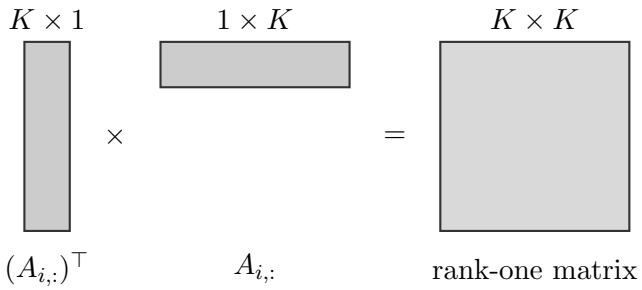
Conventionally, this product is regarded as a collection of row-column inner products:

$$B_{ik} = \left(A^\top A\right)_{ik} = \sum_{j=1}^{N} (A^\top)_{ij} A_{jk} = (A^\top)_{i,:} A_{:,k}.$$

However, to leverage the concepts of matrix sketching, it is helpful to think of matrix multiplication as a sum of independent outer products, which structurally corresponds to a sum of rank-one matrices:

$$B = A^\top A = \sum_{i=1}^{N} (A_{i,:})^\top A_{i,:},$$

where $A_{i,:} \in \mathbb{R}^{1\times K}$ is the i-th row of A. The figure below illustrates the construction of a summand formed by taking the outer product of the transposed row vector $(A_{i,:})^\top \in \mathbb{R}^{K\times 1}$ and the row vector $A_{i,:} \in \mathbb{R}^{1\times K}$, yielding a $K \times K$ rank-one matrix.



In particular, when A is a tall and narrow matrix, characterized by having significantly more rows than columns $N \gg K$, this formulation highlights a severe computational inefficiency in calculating the matrix product. A computational effort of $\mathcal{O}(NK^2)$ is spent to evaluate the sum over the entire number of rows N. The result is a matrix $B \in \mathbb{R}^{K\times K}$ that needs significantly fewer memory space.

3.2 Matrix Sketching

To reduce this computational load, we approximate B by executing randomized row-sampling. We define a matrix-valued random variable X that represents drawing rows of A based on a discrete probability distribution π on $\{1, \ldots, N\}$:

$$X_j = \frac{1}{\pi(i_j)} A_{i_j,:}^\top A_{i_j,:}, \tag{3.1}$$

where a row index $i_j \in \{1, \ldots, N\}$ is drawn with probability $\pi(i_j)$. We assume that $\pi(i) > 0$ for all i, so that (3.1) is well-defined. By drawing $c \ll N$ i.i.d. samples, the matrix product is approximated by the sample average

$$B \approx \widetilde{B} = \frac{1}{c} \sum_{j=1}^{c} X_j. \tag{3.2}$$

In the following Lemmas 2 and 3, we analyze the expectation and the component-wise variance of the estimator $\widetilde{B}$. The two Lemmas, as well as their respective proofs, are similar to [Mah16, Lemma 3].

Lemma 2 (Unbiasedness of the estimator)**.** *The randomized sketching estimator $\widetilde{B}$ is an unbiased estimator of B, satisfying:*

$$\mathbb{E}\left[\widetilde{B}\right] = B.$$

Proof. Using the linearity of the expected value $\mathbb{E}[\cdot]$ and taking into account the fact that all X_j are i.i.d. realizations of X_1, we obtain:

$$\mathbb{E}\left[\widetilde{B}\right] = \mathbb{E}\left[\frac{1}{c}\sum_{j=1}^{c} X_j\right] = \frac{1}{c}\sum_{j=1}^{c} \mathbb{E}\left[X_j\right] = \frac{1}{c} \cdot c \cdot \mathbb{E}\left[X_1\right] = \mathbb{E}\left[X_1\right].$$

Evaluating the expected value of the matrix-valued random variable X_1 by definition yields:

$$\mathbb{E}\left[X_1\right] = \sum_{i=1}^{N} \pi(i) \left(\frac{1}{\pi(i)} A_{i,:}^\top A_{i,:}\right) = \sum_{i=1}^{N} A_{i,:}^\top A_{i,:} = B. \tag{3.3}$$

Hence, $\mathbb{E}\left[\widetilde{B}\right] = B$. □

In this context, the outer product formulation is very advantageous, as defining X_1 as a rank-one matrix allows the expected value to be computed globally for the entire matrix structure at once, eliminating the need for element-wise indexing.

Lemma 3 (Component-wise variance of the estimator)**.** *For any element (g, h) of the sketched matrix $\widetilde{B}$, the variance of the estimator scales inversely with the sample size c and is explicitly given by:*

$$Var\left[\widetilde{B}_{g,h}\right] = \frac{1}{c}\sum_{i=1}^{N}\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)} - \frac{1}{c}B_{g,h}^2. \tag{3.4}$$

Proof. Since the sketched matrix $\widetilde{B}$ is computed as the sample average of i.i.d. random matrices X_j, the variance of any scalar component (g, h) satisfies:

$$\mathrm{Var}\left[\widetilde{B}_{g,h}\right] = \mathrm{Var}\left[\frac{1}{c}\sum_{j=1}^{c}(X_j)_{g,h}\right] = \frac{1}{c^2}\sum_{j=1}^{c}\mathrm{Var}\left[(X_1)_{g,h}\right] = \frac{1}{c}\mathrm{Var}\left[(X_1)_{g,h}\right].$$

Applying the identity $\mathrm{Var}[Z] = \mathbb{E}[Z^2] - (\mathbb{E}[Z])^2$ to the scalar component $(X_1)_{g,h}$ yields:

$$\mathrm{Var}\left[(X_1)_{g,h}\right] = \mathbb{E}\left[(X_1)_{g,h}^2\right] - \left(\mathbb{E}\left[(X_1)_{g,h}\right]\right)^2.$$

From Lemma 2 (3.3), we already know that $\mathbb{E}\left[(X_1)_{g,h}\right] = B_{g,h}$. Next, we evaluate the second moment of the component $(X_1)_{g,h}$ by using its discrete probability distribution:

$$\mathbb{E}\left[(X_1)_{g,h}^2\right] = \sum_{i=1}^{N}\pi(i)\left(\frac{1}{\pi(i)}A_{i,g}A_{i,h}\right)^2 = \sum_{i=1}^{N}\pi(i)\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)^2} = \sum_{i=1}^{N}\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)}.$$

Combining these results in the variance expression yields:

$$\mathrm{Var}\left[(X_1)_{g,h}\right] = \sum_{i=1}^{N}\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)} - B_{g,h}^2.$$

By multiplying the factor $\frac{1}{c}$, we obtain the component-wise variance of the estimator $\widetilde{B}$:

$$\mathrm{Var}\left[\widetilde{B}_{g,h}\right] = \frac{1}{c}\sum_{i=1}^{N}\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)} - \frac{1}{c}B_{g,h}^2.$$

□

We observe that the variance scales inversely with the sample size c. This confirms that drawing a larger number of samples enhances the approximation accuracy of the matrix estimator. Consequently, the standard error decreases at a rate of $\frac{1}{\sqrt{c}}$.

Furthermore, we consider the sum $\sum_{i=1}^{N}\frac{A_{i,g}^2 A_{i,h}^2}{\pi(i)}$ within the variance term. Since all terms in this sum are strictly positive or zero for all N rows, they cannot cancel each other out. This carries the risk that the variance will increase significantly if, even for

a single row i, the matrix entries $A_{i,g}$ and $A_{i,h}$ are large in terms of amount and the row is assigned a small sampling probability close to zero. To prevent this, we will take a closer look at how to choose the sampling probabilities $\pi(i)$ in Lemma 5.

To execute the sketching, we formulate the procedure in Algorithm 1 taken from [HKL26]. In the algorithm we replace the sum-notation (3.2) by an equivalent, compact sketched matrix multiplication $C^\top C$.

Algorithm 1: BasicMatrixMultiplication

Input: Matrix $A \in \mathbb{R}^{N\times K}$, sample size $c \ll N$, probability measure π on $\{1,\dots,N\}$

Output: Matrix $C \in \mathbb{R}^{c\times K}$ such that $C^\top C \approx A^\top A$

1 **for** $j = 1,\dots,c$ **do**
2 Sample $i_j \sim \pi$ i.i.d.;
3 Set the j-th row of C as $C_{j,:} = \frac{A_{i_j,:}}{\sqrt{c\pi(i_j)}}$.
4 **return** Matrix C and $C^\top C$

Since we sum up significantly fewer rows during the sketching process, the drawn rows must be reweighted by the factor $\frac{1}{\sqrt{c\pi(i_j)}}$ so that the final result is not systematically underestimated. This scaling ensures that the estimator is unbiased.

The equivalence between the matrix multiplication $C^\top C$ and the sample average is verified by analyzing the outer product structure of the sketched matrix:

$$\begin{aligned} C^\top C &= \sum_{j=1}^{c} C_{j,:}^\top C_{j,:} \\ &= \sum_{j=1}^{c} \left(\frac{A_{i_j,:}}{\sqrt{c\pi(i_j)}}\right)^\top \left(\frac{A_{i_j,:}}{\sqrt{c\pi(i_j)}}\right) \\ &= \sum_{j=1}^{c} \frac{1}{c\pi(i_j)} A_{i_j,:}^\top A_{i_j,:} \qquad (3.5) \\ &\overset{(3.1)}{=} \frac{1}{c}\sum_{j=1}^{c} X_j \\ &\approx A^\top A. \end{aligned}$$

3.3 Theoretical Guarantees

To clarify the meaning of the heuristic approximation symbol, we quantify the expected discrepancy between the true matrix product $A^\top A$ and its sketched counterpart $C^\top C$

in terms of the Frobenius norm. The following Lemma 4 and its proof are based on [Mah16, Lemma 4].

Lemma 4 (Expected approximation error)**.** *Let a matrix C be constructed from a matrix A by Algorithm 1. Then,*

$$\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] = \frac{1}{c}\sum_{i=1}^{N}\frac{\|A_{i,:}\|_2^4}{\pi(i)} - \frac{1}{c}\|A^\top A\|_F^2. \tag{3.6}$$

Proof. Using the definition of the Frobenius norm

$$\|A\|_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} A_{ij}^2},$$

the expected approximation error can be decomposed component-wise into the scalar entries of the respective $K \times K$ matrices:

$$\begin{aligned}\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] &= \mathbb{E}\left[\sum_{k=1}^{K}\sum_{l=1}^{K}\left((A^\top A) - (C^\top C)\right)_{k,l}^2\right] \\ &= \sum_{k=1}^{K}\sum_{l=1}^{K}\mathbb{E}\left[\left((A^\top A)_{k,l} - (C^\top C)_{k,l}\right)^2\right].\end{aligned}$$

Since the sketching estimator $C^\top C$ is unbiased according to Lemma 2, the entry $(A^\top A)_{k,l}$ represents the expected value of $(C^\top C)_{k,l}$. Using the identity $\mathrm{Var}(Z) = \mathbb{E}[(Z - \mathbb{E}[Z])^2]$ with $Z = (C^\top C)_{k,l}$ yields,

$$\mathrm{Var}((C^\top C)_{k,l}) = \mathbb{E}[((C^\top C)_{k,l} - \mathbb{E}[(C^\top C)_{k,l}])^2] = \mathbb{E}[((C^\top C)_{k,l} - (A^\top A)_{k,l})^2].$$

Since $(x-y)^2 = (y-x)^2$ it follows

$$\mathbb{E}\left[\left((A^\top A)_{k,l} - (C^\top C)_{k,l}\right)^2\right] = \mathrm{Var}\left((C^\top C)_{k,l}\right).$$

Substituting the component-wise variance from Lemma 3 (3.4)

$$\mathrm{Var}\left((C^\top C)_{k,l}\right) = \frac{1}{c}\sum_{i=1}^{N}\frac{A_{i,k}^2 A_{i,l}^2}{\pi(i)} - \frac{1}{c}(A^\top A)_{k,l}^2$$

back into the double summation allows us to pull out the row-dependent terms:

$$\begin{aligned}\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] &= \sum_{k=1}^{K}\sum_{l=1}^{K}\left(\frac{1}{c}\sum_{i=1}^{N}\frac{A_{i,k}^2 A_{i,l}^2}{\pi(i)} - \frac{1}{c}(A^\top A)_{k,l}^2\right) \\ &= \frac{1}{c}\sum_{i=1}^{N}\frac{1}{\pi(i)}\left(\sum_{k=1}^{K}A_{i,k}^2\right)\left(\sum_{l=1}^{K}A_{i,l}^2\right) - \frac{1}{c}\sum_{k=1}^{K}\sum_{l=1}^{K}(A^\top A)_{k,l}^2.\end{aligned}$$

Recognizing that $\sum_{k=1}^{K} A_{i,k}^2 = \sum_{l=1}^{K} A_{i,l}^2 = \|A_{i,:}\|_2^2$ is the squared Euclidean norm of the i-th row of A, and that the final term defines the squared Frobenius norm $\|A^\top A\|_F^2$, the total expected error condenses to

$$\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] = \frac{1}{c}\sum_{i=1}^{N}\frac{\|A_{i,:}\|_2^4}{\pi(i)} - \frac{1}{c}\|A^\top A\|_F^2.$$

□

As one might expect from our previous observation regarding the variance in Lemma 3, we observe that the expected approximation error depends on the probability distribution π we use.

A baseline approach proposed in [Mah16] is uniform sampling, where every row is assigned the same probability, i.e., $\pi_{\text{uni}}(i) = \frac{1}{N}$ for all $i \in \{1, \ldots, N\}$. Since it requires no prior inspection of the matrix A, uniform sampling is very fast and trivial to implement. According to (3.6), the expected squared approximation error decays like $\frac{1}{c}$, so the error itself decays like $\frac{1}{\sqrt{c}}$. Hence, an accurate approximation can be achieved if the sample size c is chosen large enough. The downside is that uniform sampling may require significantly more samples than sampling with a distribution that takes the structural importance of each row of A into account. Especially if a few rows carry most of the information, they are likely to be overlooked in a uniform sampling approach.

Consequently, the choice of π involves a trade-off between two costs. The cost of drawing more samples and the cost of constructing and evaluating a more sophisticated distribution, which requires knowledge of the row norms of A. An informed sampling distribution lowers the expected approximation error for a given c, and thus allows a smaller c for the same accuracy. We therefore look for the distribution that minimizes the expected approximation error in (3.6). The following Lemma 5 adapted from [Mah16, Lemma 5] identifies it.

Lemma 5 (Optimal sampling distribution)**.** *The optimal sampling distribution mini-*

mizing $\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right]$ *is given by:*

$$\tilde{\pi}(i) = \frac{\|A_{i,:}\|_2^2}{\|A\|_F^2}.$$

Proof. To determine the optimal sampling distribution $\tilde{\pi}$, we minimize the objective function (3.6) under the constraint that π forms a valid probability measure, i.e., $\sum_{i=1}^N \pi(i) = 1$. Since the second term $\frac{1}{c}\|A^\top A\|_F^2$ is a constant independent of π, the error depends solely on the function

$$f(\pi(1), \ldots, \pi(N)) = \sum_{i=1}^N \frac{\|A_{i,:}\|_2^4}{\pi(i)}.$$

We formulate the optimization problem via the Lagrange multiplier λ:

$$\mathcal{L}(\pi, \lambda) = \sum_{i=1}^N \frac{\|A_{i,:}\|_2^4}{\pi(i)} + \lambda \left(\sum_{i=1}^N \pi(i) - 1 \right).$$

Taking the partial derivative with respect to a single probability entry $\pi(i)$ and setting it to zero yields the optimality condition:

$$\frac{\partial \mathcal{L}}{\partial \pi(i)} = -\frac{\|A_{i,:}\|_2^4}{\pi(i)^2} + \lambda = 0 \implies \pi(i) = \frac{\|A_{i,:}\|_2^2}{\sqrt{\lambda}}.$$

To eliminate the unknown multiplier λ, we use the constraint $\sum_{i=1}^N \pi(i) = 1$:

$$\sum_{i=1}^N \pi(i) = \sum_{i=1}^N \frac{\|A_{i,:}\|_2^2}{\sqrt{\lambda}} = 1 \implies \sqrt{\lambda} = \sum_{k=1}^N \|A_{k,:}\|_2^2 = \|A\|_F^2.$$

Substituting $\sqrt{\lambda}$ back into the expression for $\pi(i)$ results in the optimal sampling distribution:

$$\tilde{\pi}(i) = \frac{\|A_{i,:}\|_2^2}{\|A\|_F^2}.$$

To verify that this probability is in fact a minimum of the objective function, we evaluate the second partial derivative. Since $\pi(i) > 0$ it follows

$$\frac{\partial^2 \mathcal{L}}{\partial \pi(i)^2} = 2\frac{\|A_{i,:}\|_2^4}{\pi(i)^3} > 0.$$

This confirms the strict convexity of the objective function, ensuring that $\tilde{\pi}(i)$ is the unique global minimizer. □

The following theorem taken from [Mah16, Theorem 7] guarantees that the approx-

imation error of the sketched matrix product can be bounded with high probability, given that the sampling distribution does not deviate too far from the optimal sampling distribution $\tilde{\pi}$.

Theorem 1. *Let $A \in \mathbb{R}^{N\times K}$ with $N \gg K$. Fix a small, finite number $K \leq c \ll N$ and consider a probability density π on the row space $\{1,\dots,N\}$, for which there exists a $\beta \in (0,1]$ with*

$$\pi(i) \geq \beta \frac{\|A_{i,:}\|_2^2}{\|A\|_F^2}, \tag{3.7}$$

and let the matrix $C \in \mathbb{R}^{c\times K}$ be constructed by Algorithm 1. Then for any failure probability rate $\delta \in (0,1)$, one has the error control

$$\mathbb{P}\left(\left\| A^\top A - C^\top C \right\|_F \leq \frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}} \|A\|_F^2 \right) \geq 1-\delta,$$

where the probability $\mathbb{P}$ is taken over all realizations of C.

The Theorem allows π to deviate from the optimal distribution. The parameter β serves as a quality indicator of the sampling distribution. If $\beta = 1$, the distribution is optimal and the error reaches its minimum. If $\beta < 1$, it means that the distribution π shifts away from the optimal weights. For $\beta \to 0$ the error bound diverges at the rate β^{-1}.

A key point is that the dimension N of the design space does not appear explicitly in the error bound. Relative to $\|A\|_F^2$, the required sample size c depends only on the desired accuracy and the confidence level and not on the size of the design space from which the rows are drawn. The error decreases by a factor of $\frac{1}{\sqrt{c}}$. This means that if we want to double the accuracy, we need to take four times as many samples.

The inequality inside the probability statement specifies the accuracy of the sketched matrix product $C^\top C$ compared to the true, deterministic product $A^\top A$. The parameter δ indicates the confidence level. Since the sampling is a random process, there is always a small probability δ that the error bound will not be met. The theorem bounds the probability of a failure by δ, meaning the algorithm succeeds with a probability of at least $1-\delta$. Notice that δ enters the error bound through $\sqrt{\log(\delta^{-1})}$. This implies that we can demand a much higher level of confidence without the error bound increasing significantly.

Proof. To establish the desired probability guarantee, the proof proceeds in three steps. First, we derive an upper bound for the expected squared approximation error. Second, we show that substituting a single sampled index i_k alters the total error by at most a

deterministic threshold Δ. Third, we construct a Doob martingale sequence, enabling the application of the Azuma-Hoeffding inequality to achieve high-probability error control. Doob martingales and the Azuma-Hoeffding inequality were introduced in Section 2.

In the first step, by Lemma 4, the expected squared Frobenius norm of the matrix approximation error is given by

$$\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] = \frac{1}{c}\sum_{i=1}^{N}\frac{\|A_{i,:}\|_2^4}{\pi(i)} - \frac{1}{c}\|A^\top A\|_F^2.$$

Incorporating the nearly-optimal sampling probabilities (3.7) the expected approximation error satisfies

$$\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] \leq \frac{1}{c}\sum_{i=1}^{N}\frac{\|A_{i,:}\|_2^4}{\beta\frac{\|A_{i,:}\|_2^2}{\|A\|_F^2}} = \frac{\|A\|_F^2}{\beta c}\sum_{i=1}^{N}\|A_{i,:}\|_2^2.$$

Noting that $\|A\|_F^2 = \sum_{i=1}^{N}\|A_{i,:}\|_2^2$, this yields

$$\mathbb{E}\left[\|A^\top A - C^\top C\|_F^2\right] \leq \frac{1}{\beta c}\|A\|_F^4 \tag{3.8}$$

Our objective is to define a high-probability event E, in which the approximation error is bounded by a scaled threshold:

$$\|A^\top A - C^\top C\|_F \leq \frac{\eta}{\sqrt{\beta c}}\|A\|_F^2$$

with a corresponding confidence level of

$$\mathbb{P}(E) \geq 1 - \delta.$$

We define the matrix discrepancy as a scalar error function

$$F(i_1, \ldots, i_c) = \|A^\top A - C^\top C\|_F,$$

which depends on the sequence of drawn row indices.

To apply the Azuma-Hoeffding inequality we need to demonstrate that changing one drawn index i_k at a time the difference in the total error F is bounded, i.e.,

$$|F(i_1, \ldots, i_k, \ldots, i_c) - F(i_1, \ldots, i_k', \ldots, i_c)| \leq \Delta.$$

To this end, let us assume two sketched configurations, C and C', that differ in exactly one coordinate k such that $i_j = i'_j$ for all $j \neq k$.

To prove that the difference between the matrix products $C^\top C$ and $(C')^\top C'$ is bounded, we exploit its structure shown in (3.5):

$$C^\top C = \sum_{j=1}^{c} \frac{1}{c\pi(i_j)} A_{i_j,:}^\top A_{i_j,:}.$$

Subtraction cancels out all components except the k-th component. Subsequently applying the triangle inequality $\|X - Y\|_F \leq \|X\|_F + \|Y\|_F$ yields

$$\|C^\top C - (C')^\top C'\|_F = \left\| \frac{A_{i_k,:}^\top A_{i_k,:}}{c\pi(i_k)} - \frac{A_{i'_k,:}^\top A_{i'_k,:}}{c\pi(i'_k)} \right\|_F \leq \frac{1}{c\pi(i_k)} \left\| A_{i_k,:}^\top A_{i_k,:} \right\|_F + \frac{1}{c\pi(i'_k)} \left\| A_{i'_k,:}^\top A_{i'_k,:} \right\|_F.$$

Recalling that the Frobenius norm of a symmetric outer product equals its squared Euclidean vector norm, $\|u^\top u\|_F = \|u\|_2^2$, we simplify the expression to

$$\|C^\top C - (C')^\top C'\|_F \leq \frac{1}{c} \frac{\|A_{i_k,:}\|_2^2}{\pi(i_k)} + \frac{1}{c} \frac{\|A_{i'_k,:}\|_2^2}{\pi(i'_k)}.$$

To obtain a uniform bound that holds independently of the specific random realization, we replace both terms with the maximum over all possible row indices

$$\|C^\top C - (C')^\top C'\|_F \leq \frac{1}{c} \max_\alpha \frac{\|A_{\alpha,:}\|_2^2}{\pi(\alpha)} + \frac{1}{c} \max_\alpha \frac{\|A_{\alpha,:}\|_2^2}{\pi(\alpha)} = \frac{2}{c} \max_\alpha \frac{\|A_{\alpha,:}\|_2^2}{\pi(\alpha)}$$

By further leveraging the constraint (3.7) on the sampling probabilities $\pi(\alpha)$, this uniform bound is upper-bounded by a constant scaling factor:

$$\|C^\top C - (C')^\top C'\|_F \leq \frac{2}{\beta c} \|A\|_F^2.$$

Using the reverse triangle inequality $|\|X\|_F - \|Y\|_F| \leq \|X - Y\|_F$, we link this result back to the error function F:

$$\begin{aligned} |F(i_1, \ldots, i_k, \ldots, i_c) - F(i_1, \ldots, i'_k, \ldots, i_c)| &= \left| \|A^\top A - C^\top C\|_F - \|A^\top A - (C')^\top C'\|_F \right| \\ &\leq \|C^\top C - (C')^\top C'\|_F \leq \frac{2}{\beta c} \|A\|_F^2. \end{aligned}$$

This successfully verifies the bounded differences condition with the maximum impact

of a single row substitution defined as $\Delta := \frac{2}{\beta c}\|A\|_F^2$.

In the third and final step, we translate this deterministic bound into probabilistic error control by constructing a Doob martingale sequence associated with F. Following the Doob martingale construction of Definition 4, let $(V_t)_{t=0}^c$ with $V_t := \mathbb{E}[F|\mathcal{F}_t]$ denote the Doob martingale with respect to the filtration $\mathcal{F}_t = \sigma(i_1, \ldots, i_t)$ generated by the sequentially revealed row indices. That $\mathbb{E}[F] < \infty$ follows from the bound established in Lemma 4. By construction, this sequence originates at the unconditional expected value, $V_0 = \mathbb{E}[F \mid \mathcal{F}_0] = \mathbb{E}[F]$ with $\mathcal{F}_0 = \{\emptyset, \Omega\}$ being the trivial σ-algebra, where no indices are yet revealed, and terminates at the actual realization of the sketching error, $V_c = \mathbb{E}[F \mid \mathcal{F}_c] = F(i_1, \ldots, i_c)$.

Because F satisfies the bounded differences condition, the increments of this martingale are almost surely bounded by $|V_t - V_{t-1}| \leq \Delta$ for all $t \in \{1, \ldots, c\}$. This property allows us to invoke the Azuma-Hoeffding inequality. For any positive real parameter $\gamma > 0$, the tail probability satisfies:

$$\mathbb{P}\left(V_c - V_0 \geq \gamma\right) \leq \exp\left(-\frac{\gamma^2}{2\sum_{t=1}^c \Delta^2}\right) = \exp\left(-\frac{\gamma^2}{2c\Delta^2}\right).$$

By setting this upper bound equal to the targeted failure probability δ, we can explicitly solve for γ:

$$\exp\left(-\frac{\gamma^2}{2c\Delta^2}\right) = \delta \implies \gamma = \sqrt{2c\log(\delta^{-1})}\Delta.$$

Substituting the explicit expression for Δ from step two yields:

$$\gamma = \sqrt{2c\log(\delta^{-1})}\left(\frac{2}{\beta c}\|A\|_F^2\right) = \frac{\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}}\|A\|_F^2.$$

Expressing this relation in terms of the complementary success probability, we obtain

$$\mathbb{P}\left(V_c - V_0 \leq \gamma\right) \geq 1 - \delta.$$

To bound the initial condition V_0, we apply Jensen's inequality to the convex function $g(x) = x^2$ along with the expectation bound established in (3.8) in the first step:

$$V_0 = \mathbb{E}\left[F\right] \leq \sqrt{\mathbb{E}\left[F^2\right]} \leq \frac{1}{\sqrt{\beta c}}\|A\|_F^2.$$

Finally, we combine the bound for V_0 and γ such that

$$V_0 + \gamma \;\leq\; \frac{1 + \sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}}\,\|A\|_F^2,$$

and therefore $\{V_c \leq V_0+\gamma\} \subseteq \left\{\|A^\top A - C^\top C\|_F \leq \frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}}\|A\|_F^2\right\}$. This yields,

$$\mathbb{P}\left(\|A^\top A - C^\top C\|_F \leq \frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}}\|A\|_F^2\right) \geq \mathbb{P}\left(V_c \leq V_0+\gamma\right) \geq 1-\delta.$$

□

4 Translation to Experimental Design

The aim of this chapter is to apply the abstract theory of matrix sketching developed in Chapter 3 to the scenario of selecting an informative, compressed set of experimental designs $\Xi_c \subset \Xi$ from a large design space Ξ. Section 4.1 establishes the connection between the sensitivity matrix G and the sketching framework of Chapter 3 under the simplifying assumption of uncorrelated noise. Section 4.2 extends the sketching theory to the weighted setting required for general, correlated noise covariances. Section 4.3 applies this weighted framework to construct a concrete down-sampling strategy for the FIM, and Section 4.4 translates the resulting approximation guarantee into a statement about the well-posedness and spectral stability of the down-sampled inverse problem.

4.1 Sketching Setup

Recall from Chapter 3 that Algorithm 1 approximates a matrix product $B = A^\top A \in \mathbb{R}^{K\times K}$, for a tall matrix $A \in \mathbb{R}^{N\times K}$, by randomly sampling and reweighting its rows. To apply this result to our design problem, we establish the relationship between the sensitivity matrix G of the forward model and the matrix A in the RNLA setting. Assuming for the moment that the measurement noise is uncorrelated, meaning that the noise covariance matrix Γ is diagonal, we may define the noise-weighted sensitivity matrix

$$A := \Gamma^{-\frac{1}{2}} G \in \mathbb{R}^{N\times K}.$$

Since Γ is diagonal and positive definite, $\Gamma^{-\frac{1}{2}}$ is well-defined and symmetric. Thus, the matrix product $B = A^\top A$ reproduces the full FIM introduced in (1.3):

$$B = A^\top A = \left(\Gamma^{-\frac{1}{2}} G\right)^\top \left(\Gamma^{-\frac{1}{2}} G\right) = G^\top \Gamma^{-1} G = \mathcal{I}(\Xi).$$

Under this formulation, each of the N rows $A_{i,:} = \Gamma_{ii}^{-\frac{1}{2}} G_{i,:}$ represents the noise-weighted local sensitivity of a single, potential design point $\xi_i \in \Xi$ with respect to the unknown parameter vector p.

Consequently, applying Algorithm 1 to A amounts to randomly selecting c individual

design points $\xi_i \in \Xi$ according to a probability distribution π on $\{1, \dots, N\}$, and reweighting the corresponding rows of A by $\frac{1}{\sqrt{c\pi(i)}}$. Writing $A_c \in \mathbb{R}^{c\times K}$ for the resulting sketch, the sketched matrix product reproduces the compressed FIM on the subsampled design set $\Xi_c = \{\xi_{i_1}, \dots, \xi_{i_c}\}$:

$$A_c^\top A_c = \tilde{\mathcal{I}}(\Xi_c).$$

Optimal sampling distribution. By Lemma 5, the sampling distribution that minimizes the expected Frobenius approximation error is proportional to the squared row norm of A. Translated into the notation of the design problem, this is

$$\tilde{\pi}(i) = \frac{\left\|(\Gamma^{-1/2}G)_{i,:}\right\|_2^2}{\|\Gamma^{-1/2}G\|_F^2}, \qquad i = 1, \dots, N. \tag{4.1}$$

This probability admits a direct physical interpretation. In Lemma 3 we observed that the variance of the sketching estimator increases whenever a row with large entries is assigned a small sampling probability. Translated to experimental design, a row $A_{i,:}$ with a large norm corresponds to a design point ξ_i at which the forward operator $\mathcal{F}$ is highly sensitive to a perturbation of the parameter p and is relatively free of measurement noise. This means this sensor location carries a significant amount of information. Formula (4.1) therefore suggests placing sensors where the model is most sensitive to the unknown parameter.

According to Theorem 1, this sampling strategy comes with the following guarantee for the approximation error. If π is chosen β-close to the optimal distribution $\tilde{\pi}$ in the sense of (3.7), then, with probability at least $1 - \delta$,

$$\left\|\mathcal{I}(\Xi) - \tilde{\mathcal{I}}(\Xi_c)\right\|_F \le \frac{1 + \sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}} \left\|\Gamma^{-\frac{1}{2}}G\right\|_F^2.$$

Hence, a small number $c \ll N$ of well-chosen sensors yields, with high probability, a compressed FIM $\tilde{\mathcal{I}}(\Xi_c)$ close to the full-data FIM $\mathcal{I}(\Xi)$.

Limitation. The construction above relies on the assumption that Γ is diagonal, since only then the sketching problem reduces to independent draws of single design points ξ_i. In realistic experimental settings, however, measurements are often correlated, and Γ, and hence Γ^{-1}, are in general non-diagonal matrices. In this case, it is not enough to consider the rows of A in isolation, as crucial parameter information is contained in the interactions between two points. To carry the sketching idea over to this more general and more realistic setting, we must therefore first extend the abstract matrix sketching

theory of Chapter 3 from the unweighted product $A^\top A$ to the weighted product $A^\top W A$ for a general symmetric positive definite matrix W. This extension is developed in the following section.

4.2 Weighted Matrix Sketching

In realistic experimental scenarios, measurements are often correlated with one another. Following [HKL26], we express this by a symmetric, positive definite, generally non-diagonal covariance matrix $W \in \mathbb{R}^{N\times N}$. Since W is not diagonal, it is no longer sufficient to consider the rows of A in isolation, like in Section 3.2. Now, crucial parameter information is also contained in the interaction between two distinct rows k and l. In contrast to the independent setting of Algorithm 1, where we consider $k = l$, we therefore sum over all combinations of rows (k, l) in order to adapt the sketching matrix to this correlated setting:

$$A^\top W A = \sum_{k,l=1}^{N} A_{k,:}^\top W_{k,l} A_{l,:} = \sum_{k,l=1}^{N} \frac{1}{2}\left(A_{k,:}^\top W_{k,l} A_{l,:} + A_{l,:}^\top W_{l,k} A_{k,:}\right),$$

where we define

$$Y_{k,l} := \frac{1}{2}\left(A_{k,:}^\top W_{k,l} A_{l,:} + A_{l,:}^\top W_{l,k} A_{k,:}\right) \in \mathbb{R}^{K\times K}. \tag{4.2}$$

The matrix $Y_{k,l}$ quantifies the joint contribution of the pair of rows $(A_{k,:}, A_{l,:})$ to the total matrix product for every index pair (k, l). Because W is symmetric, i.e. $W_{k,l} = W_{l,k}$, it is

$$\left(A_{k,:}^\top W_{k,l} A_{l,:}\right)^\top = A_{l,:}^\top W_{k,l}^\top A_{k,:} = A_{l,:}^\top W_{l,k} A_{k,:}$$

and each $Y_{k,l}$ is itself symmetric:

$$\begin{aligned} Y_{k,l}^\top &= \frac{1}{2}\left(\left(A_{k,:}^\top W_{k,l} A_{l,:}\right)^\top + \left(A_{l,:}^\top W_{l,k} A_{k,:}\right)^\top\right) \\ &= \frac{1}{2}\left(A_{l,:}^\top W_{l,k} A_{k,:} + A_{k,:}^\top W_{k,l} A_{l,:}\right) \\ &= Y_{k,l}. \end{aligned}$$

So the construction preserves the symmetry of $A^\top W A$.

Similarly to (3.1), we approximate the weighted product $A^\top W A$ by drawing $c \ll N$ i.i.d. index pairs $(k_j, l_j) \sim \pi$ from a probability distribution π on the index space

$\{(k,l)\}_{k,l=1}^{N}$ and defining a sequence of i.i.d. random symmetric matrices

$$\mathbb{P}\left[X_j = \frac{1}{\pi(k,l)} Y_{k,l}\right] = \pi(k,l), \quad j = 1, \dots, c,$$

so that

$$A^\top W A \approx \frac{1}{c} \sum_{j=1}^{c} X_j. \tag{4.3}$$

This pair-sampling framework is formalized in the following algorithm, taken from [HKL26]:

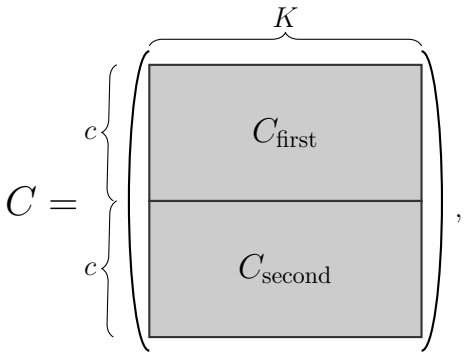
Algorithm 2: WeightedMatrixMultiplication

Input: Matrix $A \in \mathbb{R}^{N \times K}$, Matrix $W \in \mathbb{R}^{N \times N}$, sample size $c \ll N$, probability measure π on $\{(k,l)\}_{k,l=1}^{N}$

Output: Matrix $C \in \mathbb{R}^{2c \times K}$ and $V \in \mathbb{R}^{2c \times 2c}$ such that $C^\top V C \approx A^\top W A$

Set $V = 0 \in \mathbb{R}^{2c \times 2c}$
for $j = 1, \dots, c$ **do**
  Sample pair $(k_j, l_j) \sim \pi$ i.i.d.
  Set the j-th and $(j+c)$-th row of C as $C_j = A_{k_j,:}$ and $C_{(j+c)} = A_{l_j,:}$.
  Set the respective entries of V to $V_{j,j+c} = V_{j+c,j} = \frac{W_{k_j,l_j}}{2c\pi(k_j,l_j)}$.
return Matrix C, V and $C^\top V C$

Because each drawn pair contributes two rows to the sketch, the number of rows of $C \in \mathbb{R}^{2c \times K}$ doubles compared to Algorithm 1. The matrix is partitioned vertically into two blocks



where $C_{\text{first}} \in \mathbb{R}^{c \times K}$ contains the first components k_j of the sampled pairs and $C_{\text{second}} \in \mathbb{R}^{c \times K}$ the second components l_j. The rows of C remain unscaled copies of the rows of A. The entire probabilistic scaling and correlation weighting are shifted into the matrix $V \in \mathbb{R}^{2c \times 2c}$. Since V is initialized as a zero matrix and only assigned values at the positions $(j, j+c)$ and $(j+c, j)$ for $j = 1, \dots c$, it has a sparse and symmetric

block structure:

$$V = \begin{pmatrix} 0_{c\times c} & D \\ D & 0_{c\times c} \end{pmatrix}, \quad \text{where} \quad D = \operatorname{diag}(D_{1,1}, \dots, D_{c,c}) \in \mathbb{R}^{c\times c}$$

with the diagonal entries $D_{j,j} := \frac{W_{k_j,l_j}}{2c\,\pi(k_j,l_j)}$.

Using this block structure, the sketched product simplifies to

$$C^\top V C = C_{\text{first}}^\top D\, C_{\text{second}} + C_{\text{second}}^\top D\, C_{\text{first}}.$$

Since D is diagonal and $C_{\text{first}}, C_{\text{second}}$ contain the unscaled rows $A_{k_j,:}$ and $A_{l_j,:}$, both matrix products on the right reduce to a single sum over the c sampled pairs, with the scaling factor $D_{j,j} = \frac{W_{k_j,l_j}}{2c\,\pi(k_j,l_j)}$:

$$\begin{aligned} C_{\text{first}}^\top D\, C_{\text{second}} + C_{\text{second}}^\top D\, C_{\text{first}} &= \sum_{j=1}^{c} D_{j,j} \left(A_{k_j,:}^\top A_{l_j,:} + A_{l_j,:}^\top A_{k_j,:} \right) \\ &= \sum_{j=1}^{c} \frac{1}{2c\,\pi(k_j,l_j)} \left(A_{k_j,:}^\top W_{k_j,l_j} A_{l_j,:} + A_{l_j,:}^\top W_{l_j,k_j} A_{k_j,:} \right). \end{aligned}$$

Recalling the definition of $Y_{k,l}$ from (4.2), the right-hand side equals $\sum_{j=1}^{c} \frac{1}{c\,\pi(k_j,l_j)} Y_{k_j,l_j}$, so that

$$C^\top V C = \sum_{j=1}^{c} \frac{1}{c\,\pi(k_j,l_j)}\, Y_{k_j,l_j}.$$

This is exactly the sample average $\frac{1}{c}\sum_{j=1}^{c} X_j$ from (4.3), confirming that Algorithm 2 indeed produces an unbiased sketch of $A^\top W A$. Despite the doubled row count, the sparsity of V keeps the computation of the sketched matrix product $C^\top V C$ affordable.

Theoretical guarantee. A direct extension of Theorem 1 to this pairwise setting is established by the following theorem, taken from [HKL26]. The proof is analogous to the proof of Theorem 1 and a sketch is provided in [HKL26, Appendix B].

Theorem 2. *Let $W \in \mathbb{R}^{N\times N}$ be a symmetric positive definite matrix and let $A \in \mathbb{R}^{N\times K}$ with $N \gg K$. For a fixed sample size $c \ll N$, consider a probability distribution π over the index space $\{(k,l)\}_{k,l=1}^{N}$ such that for some $\beta \in (0,1]$, we have:*

$$\pi(k,l) \geq \beta \frac{\|Y_{k,l}\|_F}{\sum_{k',l'} \|Y_{k',l'}\|_F}. \tag{4.4}$$

Let C and V be the sketched matrices constructed according to Algorithm 2. Then, for

any failure rate $\delta \in (0,1)$*, the following high-probability error bound holds:*

$$\mathbb{P}\left(\|A^\top W A - C^\top V C\|_F \leq \frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}} \sum_{k,l} \|Y_{k,l}\|_F \right) \geq 1-\delta, \tag{4.5}$$

where the probability $\mathbb{P}$ *is taken over all realizations of the* c *samples.*

The following two remarks put Theorem 2 into perspective relative to its unweighted counterpart Theorem 1.

First, the sampling condition (4.4) again identifies an optimal distribution,

$$\tilde{\pi}(k,l) = \frac{\|Y_{k,l}\|_F}{\sum_{k',l'} \|Y_{k',l'}\|_F},$$

proportional not to a row norm but to the Frobenius norm of the pairwise interaction matrix $Y_{k,l}$. The quality parameter $\beta \in (0,1]$ plays exactly the same role as before. For $\beta = 1$, π is optimal and the bound (4.5) is the tightest. As $\beta \to 0$, the distribution $\pi(k,l)$ shifts away from the optimal weights and the error bound diverges.

Second, the structural dependencies of the error bound are unchanged. The dimension N of the design space does not appear explicitly. The error decays at the rate $\frac{1}{\sqrt{c}}$, confirming that the fundamental trade-off between computational cost and precision is preserved in the pairwise sampling framework. And the confidence level enters logarithmically through $\sqrt{\log(\delta^{-1})}$, so that high confidence can be demanded at negligible additional sampling cost.

With this extended theory in hand, we can now return to the physical problem and translate Theorem 2 into a concrete down-sampling strategy for the FIM that is aware of correlations.

4.3 Weighted Sketching Setup

We now apply Algorithm 2 and Theorem 2 to the experimental design problem, this time without assuming that the noise covariance Γ is diagonal. Recall that our objective is to identify a compressed design subset Ξ_c such that the sketched FIM approximates the full FIM:

$$\mathcal{I}(\Xi) = G^\top \Gamma^{-1} G \;\approx\; G_c^\top \Gamma_c^{-1} G_c = \mathcal{I}(\Xi_c).$$

To this end we identify $A := G$ and $W := \Gamma^{-1}$ in Algorithm 2 and Theorem 2. Substituting this identification into (4.2), the pairwise interaction matrix takes the

form

$$Y_{\xi,\theta} = \frac{1}{2}\left(G_{\xi,:}^{\top}\Gamma_{\xi,\theta}^{-1}G_{\theta,:} + G_{\theta,:}^{\top}\Gamma_{\theta,\xi}^{-1}G_{\xi,:}\right) \in \mathbb{R}^{K\times K}, \qquad \xi,\theta \in \Xi. \tag{4.6}$$

Drawing c i.i.d. pairs $(\xi_j, \theta_j) \sim \pi$ yields the following down-sampling strategy, adapted from Algorithm 2 to the experimental design context:

$$\begin{aligned} &\Xi_c = \{\xi_j\}_{j=1}^{c} \cup \{\theta_j\}_{j=1}^{c} \quad \text{with} \quad |\Xi_c| \le 2c, \qquad y_c = y|_{\Xi_c}, \\ &P \in \mathbb{R}^{2c\times 2c} \quad \text{with} \quad P_{k,l} = \begin{cases} \dfrac{\Gamma^{-1}_{\xi_j,\theta_j}}{2c\,\pi(\xi_j,\theta_j)}, & \text{if } |k-l| = c \text{ and } j = \min(k,l), \\ 0, & \text{else.} \end{cases} \end{aligned} \tag{4.7}$$

Because $2c$ indices are drawn, but design points may repeat, the cardinality of the resulting reduced design set satisfies $|\Xi_c| \le 2c$ instead of $|\Xi_c| = 2c$. The associated sketched FIM at the ground-truth parameter is then

$$\tilde{\mathcal{I}}(\Xi_c) = G_c^{\top} P\, G_c = \sum_{\xi,\theta\in\Xi_c} \frac{1}{2c\,\pi(\xi,\theta)}\Big(G_{\xi,:}^{\top}\Gamma_{\xi,\theta}^{-1}G_{\theta,:} + G_{\theta,:}^{\top}\Gamma_{\theta,\xi}^{-1}G_{\xi,:}\Big). \tag{4.8}$$

The procedure is summarized in the following algorithm taken from [HKL26]:

Algorithm 3: Experimental Design sampling

Input: probability distribution π on design index pairs $\{(k,l)\}_{k,l=1}^{N}$, desired number $2c \ll N$ of experiments, noise precision matrix Γ^{-1}.

Output: Down-sampled designs $\{\xi_1, \dots, \xi_c, \theta_1, \dots, \theta_c\}$.

Set $P \leftarrow 0 \in \mathbb{R}^{2c\times 2c}$
for $j = 1, \dots, c$ **do**
  Sample pair $(\xi_j, \theta_j) \sim \pi$ i.i.d.
  Add indices ξ_j, θ_j to set of indices Ξ_c.
  Store new precision value $P_{j,j+c} = P_{j+c,j} \leftarrow \frac{\Gamma^{-1}_{\xi_j,\theta_j}}{2c\pi(\xi_j,\theta_j)}$
return Down-sampled experimental designs $\{\xi_j, \theta_j\}_{j=1,\dots,c}$ and noise precision matrix P.

Optimal sampling strategy. According to Theorem 2, the pair-sampling distribution that minimizes the expected approximation error is proportional to the amount each pair contributes:

$$\tilde{\pi}(\xi,\theta) \;=\; \frac{\|Y_{\xi,\theta}\|_F}{\sum_{\xi',\theta'}\|Y_{\xi',\theta'}\|_F}. \tag{4.9}$$

If a chosen sampling distribution π satisfies $\pi \ge \beta\tilde{\pi}$ for some $\beta \in (0,1]$, then Theorem 2 guarantees that, with high probability, the down-sampled FIM $\tilde{\mathcal{I}}(\Xi_c)$ remains close to

the full FIM $\mathcal{I}(\Xi)$ in Frobenius norm, provided the sample size c is large enough.

Remark 1 (Accessibility of $\tilde{\pi}$). Evaluating the optimal sampling distribution $\tilde{\pi}(\xi, \theta)$ in (4.9) requires access to two quantities that are in general not directly available: the inverse noise covariance Γ^{-1}, and the sensitivity matrix G, whose rows are gradients of the forward model evaluated at the unknown ground truth p_*. The latter creates a classical circularity in experimental design, because the optimal design depends on the parameter that the inverse problem seeks to recover. We will return to both points and discuss how they are handled in practice in Section 5.5.

4.4 Sensitivity Guarantees under Sketching

Theorem 2 controls the distance between the full FIM $\mathcal{I}(\Xi)$ and its sketched counterpart $\tilde{\mathcal{I}}(\Xi_c)$ in the Frobenius norm. Although a small Frobenius norm error is reassuring, it does not ensure that the down-sampled inverse problem (1.4) restricted to Ξ_c is well-posed. Well-posedness in the sense of Hadamard requires the existence, uniqueness, and stability of the solution. In a noise-free setting, existence is unproblematic. Then, the ground truth p_* solves the equations $y = \mathcal{F}(p_*, \cdot)$ exactly for every subset Ξ_c. In the presence of noise, we do not treat the question of existence and instead focus on uniqueness and stability, since these properties are governed by the spectral behavior of the FIM and may be affected by down-sampling.

Well-posedness of the full problem. To have a meaningful benchmark against which to measure the sketched design, we must first assume that the full design Ξ is itself sufficiently informative. Following [HKL26], we formalize this as follows.

Assumption 1. *The experimental design Ξ is locally sensitive at p_* in the sense that the FIM defined in (1.3) is of full rank for a given positive definite covariance Γ. The inverse problem is therefore theoretically well-posed.*

This assumption is foundational for everything that follows. It states that the full dataset contains enough information in a neighborhood of the presumed ground truth p_* to uniquely and stably estimate it. If the FIM were rank-deficient, the design space Ξ would exhibit directions in the parameter space along which a perturbation of p produces no change in the expected data. Consequently, no amount of data, not even the full design space, let alone a down-sampled subset, could recover the parameter uniquely. Therefore, Assumption 1 is the minimal requirement on the problem. Since the FIM $\mathcal{I}(\Xi)$ is symmetric, its rank is equal to its number of non-zero eigenvalues. Consequently, requiring the by construction positive semidefinite FIM to be of full rank,

given a positive definite Γ, is equivalent to requiring it to be positive definite. Thus, its smallest eigenvalue is strictly positive. Consequently, the inverse problem is well-posed if the down-sampled FIM $\mathcal{I}(\Xi_c)$ retains a strictly positive smallest eigenvalue. Ideally, the smallest eigenvalue of the compressed design should not differ greatly from that of the full design because, according to the Cramér-Rao inequality mentioned in Section 1.2, the smaller the eigenvalues, the greater the loss of accuracy.

Probabilistic bound to preserve data sensitivity. To quantify how well a given down-sampled design performs relative to the full design, we analyze the impact of sketching on the spectral properties of the FIM. Like in [HKL26] we denote by $\lambda^{\Xi}_{\min}$ and $\lambda^{c}_{\min}$ the smallest eigenvalues of $\mathcal{I}(\Xi)$ and $\tilde{\mathcal{I}}(\Xi_c)$, respectively, and by $c^{\Xi}_{\text{inv}} = \frac{\lambda^{\Xi}_{\min}}{\lambda^{\Xi}_{\max}}$ and $c^{c}_{\text{inv}} = \frac{\lambda^{c}_{\min}}{\lambda^{c}_{\max}}$ their respective inverse condition numbers. Under this notation, a down-sampling scheme is considered successful if $\lambda^{c}_{\min}$ remains strictly positive, thereby preserving the well-posedness guaranteed by Assumption 1, and if c^{c}_{inv} remains close to 1, ensuring numerical stability.

We can now state the main result of this section, adapted from [HKL26], which links the sample size c to the preservation of local sensitivity and spectral stability under the pairwise down-sampling strategy of Section 4.3.

Theorem 3. *Consider an inverse problem that satisfies Assumption 1, and let the re-weighted data be constructed as in (4.7), where the sampling probability density $\pi(\xi, \theta)$ on $\Xi \times \Xi$ satisfies*

$$\pi \geq \beta \tilde{\pi}$$

for some $\beta \in (0, 1]$. Then, for any failure probability $\delta \in (0, 1)$ and any error tolerance $\varepsilon \in \left(0, \lambda^{\Xi}_{\min}\right)$ a choice of the sample size

$$c \geq \left(\sum_{\xi,\theta \in \Xi} \|Y_{\xi,\theta}\|_F \right)^2 \frac{\left(1 + \sqrt{8\beta^{-1} \log(\delta^{-1})}\right)^2}{\beta \varepsilon^2} \tag{4.10}$$

assures that with probability at least $1 - \delta$:

$$\lambda^{c}_{\min} \geq \lambda^{\Xi}_{\min} - \varepsilon > 0, \quad \textit{and} \quad c^{c}_{inv} \geq c^{\Xi}_{inv} \frac{\lambda^{\Xi}_{\min} - \varepsilon}{\lambda^{\Xi}_{\min} + \varepsilon}.$$

Theorem 3 converts the abstract sketching guarantee for the Frobenius norm into a statement about the well-posedness and numerical stability of the down-sampled inverse problem. It establishes a lower bound on the sample size c required to guarantee the positive definiteness of the reduced FIM as well as a controlled shrinkage of its

inverse condition number with high confidence at least $1 - \delta$. Its dependence on Assumption 1 is essential. Without $\lambda^{\Xi}_{\min} > 0$, the admissible error tolerance $\varepsilon \in (0, \lambda^{\Xi}_{\min})$ collapses to the empty set. In this case, no perturbation could guarantee that the reduced FIM remains positive definite. Assumption 1 is the condition that makes a meaningful stability guarantee for the sketched problem possible in the first place.

The structure of the lower bound (4.10) reveals how the experimental setup impacts the required sampling effort.

The prefactor $\left(\sum_{\xi,\theta\in\Xi} \|Y_{\xi,\theta}\|_F\right)^2$ measures the total sensitivity content of the full design space Ξ. The larger Ξ is, the more variation there is in the data. Consequently, the more variation there is to represent in the reduced design, and hence the larger c must be chosen. Because this term is squared, the sample size scales quadratically with the total sensitivity content of the design space.

The bound scales as ε^{-2} to the error tolerance. Consequently, reducing the eigenvalue perturbation ε by half requires quadrupling the sample size c. This is the familiar trade-off between accuracy and cost already encountered in Theorem 1 and Theorem 2.

The failure probability δ enters the bound through $\sqrt{\log(\delta^{-1})}$. A significantly higher confidence level can therefore be achieved with a relatively small increase in sample size, which also mirrors the corresponding observation for Theorem 1 and Theorem 2.

The parameter $\beta \in (0, 1]$ again measures how closely the chosen sampling distribution π approximates the optimal distribution $\tilde{\pi}$ from (4.9). A poor choice of π, which means a small β, penalizes the sample size requirement through β^{-2}.

Proof. To establish the lower bound for $\lambda^{c}_{\min}$, we utilize

$$\lambda^{c}_{\min} \geq \lambda^{\Xi}_{\min} - |\lambda^{\Xi}_{\min} - \lambda^{c}_{\min}|.$$

By Weyl's inequality, for symmetric matrices $A, B \in \mathbb{R}^{K\times K}$ it can be shown that the difference between eigenvalues is bounded by

$$\left|\lambda^{A}_{j} - \lambda^{B}_{j}\right| \leq \|A - B\|_F \quad \text{for } j = 1, \ldots, K.$$

Applying this property to the smallest eigenvalues of the FIMs yields

$$\left|\lambda^{\Xi}_{\min} - \lambda^{c}_{\min}\right| \leq \|I(\Xi) - I(\Xi_c)\|_F .$$

According to Theorem 2, using the assumption that $\pi \geq \beta\tilde{\pi}$, the Frobenius norm is

bounded, with probability at least $1-\delta$, by

$$\|I(\Xi) - I(\Xi_c)\|_F \leq \frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}} \sum_{\xi,\theta\in\Xi} \|Y_{\xi,\theta}\|_F$$

Requiring the right-hand side to not exceed the prescribed tolerance ε

$$\frac{1+\sqrt{8\beta^{-1}\log(\delta^{-1})}}{\sqrt{\beta c}} \sum_{\xi,\theta\in\Xi} \|Y_{\xi,\theta}\|_F \leq \varepsilon$$

and solving this inequality for the sample size c yields the lower bound (4.10):

$$c \geq \left(\sum_{\xi,\theta\in\Xi} \|Y_{\xi,\theta}\|_F\right)^2 \frac{\left(1+\sqrt{8\beta^{-1}\log(\delta^{-1})}\right)^2}{\beta\varepsilon^2}.$$

Hence, on the event of Theorem 2, which has probability at least $1-\delta$, we obtain

$$\lambda^c_{\min} \geq \lambda^{\Xi}_{\min} - \varepsilon. \tag{4.11}$$

Next, we derive the bound for the inverse condition number c^c_{inv}. Using the triangle inequality combined with the Frobenius norm bound from Theorem 2, the largest eigenvalue can be upper-bounded by

$$\lambda^c_{max} \leq \lambda^{\Xi}_{max} + |\lambda^c_{max} - \lambda^{\Xi}_{max}| \leq \lambda^{\Xi}_{max} + \|I(\Xi_c) - I(\Xi)\|_F \leq \lambda^{\Xi}_{max} + \varepsilon. \tag{4.12}$$

Substituting the lower bound for $\lambda^c_{\min}$ (4.11) and the upper bound for $\lambda^c_{\max}$ (4.12) into the definition of c^c_{inv} yields

$$c^c_{\text{inv}} = \frac{\lambda^c_{\min}}{\lambda^c_{\max}} \geq \frac{\lambda^{\Xi}_{\min} - \varepsilon}{\lambda^{\Xi}_{\max} + \varepsilon}.$$

Finally, using the definition of the inverse condition number

$$c^{\Xi}_{\text{inv}} = \frac{\lambda^{\Xi}_{\min}}{\lambda^{\Xi}_{\max}} \quad \Leftrightarrow \quad \lambda^{\Xi}_{\max} = \frac{\lambda^{\Xi}_{\min}}{c^{\Xi}_{\text{inv}}},$$

we substitute the expression for $\lambda^{\Xi}_{\max}$ into the denominator

$$c^c_{\text{inv}} \geq \frac{\lambda^{\Xi}_{\min} - \varepsilon}{\frac{\lambda^{\Xi}_{\min}}{c^{\Xi}_{\text{inv}}} + \varepsilon} = c^{\Xi}_{\text{inv}} \frac{\lambda^{\Xi}_{\min} - \varepsilon}{\lambda^{\Xi}_{\min} + \varepsilon \cdot c^{\Xi}_{\text{inv}}}.$$

Since $c_{\text{inv}}^{\Xi} \leq 1$, it holds that $\varepsilon \cdot c_{\text{inv}}^{\Xi} \leq \varepsilon$, yielding:

$$c_{\text{inv}}^{c} \geq c_{\text{inv}}^{\Xi} \frac{\lambda_{\min}^{\Xi} - \varepsilon}{\lambda_{\min}^{\Xi} + \varepsilon},$$

which completes the proof. □

5 Sampling

Achieving the guarantees of the Theorems 2 and 3 requires sampling design pairs (ξ, θ) from the optimal distribution $\tilde{\pi}(\xi, \theta)$ from (4.9). In practice, however, this is often infeasible. When Ξ is large, evaluating the normalization constant $\sum_{\xi',\theta'} \|Y_{\xi',\theta'}\|_F$ in the denominator of $\tilde{\pi}$ becomes numerically intractable. This chapter addresses this obstacle. It introduces a sampling method based on [GI+20] and [HKL26] that is able to generate representative draws from a target distribution such as $\tilde{\pi}$ without requiring the normalization constant.

Section 5.1 formulates the generic Bayesian inverse problem underlying the sampling target distribution. Section 5.2 introduces the Ensemble Kalman Sampler (EKS) as a gradient-free method for drawing from this distribution. Section 5.3 discusses the mean-field limit of the EKS dynamics and its convergence to the target distribution. Section 5.4 specializes this general sampling framework to the design-selection problem, resulting in the complete algorithmic pipeline of this thesis. Section 5.5 concludes the chapter with a discussion of the drawbacks of the sampling approach.

5.1 The Bayesian Inverse Problem

Following the general formulation in [GI+20], we first formulate a generic inverse problem similar to (1.1) because, rather than reconstructing the ground truth parameter p_*, we use the Bayesian formulation of this problem to derive a sampling target distribution over an arbitrary Bayesian variable u. Throughout this chapter, p_* remains the fixed ground truth. In Section 5.4, we specialize the variable u to the design pairs relevant to our down-sampling problem.

Consider the inverse problem of finding $u \in \mathbb{R}^d$ from observed data $y \in \mathbb{R}^N$,

$$y = f(u) + \eta, \tag{5.1}$$

where $f : \mathbb{R}^d \to \mathbb{R}^N$ is a deterministic forward operator and $\eta \sim \mathcal{N}(0, \Gamma)$ is additive Gaussian measurement noise with a known, symmetric, positive definite covariance

matrix $\Gamma \in \mathbb{R}^{N\times N}$. Note that u ranges over the continuous space $\mathbb{R}^d$, in contrast to the discrete design space Ξ of Chapter 4. The relation between the two is established in Section 5.4, where the EKS is applied on a continuous extension of Ξ and the iterates are projected back onto Ξ.

In classical approaches, when dealing with an inverse problem of the form (5.1), the goal is to find an optimal parameter u. In the Bayesian approach to inverse problems, however, the uncertainty introduced by the noise is treated by modeling the unknown parameter u as a random variable. Instead of a single optimal value, the solution is formulated as a posterior probability distribution of all possible parameters u, given the observed data y.

Given a prior density $\pi_{\text{pr}}(u)$ for the unknown u, Bayes' theorem yields the posterior density

$$\pi_{\text{pos}}(u \mid y) \;\propto\; \ell(y \mid u)\, \pi_{\text{pr}}(u),$$

where $\ell(y \mid u)$ is the likelihood of observing y given u. Since our focus is on the sensitivity of the data to the parameter, and not on the additional regularizing effect of a prior, we assume throughout this chapter that the inverse problem is well-posed and π_{pr} is proportional to a constant. Under the Gaussian noise model with $\eta = y - f(u)$, the likelihood is given by

$$\ell(y \mid u) = \frac{1}{\sqrt{(2\pi)^N \det(\Gamma)}} \exp\left(-\frac{1}{2}(y - f(u))^\top \Gamma^{-1}(y - f(u))\right).$$

Since the normalization factor $\frac{1}{\sqrt{(2\pi)^N \det(\Gamma)}}$ is independent of the variable u, it can be incorporated into the proportionality constant. Defining the nonlinear least squares functional, sometimes referred to as the potential,

$$\Phi(u) \coloneqq \frac{1}{2}\|y - f(u)\|_\Gamma^2 = \frac{1}{2}(y - f(u))^\top \Gamma^{-1}(y - f(u)), \tag{5.2}$$

the target distribution takes the exponential form

$$\tilde{\pi}(u) \;\propto\; e^{-\Phi(u)}. \tag{5.3}$$

This probability distribution satisfies two key criteria. First, the exponential form ensures that the probability density is strictly positive. Second, $\tilde{\pi}$ is known only up to a normalization constant, which is the practical obstacle.

5.2 Ensemble Kalman Sampling

Why not Markov Chain Monte Carlo? Markov Chain Monte Carlo (MCMC) methods are the classical tools for drawing samples from a distribution known only up to a normalization constant. As described in [HKL26], the underlying idea is to construct a Markov chain on the parameter space whose stationary distribution is the target distribution $\tilde{\pi}(u) \propto e^{-\Phi(u)}$ from (5.3). As a sample propagates through this Markov chain, its distribution converges to the stationary distribution over time, and the normalization constant of $\tilde{\pi}$ is never evaluated.

Many MCMC methods exploit the mathematical structure of the target distribution. An example for this are MCMC methods in [LMW18], that use systems of interacting particles derived from coupled Langevin-type equations. As [GI+20] points out, these MCMC methods require the calculation of gradients of the log likelihood or the potential Φ. However, in many practical applications, these derivatives are either unavailable or expensive to compute. The EKS circumvents this gradient dependency by building upon the framework of coupled Langevin dynamics while approximating exact gradients by ensemble differences. Yet, for ensemble methods, the theory on non-asymptotic convergence rates, particularly for non-Gaussian target distributions, is still under development. The choice between the two is, therefore, a compromise between accuracy guarantees and computational cost. For qualitative experimental design precise sampling from $\tilde{\pi}$ is not necessary, since the main goal is to improve the conditioning of the down-sampled FIM. As for our use case, the gradient-free EKS is typically cheaper than MCMC, its computational advantage outweighs the weaker theoretical guarantees.

Langevin dynamics as a conceptual foundation. According to [GI+20], the conceptual foundation of the EKS is the overdamped Langevin equation,

$$\mathrm{d}u = -\nabla\Phi(u)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}W, \tag{5.4}$$

where W denotes a standard Brownian motion in $\mathbb{R}^d$.

This is an SDE, which we introduced in Section 2.4. The term $-\nabla\Phi(u(t))\,\mathrm{d}t$ describes the deterministic drift. It pulls the particle down the potential landscape toward its minima, which correspond to the maxima of the density $\tilde{\pi}(u) \propto e^{-\Phi(u)}$, acting analogously to a gradient-based optimization step. Conversely, the term $\sqrt{2}\,\mathrm{d}W$ represents the diffusion component, supplying stochastic noise. This noise allows the particle to leave local minima and to explore the probability landscape.

The target distribution (5.3) is specified only up to a multiplicative constant. To see that this constant does not affect the dynamics, we make it explicit:

$$\tilde{\pi}(u) = \frac{1}{Z}\, e^{-\Phi(u)}, \qquad Z := \int_{\mathbb{R}^d} e^{-\Phi(u)}\, \mathrm{d}u,$$

where Z is determined by the requirement $\int_{\mathbb{R}^d} \tilde{\pi}(u)\, \mathrm{d}u = 1$ and is finite and positive whenever $e^{-\Phi}$ is integrable. The constant Z is generally intractable, since it requires integrating over the entire state space. However, it does not need to be computed, because $\log(Z)$ has zero gradient and cancels during differentiation:

$$\nabla \log \tilde{\pi}(u) = \nabla \log\left(\frac{1}{Z}\, e^{-\Phi(u)}\right) = -\nabla \log Z - \nabla\Phi(u) = -\nabla\Phi(u). \tag{5.5}$$

Consequently, only $\nabla\Phi$ enters the dynamics, and the normalization constant is never evaluated.

The time t introduced in (5.4) is an artificial pseudo-time. This pseudo-time serves no physical purpose. Its role is to let the particles flow, step by step, from an arbitrary initial ensemble towards the target distribution $\tilde{\pi}(u) \propto e^{-\Phi(u)}$.

To speed up the process, a symmetric, positive definite preconditioning matrix $C \in \mathbb{R}^{d\times d}$ is introduced into the drift and diffusion terms:

$$\mathrm{d}u = -C\, \nabla\Phi(u)\, \mathrm{d}t + \sqrt{2C}\, \mathrm{d}W.$$

A central algorithmic idea underlying the EKS is to select C automatically and adaptively, using an interacting ensemble of particles rather than a fixed matrix. Concretely, we evolve a set of J interacting particles $U = \{u_j\}_{j=1}^J$ according to

$$\mathrm{d}u_j = -C(U)\, \nabla\Phi\big(u_j\big) + \sqrt{2\, C(U)}\, \mathrm{d}W_j, \qquad j = 1, \dots, J, \tag{5.6}$$

where $\{W_j\}_{j=1}^J$ are i.i.d. standard Brownian motions in $\mathbb{R}^d$, and $C(U)$ is chosen to be the empirical covariance of the ensemble,

$$C(U) = \frac{1}{J}\sum_{k=1}^{J} \big(u_k - \bar{u}\big) \otimes \big(u_k - \bar{u}\big) \in \mathbb{R}^{d\times d}, \qquad \bar{u} = \frac{1}{J}\sum_{j=1}^{J} u_j. \tag{5.7}$$

Because $C(U)$ depends nonlinearly on the entire ensemble, the drift and diffusion of every particle are influenced by the current spread of the whole ensemble around its mean.

Gradient approximation via ensemble differences. A common obstacle to using (5.6) in practice is that the gradient $\nabla\Phi$ is either unavailable or prohibitively expensive to evaluate. The EKS offers a solution to that problem. Instead of computing the gradient exactly, it approximates it using ensemble differences, yielding a fully gradient-free algorithm. We will do that similarly to [HKL26] and [GI+20].

Consider Φ of the quadratic, least-squares form (5.2), whose gradient is

$$\nabla\Phi(u) = Df(u)^\top \Gamma^{-1}\big(f(u) - y\big).$$

Substituting this into the drift term of (5.6) and using the definition (5.7) of the empirical covariance, the drift acting on particle u_j can be rewritten as

$$\begin{aligned} C(U)\nabla\Phi\big(u_j\big) &= \frac{1}{J}\sum_{k=1}^{J}\big(u_k - \bar{u}\big) \otimes \big(u_k - \bar{u}\big) \cdot Df\big(u_j\big)^\top \Gamma^{-1}\big(f(u_j) - y\big) \\ &= \frac{1}{J}\sum_{k=1}^{J}\big(u_k - \bar{u}\big)\big(u_k - \bar{u}\big)^\top \cdot Df\big(u_j\big)^\top \Gamma^{-1}\big(f(u_j) - y\big) \\ &= \frac{1}{J}\sum_{k=1}^{J}\big(u_k - \bar{u}\big)\big(Df\big(u_j\big)\big(u_k - \bar{u}\big)\big)^\top \Gamma^{-1}\big(f(u_j) - y\big) \\ &= \frac{1}{J}\sum_{k=1}^{J}\big(u_k - \bar{u}\big)\,\big\langle Df\big(u_j\big)\big(u_k - \bar{u}\big), \big(f(u_j) - y\big)\big\rangle_\Gamma\,. \end{aligned} \tag{5.8}$$

Here every term is directly computable from the ensemble except the directional derivative $Df\big(u_j\big)\big(u_k - \bar{u}\big)$, which still requires differentiating f.

To solve this problem, we start from the first-order Taylor expansion of f around u_j,

$$f\big(u_k\big) \approx f\big(u_j\big) + Df\big(u_j\big)\big(u_k - u_j\big),$$

which rearranges to

$$Df\big(u_j\big)\big(u_k - u_j\big) \;\approx\; f\big(u_k\big) - f\big(u_j\big). \tag{5.9}$$

Because the ensemble consists of J particles rather than a single point, we relate it to the ensemble mean. Averaging (5.9) over all $k = 1, \dots, J$ for fixed j gives

$$\frac{1}{J}\sum_{k=1}^{J} Df\big(u_j\big)\big(u_k - u_j\big) \;\approx\; \frac{1}{J}\sum_{k=1}^{J}\Big(f\big(u_k\big) - f\big(u_j\big)\Big),$$

and using the ensemble mean $\bar{u}$ from (5.7) and $\bar{f} := \frac{1}{J}\sum_{k=1}^{J} f\big(u_k\big)$ yields

$$Df\big(u_j\big)\big(\bar{u} - u_j\big) \;\approx\; \bar{f} - f\big(u_j\big). \tag{5.10}$$

Subtracting (5.10) from (5.9)

$$Df(u_j)(u_k - u_j) - Df(u_j)(\bar{u} - u_j) \;\approx\; f(u_k) - f(u_j) - (\bar{f} - f(u_j))$$

and using the linearity of $Df(u_j)$ we establish an approximation for the directional derivative

$$Df(u_j)(u_k - \bar{u}) \;\approx\; f(u_k) - \bar{f}. \tag{5.11}$$

The gradient-free EKS dynamics. Substituting the approximation (5.11) into (5.8), and hence into the drift of (5.6), yields the gradient-free EKS dynamics

$$\mathrm{d}u_j = -\frac{1}{J}\sum_{k=1}^{J} \Big\langle f(u_k) - \bar{f},\; f(u_j) - y \Big\rangle_{\Gamma} u_k \;+\; \sqrt{2\,C(U)}\,\mathrm{d}W_j. \tag{5.12}$$

The system requires only evaluations of f at the current ensemble locations $u_1, \ldots, u_J$, and never its derivative. That is precisely the property that makes it viable for the design-selection application.

Two remarks on this approximation are worth noting. First, whenever f is affine, the approximation (5.11) is exact, since the Taylor expansion has no remainder. Consequently, (5.12) reproduces the exact dynamics (5.6). Second, if f is mildly nonlinear within the area of the ensemble, the Jacobi matrix $Df(u)$ varies only slightly locally. Therefore, the approximation also holds in that case.

5.3 Mean-Field Limit

For large ensemble $J \to \infty$, a single particle $u_j(t)$ behaves approximately like a sample drawn from a distribution $\rho(t, \cdot)$, which itself no longer depends on the randomness of the individual realization but corresponds deterministically to the probability density of the system at time t. To formalize this, we replace the empirical covariance in (5.7) with the corresponding covariance of the distribution $\rho(t, \cdot)$ following [GI+20]:

$$C(\rho) := \int (u - \bar{u}) \otimes (u - \bar{u})\, \rho(t, u)\, \mathrm{d}u \quad \text{with} \quad \bar{u} := \int u\, \rho(t, u)\, \mathrm{d}u. \tag{5.13}$$

By replacing the empirical covariance $C(U)$ of the ensemble in (5.6) by the covariance of the limiting distribution $C(\rho)$ as $J \to \infty$, a single particle $u(t)$ then follows the mean-field SDE

$$\mathrm{d}u(t) = -C(\rho(t))\, \nabla\Phi\big(u\big)\mathrm{d}t + \sqrt{2\,C(\rho(t))}\,\mathrm{d}W(t) \tag{5.14}$$

provided that Φ is Lipschitz-smooth, as it was shown in [DL21] and [Vae24]. Correspondingly, the probability density ρ_t evolves according to the nonlinear Fokker-Planck equation

$$\partial_t \rho = \nabla \cdot \big(\rho\, \mathcal{C}(\rho) \nabla \Phi\big) + \mathcal{C}(\rho) : D^2 \rho, \tag{5.15}$$

where $D^2\rho$ is the Hessian of ρ and $\mathcal{C}(\rho) : D^2\rho = \mathrm{tr}(\mathcal{C}(\rho) D^2\rho)$ is the Frobenius inner product. This is the special case of the general Fokker-Planck equation (2.4) for $b = -\mathcal{C}(\rho)\nabla\Phi$ and $\Sigma = 2\,\mathcal{C}(\rho)$. Since $C(\rho)$ does not depend on u it holds $\nabla \cdot \nabla \cdot (C(\rho)\rho) = \sum_{i,j} \partial_i \partial_j (C(\rho)_{ij} \rho) = \sum_{i,j} C(\rho)_{ij}\, \partial_i \partial_j \rho = C(\rho) : D^2\rho$. The equation is nonlinear because $\mathcal{C}(\rho)$ is itself computed from ρ via (5.13). The rigorous derivation of the mean-field equations (5.14) and (5.15) would go beyond the scope of this thesis.

Instead, we record the qualitative fact we rely on in Section 5.4. It can be checked directly by substitution that the target distribution $\rho \propto e^{-\Phi}$ from (5.3) is an invariant measure of (5.15). Writing $\rho = \frac{1}{Z} e^{-\Phi}$ and $C = C(\rho)$ for the mean-field covariance. The chain rule gives $\nabla\rho = -\rho\,\nabla\Phi$. Substituting into the right-hand side of (5.15) and expanding both terms by the product rule, leads to

$$\nabla \cdot (\rho\, C \nabla\Phi) = (\nabla\rho)^\top C \nabla\Phi + \rho\, \nabla \cdot (C\nabla\Phi) = -\rho\, \nabla\Phi^\top C \nabla\Phi + \rho\, \nabla \cdot (C\nabla\Phi),$$

and

$$D^2\rho = D(\nabla\rho) = D(-\rho\nabla\Phi) = -(\nabla\rho) \otimes \nabla\Phi - \rho D^2\Phi = \rho\,(\nabla\Phi \otimes \nabla\Phi) - \rho\, D^2\Phi.$$

Using the symmetry of C and $C : (a \otimes a) = a^\top C a$ for any vector a,

$$C : D^2\rho = \rho\, \nabla\Phi^\top C \nabla\Phi - \rho\, C : D^2\Phi = \rho\, \nabla\Phi^\top C \nabla\Phi - \rho\, \nabla \cdot (C\nabla\Phi),$$

where the last equality uses $C : D^2\Phi = \nabla \cdot (C\nabla\Phi)$ for constant C. Adding the two terms, yields

$$\partial_t \rho \Big|_{\rho \propto e^{-\Phi}} = \underbrace{-\rho\, \nabla\Phi^\top C \nabla\Phi + \rho\, \nabla \cdot (C\nabla\Phi)}_{\nabla \cdot (\rho C \nabla\Phi)} + \underbrace{\rho\, \nabla\Phi^\top C \nabla\Phi - \rho\, \nabla \cdot (C\nabla\Phi)}_{C : D^2\rho} = 0.$$

Both terms cancel exactly, confirming that $\rho \propto e^{-\Phi}$ is a stationary solution of (5.15), i.e., an invariant measure. In fact, as shown by [GI+20], $\rho \propto e^{-\Phi}$ is the unique stationary solution of the mean-field PDE that is not concentrated on a single point. The only other steady states of (5.15) are Dirac measures, which correspond to a vanishing ensemble covariance $C(\rho) = 0$. A collapse onto some Dirac measure would correspond to the ensemble contracting onto a single point, after which the dynamics would stagnate and could no longer explore the design space.

When Φ is strongly convex and the ensemble covariance matrix is uniformly bounded from below away from zero, [GI+20] demonstrates that the PDE (5.15) converges exponentially fast to $\rho \propto e^{-\Phi}$ as $t \to \infty$. Together with the shown invariant measure, this is what justifies treating the ensemble as samples drawn from the target distribution $\tilde{\pi} \propto e^{-\Phi}$, after running the EKS dynamics for a sufficiently large ensemble J and a sufficiently long time t.

5.4 Experimental Design through Sampling

We now apply the sampling method developed in the Section 5.2 to the target distribution derived in Section 4.3. Identifying the generic Bayesian variable u from the previous sections with the design pair $\chi = (\xi, \theta) \in \Xi \times \Xi$ specializes the general EKS dynamics (5.12) to the design-selection setting of [HKL26].

Recall from (4.9) that the sampling distribution that best preserves the sensitivity of the FIM under down-sampling is

$$\tilde{\pi}(\chi) \;\propto\; \|Y_\chi\|_F,$$

with the pairwise interaction matrix $Y_\chi = Y_{\xi,\theta}$ from (4.6). We can write it in the exponential form (5.3) as

$$\tilde{\pi}(\chi) = \frac{1}{Z}\, e^{-\Phi(\chi)}, \qquad \Phi(\chi) := -\log\big(\|Y_\chi\|_F\big), \tag{5.16}$$

where $Z = \sum_{\xi',\theta'} \|Y_{\xi',\theta'}\|_F$ is the generally intractable normalization constant.

To apply the EKS dynamic, we work on a continuous extension $\tilde{\Xi} \supset \Xi$ of the design space. For Φ and $\tilde{\pi}$ on $\tilde{\Xi} \times \tilde{\Xi}$ to be well-defined, we assume that $\tilde{\Xi}$ is of finite Lebesgue measure and that $G_{\xi,:}$ and $\Gamma^{-1}_{\xi,\theta}$ are uniformly bounded and non-zero for all $\xi, \theta \in \tilde{\Xi}$. It follows that $\|Y_\chi\|_F$ is also bounded and non-zero. This ensures that $\Phi(\chi) = -\log\big(\|Y_\chi\|_F\big)$ remains finite. Together with the finite measure of $\tilde{\Xi}$, this guarantees that the normalization constant Z is finite and strictly positive, so that $\tilde{\pi}$ is a well-defined probability density.

The EKS only requires evaluations of $\|Y_\chi\|_F$, and never Z itself, since Z cancels in the drift term as shown in (5.5). Moreover, since Ξ is discrete, the gradient $\nabla\Phi$ appearing in (5.6) is generally ill-defined. The EKS is gradient-free by construction as established in (5.12), and remains applicable even when the gradient is not available. Like in [HKL26] it should be noted that while the EKS was developed for continuous state spaces, we employ a nearest neighbor projection for its numerical realization on

our discrete design set.

Gradient approximation. Unlike the generic setting of Section 5.2, the potential Φ in (5.16) is not of the least-squares form we defined (5.2). Instead, it is a functional of a scalar quantity $h(\chi) := \|Y_\chi\|_F$, namely $\Phi(\chi) = -\log(h(\chi))$. Since Φ is nonetheless smooth, we can derive a gradient-free approximation of the drift term $C(U)\nabla\Phi(\chi_j)$ analogous to Section 5.2.

Writing $\chi_j := (\xi_j, \theta_j) \in \Xi \times \Xi$ for the j-th element out of the ensemble $S = \{\chi_j = (\xi_j, \theta_j)\}_{j=1,\dots,c}$. By the chain rule and definition (5.7) we have

$$C(S)\nabla_\chi\Phi(\chi_j) = -\left(\frac{1}{c}\sum_{k=1}^{c}\big(\chi_k - \bar{\chi}\big)\big(\chi_k - \bar{\chi}\big)^\top\right)\frac{\nabla_\chi h(\chi_j)}{h(\chi_j)}.$$

Using the fact that $h(\chi_j) = \|Y_{\chi_j}\|_F = \sqrt{\langle Y_{\chi_j}, Y_{\chi_j}\rangle_F}$ and the chain rule once more it follows

$$C(S)\nabla_\chi\Phi(\chi_j) = -\frac{1}{2c\|Y_{\chi_j}\|_F^2}\sum_{k=1}^{c}\big(\chi_k - \bar{\chi}\big)\big(\chi_k - \bar{\chi}\big)^\top\nabla_\chi\langle Y_{\chi_j}, Y_{\chi_j}\rangle_F.$$

Since the Frobenius scalar product is symmetric and bilinear by the product rule it is

$$\nabla_\chi\langle Y_{\chi_j}, Y_{\chi_j}\rangle_F = 2\langle\nabla_\chi Y_{\chi_j}, Y_{\chi_j}\rangle_F.$$

Consequently,

$$C(S)\nabla_\chi\Phi(\chi_j) = -\frac{1}{c\|Y_{\chi_j}\|_F^2}\sum_{k=1}^{c}\big(\chi_k - \bar{\chi}\big)\,\langle\nabla_\chi Y_{\chi_j}\big(\chi_k - \bar{\chi}\big), Y_{\chi_j}\rangle_F.$$

This still contains the unknown directional derivative $\nabla_\chi Y_{\chi_j}(\chi_k - \bar{\chi})$. Substituting the matrix-valued map Y_χ for f in the established approximation (5.11) leads to the ensemble-difference approximation

$$\nabla_\chi Y_{\chi_j}(\chi_k - \bar{\chi}) \approx Y_{\chi_k} - \bar{Y_\chi}, \qquad \bar{Y_\chi} := \frac{1}{c}\sum_{k=1}^{c} Y_{\chi_k}$$

and

$$C(S)\nabla_\chi\Phi(\chi_j) \approx -\frac{1}{c\|Y_{\chi_j}\|_F^2}\sum_{k=1}^{c}\big(\chi_k - \bar{\chi}\big)\,\langle Y_{\chi_k} - \bar{Y_\chi}, Y_{\chi_j}\rangle_F.$$

We define the gradient-free interaction kernel

$$D_{j,k} := \frac{1}{c\,\|Y_{\chi_j}\|_F^2}\langle Y_{\chi_k} - \bar{Y}_\chi, Y_{\chi_j}\rangle_F, \tag{5.17}$$

which satisfies

$$\sum_{k=1}^{c} D_{j,k}\,\chi_k \approx -C(S)\nabla_\chi \Phi(\chi_j). \tag{5.18}$$

As in Section 5.2, the approximation in (5.18) becomes an equality when Y_χ is affine in χ and remains applicable when Y_χ is mildly nonlinear. Furthermore, the left side of (5.18) is gradient-free. The approximation is built entirely from the interaction matrices Y_{χ_j}, which are directly computable once G and Γ^{-1} are known.

Discretization. Substituting (5.17) for the drift in (5.6) yields the gradient-free, design-specific EKS dynamics

$$\mathrm{d}\chi_j = \sum_{k=1}^{c} D_{j,k}\,\chi_k\,\mathrm{d}t + \sqrt{2\,C(S)}\,\mathrm{d}W_j, \qquad j = 1,\dots,c, \tag{5.19}$$

where $C(S) = \frac{1}{c}\sum_k(\chi_k-\bar{\chi})\otimes(\chi_k-\bar{\chi})$ is the empirical covariance of the current ensemble $S = \{\chi_j\}_{j=1,\dots,c}$. Discretizing (5.19) with the Euler–Maruyama method introduced in Section 2.4 gives the update samples

$$\chi_j^{t_{n+1}} = \chi_j^{t_n} + \Delta t_n \sum_k D_{j,k}^{t_n}\,\chi_k^{t_n} + \sqrt{2\,\Delta t_n\, C\big(S^{t_n}\big)}\,\eta_j^{t_n}, \qquad \eta_j^{t_n} \sim \mathcal{N}(0, I). \tag{5.20}$$

Here, the time step Δt_n can be chosen adaptively to enhance numerical stability as the ensemble evolves, for instance as proposed in [KS19]. Because Ξ is discrete, the resulting iterates $\chi_j^{t_{n+1}}$ produced by (5.20) must be projected back onto the design space after every step. Based on [HKL26] we do this by nearest-neighbor projection onto $\Xi \times \Xi$.

Early stopping. Since the continuous-time EKS dynamics (5.19) are only realized numerically through a time-discretization with the Euler–Maruyama method (5.20), the resulting discrete-time ensemble does not converge exactly to $\tilde{\pi}$. This means the discretization introduces a bias into the system. A common improvement is to integrate a Metropolis–Hastings step as post-processing, following [SWZ25]. However, in our setting this is not necessary. Our goal is not to obtain high-fidelity samples from $\tilde{\pi}$, but rather to improve the conditioning and ensure the positivity of the down-sampled FIM $\tilde{\mathcal{I}}(\Xi_c)$ constructed from the current ensemble via (4.8). In our setting, natural choices

for acceptance criteria are the smallest eigenvalue $\lambda_{\min}^c$ or the inverse condition number c_{inv}^c of the down-sampled FIM $\tilde{\mathcal{I}}(\Xi_c)$. Both are motivated by Theorem 3. Furthermore, running the EKS to full convergence would be computationally prohibitive, as each iteration requires recomputing the sensitivity matrix G_c and the sketched FIM $\tilde{\mathcal{I}}(\Xi_c)$. Since sampling accuracy is not the main objective, it is sufficient to track the ensemble evolution until a favorable matrix configuration is reached. We refer to truncating the underlying EKS dynamics as early stopping.

This idea is implemented in the following Algorithm 4 like in [HKL26]. Given an evaluation criterion Q to be maximized, the ensemble $\{\xi_j\}_{j=1,\dots,c}$ is repeatedly updated via the sampling rule R for a fixed number I of iterations. The evolving ensemble is updated only when the current ensemble improves upon the evaluation criterion Q. This greedy strategy preserves favorable ensemble configurations by stopping the sampling process before the desired matrix properties get worse.

Algorithm 4: Sampling with early stopping

Input: Initial sample $\{\xi_j\}_{j=1,\dots,c}$, sample update rule $R\colon \Xi^c \to \Xi^c$, number of iterations $I > 0$, evaluation criterion $\mathsf{Q}\colon \Xi^c \to \mathbb{R}$

Output: Updated sample $\{\xi_j\}_{j=1,\dots,c}$ with improved evaluation criterion

```
for i = 1, ..., I do
    Generate candidate update: {θ_j}_{j=1,...,c} = R({ξ_j}_{j=1,...,c})
    if Q({θ_j}_{j=1,...,c}) > Q({ξ_j}_{j=1,...,c}) then
        Update sample: {ξ_j}_{j=1,...,c} ← {θ_j}_{j=1,...,c}
return Sample {ξ_j}_{j=1,...,c}
```

Pipeline for sensitivity-preserving experimental design. Combining the gradient-free EKS dynamics (5.20) with Algorithm 4 for early stopping and the FIM construction (4.8) yields the complete pipeline for sensitivity-preserving experimental design, summarized in Algorithm 5, as proposed in [HKL26]. This algorithm brings together the central concepts of this thesis. The matrix sketching theory of Chapter 3 and its translation to the FIM of Chapter 4 determine $\tilde{\mathcal{I}}(\Xi_c)$, which is being optimized. The sampling algorithm of Chapter 5 determines how the design space is explored.

Algorithm 5: Experimental Design sampling with EKS and early stopping

Input: Initial data distribution π_0, algorithm to compute $G_{\xi,:}$ for given ξ, desired number $c \ll N$ of experiments, number of iterations $I > 0$, evaluation criterion: scalarization of $\mathcal{I}(\Xi_c)$ to be maximized $q : \mathbb{R}^{K\times K} \mapsto [0, \infty)$.

Output: Ensemble $\{\xi_j\}_{j=1,\dots,c}$ with best seen evaluation criterion.

Initialize ensemble $\Xi_c = \{\theta_j\}_{j=1,\dots,c}$ with $\theta_j \sim \pi_0, j = 1, \dots, c$
Initialize best seen ensemble $\Xi_c^* \leftarrow \Xi_c$
Compute $G_c = (G_{\theta_j,:})_{\theta_j \in \Xi_c}$
Build $\tilde{\mathcal{I}}(\Xi_c)$ according to (4.8) with $\pi|_{\Xi_c} = \pi_0|_{\Xi_c}$
Initialize best evaluation criterion $q^* \leftarrow q(\tilde{\mathcal{I}}(\Xi_c))$
for *i = 1, . . . , I* **do**
    Update Ξ_c, by (5.20), using G_c
    Update $G_c = (G_{\theta_j,:})_{\theta_j \in \Xi_c}$ on the new ensemble
    Compute $\tilde{\mathcal{I}}(\Xi_c)$ according to (4.8)
    if $q(\tilde{\mathcal{I}}(\Xi_c)) > q^*$ **then**
        Update $\Xi_c^* \leftarrow \Xi_c$ and $q^* \leftarrow q(\tilde{\mathcal{I}}(\Xi_c))$
return *Best seen ensemble* Ξ_c^*

Algorithm 5 constitutes the practically executable pipeline for sensitivity-preserving experimental design. While a numerical evaluation of this framework lies beyond the scope of this thesis, it is instructive to relate it to the numerical experiments presented in [HKL26], where the same algorithm is applied to a Schrödinger potential reconstruction problem. Interestingly, the numerical experiments suggest that the down-sampling does not need to come with a practical cost. As summarized in Table 2 in [HKL26], the down-sampled FIM obtained via EKS exhibited a larger minimum eigenvalue and a better inverse condition number than the FIM computed from the full dataset. This is a somewhat counterintuitive finding. Intuitively, one would expect that using more data can only improve or at worst not harm the conditioning of the FIM. The authors attribute this effect to data dilution. Incorporating a large number of weakly informative or redundant measurements can average out the contribution of a few highly sensitive sensors, so that a carefully selected, smaller subset of informative points may in fact carry more concentrated information than the full dataset.

However, this result should be interpreted with some caution. In the synthetic benchmark of [HKL26], the ground-truth parameter p_* is known by construction, so the sensitivity matrix G and hence the interaction matrices $Y_{\xi,\theta}$ from (4.6) can be evaluated exactly at p_*. As discussed in Remark 1, this is precisely the quantity that is unavailable in a real reconstruction task, where p_* is the unknown to be inferred in

the first place. The conditioning reported in Table 2 of [HKL26] therefore reflects the performance of the method under an idealized, best-case sensitivity matrix.

5.5 Drawbacks of the Sampling Approach

In conclusion, we collect the main limitations of the sampling method developed in this thesis and explain how they are compensated for.

Dependence on unknown quantities. As noted in Remark 1, evaluating the optimal sampling distribution $\tilde{\pi}$ requires the precision matrix Γ^{-1} and the sensitivity matrix G at the unknown ground truth p_*. If the noise model is defined by a dense covariance matrix Γ, computing Γ^{-1} can be computationally prohibitive for large datasets. The presented framework relies on Γ^{-1} being either directly available or efficiently computable. The rows of G are the gradients of the forward model evaluated at the unknown ground truth p_*. Since the goal of the inverse problem is to determine p_*, this creates the classical circularity of experimental design. Determining the optimal measurement points would require already knowing the outcome of the reconstruction. Since $\tilde{\pi}$ is therefore inaccessible in practice, we require a computable proxy π_0. According to [HKL26], assuming that $G_{\xi,:}$ and $\Gamma^{-1}_{\xi,\theta}$ are uniformly bounded, a simple uniform distribution can be used as a sufficient π_0 for $\tilde{\pi}$. By Theorem 3, this choice is justified. Since $\tilde{\pi}$ is bounded by construction, for a uniform π_0 there exists a $\beta \in (0, 1]$ such that $\pi_0 \geq \beta\tilde{\pi}$. Algorithm 5 starts with this neutral distribution and uses early stopping to optimize the conditioning of the down-sampled FIM. The numerical experiments in Figure 6 in [HKL26] show that this heuristic yields robust results. For strongly nonlinear problems, [HKL26] refers to adaptive sequential strategies, in which parameter reconstruction and experimental design are updated alternately in an iterative process, see [Ale21; HJM24] for concrete instances of such schemes.

Another dependency concerns the smallest eigenvalue of the full FIM $\lambda^{\Xi}_{\min}$, which quantifies the sensitivity of the full design and enters Theorem 3 as an upper bound on the admissible error tolerance $\varepsilon \in \left(0, \lambda^{\Xi}_{\min}\right)$. Only for ε chosen within this range does Theorem 3 guarantee positive definiteness of the down-sampled FIM, and hence stable and unique recovery of p_* from the reduced measurements. Without knowledge of $\lambda^{\Xi}_{\min}$, the sample size c prescribed by (4.10) in Theorem 3 cannot be computed a priori. A second, structurally similar dependency arises through the normalization constant of $\tilde{\pi}$, $\sum_{\xi,\theta\in\Xi} \|Y_{\xi,\theta}\|_F$, which enters (4.10) quadratically. Since this sum ranges over all combinations of design pairs, evaluating it exactly becomes prohibitively expensive for fine measurement grids. Instead of computing the theoretical sample size c in advance,

in practice c is not chosen by the theoretical bound (4.10), but rather based on the parameter dimension K. It is essential that $c \geq \frac{K}{2}$ in the case of pairwise sampling, since otherwise the FIM of dimension $K \times K$ is rank-deficient and remains singular. In [HKL26], the numerical simulations use $c = 2K = 18$ sensors for $K = 9$ parameters to be estimated. Note that satisfying $c \geq \frac{K}{2}$ alone does not guarantee well-posedness. If too many of the drawn design points coincide or are linearly dependent, G_c may become rank deficient and $\tilde{\mathcal{I}}(\Xi_c)$ remains singular even when $c \geq \frac{K}{2}$. Since it is not possible to calculate in advance whether c is large enough for a stable FIM, Algorithm 5 performs a quality check through the evaluation criterion q_*. If the procedure reaches a satisfactory target value via early stopping, we can conclude that the chosen sample size c was sufficient. If the value remains low throughout all iterations, it suggests that c was chosen too small and must be increased for a new run.

Limitations of the convergence theory. As discussed in Section 5.3, whenever the potential Φ is Lipschitz-smooth, the target distribution $\tilde{\pi} \propto e^{-\Phi}$ is the unique invariant measure of the mean-field EKS dynamics that is not concentrated on a single point. The only other stationary states of the mean-field Fokker-Planck equation are Dirac measures, where the covariance matrix vanishes. This uniqueness implies that $\tilde{\pi}$ is the only smooth distribution consistent with the EKS dynamics, in the sense that the ensemble does not systematically drift away from it once it is reached. However, it does not guarantee that the ensemble converges to $\tilde{\pi}$, nor does it prevent a collapse to a Dirac state. The stronger result is exponential convergence of the mean-field PDE to the target distribution, which requires the potential Φ to be strongly convex, as stated in Section 5.3. This assumption fails for the potential $\Phi(\chi) = -\log\left(\|Y_\chi\|_F\right)$ relevant to our setting, since it is generally non-convex. Without a theoretical convergence guarantee or lower bounds on the ensemble covariance under non-convexity, it cannot be ruled out that the ensemble stagnates or fails to move toward favorable regions of the design space.

A second gap separates the convergence theory from practice. Even when the exponential convergence result applies, it is a statement about the mean-field PDE (5.15), that is, about the limiting behavior of infinitely many interacting particles $c \to \infty$. In practice, however, Algorithm 5 is run with a small, finite ensemble, such as $c = 18$ in the numerical experiments in [HKL26]. The transition from the infinite mean-field limit to a system with a few interacting particles remains an active field of research, particularly for non-Gaussian target distributions. Since non-asymptotic convergence rates are unavailable, there is no theory on how close an EKS run of Algorithm 5 comes to the target distribution after a finite number of iterations I and a finite ensemble size

c.

A third gap arises from the discrete design space Ξ. The EKS dynamics (5.19), and the invariance and convergence results established in Section 5.3, are formulated on the continuous extension $\tilde{\Xi} \supset \Xi$. However, after each Euler-Maruyama step (5.20) the updated particles $\chi_j^{t_{n+1}}$ generally do not lie on the discrete design space and must therefore be projected onto their nearest neighbor in $\Xi \times \Xi$ as mentioned in Section 5.4. This projection introduces an additional error at every iteration that is not covered by the mean-field convergence theory. The projection step could map distinct nearby particles onto the same discrete design point, which can reduce the effective diversity of the ensemble Ξ_c and degrade the empirical covariance $C(S)$ that drives the exploration of the design space.

Because of these three gaps, the algorithm should not be understood as a sampling method that generates statistically faithful samples of $\tilde{\pi}$. Instead, its primary objective is the positivity and favorable conditioning of the down-sampled FIM $\tilde{\mathcal{I}}(\Xi_c)$, not distributional fidelity to $\tilde{\pi}$. This is what motivates the early stopping strategy of Section 5.4. Rather than running the EKS dynamics to convergence, which would be computationally prohibitive and theoretically unjustified, the procedure is terminated as soon as an acceptable value of the evaluation criterion is observed.

Nevertheless, the invariance property still plays an important role in justifying the approach. That the particles are qualitatively driven toward the sensitive regions of the design space is sufficient for Theorem 3. For the local sensitivity guarantee to hold, it is sufficient that the sampling distribution satisfies $\pi \geq \beta\tilde{\pi}$ for some $\beta \in (0, 1]$, which is an incomplete or disturbed approximation of the kind produced by a finite, early-stopped EKS run.

Scope of the numerical validation in [HKL26]. As mentioned in Section 5.4 the simulations in [HKL26] make use of a known ground truth parameter p_*. In addition to the previously discussed experiment, the authors conducted several encouraging, but partial, robustness checks with respect to the sampling distribution and its normalization. In one experiment, the sensors are initialized uniformly over the design space as an information-free proxy for $\tilde{\pi}$. Another experiment uses a deliberately poor initial ensemble to test whether EKS can effectively improve the conditioning of the FIM from an unfavorable starting configuration. Furthermore, they examine a setting, where the fine discretization of the design space makes the computation of the exact global normalization constant of $\tilde{\pi}$ intractable.

However, this is not sufficient for a full validation of the method under a genuine

unknown p_*, and some aspects remain unaddressed. Most importantly the sensitivity matrix G itself is still evaluated at the true p_*. In the experiments the sampling distribution or the initial ensemble is varied, while the derivatives entering G continue to rely on exact knowledge of the ground truth p_*.

Additionally, experiments that quantify the propagation of estimation error of p_* could be conducted. Meaning how far an estimate $\hat{p}$ may deviate from p_* before the EKS dynamics drift toward the wrong region of the design space. This would demonstrate the robustness of the method in the case of a misspecified sensitivity matrix.

The authors themselves explicitly frame their simulations as a proof of concept under a known ground truth parameter. They leave the investigation of the dependence on p_* as an open research question for the future. They refer to strategies in the classical OED literature that have been developed to mitigate this drawback and are summarized under the term sequential experimental design, e.g. see [Ale21; HJM24].

6 Conclusion and Outlook

Summary. This thesis examined the problem of qualitative experimental design for inverse problems following the concept of [HKL26]. Given a large design space Ξ, we looked for a smaller subset $\Xi_c \subset \Xi$ that preserves the sensitivity properties of the full design, as quantified by the FIM. Reformulating this task as a matrix sketching problem allowed us to leverage tools from RNLA. We derived probabilistic error bounds on the approximation of the full FIM by its down-sampled counterpart in Theorem 1 and Theorem 2, and translated these bounds into explicit guarantees for the well-posedness and spectral stability of the reduced inverse problem in Theorem 3. Since the resulting optimal sampling distribution $\tilde{\pi}$ is generally intractable to evaluate exactly, we further employed a gradient-free EKS. The EKS is adapted to the discrete design setting by a nearest-neighbor projection onto $\Xi \times \Xi$. Using early stopping the conditioning of the down-sampled FIM is prioritized over distributional fidelity to $\tilde{\pi}$. Interestingly, the numerical results reported in [HKL26] show that this relaxation does not necessarily come at a practical cost. In an idealized setting, the down-sampled FIM obtained via EKS exhibited both a larger minimum eigenvalue and a better inverse condition number than the FIM computed from the full dataset. This suggests that including a large number of redundant or weakly informative data points can, in fact, dilute the information content of the full dataset. A qualitative subset sampling strategy focused on informative points counteracts this effect at a feasible computational cost.

Outlook. Directions for future work follow directly from the limitations discussed in Section 5.5.

The one with the biggest impact concerns the dependence of the sensitivity matrix G, and hence of the optimal sampling distribution $\tilde{\pi}$, on the unknown ground truth p_*, as discussed in Remark 1 and Section 5.5. A natural extension would be to embed the down-sampling pipeline of Algorithm 5 into a sequential experimental design loop following the general strategies provided in [Ale21; HJM24]. In this context, it would also be valuable to numerically quantify how sensitive the resulting FIM conditioning is to errors in the estimate $\hat{p}$, complementing the partial robustness checks already

reported in [HKL26].

Another direction concerns the theoretical guarantees for the sampling step. As discussed in Section 5.5, the convergence theory for the EKS is currently restricted to the mean-field limit $c \to \infty$ and to strongly convex potentials Φ. Establishing non-asymptotic convergence rates for finite ensembles, or extending the convergence theory to non-convex potentials, remains an open challenge, but would provide a rigorous theoretical foundation for the sampling step based on the EKS. A further open question concerns the nearest-neighbor projection used to map the continuous EKS iterates back onto the discrete design space Ξ. Quantifying the resulting projection error and its interaction with the mean-field convergence rate, would close the mentioned gap between the continuous-space theory and the application on discrete design spaces.

Finally, applying the theoretical framework presented here to other classes of inverse problems, such as those arising in medical imaging or geophysical exploration, would help demonstrate its broader applicability beyond the Schrödinger potential reconstruction setting considered in [HKL26].

Bibliography

[Ale21] Alen Alexanderian. "Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: A review". In: *Inverse Problems* 37.4 (2021), p. 043001.

[DL21] Zhiyan Ding and Qin Li. "Ensemble Kalman Sampler: Mean-Field Limit and Convergence Analysis". In: *SIAM Journal on Mathematical Analysis* 53.2 (2021), pp. 1546–1578.

[FSU19] Joel Feldman, Mikko Salo, and Gunther Uhlmann. "The Calderón problem—an introduction to inverse problems". In: *Preliminary notes on the book in preparation* (2019).

[GI+20] Alfredo Garbuno-Iñigo et al. "Interacting Langevin Diffusions: Gradient Structure and Ensemble Kalman Sampler". In: *SIAM Journal on Applied Dynamical Systems* 19.1 (2020), pp. 412–441.

[HJM24] Xun Huan, Jayanth Jagalur, and Youssef Marzouk. "Optimal experimental design: Formulations and computations". In: *Acta Numerica* 33 (2024), pp. 715–840.

[HKL26] Kathrin Hellmuth, Christian Klingenberg, and Qin Li. "Local sensitivity-preserving random data down-sampling for experimental design". In: *arXiv preprint arXiv:2409.15906* (2026).

[KS19] Nikola B. Kovachki and Andrew M. Stuart. "Ensemble Kalman Inversion: A Derivative-Free Technique for Machine Learning Tasks". In: *Inverse Problems* 35.9 (2019), p. 095005.

[LMW18] Benedict Leimkuhler, Charles Matthews, and Jonathan Weare. "Ensemble Preconditioning for Markov Chain Monte Carlo Simulation". In: *Statistics and Computing* 28.2 (2018), pp. 277–290.

[Mah16] Michael W. Mahoney. *Lecture Notes on Randomized Linear Algebra*. 2016. arXiv: 1608.04481.

[Pav14] Grigorios A. Pavliotis. *Stochastic Processes and Applications: Diffusion Processes, the Fokker–Planck and Langevin Equations*. Vol. 60. Texts in Applied Mathematics. New York: Springer, 2014.

[SWZ25] Björn Sprungk, Simon Weissmann, and Jakob Zech. “Metropolis-Adjusted Interacting Particle Sampling”. In: *Statistics and Computing* 35.3 (2025), p. 64.

[Vae24] Urbain Vaes. “Sharp Propagation of Chaos for the Ensemble Langevin Sampler”. In: *Journal of the London Mathematical Society* 110.5 (2024), e13008.