

Probabilistic Linear Algebra

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von
Simon Bartels
aus Pasewalk

Tübingen
2019

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation:

20.10.2020

Stellvertretender Dekan:

Prof. Dr. József Fortágh

1. Berichterstatter:

Prof. Dr. Philipp Hennig

2. Berichterstatter:

Prof. Dr. Matthias Hein

Probabilistic Linear Algebra

Simon Bartels
Eberhard Karls Universität Tübingen

Eidesstattliche Erklärung

Hiermit erkläre ich, dass ich, Simon Bartels, die vorliegende Arbeit selbstständig angefertigt, keine anderen als die angegebenen Hilfsmittel benutzt und alle Stellen, die dem Wortlaut oder Sinne nach anderen Werken entnommen sind, durch Angabe der Quellen als Entlehnung kenntlich gemacht habe. Die Arbeit hat in gleicher oder ähnlicher Form noch keiner anderen Prüfungsbehörde vorgelegen.

Ort, Datum

Unterschrift

Abstract

Linear algebra operations are at the core of many computational tasks. For example, evaluating the density of a multivariate normal distribution requires the solution of a linear equation system and the determinant of a square matrix. Frequently, and in particular in machine learning, the size of the involved matrices is too large to compute exact solutions, and necessitate approximation. Building upon recent work (Hennig and Kiefel 2012; Hennig 2015) this thesis considers numerical linear algebra from a probabilistic perspective.

Part [iii](#) establishes connections between approximate linear solvers and Gaussian inference, with a focus on projection methods. One result is the observation that solution-based inference (Cockayne, Oates, Ipsen, and Girolami 2018) is subsumed in the matrix-based inference perspective (Hennig and Kiefel 2012). Part [iv](#) shows how the probabilistic viewpoint leads to a novel algorithm for kernel least-squares problems. A Gaussian model over kernel functions is proposed that uses matrix-multiplications computed by conjugate gradients to obtain a low-rank approximation of the kernel. The derived algorithm *kernel machine conjugate gradients* provides empirically better approximations than conjugate gradients and, when used for Gaussian process regression, additionally provides estimates for posterior variance and log-marginal likelihood, without the need to rerun. Part [v](#) is concerned with the approximation of kernel matrix determinants. Assuming the inputs to the kernel are independent and identically distributed, a stopping condition for the Cholesky decomposition is presented that provides probably approximately correct (PAC) estimates of the log-determinant with only little overhead.

Zusammenfassung

Operationen der linearen Algebra bilden häufig die Basis vieler anderer Algorithmen und mathematischen Probleme. Beispielsweise, um die Dichte einer multivariaten Gauß-Verteilung zu berechnen, benötigt man die Lösung eines lineares Gleichungssystem und die Determinante einer quadratischen Matrix. Häufig, und insbesondere in Anwendungen des Maschinellen Lernens, ist die Größe der involvierten Matrizen zu groß, um exakte Lösungen berechnen zu können und man muss auf Approximationsmethoden zurückgreifen. Aufbauend auf kürzlich veröffentlichten Arbeiten (Hennig und Kiefel 2012; Hennig 2015), betrachtet diese Arbeit, Approximationsmethoden der linearen Algebra aus einer probabilistischen Perspektive.

Part [iii](#) zeigt Verbindungen auf zwischen Approximationsalgorithmen für lineare Gleichungssysteme und Gaußscher Inferenz, mit einem Fokus auf Projektionsmethoden. Ein Resultat ist die Beobachtung, dass lösungsfokussierte Inferenz (Cockayne, Oates, Ipsen und Girolami 2018) enthalten ist in der matrixfokussierten Perspektive (Hennig und Kiefel 2012). Part [iv](#) zeigt, wie sich die probabilistische Perspektive nutzen lässt, um neue Approximationsalgorithmen zu entwickeln, hier für die Lösungen von Normalgleichungssystemen für Kernmethoden. Mit einem speziellen Gauß-Prozess Modell über die Kernfunktion und unter Verwendung von Matrixmultiplikationen mit der Methode der konjugierten Gradienten erhält man eine näherungsweise Kernfunktion von niedrigem Rang. Der resultierende Algorithmus *kernel machine conjugate gradients* gibt empirisch bessere Approximationen als die Methode der konjugierten Gradienten, und, im Fall von Gauß-Prozess Regression, gibt zudem Schätzungen für Unsicherheit und Evidenz, ohne Neustart des Algorithmus. Part [v](#) beschäftigt sich der Approximation von Determinanten für Kernmatrizen. Unter der Annahme, dass die Argumente der Kernfunktion unabhängig und identisch verteilt sind, beschreibt dieser Teil eine Stopstrategie für die Cholesky-Dekomposition, die Näherungslösungen der Log-Determinante mit gewünschter Präzision und Wahrscheinlichkeit ausgibt, mit nur geringem Mehraufwand.

Acknowledgments

I am most grateful for the endless patience, guidance and trust of my adviser Prof. Dr. Philipp Hennig, for allowing me to deal with the ups and downs of working towards a Ph.D. in my own way. As a colleague once put it: discussing research with Philipp is like pursuing a Ferrari on a bicycle. It was an honor to write my dissertation under the supervision of such a brilliant and at the same time humble and understanding researcher.

I am thanking Prof. Dr. Matthias Hein for the time and effort spent to evaluate this thesis. Further, Prof. Dr. Kay Nieselt and Prof. Dr. Philipp Behrens have my gratitude for sparing valuable time for my defense.

I am grateful for inspiration and support at all times from the probabilistic numerics group: thank you, Edgar Klenske, Maren Mahsereci, Michael Schober, Hans Kersting, Alexandra Gessner, Lukas Balles, Filip de Roos, Motonobu Kanagawa, Frank Schneider, Matthias Werner, Felix Dangel, Susanne Zabel, Frederik Künstner, Agustinus Kristiadi, Jonathan Wenger and Nicholas Krämer.

Many people have patiently listened to me while I was rambling about the mathematical problems on the way to Part [v](#). Thank you, Gabriele Abbati, Damien Garreau, Motonobu Kanagawa, Hans Kersting, Jonas Kübler, Simon Julien-Lacoste, Krikamol Muandet, Alexander Neitz, Giambatista Parascandolo, Michael Perrot, Carl Rasmussen, Luca Rendsburg, Michael Schober, Sebastian Weichwald and everyone else I may have forgotten.

Special thanks go to Inna Zeitler for listening when I had a first idea and because it was most inconvenient at the time. Special thanks go to Maja Rudolph for pointing me to the works of David Freedman which eventually lead to discovering the theorem by Fan, Grama, and Liu (2012). Special thanks go to Damien Garreau for the most difficult first proof-read of Part [v](#).

Thank you Jon Cockayne and Ilse Ipsen for inspiring discussions about probabilistic views on linear solvers.

Thank you Pablo Garcia Moreno, Javier Gonzalez and Neil Lawrence for your supervision at Amazon Cambridge.

On my path towards this degree, I am (in chronological order) grateful for the time and effort which Prof. Karsten Wolf, Prof. Dietlinde Lau, Dr. Roland Ewald, Dr. Jan Himmelspach, Prof. Dr. Clemens Căp, Prof. Dr. Frank Hutter, Manuel Blum and Prof. Hans Rudolf Lerche invested into teaching and supervision.

Thank you Jonas Schöley for pushing me just the right way while writing this thesis.

Thank you Dr. Simon Donné for your couch the night before my defense.

Contents

i	Prologue	1
1	Introduction	3
2	Publications	5
ii	Preliminaries	7
3	Introduction	9
4	Gaussian Process Regression and Regularized Least-Squares	11
5	Classic Linear Solvers	13
5.1	Arnoldi's Procedure	13
5.2	Conjugate Gradients	14
5.3	Generalized Minimal Residual	16
6	Kronecker Calculus	19
iii	Probabilistic Linear Solvers	21
7	Preliminaries	23
7.1	Introduction	23
7.2	Projection Methods	23
7.3	Constructing Probabilistic Linear Solvers	24
8	Probabilistic Interpretation of Projection Methods	27
8.1	Matrix-based Inference and Solution-based Inference	27
8.2	Left-information Views	28
8.3	Right-information Views	35
8.4	Preconditioning	37
9	Discussion	41
9.1	Summary	41
9.2	Future Directions	41
iv	Probabilistic Solvers for Kernel Least-Squares Problems	43
10	Preliminaries	45
10.1	Introduction	45
10.2	Finite-rank Kernel	46
11	Model	47
11.1	Prior	47
11.2	Likelihood	49
11.3	Posterior	50
12	Conjugate Gradients for Kernel Machines	53
12.1	Properties	54
12.2	Related Work	56
13	Empirical Comparison of CG and KMCG	57

13.1	Common Regression Datasets	58
13.2	Grid-structured Datasets	59
14	Discussion	67
14.1	Summary	67
14.2	Future Directions	67
v	Probabilistic Kernel-Matrix Determinant Estimation	69
15	Preliminaries	71
15.1	Introduction	71
15.2	The Cholesky Decomposition	73
15.3	Martingales and Stopping Times	74
16	A Probably Approximately Correct Bound	75
16.1	Problem Definition	75
16.2	Stopping Condition	75
16.3	Discussion	78
16.4	Related Work	80
17	Proof of Theorem 27	81
17.1	Using the Stopping Conditions	83
17.2	D_N is probably close to its Expected Value	84
17.3	\mathcal{U}_τ is probably large enough	86
18	Application to Kernel Matrix Determinant Estimation	89
19	Discussion	93
19.1	Summary	93
19.2	Future Directions	93
vi	Epilogue	97
20	Conclusions	99
20.1	Discussion	99
20.2	Future Work	100
vii	Appendix	103
A	Benchmark Datasets	105
B	Gaussian Processes	107
C	Additional Material for Part ii	109
D	Additional Material for Part iii	111
D.1	Proof of Proposition 5	111
D.2	Proof of Theorem 15	112
E	Additional Material for Part iv	117
E.1	Sampling from a Gaussian with Symmetric Kronecker Covariance	117
E.2	Proofs	117
E.3	Additional Experiments and Results	120
F	Additional Material for Part v	129
	Bibliography	133

List of Figures

- 1 convergence of posterior mean (top plot) and trace of the posterior covariance matrix (bottom plot) of the probabilistic interpretation of GMRES from Proposition 11. 33
- 2 assessment of the uncertainty quantification. Plotted are kernel density estimates for the statistic Z based on 500 randomly sampled test problems for steps $m = \{1, 3, 5, 8, 10\}$. These are compared with the theoretical distribution of Z when the posterior distribution is well-calibrated. 34
- 3 the algorithm KMCG in comparison to CG on a toy setup. The dataset consists of one hundred data-points where the targets are a draw from a zero-mean Gaussian process with squared exponential kernel (Eq. (43) with $\Lambda = 0.25$ and $\theta_f = 2$). The thin, black line is the posterior mean of that Gaussian process (Eq. (1)). The light-green line is the mean prediction produced by conjugate gradients after $P = 7$ steps and the dark-red line is the mean prediction of KMCG (where the number of inducing inputs $M = N$). 46
- 4 schematic summary of the proposed kernel approximation method. 48
- 5 progression of the posterior (Eq. (37)) for KMCG on the toy example from Fig. 3 for $P = 2, 4$ and 8 conjugate gradients steps. The columns show from left to right: mean, standard deviation, standardized error (white refers to perfect calibration, green to overconfidence and red to underconfidence) and a sample. 51
- 6 progression of the relative error ϵ_f as a function of the number of iterations of CG and KMCG for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 60
- 7 progression of the relative error of the variance ϵ_{var} as a function of the number of iterations of KMCG and baseline for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 61

- 8 progression of the relative error of the evidence ϵ_{ev} as a function of the number of iterations of baseline and KMCG for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. The small spikes in the plots where KMCG appears to be close to the solution correspond to changes of the estimate from too small to too large. 62
- 9 comparison of baseline and KMCG on grid-structured datasets using the squared exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. 64
- 10 comparison of KMCG and CG on the SOUND dataset using the squared exponential kernel (Eq. (43)) with the hyper-parameters from Dong, Eriksson, Nickisch, Bindel, and Wilson (2017). The shaded area visualizes minimum and maximum over all baseline runs. 65
- 11 evolution of the log determinant $\ln |\mathbf{K}_n + 0.001\mathbf{I}|$ for common benchmark problems using the squared-exponential kernel Eq. (43). The datasets have been centered and standardized but are otherwise unprocessed. Note the different magnitudes of the y-axis. The figure demonstrates that depending on the kernel and the distribution of the dataset, there are cases in which a linear extrapolation from some $\ln |\mathbf{K}_n|$ is sufficient to estimate $\ln |\mathbf{K}_N|$ with high relative precision. More details about the datasets can be found in Appendix A. 72
- 12 the function $H_N(x)$ for different N . Increasing N requires a larger ϵ for $H_N(\epsilon)$ to fall below a certain threshold. However, in relation to N the increase is small. The thresholds $\delta/2$ and $\delta/2|S|$ become relevant for Fig. 14. 77
- 13 relative overhead of the OpenBLAS Cholesky decomposition with early stopping against without. Both algorithms compute the Cholesky decomposition of identity matrices. The nuisance parameters were set to $m := 256$ and $S := m \lfloor \{m, \dots, N - 1\} / m \rfloor$. The parameters r and δ are in this case irrelevant for the performance, since early stopping is impossible by design. With increasing matrix size, the measurement noise is decreasing and more runs remain below the one percent overhead mark. The experiments were executed on an Intel i7 CPU with 32 Gigabytes of RAM running Ubuntu 18.04. 78

- 14 visualization of the quantities of Theorem 28 for one random shuffling of the SARCOS dataset ($N = 40000$), the squared-exponential kernel (Eq. (43)) and $\sigma^2 := 0.001$. The parameters have been set to $r := 0.1, \delta := 0.25, m := 100, 1000$ and $S := \{0.75N, 0.9N\}$. The deterministic upper bound is $\mathcal{U}_n := \mathcal{L}_n + (N - n)C^+$. It may seem surprising that the upper bound for $m = 100$ is higher than for $m = 1000$. However, recall that ϵ_μ has a term C/m . In the last picture, both stopping conditions are fulfilled for $\tau = 0.9N$, for both m . This figure demonstrates that Theorem 28 is not trivial. Arguably, it is desirable to achieve early stopping for less conservative parameters than $|S| = 2$ and $r = 0.1$. Section 19.2 discusses possibilities how to obtain more practical versions of Theorem 27. 79
- 15 progression of the relative error ϵ_f over training time for different datasets using the squared-exponential kernel (Eq. 43). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 121
- 16 comparison of baseline and KMCG on grid-structured datasets using the squared exponential kernel (Eq. 43). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. It may seem surprising that the runs on the $100 \times 100 \times 100$ dataset take more than twice as long. By chance the dataset contains more extreme values in the kernel matrix, *i.e.* smaller than $1e^{-50}$. Multiplication with these elements takes more time. 122
- 17 progression of the relative error ϵ_f as a function of the number of iterations of CG and KMCG for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 123
- 18 progression of the relative error of the variance as a function of the number of iterations of KMCG and baseline for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 124
- 19 progression of the relative error of the evidence as a function of the number of iterations of baseline and KMCG for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 125
- 20 progression of the relative error ϵ_f over training time for different datasets using the Matérn $5/2$ kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. 126

progression of the relative error ϵ_f over 100 CG-steps for different datasets using the squared exponential kernel (Eq. 43), comparing CG and FOM. The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

127

List of Tables

1	overview over frequently used letters and their context	xxi
2	compact overview of all the definitions of Part v.	82
3	descriptions and sources for all datasets considered in this work.	106

List of Algorithms

- 1 Arnoldi's procedure (Saad 2003, Algorithm 6.2) 14
- 2 the conjugate gradients algorithm (Hestenes and Stiefel 1952;
Saad 2003, p. 199f), adapted in notation and presentation to this
dissertation. 15
- 3 The algorithm referred to by Proposition 15, which reproduces the
search directions from CG. 37
- 4 Kernel Machine Conjugate Gradients 54
- 5 Cholesky decomposition according to Meister (2015, p. 48). The
letters of the indices have been adapted to fit with the notation
here. A blocked version of the algorithm can be found in Golub
and Van Loan (2013, p. 170). 73

Notation

i, j, k, l, m	discrete index variables
d, D	number of dimensions
n, N	number of datapoints
$\boldsymbol{\mu}$	mean vector
$\boldsymbol{\Sigma}$	covariance matrix
\mathbf{I}	identity matrix, $(\mathbf{I})_{ij} = \delta_{ij}$
$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	\mathbf{x} is distributed multivariate normal with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$
$K_m(\mathbf{A}, \mathbf{b})$	Krylov space of order m generated by the matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ and the vector $\mathbf{b} \in \mathbb{R}^d$
$f^{-1}(\delta)$	$\arg \sup_{\epsilon \in \mathbb{R}} \{f(\epsilon) \leq \delta\}$
$\sigma(\mathbf{x}_1, \dots, \mathbf{x}_N)$	the σ -Algebra generated from random variables $\mathbf{x}_1, \dots, \mathbf{x}_N$
$\mathbf{1}_A$	indicator function for a set A

Table 1: overview over frequently used letters and their context

Vectors are denoted with small bold letters \mathbf{x} , matrices with capital bold letters \mathbf{C} . For a symmetric positive-definite (s.p.d.) matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$ and two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^d$, I write $\langle \mathbf{v}, \mathbf{w} \rangle_M = \mathbf{v}^\top \mathbf{M} \mathbf{w}$ for the inner product induced by \mathbf{M} , and $\|\mathbf{v}\|_M^2 = \langle \mathbf{v}, \mathbf{v} \rangle_M$ for the corresponding norm. If $\langle \mathbf{v}, \mathbf{w} \rangle_M = 0$ these vectors are called \mathbf{M} -conjugate.

I will slightly abuse notation to describe shifted and scaled subspaces of \mathbb{R}^d : Let \mathbf{S} be an m -dimensional linear subspace of \mathbb{R}^d with basis $\{\mathbf{s}_1, \dots, \mathbf{s}_m\}$. Then for a vector $\mathbf{v} \in \mathbb{R}^d$ and a matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, let

$$\mathbf{v} + \mathbf{M}\mathbf{S} = \text{span}(\mathbf{v} + \mathbf{M}\mathbf{s}_1, \dots, \mathbf{v} + \mathbf{M}\mathbf{s}_m).$$

Part I

Prologue

“I have heard the heartbeat of the universe. I know the answers to many questions. Ask me.” The apprentice gave him a bleary look. It was too early in the morning for it to be early in the morning. That was the only thing he currently knew for sure. “Er... what does master want for breakfast?” he said. Wen looked down on their camp and across the snowfields and purple mountains to the golden daylight creating the world, and mused upon certain aspects of humanity. “Ah,” he said. “One of the difficult ones.”

—Terry Pratchett, *Thief of Time*

Introduction

“What is the chance of an earthquake?” Stark and Freedman (2003)

In their article, Stark and Freedman (2003) discuss how to interpret the probability predicted by the U.S. Geological Survey that a large earthquake will occur within the next 30 years in the San Francisco Bay area. One concern (among many) raised in that article is that the predicted probability of 0.7 ± 0.1 relies on simulations involving numerical approximation techniques. The complaint is that the parameters for the approximation algorithms have been set *ad hoc*, such that it is not clear how reliable the outcomes of the simulations actually are.

There is no straightforward interpretation of the USGS probability forecast. Many steps involve models that are largely untestable; modeling choices often seem arbitrary. (Stark and Freedman 2003, p. 9)

Computationally intricate models for which numerical approximation is inevitable are prevalent across sciences¹. Ideally, one would like to run approximation algorithms only as long as is necessary for sufficiently accurate solutions, since computation costs time, energy and hence, money.

The problem translates into the two questions: how sensitive is the model to approximation and how stable is the approximation. To reason about the second question, the field of probabilistic numerics (Hennig, Osborne, and Girolami 2015) investigates *probabilistic numerical methods* (PNM) which return probability distributions, instead of only point-estimates.

Probabilistic interpretations provide an alternative perspective on numerical algorithms, and can also provide extensions such as the ability to exploit noisy or corrupted observations (Hennig, Osborne, and Girolami 2015; Cockayne, Oates, Sullivan, and Girolami 2017). Of particular interest are those PNM whose estimate coincides with approximation methods from classic numerical analysis. The relationship between PNM and classic methods has been explored for integration (e.g. Karvonen and Sarkka 2017), ODE-solvers (Schober, Duvenaud, and Hennig 2014; Schober, Särkkä, and Hennig 2018; Kersting, Sullivan, and Hennig 2018; Tronarp, Kersting, Särkkä, and Hennig 2019) and PDE solvers (Cockayne, Oates, Sullivan, and Girolami 2016) in some generality. Concerning linear algebra, attention has thus far been restricted to the conjugate gradient (CG) method (Hennig 2015; Cockayne, Oates, Ipsen, and Girolami 2018). CG is but a single member of

¹Roeckner, Bäuml, Bonaventura, Brokopf, Esch, and Giorgetta 2003; Arras, Knollmüller, Junklewitz, and Enßlin 2018; Nille, Toussaint, Sieglin, and Faitsch 2018.

a larger class of iterative solvers, and applicable only if the matrix A is symmetric and positive-definite. Broadening the understanding of the connections between PNMs and classic linear solvers is the concern of Part [iii](#).

Linear algebra operations appear frequently across mathematical problems: the omnipresent multivariate Gaussian probability distribution requires the evaluations of a determinant and a quadratic form, Newton’s method for root finding and second order optimization requires solving linear equation systems, to name two examples. In particular, for Gaussian process inference, the size of modern machine learning datasets makes approximation necessary, and therefore Gaussian process inference will be the main application studied in this thesis. Part [iv](#) describes a novel method arising from the probabilistic perspective, particularly tailored for kernel least-squares problems of which Gaussian process regression is an instance. Part [v](#) addresses the question how statistical structure in numerical problems can be exploited in classic numerical algorithms. The particular problem under consideration is the computation of kernel-matrix determinants.

Publications

Some of the results presented in this thesis have been devised, developed, and published in collaboration. Of relevance for this thesis are the publications

- S. Bartels and P. Hennig (2016). “Probabilistic Approximate Least-Squares.” In: *Proceedings of Artificial Intelligence and Statistics (AISTATS)* for Part [iii](#),
- S. Bartels, J. Cockayne, I. C. F. Ipsen, and P. Hennig (2019). “Probabilistic Linear Solvers: A Unifying View.” In: *Statistics and Computing* 29.6, pp. 1249–1263 also for Part [iii](#), and
- S. Bartels and P. Hennig (2019). “Conjugate Gradients for Kernel Machines.” In: *ArXiv e-prints* 1911.06048. [arXiv:1911.06048](#) for Part [iv](#).

The results of Part [v](#) have not yet been published.

The most credit for the publication “Probabilistic Approximate Least-Squares” goes to Philipp Hennig. The scientific ideas, the writing, as well as the analysis are mainly his work. My main contribution was the data generation.

“Probabilistic Linear Solvers: A Unifying View” has been initiated by Jon Cockayne and me. Scientific ideas have been contributed equally by all four authors. In particular, Propositions [8](#), [15](#) and [16](#) and Corollary [9](#) in this thesis (which are taken from the *op. cit.*) are contributions by my collaborators. In due place, I will be giving exact credit. The figures in Section [8.2.3](#) (also taken from *op. cit.*) have been generated by me, based on Jon Cockayne’s simulation study from Cockayne, Oates, Ipsen, and Girolami ([2018](#)). The analysis has been performed by the two of us. All four authors have been contributing equally to the scientific writing, everyone working in every part. The final wording is more due to Jon Cockayne and Philipp Hennig.

Part [iv](#) corresponds to the publication Bartels and Hennig ([2019](#)) which is currently under its second revision at the *Journal of Machine Learning Research* (JMLR). I was the primary author and performed the principal analysis and work. Philipp Hennig provided initial ideas, direction and supervision.

For Part [v](#), ideas, data generation, analysis and writing are my own. The results are not yet published but will be submitted to the *International Conference on Machine Learning* (ICML).

Part II

Preliminaries

Bevor ich fortfahre, will ich feierlich versichern, daß von jetzt ab das Fäkalthema so erledigt ist, wie Chopin schon auf Seite 189 war. Auch mit der Schilderung erzieherischer Maßnahmen bei militärischen Organisationen bin ich am Ende. Es könnte zu leicht der Verdacht entstehen, dieses Erzählwerk wäre antimilitaristisch oder gar abrüstungsfreundlich bzw. aufrüstungsfreundlich. O nein, es geht um Höheres, um – wie jeder unvoreingenommene Leser längst weiß – um die Liebe und um die Unschuld. Daß die Umstände, unter, die Details, mit denen ich beides hier zu schildern versuche, die Erwähnung gewisser Formationen, Organisationen, Institutionen notwendig macht, ist nicht meine Schuld, sondern die eines Schicksals, mit dem jeder hadern mag, soviel er Lust hat.

—Heinrich Böll, *Entfernung von der Truppe*

Introduction

Each part in this thesis has its own “Preliminaries” chapter containing background material specific to that part. This part contains background material that is of relevance for all or several of the parts to follow. The following elaborations assume the reader to be familiar with basic concepts of probability, such as σ -Algebra, Bayes rule and expectation. Otherwise, the work by DeGroot and Schervish (2012) is an excellent starting point. For an overview on (numerical) linear algebra, Golub and Van Loan (2013) contains most information relevant for this thesis. The definitions and properties of multivariate Gaussian random variables and Gaussian processes are fundamental to this dissertation, but likely familiar to the reader and are therefore in Appendix B.

The exposition starts with Gaussian process regression and its connection to kernel least-squares. The purpose of that chapter is to provide the reader with an example application that all parts of this thesis contribute to. Thereafter, two popular solvers for linear equation systems are introduced: the generalized minimal residual (GMRES) method and the method of conjugate gradients (CG). Part iii will provide probabilistic interpretations of these solvers and Part iv shows how to use them in combination with certain approximation methods for kernel machines. The last chapter, Chapter 6, presents the Kronecker product and its symmetric sibling which will be a major tool for Parts iii and iv.

Gaussian Process Regression and Regularized Least-Squares

One of the main applications, for all parts of this work, are machine learning algorithms involving operations with kernel matrices. Exemplary we will consider here regularized least-squares regression. Regularized least-squares is known under a variety of names such as kernel ridge regression (Hoerl and Kennard 1970), spline regression (e.g. Wahba (1990)), Kriging (e.g. Matheron (1973)) and Gaussian process (GP) regression (e.g. Rasmussen and Williams (2006)). The common principle is the estimation of a regression function from a reproducing kernel Hilbert space (RKHS) $f : \mathbb{X} \rightarrow \mathbb{R}$ over some domain \mathbb{X} that minimizes the regularized loss (Rasmussen and Williams 2006, Eq. (6.19))

$$\mathcal{L}(f) = \frac{1}{2} \|f\|_k^2 + \frac{1}{2\sigma^2} \sum_{i=1}^N (y_i - f(x_i))^2,$$

where $(x_i, y_i) \in \mathbb{X} \times \mathbb{R}$, $i = 1, \dots, N$ are observations, $\sigma^2 \in \mathbb{R}^+$ is a regularization parameter, k is the corresponding kernel and $\|\cdot\|_k$ is the RKHS norm of f .

The minimizer of this loss has a closed-form solution that coincides with the posterior mean of the Gaussian process $p(f | \mathbf{X}, \mathbf{y}) = \mathcal{GP}(f; \bar{f}, \bar{c})$ under a zero-mean prior $p(f) = \mathcal{GP}(f; 0, k)$ and likelihood $p(\mathbf{y} | f(\mathbf{X})) = \mathcal{N}(\mathbf{y}; f(\mathbf{X}), \sigma^2 \mathbf{I})$ (Kimeldorf and Wahba 1970; Wahba 1990; Rasmussen and Williams 2006):

$$\bar{f}(\mathbf{x}_*) = \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \quad (1)$$

$$\bar{c}(\mathbf{x}_*, \mathbf{x}_{**}) = k(\mathbf{x}_*, \mathbf{x}_{**}) - \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{**} \quad (2)$$

where $\mathbf{K}_{ij} = k(x_i, x_j)$, and $\mathbf{k}_{*,i} = k(\mathbf{x}_*, \mathbf{x}_i)$. For medium-size datasets,¹ the standard approach to solve Equations (1) and (2) is to compute a Cholesky decomposition \mathbf{C} (Benoit 1924) of the symmetric and positive definite (s. p. d.) $\mathbf{K} + \sigma^2 \mathbf{I}$ at a cubic cost $\mathcal{O}(N^3)$. For larger datasets, a number of approximate algorithms have been proposed that yield an approximation \hat{f} to \bar{f} in linear time.² Comparative empirical studies like those of Chalupka, Williams, and Murray (2013) or Quiñero-Candela and Rasmussen (2005) indicate that some of these methods can provide good approximations in a reasonable amount of time, although there is no conclusive ‘best practice’ among these choices.

The machine learning community tends to prefer the methods above over the linear solvers that will be introduced in Chapter 5. In compar-

¹ For currently modern machines this means roughly $N \sim 5 \cdot 10^4$ observations.

² Zhu, Williams, Rohwer, and Morciniec 1998; Csató and Opper 2002; Snelson and Ghahramani 2007; Walder, Kim, and Schölkopf 2008; Rahimi and Recht 2009; Titsias 2009; Lázaro-Gredilla, Quiñero-Candela, Rasmussen, and Figueiras-Vidal 2010; Yan and Qi 2010; Le, Sarlos, and Smola 2013; Solin and Särkkä 2014; Wilson and Nickisch 2015; Hensman, Durrande, and Solin 2018.

ision the latter are computationally more expensive, and a linear solver needs to run again for new test inputs when computing the posterior uncertainty (Eq. (2)). Furthermore, Gaussian process regression often requires the evaluation of the log marginal likelihood

$$\ln p(\mathbf{y}) = -\frac{1}{2}\mathbf{y}^\top(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y} + \frac{1}{2}\ln |2\pi(\mathbf{K} + \sigma^2\mathbf{I})|^{-1}, \quad (3)$$

e.g. for model selection and linear solvers do not provide estimates for the determinant, off-the-shelf.³

³ Conjugate gradients can be used to estimate $|\mathbf{K}|$ (Filippone and Engler 2015), yet also requiring several runs.

Classic Linear Solvers

Consider the linear equation system

$$\mathbf{A}\mathbf{x}^* = \mathbf{b} \quad (\text{LES})$$

where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is no longer s.p.d. but more generally, an invertible matrix, $\mathbf{b} \in \mathbb{R}^d$ is a given vector and $\mathbf{x}^* \in \mathbb{R}^d$ is an unknown to be determined. The problem can be solved exactly using for example Gaussian elimination. The costs for solving such systems in general scale as $\mathcal{O}(d^3)$. This thesis is concerned with linear equation systems so large that running exact algorithms is precluded.

Two approximation algorithms will be of particular interest: the conjugate gradient (CG) method (Hestenes and Stiefel 1952) and the generalized minimal residual (GMRES) method (Saad and Schultz 1986). The following presentation follows Saad (2003). Both methods can be derived from Arnoldi's procedure as outlined below.

5.1 Arnoldi's Procedure

Arnoldi's procedure (Saad 2003, Section 6.3) is used to construct orthonormal bases for Krylov spaces $K_m(\mathbf{A}, \mathbf{r}_0) = \text{span}(\mathbf{r}_0, \mathbf{A}\mathbf{r}_0, \dots, \mathbf{A}^{m-1}\mathbf{r}_0)$ of general, nonsingular matrices \mathbf{A} and vectors \mathbf{r}_0 . Starting with $\mathbf{q}_1 = \mathbf{r}_0 / \|\mathbf{r}_0\|_2$, for some vector \mathbf{r}_0 , at each iteration j , the algorithm multiplies the previous \mathbf{q}_{j-1} by \mathbf{A} and then orthonormalizes the resulting vector $\mathbf{w}_j = \mathbf{A}\mathbf{q}_{j-1}$ against all previous \mathbf{q}_i , using the Gram-Schmidt procedure (cf. Algorithm 1).

Define

$$\mathbf{Q}_m := [\mathbf{q}_1 \quad \dots \quad \mathbf{q}_m] \in \mathbb{R}^{d \times m},$$

$$h_{ij} := \langle \mathbf{A}\mathbf{q}_j, \mathbf{q}_i \rangle \quad 1 \leq i \leq j \leq m \text{ and}$$

$$h_{j+1,j} := \|\mathbf{w}_j\|.$$

The basis vectors satisfy the relations

$$\mathbf{Q}_{m+1}\tilde{\mathbf{H}}_m = \mathbf{A}\mathbf{Q}_m = \mathbf{Q}_m\mathbf{H}_m + h_{m+1,m}\mathbf{q}_{m+1}\mathbf{e}_m^\top \text{ and} \quad (4)$$

$$\mathbf{Q}_m^\top \mathbf{A}\mathbf{Q}_m = \mathbf{H}_m, \quad (5)$$

where the *upper Hessenberg* matrix \mathbf{H}_m is defined as

$$\mathbf{H}_m = \begin{bmatrix} h_{11} & h_{12} & h_{13} & \dots & h_{1,m-1} & h_{1m} \\ h_{21} & h_{22} & h_{23} & \dots & h_{2,m-1} & h_{2m} \\ 0 & h_{32} & h_{33} & \dots & h_{3,m-1} & h_{3m} \\ \vdots & 0 & h_{43} & \dots & h_{4,m-1} & h_{4m} \\ \vdots & & \ddots & \ddots & \vdots & \vdots \\ 0 & \dots & \dots & 0 & h_{m,m-1} & h_{mm} \end{bmatrix} \in \mathbb{R}^{m \times m}$$

and

$$\tilde{\mathbf{H}}_m = \begin{bmatrix} \mathbf{H}_m \\ h_{m+1,m} \mathbf{e}_m^\top \end{bmatrix} \in \mathbb{R}^{(m+1) \times m}.$$

Algorithm 1: Arnoldi's procedure
(Saad 2003, Algorithm 6.2)

```

1  $\mathbf{r}_0 \leftarrow \mathbf{b} - \mathbf{A}\mathbf{x}_0$ ,  $\beta \leftarrow \|\mathbf{r}_0\|_2$ ,  $\mathbf{q}_1 \leftarrow \mathbf{r}_0/\beta$ 
2 for  $j = 1, \dots, m$  do
3    $h_{ij} \leftarrow \langle \mathbf{A}\mathbf{q}_j, \mathbf{q}_i \rangle$ 
4    $\mathbf{w}_j \leftarrow \mathbf{A}\mathbf{q}_j - \sum_{i=1}^j h_{ij}\mathbf{q}_i$ 
5    $h_{j+1,j} \leftarrow \|\mathbf{w}_j\|_2$ 
6   if  $h_{j+1,j} = 0$  then
7     | Stop
8   end if
9    $\mathbf{q}_{j+1} \leftarrow \mathbf{w}_j/h_{j+1,j}$ 
10 end for
11 Define  $\tilde{\mathbf{H}}_m \in \mathbb{R}^{(m+1) \times m}$  with elements  $h_{ij}$ 

```

5.2 Conjugate Gradients

For general matrices \mathbf{A} , conjugate gradients is known as Arnoldi's method for linear systems, or Full Orthogonalization Method (FOM) (Saad 2003, p. 165). For some initial guess \mathbf{x}_0 , FOM computes the iterate

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{Q}_m \mathbf{c}_m$$

where \mathbf{c}_m is chosen s.t. the residual $\mathbf{r}_m := \mathbf{b} - \mathbf{A}\mathbf{x}_m$ is orthogonal to \mathbf{Q}_m . This implies

$$\begin{aligned} \mathbf{0} &= \mathbf{r}_m^\top \mathbf{Q}_m \\ &= \mathbf{b}^\top \mathbf{Q}_m - \mathbf{x}_m^\top \mathbf{A}^\top \mathbf{Q}_m \\ &= \mathbf{b}^\top \mathbf{Q}_m - \mathbf{x}_0^\top \mathbf{A}^\top \mathbf{Q}_m - \mathbf{c}_m^\top \mathbf{Q}_m^\top \mathbf{A}^\top \mathbf{Q}_m \\ &\quad // \text{definition of } \mathbf{x}_m \end{aligned} \tag{6}$$

$$\begin{aligned}
&= \mathbf{r}_0^\top \mathbf{Q}_m - \mathbf{c}_m^\top \mathbf{H}_m^\top \\
&\quad // \text{definition of } \mathbf{H}_m \\
&= \|\mathbf{r}_0\|_2 \mathbf{e}_1 - \mathbf{c}_m^\top \mathbf{H}_m^\top \\
&\quad // \text{by choice of } \mathbf{q}_1
\end{aligned}$$

and hence, $\mathbf{c}_m = \|\mathbf{r}_0\|_2 \mathbf{H}_m^{-1} \mathbf{e}_1$. For \mathbf{x}_m this implies

$$\begin{aligned}
\mathbf{x}_m &= \mathbf{x}_0 + \|\mathbf{r}_0\|_2 \mathbf{Q}_m^\top \mathbf{H}_m^{-1} \mathbf{e}_1, \text{ or} \\
&= \mathbf{x}_0 + \mathbf{Q}_m (\mathbf{Q}_m^\top \mathbf{A} \mathbf{Q}_m)^{-1} \mathbf{Q}_m^\top \mathbf{r}_0 \quad (\text{FOM}) \\
&\quad // \text{when stopping simplification at Eq. (6)}
\end{aligned}$$

where Eq. (FOM) will be relevant in Section 7.2. After running m steps of Arnoldi's procedure, solving the $m \times m$ linear equation system is possible in $\mathcal{O}(m^2)$ (Golub and Van Loan 2013, p. 179) s.t. \mathbf{x}_m can be evaluated in $\mathcal{O}(d + m^2)$.

If the matrix \mathbf{A} is symmetric and positive definite, the Hessenberg matrix $\mathbf{H}_m = \mathbf{Q}_m^\top \mathbf{A} \mathbf{Q}_m$ is tridiagonal which allows to simplify Arnoldi's procedure, then called symmetric Lanczos algorithm, and FOM can be simplified to conjugate gradients.

Conjugate gradients is the method that produces in each step the minimizer \mathbf{x}_m of the function $\phi(\mathbf{x}) := 1/2 \mathbf{x}^\top \mathbf{A} \mathbf{x} - \mathbf{x}^\top \mathbf{b}$ where $\mathbf{x} \in \mathbf{x}_0 + K_m(\mathbf{A}, \mathbf{r}_0)$ (Nocedal and Wright 1999, Section 5) and shown in Algorithm 2. Alternatively \mathbf{x}_m can be written as

$$\mathbf{x}_m = \arg \min_{\mathbf{x} \in \mathbf{x}_0 + K_m(\mathbf{A}, \mathbf{r}_0)} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_{\mathbf{A}},$$

which is interesting for a comparison to GMRES, following soon (*c.f.* Eq. (8)).

```

1   $\mathbf{r}_0 \leftarrow \mathbf{b} - \mathbf{A} \mathbf{x}_0$  // The initial residual ...
2   $\mathbf{s}_0 \leftarrow \mathbf{r}_0$  // ... is the first search direction.
3   $i \leftarrow 0$ 
4  while  $\|\mathbf{r}_i\|_2 > \epsilon$  do
5  |    $\mathbf{z}_i \leftarrow \mathbf{A} \mathbf{s}_i$  // the most expensive step:  $\mathcal{O}(d^2)$  matrix-multiplication
6  |    $\alpha_i \leftarrow \frac{\mathbf{r}_i^\top \mathbf{r}_i}{\mathbf{s}_i^\top \mathbf{z}_i}$  // optimal linesearch along  $\mathbf{s}_i$  for  $\phi(\mathbf{x}) := \mathbf{x}^\top \mathbf{A} \mathbf{x} - 2 \mathbf{x}^\top \mathbf{b}$ 
7  |    $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i + \alpha_i \mathbf{s}_i$  // update to the solution
8  |    $\mathbf{r}_{i+1} \leftarrow \mathbf{r}_i - \alpha_i \mathbf{z}_i$  // analogue update to the residual
9  |    $\mathbf{s}_{i+1} \leftarrow \mathbf{r}_{i+1} + \frac{\mathbf{r}_{i+1}^\top \mathbf{r}_i}{\mathbf{r}_i^\top \mathbf{r}_i} \mathbf{s}_i$  // Gram-Schmidt applied to the new residual
10 |    $i \leftarrow i + 1$ 
11 end while
12 return  $\mathbf{x}_i$ 

```

Algorithm 2: the conjugate gradients algorithm (Hestenes and Stiefel 1952; Saad 2003, p. 199f), adapted in notation and presentation to this dissertation.

The transition from FOM to CG takes the pages 196 to 200 in the book by Saad (2003). The process itself is not relevant for this work. Relevant is the fact that CG is a special case of FOM, when considering

probabilistic interpretations of conjugate gradients in Sections 8.2.1 and 8.3.2. Here, we will just show that the two algorithms produce the same solutions in each step.

Proposition 1. *If A is symmetric and positive definite, FOM and CG produce the same solutions in each step.*

Proof. Denote with $\mathbf{s}_0, \dots, \mathbf{s}_{m-1}$ the conjugate gradients directions which span the vector space $K_m(\mathbf{A}, \mathbf{r}_0)$ (Nocedal and Wright 1999, Theorem 5.3). The proof of Theorem 5.2 by Nocedal and Wright (1999, p. 106) states that $\mathbf{x}_m = \arg \min_{\mathbf{x} \in \mathbf{x}_0 + K_m(\mathbf{A}, \mathbf{r}_0)} \phi(\mathbf{x})$ iff $\mathbf{r}_m^\top \mathbf{s}_i$ for $i = 0, \dots, m-1$. By construction through Arnoldi's procedure, \mathbf{Q}_m forms a basis of $K_m(\mathbf{A}, \mathbf{r}_0)$, and by definition of FOM $\mathbf{Q}_m^\top \mathbf{r}_m = \mathbf{0}$. \square

5.3 Generalized Minimal Residual

Another method derived from the Arnoldi procedure is the *Generalized Minimal Residual Method* (Saad 2003, Section 6.5). GMRES computes the iterate

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{Q}_m \mathbf{c}_m$$

where \mathbf{c}_m is chosen to satisfy the optimality condition in Eq. (7).

$$\|\mathbf{r}_m\|_2 = \min_{\mathbf{x} \in K_m(\mathbf{A}, \mathbf{r}_0)} \|\mathbf{A}\mathbf{x} - \mathbf{r}_0\|_2 \quad (7)$$

$$= \min_{\mathbf{x} \in \mathbf{x}_0 + K_m(\mathbf{A}, \mathbf{r}_0)} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2. \quad (8)$$

That is, at iteration m , GMRES minimises the residual over the vector space $\mathbf{x}_0 + K_m(\mathbf{A}, \mathbf{r}_0)$. Hence,

$$\begin{aligned} \mathbf{c}_m &= \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|\mathbf{A}\mathbf{Q}_m \mathbf{c} - \mathbf{r}_0\|_2 \quad (9) \\ &= ((\mathbf{A}\mathbf{Q}_m)^\top (\mathbf{A}\mathbf{Q}_m))^{-1} (\mathbf{A}\mathbf{Q}_m)^\top \mathbf{r}_0. \end{aligned}$$

and

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{Q}_m (\mathbf{Q}_m^\top \mathbf{A}^\top \mathbf{A} \mathbf{Q}_m)^{-1} \mathbf{Q}_m^\top \mathbf{A}^\top \mathbf{r}_0, \quad (\text{GMR})$$

where the above equation will be relevant in Section 7.2. The least-squares problem in Eq. (9) can be solved exactly in $\mathcal{O}(d + m^3)$ but GMRES solves it more efficiently via Arnoldi's method. To this end, express the starting vector in the Krylov basis,

$$\mathbf{r}_0 = \|\mathbf{r}_0\|_2 \mathbf{q}_1 = \|\mathbf{r}_0\|_2 \mathbf{Q}_{m+1} \mathbf{e}_1,$$

and exploit the Arnoldi recursion from Eq. (4),

$$\mathbf{A}\mathbf{Q}_m \mathbf{c} - \mathbf{r}_0 = \mathbf{Q}_{m+1} (\tilde{\mathbf{H}}_{m+1} \mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1),$$

followed by the unitary invariance of the two-norm,

$$\|\mathbf{A}\mathbf{Q}_m\mathbf{c} - \mathbf{r}_0\|_2 = \|\tilde{\mathbf{H}}_m\mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1\|_2.$$

Thus, instead of solving the least squares problem Equation (9), GMRES solves

$$\mathbf{c}_m = \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|\tilde{\mathbf{H}}_m\mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1\|_2. \quad (10)$$

Solving above systems costs only $\mathcal{O}(m^2)$ (Golub and Van Loan 2013, p. 179) instead of the $\mathcal{O}(m^3)$ for the naive solution of Eq. (9).

6

Kronecker Calculus

The Kronecker product will be a vital component throughout Parts [iii](#) and [iv](#). The Kronecker product and its symmetric version have been studied, among others, by Loan (2000) and Magnus and Neudecker (1980). The definitions used in this work slightly differ from the authors above and instead follow Hennig (2015).

The Kronecker product for two arbitrary matrices $\mathbf{A} \in \mathbb{R}^{N_1 \times N_2}$, $\mathbf{B} \in \mathbb{R}^{N_3 \times N_4}$ is defined as

$$[\mathbf{A} \otimes \mathbf{B}]_{ij,kl} := \mathbf{A}_{ik} \mathbf{B}_{jl}$$

where $i \in \{1, \dots, N_1\}$, $j \in \{1, \dots, N_3\}$, $k \in \{1, \dots, N_2\}$ and $l \in \{1, \dots, N_4\}$, and ij is not a product but a double-index. The following identities about Kronecker products and the vectorization operator can be found in Hennig and Kiefel (2013), and are restated here for the convenience of the reader:

$$(\mathbf{A} \otimes \mathbf{B}) \text{vec}(\mathbf{C}) = \text{vec}(\mathbf{ACB}^\top) \quad (\text{K1})$$

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD}) \quad (\text{K2})$$

$$(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1} \quad (\text{K3})$$

$$(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top \quad (\text{K4})$$

$$(\mathbf{A} + \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes \mathbf{C} + \mathbf{B} \otimes \mathbf{C} \quad (\text{K5})$$

where¹ $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D} \in \mathbb{R}^{N \times N}$, and \mathbf{A} and \mathbf{B} are assumed to be invertible. An appealing property of Kronecker-structured matrices is their interaction with vectorized matrices. For a square matrix $\mathbf{A} = \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_N \end{bmatrix}^\top \in \mathbb{R}^{N \times N}$, the operator $\text{vec}(\cdot) : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N^2}$ stacks the rows of \mathbf{A} into one vector:

$$\text{vec}(\mathbf{A}) = \begin{bmatrix} \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_N \end{bmatrix}, \quad \text{with} \quad [\text{vec}(\mathbf{A})]_{(ij)} = [\mathbf{A}]_{ij}$$

and $\text{mat}(\cdot) : \mathbb{R}^{N^2} \rightarrow \mathbb{R}^N \times \mathbb{R}^N$ transforms an N^2 vector into an $N \times N$ matrix, s.t. $\text{mat}(\text{vec}(\mathbf{A})) = \mathbf{A}$. A vector product of vectorized matrices corresponds to the trace of their product:

$$\text{vec}(\mathbf{A})^\top \text{vec}(\mathbf{B}) = \text{tr} \mathbf{AB}^\top. \quad (\text{V1})$$

¹The conditions can be more general but for ease of exposition, we assume all matrices are square and of equal size.

Proof.

$$\operatorname{tr} \mathbf{A} \mathbf{B}^\top = \sum_i [\mathbf{A} \mathbf{B}^\top]_{ii} = \sum_{i,j} \mathbf{A}_{ij} \mathbf{B}_{ji}^\top = \sum_{i,j} \mathbf{A}_{ij} \mathbf{B}_{ij} = \operatorname{vec}(\mathbf{A})^\top \operatorname{vec}(\mathbf{B})$$

□

The *symmetric* Kronecker product for two square matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{N \times N}$ of equal size is defined as

$$\mathbf{A} \otimes \mathbf{B} := \mathbf{\Gamma}_N(\mathbf{A} \otimes \mathbf{B}) \mathbf{\Gamma}_N$$

where $[\mathbf{\Gamma}_N]_{ij,kl} := 1/2 \delta_{ik} \delta_{jl} + 1/2 \delta_{il} \delta_{jk}$ satisfies

$$\mathbf{\Gamma} \operatorname{vec}(\mathbf{C}) = 1/2 \operatorname{vec}(\mathbf{C}) + 1/2 \operatorname{vec}(\mathbf{C}^\top)$$

for all square-matrices $\mathbf{C} \in \mathbb{R}^{N \times N}$. Equivalently, one can write

$$(\mathbf{A} \otimes \mathbf{B})_{ij,kl} = \frac{1}{4} \left(\mathbf{A}_{ik} \mathbf{B}_{jl} + \mathbf{A}_{il} \mathbf{B}_{jk} + \mathbf{B}_{ik} \mathbf{A}_{jl} + \mathbf{B}_{il} \mathbf{A}_{jk} \right).$$

The symmetric Kronecker product inherits some of the desirable properties of the Kronecker product. Some of the following identities can, again, be found in Hennig (2015), some are due to Loan (2000) and Magnus and Neudecker (1980) and some are novel. The proof gives exact credit.

Proposition 2. *Let $\mathbf{V}, \mathbf{W} \in \mathbb{R}^{N \times N}$ be square matrices and $\mathbf{A}^\top, \mathbf{B} \in \mathbb{R}^{N \times M}$ be rectangular.*

$$\mathbf{W} \otimes \mathbf{W} = \mathbf{\Gamma}_N(\mathbf{W} \otimes \mathbf{W}) \quad (\text{SK1})$$

$$\mathbf{\Gamma}_M(\mathbf{A} \otimes \mathbf{A}) = (\mathbf{A} \otimes \mathbf{A}) \mathbf{\Gamma}_N \quad (\text{SK2})$$

$$\mathbf{V} \otimes \mathbf{W} = \mathbf{W} \otimes \mathbf{V} \quad (\text{SK3})$$

$$(\mathbf{A} \otimes \mathbf{A})(\mathbf{W} \otimes \mathbf{W})(\mathbf{B} \otimes \mathbf{B}) = (\mathbf{A} \mathbf{W} \mathbf{B}) \otimes (\mathbf{A} \mathbf{W} \mathbf{B}) \quad (\text{SK4})$$

$$\mathbf{W} \otimes \mathbf{W} - \mathbf{V} \otimes \mathbf{V} = (\mathbf{W} + \mathbf{V}) \otimes (\mathbf{W} - \mathbf{V}) \quad (\text{SK5})$$

$$(\mathbf{W} \otimes \mathbf{W})^{-1} = (\mathbf{W}^{-1} \otimes \mathbf{W}^{-1}). \quad (\text{SK6})$$

The interpretation of Eq. (SK6) requires some care: symmetric Kronecker product matrices are rank deficient. Eq. (SK6) is to be read in the sense that for symmetric $\mathbf{Y} \in \mathbb{R}^{N \times N}$, i.e. $\mathbf{Y} = \mathbf{Y}^\top$, $\mathbf{X} := \operatorname{mat} \left((\mathbf{W}^{-1} \otimes \mathbf{W}^{-1}) \operatorname{vec}(\mathbf{Y}) \right)$ satisfies $\operatorname{vec}(\mathbf{Y}) = (\mathbf{W} \otimes \mathbf{W}) \operatorname{vec}(\mathbf{X})$ and \mathbf{X} is the unique symmetric solution.

The proof is part of Appendix C.

Part III

Probabilistic Linear Solvers

Spätestens hier wird der kluge Leser wissen, was wir, er und ich, nun auch dem weniger klugen Leser nicht länger vorenthalten sollten: daß dieses Erzählwerk wirklich eine reine Idylle werden soll, in der Kloakendüfte dieselbe Funktion haben wie anderswo Rosendüfte, in der die Auseinandersetzung mit dem Krieg vermieden oder zumindest sehr reduziert wird, die Nazi-Angelegenheit wie etwas zwischen Schnupfen und Schwefelregen abgetan werden soll, [...]

—Heinrich Böll, *Entfernung von der Truppe*

Preliminaries

The first part of this thesis is concerned with the question how classic linear solvers are connected to probabilistic numerical methods (PNM). The motivation to answer this question is to gain a deeper understanding of linear solvers; which assumptions are implicitly encoded and under which circumstances, which solvers are preferable.

7.1 Introduction

The first publications, that show how to construct probabilistic linear solvers are those by Hennig and Kiefel (2012), Hennig and Kiefel (2013) and Hennig (2015). Their focus is mainly on the connection between Gaussian inference and optimization—the only linear solver analyzed is the conjugate gradient (CG) method. Later Cockayne, Oates, Ipsen, and Girolami (2018) proposed a (seemingly) different approach, also with a focus on CG. CG is but a single member of a larger class of iterative solvers, and applicable only if the matrix \mathbf{A} is symmetric and positive-definite. This part is based on the publication “Probabilistic Linear Solvers: A Unifying View” (Bartels, Cockayne, Ipsen, and Hennig 2019) which explored the connection for a larger class of solvers, in particular, projection methods. As a step towards a probabilistic understanding of projection methods, the goal of that publication was, for a given projection method, to find possible prior assumptions and information¹ s.t. the posterior mean estimate coincides with the solution of the projection method.

Projection methods will be introduced in Section 7.2. Section 7.3 presents existing approaches to construct probabilistic linear solvers. Thereafter, Chapter 8 shows connections between projection methods and probabilistic linear solvers.

“Probabilistic Linear Solvers: A Unifying View” is a joint publication, and Propositions 8, 15 and 16 and Corollary 9 are contributions by my collaborators, which I will point out again in due time.

7.2 Projection Methods

Recall Eq. (LES) which describes the system of linear equations

$$\mathbf{Ax}^* = \mathbf{b} \quad (\text{LES})$$

¹The meaning of “information” will be defined properly Section 7.3.

where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is an invertible matrix, $\mathbf{b} \in \mathbb{R}^d$ is a given vector and $\mathbf{x}^* \in \mathbb{R}^d$ is an unknown to be determined. Denote with \mathbf{x}_0 an initial guess and with $\mathbf{r}_m := \mathbf{b} - \mathbf{A}\mathbf{x}_m$ the residual for an approximation \mathbf{x}_m .

Many iterative methods for linear systems belong to the class of projection methods (Saad 2003, p. 130f.), including popular linear solvers such as the conjugate gradients (CG) (see Section 5.2) method and the generalized minimal residual (GMRES) method (see Section 5.3). Saad describes a projection method as an iterative scheme in which, at each iteration, a solution vector \mathbf{x}_m is constructed by projecting \mathbf{x}^* into a *solution space* $\mathbb{X}_m \subset \mathbb{R}^d$, subject to the restriction that the residual $\mathbf{r}_m = \mathbf{b} - \mathbf{A}\mathbf{x}_m$ is orthogonal to a *constraint space* $\mathbb{U}_m \subset \mathbb{R}^d$.

More formally, each iteration of a projection method is defined by two matrices $\mathbf{X}_m, \mathbf{U}_m \in \mathbb{R}^{d \times m}$, and by a starting point \mathbf{x}_0 . The matrices \mathbf{X}_m and \mathbf{U}_m each encode the solution and constraint spaces as $\mathbb{X}_m = \text{range}(\mathbf{X}_m)$ and $\mathbb{U}_m = \text{range}(\mathbf{U}_m)$. The projection method then constructs \mathbf{x}_m as $\mathbf{x}_m = \mathbf{x}_0 + \mathbf{X}_m\boldsymbol{\alpha}_m$ with $\boldsymbol{\alpha}_m \in \mathbb{R}^m$ determined by the constraint $\mathbf{U}_m^\top \mathbf{r}_m = \mathbf{0}$. If $\mathbf{U}_m^\top \mathbf{A}\mathbf{X}_m$ is nonsingular, one obtains

$$\begin{aligned} \boldsymbol{\alpha}_m &= (\mathbf{U}_m^\top \mathbf{A}\mathbf{X}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0, \text{ and thus} \\ \mathbf{x}_m &= \mathbf{x}_0 + \mathbf{X}_m (\mathbf{U}_m^\top \mathbf{A}\mathbf{X}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0. \end{aligned} \quad (\mathbf{P})$$

Eq. (P) is tagged with a special letter as it will be referenced frequently in Chapter 8 when examining probabilistic interpretations of projection methods. Observe that another way to express Eq. (P) is

$$\begin{aligned} \mathbf{x}_m &= \mathbf{x}_0 + \mathbf{P}(\mathbf{x}^* - \mathbf{x}_0), \text{ where} \\ \mathbf{P} &:= \mathbf{X}_m (\mathbf{U}_m^\top \mathbf{A}\mathbf{X}_m)^{-1} \mathbf{U}_m^\top \mathbf{A}, \end{aligned}$$

and note further that \mathbf{P} is idempotent, *i.e.* $\mathbf{P}\mathbf{P} = \mathbf{P}$. Thus \mathbf{P} is a linear projection, hence the name, projection method.

Differing from the presentation of FOM/CG and GMRES in Chapter 5, in the projection method perspective, both algorithms perform only a single step with the size of the Krylov subspace m fixed and determined in advance. For FOM the spaces are $\mathbb{U}_m = \mathbb{X}_m = K_m(\mathbf{A}, \mathbf{b})$, while for GMRES they are $\mathbb{X}_m = K_m(\mathbf{A}, \mathbf{b})$ and $\mathbb{U}_m = \mathbf{A}K_m(\mathbf{A}, \mathbf{b})$ (compare Eq. (P) with Eqs. (FOM) and (GMR) on page 15 and page 16, respectively).

7.3 Constructing Probabilistic Linear Solvers

Hennig (2015) and Cockayne, Oates, Ipsen, and Girolami (2018) proposed two seemingly different approaches to probabilistic linear solvers. In the matrix-based inference (MBI) approach of Hennig (2015), a probability measure is constructed on the matrix \mathbf{A}^{-1} , while the solution-

based inference (SBI) method of Cockayne, Oates, Ipsen, and Girolami (2018) constructs a measure on the solution vector \mathbf{x}^* .

7.3.1 Solution-Based Inference

To phrase the solution of Eq. (LES) as a form of probabilistic inference, Cockayne, Oates, Ipsen, and Girolami (2018) consider a Gaussian prior over the solution \mathbf{x}^* , and condition on observations provided by a set of *search directions* $\mathbf{s}_1, \dots, \mathbf{s}_m$, $m < d$. Let $\mathbf{S}_m \in \mathbb{R}^{d \times m}$ be given by $\mathbf{S}_m := [\mathbf{s}_1, \dots, \mathbf{s}_m]$, and let information be given by $\mathbf{y}_m := \mathbf{S}_m^\top \mathbf{A} \mathbf{x}^* = \mathbf{S}_m^\top \mathbf{b}$. Since the information is a linear projection of \mathbf{x}^* , the posterior distribution is a Gaussian distribution on \mathbf{x}^* (*c.f.* Lemma 35 in Appendix B):

Lemma 3 (Cockayne, Oates, Ipsen, and Girolami (2018)). *Assume that the columns of \mathbf{S}_m are linearly independent. Consider the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0).$$

The posterior from SBI is then given by

$$p(\mathbf{x} \mid \mathbf{y}_m) = \mathcal{N}(\mathbf{x}; \mathbf{x}_m, \boldsymbol{\Sigma}_m)$$

where

$$\begin{aligned} \mathbf{x}_m &= \mathbf{x}_0 + \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m (\mathbf{S}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \mathbf{r}_0 \\ \boldsymbol{\Sigma}_m &= \boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m (\mathbf{S}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \boldsymbol{\Sigma}_0, \end{aligned} \quad (\text{SBI})$$

and $\mathbf{r}_0 = \mathbf{b} - \mathbf{A} \mathbf{x}_0$.

The naive evaluation of Eq. (SBI) requires $\mathcal{O}(m^3)$ floating point operations, due to the inversion of $\mathbf{Y}_m^\top \mathbf{W}_0 \mathbf{Y}_m$, which is still less than for the original problem Eq. (LES). If \mathbf{S}_m is chosen $\mathbf{A}^\top \boldsymbol{\Sigma}_0 \mathbf{A}$ -orthogonal, *e.g.* by applying Gram-Schmidt orthogonalization, the solution can be updated progressively from the last guess, avoiding the $\mathcal{O}(m^3)$ costs in each step.

7.3.2 Matrix-Based Inference

In contrast to SBI, the MBI approach of Hennig (2015) treats the matrix inverse \mathbf{A}^{-1} as the unknown in the inference procedure. As in the previous section, search directions \mathbf{S}_m yield matrix-vector products $\mathbf{Y}_m \in \mathbb{R}^{d \times m}$. In Hennig (2015) these arise from *right-multiplying*² \mathbf{A} with \mathbf{S}_m , *i.e.* $\mathbf{Y}_m = \mathbf{A} \mathbf{S}_m$. Note that

$$\mathbf{S}_m = \mathbf{A}^{-1} \mathbf{Y}_m, \text{ or, equivalently } \text{vec}(\mathbf{S}_m) = (\mathbf{I} \otimes \mathbf{Y}_m^\top) \text{vec}(\mathbf{A}^{-1}). \quad (11)$$

²This work also considers a model class that explicitly encodes *symmetry* of \mathbf{A} , such that the distinction between left- and right- multiplication vanishes. Proposition 15 in Section 8.3.2 will make use of this model class.

Thus \mathbf{S}_m is a linear transformation of \mathbf{A}^{-1} and the posterior is again distributed Gaussian:

Lemma 4 (Lemma 2.1 in Hennig (2015)). *Consider the prior*

$$p\left(\text{vec}\left(\mathbf{A}^{-1}\right)\right) = \mathcal{N}\left(\text{vec}\left(\mathbf{A}^{-1}\right); \text{vec}\left(\mathbf{A}_0^{-1}\right), \boldsymbol{\Sigma}_0 \otimes \mathbf{W}_0\right). \quad (12)$$

Then the posterior given the observations $\text{vec}\left(\mathbf{S}_m\right) = \mathbf{A}^{-1}\mathbf{Y}_m$ is given by

$$p\left(\text{vec}\left(\mathbf{A}^{-1}\right) \mid \text{vec}\left(\mathbf{S}_m\right)\right) = \mathcal{N}\left(\text{vec}\left(\mathbf{A}^{-1}\right); \text{vec}\left(\mathbf{A}_m^{-1}\right), \boldsymbol{\Sigma}_0 \otimes \mathbf{W}_m\right)$$

with

$$\begin{aligned} \mathbf{A}_m^{-1} &= \mathbf{A}_0^{-1} + (\mathbf{S}_m - \mathbf{A}_0^{-1}\mathbf{Y}_m)(\mathbf{Y}_m^\top \mathbf{W}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \mathbf{W}_0 \\ \mathbf{W}_m &= \mathbf{W}_0 - \mathbf{W}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \mathbf{W}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \mathbf{W}_0. \end{aligned}$$

For linear solvers, the object of interest is $\mathbf{x}^* = \mathbf{A}^{-1}\mathbf{b}$. Writing $\mathbf{A}^{-1}\mathbf{b} = (\mathbf{I} \otimes \mathbf{b}^\top) \text{vec}\left(\mathbf{A}^{-1}\right)$, and using Lemma 34 in Appendix B, the associated marginal is also Gaussian, and given by

$$p(\mathbf{x} \mid \mathbf{S}, \mathbf{Y}) = \mathcal{N}(\mathbf{x}; \mathbf{A}_m^{-1}\mathbf{b}, \mathbf{b}^\top \mathbf{W}_m \mathbf{b} \cdot \boldsymbol{\Sigma}_0). \quad (\text{MBI})$$

Again, the naive evaluation of Eq. (MBI) requires $\mathcal{O}(m^3)$ floating point operations, which can be avoided with Gram-Schmidt orthogonalization.

The Kronecker structure of the prior covariance matrix in Eq. (12) is by no means the only option that facilitates tractable inference. For example, a diagonal covariance matrix would allow efficient inference, as well. However, in absence of literature exploring other approaches within MBI, throughout MBI will refer to the use of a prior covariance matrix with Kronecker structure.

Since the typical approximate linear solver does not construct a matrix inverse, the SBI view appears to be preferable. Also, MBI has more model parameters than SBI. However, unlike in SBI, the information obtained in MBI need not be specific to a particular solution vector \mathbf{x}^* and thus can be propagated and recycled over several linear problems, similar to the notion of subspace recycling (Soodhalter, Szyld, and Xue 2014). Furthermore, MBI is able to utilize the information $\mathbf{y}_m := \mathbf{S}_m^\top \mathbf{b}$, as well. In fact, this observation will be the key in the following chapter, where it is shown that SBI is subsumed in MBI.

Probabilistic Interpretation of Projection Methods

The insights presented in Bartels, Cockayne, Ipsen, and Hennig (2019) can be summarized into three contributions. First, solution-based inference is a special case of matrix-based inference (Section 8.1). Second, the derivations in that section reveal that two points of view are useful to reason about probabilistic interpretations of projection methods: whether information stems from left-multiplication (Section 8.2)

$$\mathbf{Y}^\top := \mathbf{S}^\top \mathbf{A}, \quad (\text{L})$$

or right-multiplication (Section 8.3)

$$\mathbf{Y} := \mathbf{A} \mathbf{S}, \quad (\text{R})$$

with \mathbf{A} .

To evaluate the expression for a projection method approximation (Eq. (P)) both forms of multiplication are possible. In practice, the implementations of GMRES and CG only perform right-multiplications, yet, astonishingly, it was easier to find results presented in the right-multiplied perspective and these are more general than for left-multiplication. The third contribution is a probabilistic understanding of preconditioning (Section 8.4).

8.1 Matrix-based Inference and Solution-based Inference

One might suspect SBI and MBI to be equivalent, yet the posterior from Lemma 4 is structurally different to the posterior in Lemma 3. For the former, $\mathbf{x}_m \in \mathbf{x}_0 + \text{span}(\mathbf{S}_m - \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S}_m)$ whereas for the latter $\mathbf{x}_m \in \mathbf{x}_0 + \text{span}(\mathbf{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m)$. However, the posterior over the solution vector from MBI can be made to coincide with the posterior from SBI, if one considers observations in MBI as $\mathbf{S}_m^\top = \mathbf{Y}_m^\top \mathbf{A}^{-1}$.

Proposition 5. *Consider a Gaussian MBI prior*

$$p(\mathbf{A}^{-1}) = \mathcal{N}(\mathbf{A}^{-1}; \text{vec}(\mathbf{A}_0^{-1}), \mathbf{\Sigma}_0 \otimes \mathbf{W}_0),$$

conditioned on the left-multiplied information of Eq. (L). The associated marginal on \mathbf{x} (Eq. (MBI)) is identical to the SBI posterior on \mathbf{x} arising in Lemma 3 from $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \mathbf{\Sigma}_0)$, under the conditions

$$\mathbf{A}_0^{-1} \mathbf{b} = \mathbf{x}_0 \quad \text{and} \quad \mathbf{b}^\top \mathbf{W}_0 \mathbf{b} = 1.$$

The proof can be found in Appendix D. The first condition can be fulfilled for arbitrary $\mathbf{x}_0 \neq \mathbf{0}$ by defining $\mathbf{A}_0^{-1} := \mathbf{x}_0(\mathbf{x}_0^\top \mathbf{b})^{-1} \mathbf{x}_0^\top$, or if $\mathbf{x}_0 = \mathbf{0}$ by defining $\mathbf{A}_0^{-1} := \mathbf{0}$. The second condition can be enforced for an arbitrary covariance $\bar{\mathbf{W}}_0$ by setting $\mathbf{W}_0 := (\mathbf{b}^\top \bar{\mathbf{W}}_0 \mathbf{b})^{-1} \bar{\mathbf{W}}_0$. The result in Proposition 5 shows that any result proven for SBI applies immediately to MBI with left-multiplied observations.

8.2 Left-information Views

In this section we first show, in Proposition 6, that the conditional mean from SBI after m steps corresponds to some projection method. Then, in Proposition 7 we prove the converse: that each projection method is also the posterior mean of a probabilistic method, for some prior covariance and choice of information. After these general results, the focus will be on left-information views on CG and GMRES.

Proposition 6 (SBI defines projection methods). *Let the columns of \mathbf{S}_m be linearly independent. Consider SBI under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0),$$

and with observations $\mathbf{y}_m = \mathbf{S}_m^\top \mathbf{b}$. Then the posterior mean \mathbf{x}_m in Lemma 3 is identical to the iterate from a projection method defined by the matrices $\mathbf{U}_m = \mathbf{S}_m$ and $\mathbf{X}_m = \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m$, and the starting vector \mathbf{x}_0 .

Proof. Substituting $\mathbf{U}_m = \mathbf{S}_m$ and $\mathbf{X}_m = \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m$ into Lemma 3 gives Eq. (P), as required. \square

The converse to this also holds:

Proposition 7. *Consider a projection method defined by the matrices $\mathbf{X}_m, \mathbf{U}_m \in \mathbb{R}^{d \times m}$, each with linearly independent columns, and the starting vector $\mathbf{x}_0 \in \mathbb{R}^d$. Then the iterate \mathbf{x}_m in Eq. (P) is identical to the SBI posterior mean in Lemma 3 under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \mathbf{X}_m \mathbf{X}_m^\top) \tag{13}$$

when search directions $\mathbf{S}_m = \mathbf{U}_m$ are used.

Proof. Abbreviate $\mathbf{Z} = \mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m$ and write the projection method iterate from Eq. (P) as

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{X}_m \mathbf{Z}^{-\top} \mathbf{U}_m^\top \mathbf{r}_0.$$

Multiply the middle matrix by the identity,

$$\begin{aligned} \mathbf{Z}^{-\top} &= \mathbf{Z} \mathbf{Z}^{-1} \mathbf{Z}^{-\top} = \mathbf{Z} (\mathbf{Z}^\top \mathbf{Z})^{-1} \\ &= \mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m (\mathbf{U}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{U}_m)^{-1}, \end{aligned}$$

and insert this into the expression for \mathbf{x}_0 ,

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_0 \mathbf{A}^\top \mathbf{U}_m (\mathbf{U}_m^\top \mathbf{A} \Sigma_0 \mathbf{A}^\top \mathbf{U}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0.$$

Setting $\mathbf{U}_m = \mathbf{S}_m$ gives the mean in Lemma 3. \square

A direct way to enforce the posterior occupying the solution space is by placing a prior on the coefficients $\boldsymbol{\alpha}$ in $\mathbf{x} = \mathbf{x}_0 + \mathbf{X}_m \boldsymbol{\alpha}$. Under a unit Gaussian prior $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the implied prior on \mathbf{x} naturally has the form of Eq. (13).

However, this prior is unsatisfying since it requires the solution space to be specified *a-priori*, precluding adaptivity in the algorithm and perhaps more worryingly, the posterior uncertainty over the solution is a matrix of zeros even though the solution is not fully identified. Again taking $\mathbf{Z} = \mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m$:

$$\begin{aligned} \Sigma_m &= \Sigma_0 - \Sigma_0 \mathbf{A}^\top \mathbf{U}_m (\mathbf{U}_m^\top \mathbf{A} \Sigma_0 \mathbf{A}^\top \mathbf{U}_m)^{-1} \mathbf{U}_m^\top \mathbf{A} \Sigma_0 \\ &= \mathbf{X}_m \mathbf{X}_m^\top - \mathbf{X}_m \mathbf{Z} (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{X}_m^\top \\ &= \mathbf{X}_m \mathbf{X}_m^\top - \mathbf{X}_m \mathbf{X}_m^\top \\ &= \mathbf{0}. \end{aligned}$$

Concerning this issue, Hennig (2015) and Bartels and Hennig (2016) each proposed to adding additional uncertainty in the null space of \mathbf{X}_m . This empirical uncertainty calibration step has not yet been analyzed in detail. Such analysis is left for future work.

Including the solution space \mathbf{X}_m in the prior covariance matrix requires it to be specified *a-priori*. For solvers like CG and GMRES which construct \mathbf{X}_m adaptively this assumption may appear problematic—a probabilistic interpretation should use for inference only quantities that have already been computed. From a projection method perspective, the computation of \mathbf{X}_m is part of the initialization, but practically such methods choose m adaptively by examining the norm of the residual.¹ Nevertheless, the proposition provides a probabilistic view for *arbitrary* projection methods and does not involve \mathbf{A}^{-1} , unlike the results presented in Hennig (2015) and Cockayne, Oates, Sullivan, and Girolami (2017).

The above prior is not unique. The next proposition establishes probabilistic interpretations of projection methods under priors that are independent of solution- and constraint-space, albeit under more restrictive conditions. The benefit of this is that m need not be fixed *a-priori*. The following proposition is a contribution by Jon Cockayne.

Proposition 8. *Consider a projection method defined by $\mathbf{X}_m, \mathbf{U}_m \in \mathbb{R}^{d \times m}$ and the starting vector \mathbf{x}_0 . Further suppose that $\mathbf{U}_m = \mathbf{R} \mathbf{X}_m$ for*

¹ Sometimes m is fixed *a-priori*, due to memory or computation time limits.

some invertible $\mathbf{R} \in \mathbb{R}^{d \times d}$, and that $\mathbf{A}^\top \mathbf{R}$ is symmetric positive-definite. Then under the prior

$$p(\mathbf{x}) = \mathcal{N}\left(\mathbf{x}; \mathbf{x}_0, (\mathbf{A}^\top \mathbf{R})^{-1}\right)$$

and the search directions $\mathbf{S}_m = \mathbf{U}_m = \mathbf{R}\mathbf{X}_m$, the iterate in the projection method is identical to the posterior mean in Lemma 3.

Proof. First substitute $\mathbf{X}_m = \mathbf{R}^{-1}\mathbf{U}_m$ into Eq. (P) to obtain

$$\begin{aligned} \mathbf{x}_m &= \mathbf{x}_0 + \mathbf{R}^{-1}\mathbf{U}_m(\mathbf{U}_m^\top \mathbf{A} \mathbf{R}^{-1} \mathbf{U}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0 \\ &= \mathbf{x}_0 + \mathbf{R}^{-1} \mathbf{A}^{-\top} \mathbf{A}^\top \mathbf{U}_m (\mathbf{U}_m^\top \mathbf{A} \mathbf{R}^{-1} \mathbf{A}^{-\top} \mathbf{A}^\top \mathbf{U}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0 \\ &= \mathbf{x}_0 + \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{U}_m (\mathbf{U}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{U}_m)^{-1} \mathbf{U}_m^\top \mathbf{r}_0. \end{aligned}$$

The third line uses $\boldsymbol{\Sigma}_0 = (\mathbf{A}^\top \mathbf{R})^{-1} = \mathbf{R}^{-1} \mathbf{A}^{-\top}$. This is equivalent to the posterior mean in Eq. (SBI) with $\mathbf{S}_m = \mathbf{U}_m$. \square

A corollary which provides further insight arises when one considers the *polar decomposition* of \mathbf{A} . Recall that an invertible matrix \mathbf{A} has a unique polar decomposition $\mathbf{A} = \mathbf{P}\mathbf{H}$, where $\mathbf{P} \in \mathbb{R}^{d \times d}$ is orthogonal and $\mathbf{H} \in \mathbb{R}^{d \times d}$ is symmetric positive-definite. The following corollary is a contribution by Ilse Ipsen.

Corollary 9. Consider a projection method defined by $\mathbf{X}_m, \mathbf{U}_m \in \mathbb{R}^{d \times m}$ and the starting vector \mathbf{x}_0 , and suppose that $\mathbf{U}_m = \mathbf{P}\mathbf{X}_m$, where \mathbf{P} arises from the polar decomposition $\mathbf{A} = \mathbf{P}\mathbf{H}$. Then under the prior

$$p(\mathbf{x}) = \mathcal{N}\left(\mathbf{x}; \mathbf{x}_0, \mathbf{H}^{-1}\right)$$

and the search directions $\mathbf{S}_m = \mathbf{U}_m = \mathbf{P}\mathbf{X}_m$, the iterate in the projection method is identical to the posterior mean in Lemma 3.

Proof. This follows from Proposition 8. Setting $\mathbf{R} = \mathbf{P}$ aligns the search directions in Corollary 9 with those in Proposition 8. Since \mathbf{P} is orthogonal, $\mathbf{P}^{-1} = \mathbf{P}^\top$, and since \mathbf{H} is symmetric positive-definite, $\mathbf{A}^\top \mathbf{P} = \mathbf{P}^\top \mathbf{A} = \mathbf{H}$ by definition of the polar decomposition, which gives the prior covariance required for Proposition 8. \square

This is an intuitive analogue of similar results by Hennig (2015) and Cockayne, Oates, Sullivan, and Girolami (2017) which show that CG is recovered under certain conditions involving a prior $\boldsymbol{\Sigma}_0 = \mathbf{A}^{-1}$.

When \mathbf{A} is not symmetric and positive definite it cannot be used as a prior covariance. This corollary suggests a natural way to select a prior covariance still linked to the linear system, though this choice is still not computationally convenient. Furthermore, in the case that \mathbf{A} is symmetric positive-definite, this recovers the prior which replicates CG: note that each of \mathbf{H} and \mathbf{P} can be stated explicitly as $\mathbf{H} = (\mathbf{A}^\top \mathbf{A})^{\frac{1}{2}}$ and $\mathbf{P} = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-\frac{1}{2}}$. Thus in the case of symmetric positive-definite \mathbf{A} we

have that $\mathbf{H} = \mathbf{A}$ and $\mathbf{P} = \mathbf{I}$, so that the prior covariance $\boldsymbol{\Sigma}_0 = \mathbf{A}^{-1}$ arises naturally from this interpretation.

8.2.1 Conjugate Gradients

Recall from Section 5.2 that conjugate gradients can be seen as projection method with $\mathbf{X}_m = \mathbf{U}_m = \mathbf{Q}_m$, where \mathbf{Q}_m is a basis of $K_m(\mathbf{A}, \mathbf{r}_0)$. A left-multiplied probabilistic interpretation for CG has been presented by Cockayne, Oates, Ipsen, and Girolami (2018) for the choice $\boldsymbol{\Sigma}_0 := \mathbf{A}^{-1}$ and \mathbf{S}_m being the conjugate gradients search directions. This is now a consequence of Corollary 9. Novel is the interpretation $\boldsymbol{\Sigma}_0 := \mathbf{Q}_m \mathbf{Q}_m^\top$ following from Proposition 7.

Corollary 10. *Under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0) \quad \text{where} \quad \boldsymbol{\Sigma}_0 = \mathbf{Q}_m \mathbf{Q}_m^\top,$$

and with observations $\mathbf{y}_m = \mathbf{Q}_m^\top \mathbf{b}$, the SBI posterior mean Eq. (SBI) is identical to the CG iterate \mathbf{x}_m in Eq. (FOM).

Above prior also gives a probabilistic interpretation of FOM, the generalization of CG, when \mathbf{A} is not a valid covariance matrix.

8.2.2 Generalized Minimal Residual

Now follow probabilistic linear solvers with posterior means that coincide with the solution estimate from GMRES.

Corollary 11. *Under the SBI prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0) \quad \text{where} \quad \boldsymbol{\Sigma}_0 = (\mathbf{A}^\top \mathbf{A})^{-1}$$

and the search directions $\mathbf{U}_m = \mathbf{A} \mathbf{Q}_m$, the posterior mean is identical to the GMRES iterate \mathbf{x}_m in Eq. (GMR).

Proof. Substitute $\mathbf{R} = \mathbf{A}$ and $\mathbf{U}_m = \mathbf{A} \mathbf{Q}_m$ into Proposition 8 and compare to Eq. (GMR). \square

This interpretation exhibits an interesting duality with CG for which $\boldsymbol{\Sigma}_0 = \mathbf{A}^{-1}$ and $\mathbf{U}_m = \mathbf{Q}_m$. Another probabilistic interpretation follows from Proposition 7.

Corollary 12. *Under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0) \quad \text{where} \quad \boldsymbol{\Sigma}_0 = \mathbf{Q}_m \mathbf{Q}_m^\top, \tag{14}$$

and with observations $\mathbf{y}_m = \mathbf{Q}_m^\top \mathbf{A}^\top \mathbf{b}$, the SBI posterior mean Eq. (SBI) is identical the GMRES iterate \mathbf{x}_m in Eq. (GMR).

Note that Proposition 11 has a posterior covariance which is not practical, as it involves \mathbf{A}^{-1} . Cockayne, Oates, Sullivan, and Girolami (2017) proposed replacing \mathbf{A}^{-1} in the prior covariance with a preconditioner, which does yield a practically computable posterior, but this extension was not explored here. Furthermore, that approach yields unsatisfactorily calibrated posterior uncertainty, as described in that work. Corollary 12 does not have this drawback, but as mentioned there, the posterior covariance is a matrix of zeroes.

8.2.3 Simulation Study

In this section the simulation study of Cockayne, Oates, Ipsen, and Girolami (2018) will be replicated to demonstrate that the uncertainty produced from GMRES in Proposition 11 is similarly poorly calibrated to CG, owing to the dependence of the Arnoldi directions \mathbf{Q}_m on \mathbf{x}^* by way of its dependence on \mathbf{b} . Throughout the size of the test problems is set to $d = 100$. The eigenvalues of \mathbf{A} were drawn from an exponential distribution with parameter $\gamma = 10$, and eigenvectors were drawn uniformly from the Haar-measure over rotation-matrices (see Diaconis and Shahshahani (1987)). In contrast to Cockayne, Oates, Ipsen, and Girolami (2018) the entries of \mathbf{b} are drawn from a standard Gaussian distribution, rather than \mathbf{x}^* . In that case, \mathbf{x}^* is a draw from the GMRES prior $\mathcal{N}(\mathbf{0}, (\mathbf{A}^\top \mathbf{A})^{-1})$ by Lemma 34 in Appendix B. Hence, the prior is perfectly calibrated for this scenario, and one would expect that the posterior should be equally well-calibrated for $m \geq 1$.

Cockayne, Oates, Ipsen, and Girolami (2018) argue that if the uncertainty is well-calibrated, then \mathbf{x}^* can be considered as a draw from the posterior. Under this assumption, *i.e.* $\boldsymbol{\Sigma}_m^{-1/2}(\mathbf{x}^* - \mathbf{x}_m) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ they derive the test statistic:

$$Z(\mathbf{x}^*) := \|\boldsymbol{\Sigma}_m^{-1/2}(\mathbf{x}^* - \mathbf{x}_m)\| \sim \chi_{d-m}^2.$$

Figure 1 shows the convergence of GMRES and below, the convergence rate of the trace of the posterior covariance matrix $\boldsymbol{\Sigma}_m$. Figure 2 displays the test statistic. It can be seen that the same poor uncertainty quantification occurs; even after just 10 iterations, the empirical distribution of the test statistic exhibits a profound left-shift, indicating an overly conservative posterior distribution. Producing well-calibrated posteriors remains an open issue in the field of probabilistic linear solvers.




Figure 1: convergence of posterior mean (top plot) and trace of the posterior covariance matrix (bottom plot) of the probabilistic interpretation of GMRES from Proposition 11.

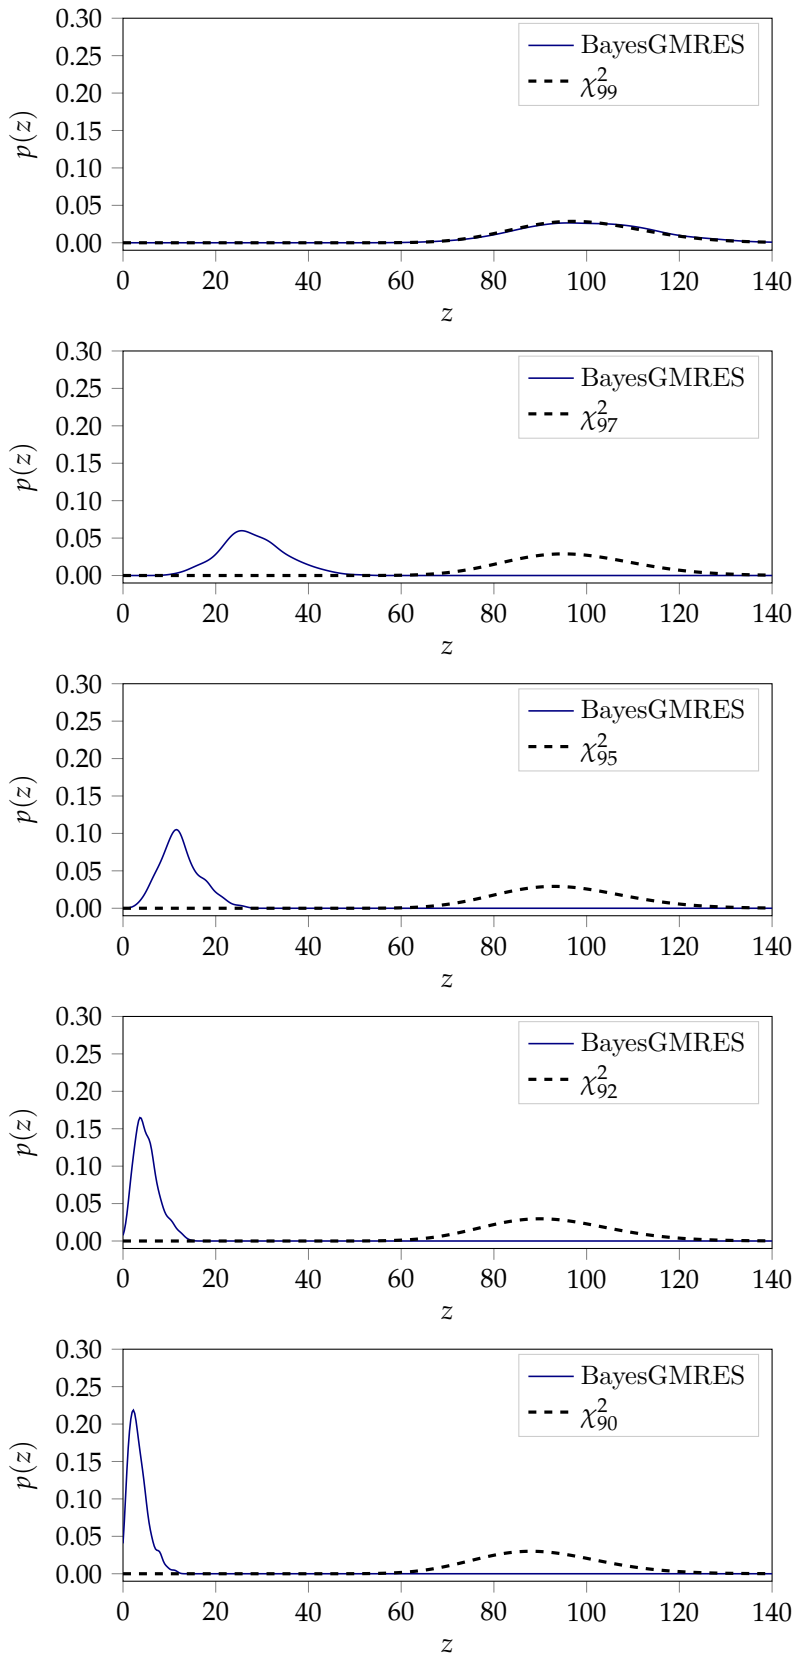


Figure 2: assessment of the uncertainty quantification. Plotted are kernel density estimates for the statistic Z based on 500 randomly sampled test problems for steps $m = \{1, 3, 5, 8, 10\}$. These are compared with the theoretical distribution of Z when the posterior distribution is well-calibrated.

8.3 Right-information Views

Surprisingly, it is harder to find probabilistic interpretations that use right-multiplied observations. A general result comparable to Proposition 7 remains yet to be found. The following proposition has not been published in Bartels, Cockayne, Ipsen, and Hennig (2019) but is a result of considerations during the writing process of this dissertation.

Proposition 13. *Consider a projection method defined by the matrices $\mathbf{X}_m, \mathbf{U}_m \in \mathbb{R}^{d \times m}$, each with linearly independent columns, and the starting vector $\mathbf{x}_0 \in \mathbb{R}^d$ and assume that $\mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m$ is invertible. Assume further that*

$$\mathbf{U}_m^\top \mathbf{A} \mathbf{x}_0 = \mathbf{0} \quad (15)$$

and if $\mathbf{x}_0 \neq \mathbf{0}$, that there exists a $\mathbf{v} \neq \mathbf{0}$, s.t.

$$\mathbf{X}_m^\top \mathbf{A}^\top \mathbf{v} = \mathbf{0}. \quad (16)$$

Then the iterate \mathbf{x}_m in Eq. (P) is identical to the projected MBI posterior mean in Eq. (MBI) under the prior

$$p(\mathbf{A}^{-1}) = \mathcal{N} \left(\mathbf{A}^{-1}; \frac{1}{\mathbf{v}^\top \mathbf{b}} \mathbf{x}_0 \mathbf{v}^\top, \mathbf{V} \otimes \mathbf{U}_m \mathbf{U}_m^\top \right) \quad (17)$$

when search directions $\mathbf{S}_m = \mathbf{X}_m$ are used and where \mathbf{V} is an arbitrary s.p.d. matrix.

Proof.

$$\begin{aligned} & \mathbf{A}_m^{-1} \mathbf{b} \\ &= \mathbf{A}_0^{-1} \mathbf{b} + (\mathbf{S}_m - \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S}_m) (\mathbf{S}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{A} \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{b} \\ & \quad // \text{ using Eq. (MBI)} \\ &= \mathbf{x}_0 + (\mathbf{S}_m - \frac{1}{\mathbf{v}^\top \mathbf{b}} \mathbf{x}_0 \mathbf{v}^\top \mathbf{A} \mathbf{S}_m) (\mathbf{S}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{A} \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{b} \\ & \quad // \text{ using } \mathbf{A}_0^{-1} = \frac{1}{\mathbf{v}^\top \mathbf{b}} \mathbf{x}_0 \mathbf{v}^\top \\ &= \mathbf{x}_0 + (\mathbf{X}_m - \frac{1}{\mathbf{v}^\top \mathbf{b}} \mathbf{x}_0 \mathbf{v}^\top \mathbf{A} \mathbf{X}_m) (\mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{A} \mathbf{X}_m)^{-1} \mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{b} \\ & \quad // \text{ using definition of } \mathbf{S}_m \\ &= \mathbf{x}_0 + \mathbf{X}_m (\mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{A} \mathbf{X}_m)^{-1} \mathbf{X}_m^\top \mathbf{A}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{b} \\ & \quad // \text{ using Eq. (16)} \\ &= \mathbf{x}_0 + \mathbf{X}_m (\mathbf{U}_m^\top \mathbf{A} \mathbf{S}_m)^{-1} \mathbf{U}_m^\top \mathbf{b} \\ & \quad // \text{ simplifying} \\ &= \mathbf{x}_0 + \mathbf{X}_m (\mathbf{U}_m^\top \mathbf{A} \mathbf{S}_m)^{-1} \mathbf{U}_m^\top (\mathbf{b} - \mathbf{A} \mathbf{x}_0) \\ & \quad // \text{ using Eq. (15)} \end{aligned}$$

□

Observe the symmetry with respect to Proposition 7 where the roles of \mathbf{X}_m and \mathbf{U}_m in prior and search directions are exchanged. Equations (15) and (16) are trivially fulfilled when $\mathbf{x}_0 = \mathbf{0}$. If $\mathbf{X}_m = \mathbf{U}_m$ and \mathbf{A} is s.p.d. one can choose $\mathbf{v} := \mathbf{x}_0$. This choice imposes the search directions to be \mathbf{A} -conjugate to the initial solution. For conjugate gradients this is generally *not* the case.

8.3.1 Generalized Minimal Residual

Now follows a probabilistic interpretation of GMRES. Again, a general result comparable to those in Section 8.2.2 remains future work.

Proposition 14. *Assume that $\mathbf{x}_0 = \mathbf{0}$. Under the prior*

$$p(\mathbf{A}^{-1}) = \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma} \otimes \mathbf{I})$$

and given $\mathbf{Y}_m = \mathbf{A}\mathbf{Q}_m$, where \mathbf{Q}_m are the Arnoldi directions, the implied posterior mean over the solution given by $\mathbf{A}_m^{-1}\mathbf{b}$ is equivalent to the GMRES solution.

Proof. Under this prior, \mathbf{b} applied to the posterior mean is

$$\begin{aligned} \mathbf{A}_m^{-1}\mathbf{b} &= \mathbf{A}_0^{-1}\mathbf{b} + (\mathbf{Q}_m - \mathbf{A}_0^{-1}\mathbf{Y}_m)(\mathbf{Y}_m^\top\mathbf{Y}_m)^{-1}\mathbf{Y}_m^\top\mathbf{b} \\ &= \mathbf{Q}_m(\mathbf{Y}_m^\top\mathbf{Y}_m)^{-1}\mathbf{Y}_m^\top\mathbf{b} \\ &= \mathbf{Q}_m(\mathbf{Q}_m^\top\mathbf{A}^\top\mathbf{A}\mathbf{Q}_m)^{-1}\mathbf{Q}_m^\top\mathbf{A}^\top\mathbf{b} \end{aligned}$$

which is the GMRES projection (Eq. (GMR)) if $\mathbf{x}_0 = \mathbf{0}$. □

If above proposition is true for $\mathbf{x}_0 \neq \mathbf{0}$, it provides structural insights into GMRES. For the choice $\boldsymbol{\Sigma} := \mathbf{I}$, one could say that running GMRES is equivalent to inference over \mathbf{A}^{-1} or \mathbf{x}^* under the assumption that all entries are independent standard normal.

8.3.2 Conjugate Gradients

The following interpretation of conjugate gradients provides an explanation of how the search directions can be motivated probabilistically. The following proposition is a contribution by Philipp Hennig.

Proposition 15. *Consider the prior*

$$p(\mathbf{A}^{-1}) = \mathcal{N}(\text{vec}(\mathbf{A}^{-1}); \text{vec}(\alpha\mathbf{I}), (\beta\mathbf{I} + \gamma\mathbf{W}) \otimes (\beta\mathbf{I} + \gamma\mathbf{W}))$$

where $\mathbf{W} := \mathbf{A}^{-1}$. For all choices $\alpha \in \mathbb{R} \setminus \{0\}$ and $\beta, \gamma \in \mathbb{R}_{+,0}$ with $\beta + \gamma > 0$, Algorithm 3 is equivalent to CG, in the sense that it produces the same search directions \mathbf{s}_i (scaled).

Proof. The proof is extensive and can be found in Appendix D. □

In general, the solution estimate \mathbf{x}_m produced by Algorithm 3 differs from the posterior over \mathbf{x} in Eq. (MBI), since the CG estimate is corrected by the step size computed in line 6. Fixing this rank-1 discrepancy would complicate the exposition of Algorithm 3 and yield a more cumbersome algorithm. Proposition 15 is remarkable since, for the case when $\gamma = 0$, it is the only result with a rank d prior matrix, that does *not* involve \mathbf{A}^{-1} .

Algorithm 3: The algorithm referred to by Proposition 15, which reproduces the search directions from CG.

```

1  $\mathbf{x}_0 \leftarrow \mathbf{A}_0^{-1} \mathbf{b}$  // initial guess
2  $\mathbf{r}_0 \leftarrow \mathbf{A} \mathbf{x}_0 - \mathbf{b}$ 
3 for  $i = 1, \dots, m$  do
4    $\mathbf{d}_i \leftarrow -\mathbf{A}_{i-1}^{-1} \mathbf{r}_{i-1}$  // compute optimization direction
5    $\mathbf{z}_i \leftarrow \mathbf{A} \mathbf{d}_i$  // observe
6    $\alpha_i \leftarrow -\frac{\mathbf{d}_i^\top \mathbf{r}_{i-1}}{\mathbf{d}_i^\top \mathbf{z}_i}$  // optimal step-size
7    $\mathbf{s}_i \leftarrow \alpha_i \mathbf{d}_i$  // re-scale step
8    $\mathbf{y}_i \leftarrow \alpha_i \mathbf{z}_i$  // re-scale observation
9    $\mathbf{x}_i \leftarrow \mathbf{x}_{i-1} + \mathbf{s}_i$  // update estimate for  $\mathbf{x}$ 
10   $\mathbf{r}_i \leftarrow \mathbf{r}_{i-1} + \mathbf{y}_i$  // new gradient at  $\mathbf{x}_i$ 
11   $\mathbf{A}_i^{-1} \leftarrow \mathbb{E}_{p(\mathbf{A}^{-1} | \mathcal{S}, \mathcal{Y})} \mathbf{A}^{-1}$  // estimate  $\mathbf{A}^{-1}$ 
12 end for
13 return  $\mathbf{x}_m$ 

```

8.4 Preconditioning

This section discusses probabilistic views on preconditioning. Preconditioning is a technique used to accelerate the convergence of iterative methods (Saad 2003, Sections 9 and 10). A preconditioner \mathbf{P} is a nonsingular matrix satisfying two requirements:

1. linear systems $\mathbf{Pz} = \mathbf{c}$ can be solved at low computational cost
2. and \mathbf{P} is “close” to \mathbf{A} in some sense.

Hence, solving systems based upon a preconditioner can be viewed as approximately inverting \mathbf{A} , and indeed many preconditioners are constructed based upon this intuition. One distinguishes between *right preconditioners* \mathbf{P}_r and *left preconditioners* \mathbf{P}_l , depending on whether they act on \mathbf{A} from the left or the right. Two-sided preconditioning with nonsingular matrices \mathbf{P}_l and \mathbf{P}_r transforms implicitly Eq. (LES) into a new linear problem

$$\mathbf{P}_l \mathbf{A} \mathbf{P}_r \mathbf{z}^* = \mathbf{P}_l \mathbf{b}, \quad \text{with } \mathbf{x}^* = \mathbf{P}_r \mathbf{z}^*. \quad (18)$$

The preconditioned system can then be solved using arbitrary projection methods as described in Section 7.2, from the starting point \mathbf{z}_0 defined by $\mathbf{x}_0 = \mathbf{P}_r \mathbf{z}_0$. The probabilistic view can be used to create

a nuanced description of preconditioning as a form of prior information. In the SBI framework, Proposition 16 below shows that solving a right-preconditioned system is equivalent to modifying the prior, while in Proposition 17 shows that left-preconditioning is equivalent to making a different choice of observations. The following proposition is a contribution by Philipp Hennig.

Proposition 16 (Right preconditioning). *Consider the right-preconditioned system*

$$\mathbf{A}\mathbf{P}_r\mathbf{z}^* = \mathbf{b} \quad \text{where} \quad \mathbf{x}^* = \mathbf{P}_r\mathbf{z}^*. \quad (19)$$

SBI on Eq. (19) under the prior

$$\mathbf{z} \sim \mathcal{N}(\mathbf{z}; \mathbf{z}_0, \boldsymbol{\Sigma}_0) \quad (20)$$

is equivalent to solving Eq. (LES) under the prior

$$\mathbf{x} \sim \mathcal{N}(\mathbf{x}; \mathbf{P}_r\mathbf{z}_0, \mathbf{P}_r\boldsymbol{\Sigma}_0\mathbf{P}_r^\top). \quad (21)$$

Proof. Define $\mathbf{B} := \mathbf{A}\mathbf{P}_r$ and $\hat{\mathbf{r}}_0 = \mathbf{b} - \mathbf{B}\mathbf{z}_0$. Consider the prior defined in Eq. (21). Lemma 3 implies that after observing information from search directions \mathbf{S}_m , the posterior mean equals

$$\mathbf{x}_m = \mathbf{P}_r\mathbf{z}_0 + \mathbf{P}_r\boldsymbol{\Sigma}_0\mathbf{B}^\top\mathbf{S}_m(\mathbf{S}_m^\top\mathbf{B}\boldsymbol{\Sigma}_0\mathbf{B}^\top\mathbf{S}_m)^{-1}\mathbf{S}_m^\top\hat{\mathbf{r}}_0.$$

Left multiplying by \mathbf{P}_r^{-1} shows that this is equivalent to

$$\begin{aligned} \mathbf{z}_m &:= \mathbf{P}_r^{-1}\mathbf{x}_m \\ &= \mathbf{z}_0 + \boldsymbol{\Sigma}_0\mathbf{B}^\top\mathbf{S}_m(\mathbf{S}_m^\top\mathbf{B}\boldsymbol{\Sigma}_0\mathbf{B}^\top\mathbf{S}_m)^{-1}\mathbf{S}_m^\top\hat{\mathbf{r}}_0. \end{aligned}$$

Now note that \mathbf{z}_m is the posterior mean from the prior Eq. (20) after observing search directions \mathbf{S}_m , when inferring the solution over of the system $\mathbf{B}\mathbf{z}^* = \mathbf{b}$. \square

Proposition 17 (Left preconditioning). *Consider the left-preconditioned system*

$$\mathbf{P}_l\mathbf{A}\mathbf{x}^* = \mathbf{P}_l\mathbf{b} \quad (22)$$

And the SBI prior

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0).$$

Then the posterior from SBI on Eq. (22) under search directions \mathbf{S}_m is equivalent to the posterior from SBI applied to the system Eq. (LES) under search directions $\mathbf{P}_l^\top\mathbf{S}_m$.

Proof. Lemma 3 implies that after observing search directions \mathbf{T}_m , the posterior mean over the solution of Eq. (LES) equals

$$\mathbf{x}_m = \mathbf{x}_0 + \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{T}_m (\mathbf{T}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{T}_m)^{-1} \mathbf{T}_m^\top \mathbf{r}_0$$

where $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$. Setting $\mathbf{T}_m = \mathbf{P}_l^\top \mathbf{S}_m$ gives

$$\mathbf{x}_m = \mathbf{x}_0 + \boldsymbol{\Sigma}_0 \mathbf{B}^\top \mathbf{S}_m (\mathbf{S}_m^\top \mathbf{B} \boldsymbol{\Sigma}_0 \mathbf{B}^\top \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \mathbf{P}_l \hat{\mathbf{r}}_0$$

where $\mathbf{B} := \mathbf{P}_l \mathbf{A}$ and $\hat{\mathbf{r}}_0 = \mathbf{P}_l \mathbf{b} - \mathbf{P}_l \mathbf{A} \mathbf{x}_0$. Thus, \mathbf{x}_m is the posterior mean of the system $\mathbf{B}\mathbf{x}^* = \mathbf{P}_l \mathbf{b}$ after observing search directions \mathbf{S}_m . \square

If a probabilistic linear solver has a posterior mean which coincides with a projection method, the Propositions 16 and 17 show how to obtain a probabilistic interpretation of the *preconditioned* version of that algorithm. With Proposition 5 these results carry over to MBI.

This interpretation of preconditioning is *not* unique. When using left-multiplied observations, the same reasoning can be used to show that left-preconditioning corresponds to a change in the prior belief, while right-preconditioning corresponds to collecting different observations.

Discussion

9.1 Summary

This part showed that solution-based inference is contained within matrix-based inference, which allows to transfer results from SBI to MBI with left-multiplied information. This correspondence motivates the classification of probabilistic interpretations as left-multiplied and right-multiplied. Connections between probabilistic linear solvers and the class of projection methods have been presented, and a probabilistic interpretation of preconditioning.

9.2 Future Directions

Posterior uncertainty calibration remains a challenge for probabilistic linear solvers. Direct probabilistic interpretations of CG and GMRES yield posterior covariance matrices which are not always computable, and even when the posterior can be computed, the uncertainty remains poorly calibrated. Mitigating this issue without sacrificing the rate of convergence provided by Krylov methods remains an important goal.

Another natural question to ask is whether other classes of linear solvers such as stationary iterative methods (Saad 2003, Chapter 4) have probabilistic interpretations. As an outlook for a non-Bayesian probabilistic interpretation consider the following relation between Jacobi-iteration (Saad 2003, pp. 105) and Gibbs sampling (Geman and Geman 1984; Bishop 2006, pp. 542).

If all components of \mathbf{x}^* are identified, except for one, say the i -th component, then one can identify this last component, by rearranging the i -th equation for \mathbf{x}_i

$$\mathbf{x}_i = \frac{1}{\mathbf{A}_{ii}} \left(\mathbf{b}_i - \sum_{\substack{j=1 \\ j \neq i}}^d \mathbf{A}_{ij} \mathbf{x}_j \right), \quad (23)$$

and has found the solution \mathbf{x}^* (assuming $\mathbf{A}_{ii} \neq 0$). The last sentence can be expressed probabilistically as

$$p(\mathbf{x}_i \mid \mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_d) = \delta \left(\mathbf{x}_i - \frac{1}{\mathbf{A}_{ii}} \left(\mathbf{b}_i - \sum_{\substack{j=1 \\ j \neq i}}^d \mathbf{A}_{ij} \mathbf{x}_j \right) \right).$$

Given an initial guess \mathbf{x}_0 , and a likelihood of the form above, Gibbs sampling cycles through all variables in order and updates each by a sample from the likelihood, where in this case the samples are deterministic and updated according to Eq. (23). At the same time, Eq. (23) is the update rule for Jacobi-iteration. Hence, when using the last sample as estimate, Gibbs sampling and Jacobi-iteration produce the same output in each step.

Part IV

Probabilistic Solvers for Kernel Least-Squares Problems

Hier soll der geduldige Leser verschmaufen. Ich schweife nicht ab, sondern zurück, verspreche feierlich: das Fäkalienthema ist noch nicht ganz erschöpft, Chopin aber endgültig erledigt, jedenfalls in seiner Qualität, als Quantität werde ich, schon aus kompositorischen Gründen, mich seiner noch einige Male bedienen müssen. Es wird nicht mehr vorkommen. Reuevoll schlage ich mir an die Brust, jene, deren Quantität bei meinem Hemdenschneider zu erfahren wäre, deren Qualität aber so schwer zu definieren ist.

—Heinrich Böll, *Entfernung von der Truppe*

Preliminaries

The previous part compared and related probabilistic numerical methods and classic approaches to approximation. The purpose of this part is to present a possibility, to use the probabilistic perspective to develop novel approximation algorithms. It is based on the publication S. Bartels and P. Hennig (2019). “Conjugate Gradients for Kernel Machines.” In: *ArXiv e-prints* 1911.06048. [arXiv:1911.06048](https://arxiv.org/abs/1911.06048), currently under revision. I was the primary author and performed the principal analysis and work, Philipp Hennig provided initial ideas and direction.

10.1 Introduction

Regularized least-squares (kernel-ridge / Gaussian process) regression is a fundamental algorithm of statistics and machine learning. Because generic algorithms for the exact solution have cubic complexity in the number of datapoints, large datasets require to resort to approximations. In the following, the computation of the least-squares prediction is itself treated as a probabilistic inference problem. The key is a structured Gaussian regression model on the kernel function that uses projections of the kernel matrix to obtain a low-rank approximation of the kernel and the matrix. This leads to an enhanced way to use the method of conjugate gradients for the specific setting of least-squares regression as encountered in machine learning dubbed *kernel machine conjugate gradients* (KMCG). KMCG improves the approximation of the kernel ridge regressor / Gaussian process posterior mean over vanilla conjugate gradients and, when used in Gaussian process models, allows computation of the posterior variance and the log marginal likelihood (evidence) without further overhead. Recall the regularized least-squares problem introduced in Chapter 4:

$$\bar{f}(\mathbf{x}_*) = \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}. \quad (1)$$

Instead of providing an approximation solely to the vector $(\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$, the approach uses the MVMs performed by CG to learn an approximation directly to the function k . This will allow to approximate Eqs. (2) and (3) (restated below) as well, without rerunning CG.

$$\bar{c}(\mathbf{x}_*, \mathbf{x}_{**}) = k(\mathbf{x}_*, \mathbf{x}_{**}) - \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{**} \quad (2)$$

$$\ln p(\mathbf{y}) = -\frac{1}{2} \mathbf{y}^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} + \frac{1}{2} \ln |2\pi(\mathbf{K} + \sigma^2 \mathbf{I})|^{-1} \quad (3)$$

Chapter 11 proposes a model template that can be used to learn finite-rank approximations to kernel functions. Then, Chapter 12 shows how conjugate gradients can be used in combination.




Figure 3: the algorithm KMCG in comparison to CG on a toy setup. The dataset consists of one hundred data-points where the targets are a draw from a zero-mean Gaussian process with squared exponential kernel (Eq. (43)) with $\Lambda = 0.25$ and $\theta_f = 2$. The thin, black line is the posterior mean of that Gaussian process (Eq. (1)). The light-green line is the mean prediction produced by conjugate gradients after $P = 7$ steps and the dark-red line is the mean prediction of KMCG (where the number of inducing inputs $M = N$).

10.2 Finite-rank Kernel

An M -rank approximation to a kernel is a factorization of the form

$$k(\mathbf{x}, \mathbf{z}) \approx \boldsymbol{\phi}(\mathbf{x})^* \boldsymbol{\Sigma}^{-1} \boldsymbol{\phi}(\mathbf{z}) \quad (24)$$

where $\boldsymbol{\phi}(\mathbf{x}) : \mathbb{X} \rightarrow \mathbb{C}^M$, $\boldsymbol{\phi}^*$ denotes the conjugate transpose, and $\boldsymbol{\Sigma}$ is an $M \times M$ Hermitian and positive definite matrix. Given such an expansion one can use the matrix-inversion, and matrix-determinant lemmata to approximate Equations (1) to (3) with the expressions below

$$\bar{f}(\mathbf{x}_*) \approx \boldsymbol{\phi}(\mathbf{x}_*)^* \mathbf{A}^{-1} \boldsymbol{\Phi} \mathbf{y} \quad (25)$$

$$\bar{c}(\mathbf{x}_*, \mathbf{z}_*) \approx \sigma^2 \boldsymbol{\phi}(\mathbf{x}_*)^* \mathbf{A}^{-1} \boldsymbol{\phi}(\mathbf{z}_*) \quad (26)$$

$$\ln p(\mathbf{y}) \approx -\frac{1}{2} \mathbf{y}^\top \boldsymbol{\Phi}^* \mathbf{A}^{-1} \boldsymbol{\Phi} \mathbf{y} - \frac{1}{2} \ln |\mathbf{A}| - \frac{N}{2} \ln(2\pi\sigma^2) \quad (27)$$

where $\boldsymbol{\phi}(\mathbf{x}_*)_j = \phi_j(\mathbf{x}_*)$, $\boldsymbol{\Phi}_{ij} = \phi_i(\mathbf{X}_j)$ and $\mathbf{A} := \boldsymbol{\Phi} \boldsymbol{\Phi}^* + \sigma^2 \boldsymbol{\Sigma}$. Typically $M \ll N$ and therefore the computational costs to evaluate Equations (25) to (27) reduce from $\mathcal{O}(N^3)$ to $\mathcal{O}(NM^2)$, *i.e.* linear in N . The dominant factor is the matrix-matrix product $\boldsymbol{\Phi} \boldsymbol{\Phi}^*$.

An example for a finite-rank kernel that will become important later is the *Subset of Regressors* (SoR) approximation (Quiñonero-Candela and Rasmussen 2005)

$$k_{\text{SoR}}(\mathbf{x}, \mathbf{z}) = k(\mathbf{x}, \mathbf{X}_U) k(\mathbf{X}_U, \mathbf{X}_U)^{-1} k(\mathbf{X}_U, \mathbf{z}) \quad (28)$$

where \mathbf{X}_U is a set of M so called inducing inputs. The method proposed in this work (KMCG) is related to SoR. Readers familiar with SoR will be aware of the associated flaws, and methods to remedy them (Quiñonero-Candela and Rasmussen 2005; Titsias 2009). For stationary kernels and tests points far away from the data, the predictive uncertainty (Eq. (26)) goes to zero. The *Deterministic Training Conditional* (DTC) approximation alleviates this issue by using the exact kernel for the prior uncertainty over the test inputs (Quiñonero-Candela and Rasmussen 2005). In effect this is a substitution of Eq. (26) for Eq. (29) below.

$$\bar{c}(\mathbf{x}_*, \mathbf{z}_*) \approx k(\mathbf{x}_*, \mathbf{x}_{**}) - \boldsymbol{\phi}(\mathbf{x}_*)^* \left(\boldsymbol{\Phi} \boldsymbol{\Phi}^* + \sigma^2 \mathbf{I} \right)^{-1} \boldsymbol{\phi}(\mathbf{z}_*) \quad (29)$$

We will apply the same substitution for our method KMCG.

Model

To approximate Equations (1) to (3), we will approximate the kernel and, to this end, present a probabilistic estimation rule for k . The idea is to treat the kernel as unknown and to choose prior and likelihood such that the posterior mean k_M is efficient to evaluate and yields a kernel of finite rank. Substituting for this finite-rank kernel in Equations (1) to (3) then allows to compute these expressions faster. The following sections describe a prior over k , possible likelihoods and resulting posteriors. Fig. 4 on page 48 shows a schematic summary of this chapter.

11.1 Prior

Consider a Gaussian process prior over bivariate functions


$$k \sim \mathcal{GP}(k_0, \gamma\psi) \quad (30)$$

where $\psi : \mathbb{X}^2 \times \mathbb{X}^2 \rightarrow \mathbb{R}$ is a covariance function over kernel and $\gamma \in \mathbb{R}^+$ is a scaling parameter. Since the posterior mean is meant to be a substitution for the exact kernel, this is an exchange of one least-squares problem for another. Without further assumptions, calculating the posterior over k is more expensive than computing the equations of interest (Eqs. (1) to (3)). Efficient inference is rendered possible by imposing the following structure on ψ

$$\psi(k(\mathbf{a}, \mathbf{b}), k(\mathbf{c}, \mathbf{d})) := \frac{1}{2}w(\mathbf{a}, \mathbf{c})w(\mathbf{b}, \mathbf{d}) + \frac{1}{2}w(\mathbf{a}, \mathbf{d})w(\mathbf{b}, \mathbf{c}) \quad (31)$$


for $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d} \in \mathbb{X}$ and where w is a covariance function on the domain \mathbb{X} . Consider the first addend. It states that the similarity between $k(\mathbf{a}, \mathbf{b})$ and $k(\mathbf{c}, \mathbf{d})$ depends on the similarity of \mathbf{a} and \mathbf{c} , and \mathbf{b} and \mathbf{d} —a natural assumption for kernel matrices. The second addend is a symmetrization of the first. Observe that each addend is a product kernel of two pairs of inputs and recall that a product kernel produces Kronecker product matrices. The sum of the two products leads to covariance matrices that have a *symmetric* Kronecker product form, *i.e.* $\forall \mathbf{A}, \mathbf{B} \in \mathbb{R}^{N \times N} : \psi(\mathbf{A}, \mathbf{B}) = \mathbf{A} \otimes \mathbf{B} \in \mathbb{R}^{N^2 \times N^2}$ (recall Chapter 6). As in the previous part, this will allow a sufficiently efficient evaluation of the posterior. Fig. 4 visualizes the variance and shows samples from this prior for the toy setup from Fig. 3.

ground-truth in \mathbb{R}^2




The kernel k , here a squared exponential (Eq. 43) is assumed to be an unknown function.

prior in \mathbb{R}^2




Section 11.1 describes a Kronecker-structured Gaussian process prior over the kernel. Above pictures show from left-to-right: prior mean (zero), prior standard deviation, the absolute error divided by the standard deviation minus one and a sample from this prior.

likelihood on \mathbb{N}^2



Observations of k stem from matrix-vector multiplications with the kernel matrix \mathbf{K} (Section 11.2), sketched using random columns of the identity matrix.

posterior in \mathbb{R}^2



The posterior is again Gaussian (Section 11.3) and similar to the top, the pictures show from left-to-right: mean, standard deviation, relative error and a sample. By design, the posterior mean k_M is an approximation of finite rank which allows to efficiently solve the original least-squares problem (Section 10.2).

Figure 4: schematic summary of the proposed kernel approximation method.

This choice of prior offers a trade-off between efficient tractable inference and the desire to encode as much prior structural information about the kernel as possible. One desirable property to encode is symmetry, and indeed, matrix-valued functions sampled from this prior distribution are symmetric (*c.f.* Fig. 4 for examples, Appendix E.1 for formal proof). Kernel functions are also positive definite. Unfortunately, since the positive definite cone is not a linear sub-space of the vector-space of real matrices, this property can not be encoded in a Gaussian prior in closed form.¹ However, it *is* possible to guarantee positive-definiteness of the posterior mean point estimate through the specific choice of prior parameters $k_0 = 0$ (proof in Appendix E.2.1). For this reason, $k_0 := 0$ for the remainder. There are other properties of certain kernels that would be desirable to encode, but which are not feasible within the chosen framework without also sacrificing fast computability at the same time. For example, stationarity of the kernel can not be represented by a prior with Kronecker structure in the covariance since \mathbf{a} and \mathbf{b} (and symmetrically \mathbf{c} and \mathbf{d}) do not appear together as arguments to w .

The question remains how to choose w . Recall that w should reflect the similarity between $k(\mathbf{a}, \mathbf{b})$ and $k(\mathbf{c}, \mathbf{d})$ which depends on the similarity of \mathbf{a} and \mathbf{c} , and \mathbf{b} and \mathbf{d} . To measure the relationship between inputs is exactly the purpose of the kernel k and we therefore set

$$w := k$$

for the remainder. Even if k fails to capture similarity between inputs, as choice for w it still captures the similarity between the kernel values. Second, samples from the approximate kernel will be a function of w and lastly, this choice is convenient computationally as expressions simplify.

11.2 Likelihood

Having specified a prior over k , we will now be concerned with how to obtain observations. We can use matrix-vector products with the kernel matrix for learning a low-rank version of the kernel by introducing the linear operator

$$\mathbf{T}_{\mathbf{p}} : k \mapsto \text{vec} \left(\left[\iint k(\mathbf{x}, \mathbf{z}) p_i(\mathbf{x}) p_j(\mathbf{z}) \, d\mathbf{x} \, d\mathbf{z} \right]_{ij} \right) \quad (32)$$

where $i, j = 1 \dots P$, $\mathbf{p} = [p_1, \dots, p_P]$ are densities or distributions.

Example 18 (Matrix-vector multiplication). *Define $\mathbf{T}_{\mathbf{p}}$ with*

$$p_i(\mathbf{x}) = \sum_{j=1}^M s_{ij} \delta(\mathbf{x} - \mathbf{x}_{u_j}). \quad (33)$$

¹ For example Hennig (2015) discusses this problem and possible solutions.

Then the evaluation of $\mathbf{T}_p k$ reduces to a matrix vector product, that is $\text{mat}(\mathbf{T}_p k) = \mathbf{S}^\top k(\mathbf{X}_U, \mathbf{X}_U) \mathbf{S}$ where $S_{ij} = s_{ij}$, $\mathbf{X}_U = [\mathbf{x}_{u_1}, \dots, \mathbf{x}_{u_M}]$.

The \mathbf{x}_{u_j} can be datapoints or arbitrary elements of the domain \mathbb{X} . In Chapter 12 we will use the conjugate gradients search directions.

Example 19 (Integrals with Eigenfunctions). *Let ϕ_i $i = 1, \dots, P$ be orthogonal Eigenfunctions of k with respect to a density v on \mathbb{X} , i.e.*

$$\int k(\mathbf{x}, \mathbf{z}) \phi_i(\mathbf{z}) v(\mathbf{z}) \, d\mathbf{z} = \lambda_i \phi_i(\mathbf{x})$$

$$\int \phi_i(\mathbf{z}) \phi_j(\mathbf{z}) v(\mathbf{z}) \, d\mathbf{z} = \delta_{ij}$$

where $\lambda_i \in \mathbb{R}_+$ and δ_{ij} is the Kronecker delta (compare Rasmussen and Williams (2006, p. 96)). Then for

$$p_i(\mathbf{x}) = \phi_i(\mathbf{x}) v(\mathbf{x})$$

the observations $[\text{mat}(\mathbf{T}_p k)]_{ij} = \delta_{ij} \lambda_i$ are spectral values of the kernel.

In essence, this example shows another possibility to express prior knowledge over the kernel. This likelihood leads to the *Projected Bayes Regressor* (Trecate, Williams, and Opper 1999), which is a historical, deterministic precursor to the more widely known random Fourier feature expansion of Rahimi and Recht 2008. For the purposes of this thesis, Example 19 is a motivation for the generality of the observation operator in Eq. (32). The example will not be considered further.

11.3 Posterior

The observation operator \mathbf{T}_p is linear, and hence transforms the Gaussian prior into an also Gaussian posterior. Given the prior (Eq. (30)) and any likelihood of the previous section, the posterior is Gaussian with:

$$p(k \mid \mathbf{Y}, \mathbf{T}_p) = \mathcal{N}(k_M, w_M)$$

$$k_M = k_0 + (\mathbf{T}_p \psi)^\top (\mathbf{T}_p (\mathbf{T}_p \psi)^\top)^{-1} (\text{vec}(\mathbf{Y}) - \mathbf{T}_p k_0) \quad (34)$$

$$\psi_M = \psi - (\mathbf{T}_p \psi)^\top (\mathbf{T}_p (\mathbf{T}_p \psi)^\top)^{-1} \mathbf{T}_p \psi \quad (35)$$

The concrete posterior depends on the choice of \mathbf{T}_p . The following propositions presents an approximation method that has a view as GP inference with low-rank kernel and how it arises in our framework.

Proposition 20 (Subset of Regressors). *Consider the prior of Eq. (30) with $k_0 := 0$ and $w := k$ and the likelihood defined in Eq. (33) with $s_{ij} = \delta_{ij}$. Then the posterior mean k_M is equivalent to that of SoR:*

$$k_M(\mathbf{x}, \mathbf{z}) = k_{\text{SoR}} = k(\mathbf{x}, \mathbf{X}_U) k(\mathbf{X}_U, \mathbf{X}_U)^{-1} k(\mathbf{X}_U, \mathbf{z})$$

where \mathbf{X}_U are inducing inputs, not necessarily part of \mathbf{X} .

The proof is part of Appendix E.2. An example of this posterior distribution is shown in Fig. 4. Fig. 5 visualizes the progression of the posterior for the KMCG algorithm, presented in the next chapter.




Figure 5: progression of the posterior (Eq. (37)) for KMCG on the toy example from Fig. 3 for $P = 2, 4$ and 8 conjugate gradients steps. The columns show from left to right: mean, standard deviation, standardized error (white refers to perfect calibration, green to overconfidence and red to underconfidence) and a sample.

Conjugate Gradients for Kernel Machines

The previous chapter introduced a probabilistic estimation rule for the kernel k . This section presents another data-collection approach using conjugate gradients that leads to a new approximation algorithm: *kernel machine conjugate gradients* (KMCG).

The interest to use conjugate gradients for kernel machines goes back to more than 25 years (Skilling 1993) and is still continuing (Davies 2015; Filippone and Engler 2015). Albeit quadratic costs per step, CG has advantages over many of the approximation methods referenced in the introduction. CG has only one parameter, the desired precision, which is more natural than *e.g.* the number of inducing inputs for inducing point methods (Quiñonero-Candela and Rasmussen 2005). Thus the computational budget of CG is not fixed in advance but varies as necessary for the problem at hand.

The approach is to run conjugate gradients for P steps on a kernel matrix of size M and to treat the matrix multiplications (\mathbf{z}_i in Algorithm 2) as observations in the model presented in Chapter 11. Formally the likelihood is defined similar to the SoR likelihood (Example 18) albeit scaled.

Definition 21 (Conjugate-gradients likelihood). *Choose a subset \mathbf{X}_M of size M from \mathbf{X} and denote as $\mathbf{y}_M \in \mathbb{R}^M$ the vector that contains the corresponding entries of \mathbf{y} . Run conjugate gradients (Algorithm 2 on p. 15) with $\mathbf{x}_0 := \mathbf{0}$, $\mathbf{A} = k(\mathbf{X}_M, \mathbf{X}_M)$, $\mathbf{b} = \mathbf{y}_M$ and $\epsilon := 0.01\|\mathbf{b}\|_2$. In Eq. (32) set*

$$p_i(\mathbf{x}) := \sum_{j=1}^M s_j \delta(\mathbf{x} - \mathbf{x}_j) \quad (36)$$

where s_j is the j -th entry of vector \mathbf{s}_i in iteration i of the CG algorithm.

Remark 22. *KMCG uses only the CG search directions $\mathbf{s}_1, \dots, \mathbf{s}_P$ and not the solution $\hat{\mathbf{x}}$.*

Using this likelihood, the resulting approximate kernel (Eq. (34)) and approximate Equations are (cf. Proposition 41):

$$\hat{k}_M(\mathbf{x}_*, \mathbf{x}_{**}) = k(\mathbf{x}_*, \mathbf{X}_M) \mathbf{S} (\mathbf{S}^\top \mathbf{K}_M \mathbf{S})^{-1} \mathbf{S}^\top k(\mathbf{X}_M, \mathbf{x}_{**}) \quad (37)$$

$$\hat{f}(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{X}_M) \mathbf{S} (\mathbf{R}^\top \mathbf{R} + \sigma^2 \mathbf{S}^\top \mathbf{K}_M \mathbf{S})^{-1} \mathbf{R}^\top \mathbf{y} \quad (38)$$

$$\begin{aligned} \hat{c}(\mathbf{x}_*, \mathbf{x}_{**}) &= k(\mathbf{x}_*, \mathbf{x}_{**}) \\ &\quad - k(\mathbf{x}_*, \mathbf{X}_M) \mathbf{S} (\mathbf{S}^\top \mathbf{K}_M \mathbf{S})^{-1} \mathbf{S}^\top k(\mathbf{X}_M, \mathbf{x}_{**}) \end{aligned} \quad (39)$$

$$\begin{aligned}
& + \sigma^2 k(\mathbf{x}_*, \mathbf{X}_M) \mathbf{S} \left(\mathbf{R}^\top \mathbf{R} + \sigma^2 \mathbf{S}^\top \mathbf{K}_M \mathbf{S} \right)^{-1} \mathbf{S}^\top k(\mathbf{X}_M, \mathbf{x}_{**}) \\
\ln \hat{Z} = & \frac{1}{2\sigma^2} (\mathbf{y}^\top \mathbf{y} - \mathbf{y}^\top \mathbf{R} (\mathbf{R}^\top \mathbf{R} + \sigma^2 \mathbf{S}^\top \mathbf{K}_M \mathbf{S})^{-1} \mathbf{R}^\top \mathbf{y}) \quad (40) \\
& + \frac{1}{2} \ln |\mathbf{R}^\top \mathbf{R} + \sigma^2 \mathbf{S}^\top \mathbf{K}_M \mathbf{S}| - \frac{1}{2} |\mathbf{S}^\top \mathbf{K}_M \mathbf{S}| \\
& + \frac{1}{2} (N - P) \ln \sigma^2 + \frac{1}{2} N \ln 2\pi
\end{aligned}$$

where $\mathbf{S} := [\mathbf{s}_1, \dots, \mathbf{s}_P]$, $\mathbf{R} := k(\mathbf{X}_M, \mathbf{X})\mathbf{S}$ and P is the number of CG iterations.

Algorithm 4: Kernel Machine Conjugate Gradients

```

1 procedure KMCG( $k, \mathbf{X}, \mathbf{y}, \sigma^2, \epsilon$ )
2   ▷ (W.l.o.g.) assume that the inducing inputs are a subset of  $\mathbf{X}$ .
3   ▷ Denote this subset by  $\mathbf{X}_M$ .
4   ▷ Let  $\mathbf{y}_M$  the be corresponding entries of  $\mathbf{y}$ .
5   Conjugate Gradients( $k(\mathbf{X}_M, \mathbf{X}_M), \mathbf{y}_M, \epsilon$ ) // ignore solution  $\hat{\mathbf{x}}$ 
6    $\mathbf{S} \leftarrow [\mathbf{s}_1, \dots, \mathbf{s}_P]$  // collect CG search directions
7    $\mathbf{Z} \leftarrow [\mathbf{z}_1, \dots, \mathbf{z}_P]$  //  $\mathbf{Z} = \mathbf{K}_M \mathbf{S}$ 
8   if  $M < N$  then
9     |  $\mathbf{R} \leftarrow k(\mathbf{X}, \mathbf{X}_M)\mathbf{S}$ 
10  else
11    |  $\mathbf{R} \leftarrow \mathbf{Z}$  // When  $\mathbf{X}_M = \mathbf{X}$  above matrix multiplication is not necessary.
12  end if
13   $\mathbf{L}_1 \leftarrow \text{chol}(\mathbf{S}^\top \mathbf{Z})$  // precompute required Choleskies
14   $\mathbf{L}_2 \leftarrow \text{chol}(\sigma^2 \mathbf{S}^\top \mathbf{Z} + \mathbf{R}^\top \mathbf{R})$ 
15  evaluate Eqs. (38) to (40)
16 end procedure

```

12.1 Properties

Fig. 5 shows how the approximation to the kernel progresses for the toy example from Fig. 3. Computing the Cholesky of $\mathbf{R}^\top \mathbf{R} + \sigma^2 \mathbf{S}^\top \mathbf{K}_M \mathbf{S}$ costs $\mathcal{O}(NMP)$. After that evaluating the mean prediction is possible in $\mathcal{O}(M)$ and the variance in $\mathcal{O}(MP)$.

In case $P = M$, KMCG reduces to SoR since all occurrences of \mathbf{S} in Eq. (37) cancel and what remains is the SoR kernel (Eq. (28)). If \mathbf{K}_M has a favorable distribution of eigenvalues such that conjugate gradients terminates in less than M steps (see Section 5.2), KMCG can be used to speed up SoR.¹ In practice, this kind of advantage can only be expected to be beneficial when realized in low-level code. The level of efficiency of existing low-level linear algebra routines makes it challenging to evaluate this area.

Recall that the computational complexity of CG for the solution of Eq. (30) in P iterations is $\mathcal{O}(N^2P)$, that of inducing point methods with M inducing inputs is $\mathcal{O}(NM^2)$, and KMCG running for P iterations

¹ The same applies to related methods such as *DTC* (Quiñero-Candela and Rasmussen 2005) and Titsias' method (Titsias 2009).

on M inducing points has complexity $\mathcal{O}(NMP)$. While the main point of the present paper is to “fix” problems of CG in kernel machines, this structure hints at an interesting side-observation: Restricting the number of steps P in advance can then allow to increase the number of inducing points M beyond what would otherwise be feasible with standard inducing input methods. The subsequent evaluation section is dedicated to the case $M = N$, *i.e.* using the whole dataset which places KMCG in direct competition to plain conjugate gradients.

12.1.1 Relationship to the Nadaraya-Watson estimator

Taking only one step ($P = 1$) implies $\mathbf{S} = \mathbf{y}_M$ and Eq. (38) takes the following form

$$\hat{f}(\mathbf{x}_*) = \alpha \sum_{m=1}^M k(\mathbf{x}_m, \mathbf{x}_*) \mathbf{y}_m$$

where $\alpha = \frac{\mathbf{y}_M^\top \mathbf{K}_M \mathbf{y}_M}{\sigma^2 \mathbf{y}_M^\top \mathbf{K}_M \mathbf{y}_M + \mathbf{y}_M^\top \mathbf{K}_M \mathbf{K}_M \mathbf{y}_M}$. The equation bears resemblance to the Nadaraya-Watson estimator (Bishop 2006, p. 301f): a sum over all training targets weighted by the similarity of the corresponding input to the test input. However, the scaling-factor α is different.

12.1.2 Uncertainty

In addition to the posterior mean k_M , the Gaussian formulation of the approximation problem also provides a posterior variance ψ_M . It is a natural question to which degree this object can be interpreted as a notion of uncertainty or, more specifically, as an estimate of the square error $(k - k_M)^2$. This section provides an analysis of this covariance for KMCG, showing it to be an outer bound on the true error. Fig. 5 visualizes this for the toy dataset from Fig. 3.

Proposition 23 (relative error bound). *The relative size of estimation error and error estimate is bounded from above by 2.*

$$\frac{(k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z}))^2}{\psi_M(k(\mathbf{x}, \mathbf{z}), k(\mathbf{x}, \mathbf{z}))} \leq 2 \quad (41)$$

Proof. Define $\mathbf{k}_x^\top := k(\mathbf{x}, \mathbf{X})$ and $\mathbf{G} := \mathbf{S}(\mathbf{S}^\top \mathbf{K} \mathbf{S})^{-1} \mathbf{S}^\top$. For KMCG posterior mean and variance evaluate to (Appendix E.2):

$$\begin{aligned} k_M(\mathbf{x}, \mathbf{z}) &= \mathbf{k}_x^\top \mathbf{G} \mathbf{k}_z, \\ \psi_M(k(\mathbf{x}, \mathbf{z}), k(\mathbf{x}, \mathbf{z})) &= \frac{1}{2} \left(k(\mathbf{x}, \mathbf{x}) k(\mathbf{z}, \mathbf{z}) + k(\mathbf{x}, \mathbf{z})^2 \right) \\ &\quad - \frac{1}{2} \left(\mathbf{k}_x^\top \mathbf{G} \mathbf{k}_x \mathbf{k}_z^\top \mathbf{G} \mathbf{k}_z + (\mathbf{k}_x^\top \mathbf{G} \mathbf{k}_z)^2 \right) \\ &= \frac{1}{2} \left(k(\mathbf{x}, \mathbf{x}) k(\mathbf{z}, \mathbf{z}) + k(\mathbf{x}, \mathbf{z})^2 \right) \end{aligned}$$

$$- \left(k_M(\mathbf{x}, \mathbf{x})k_M(\mathbf{z}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z})^2 \right).$$

As a variance $\psi_M(k(\mathbf{x}, \mathbf{x}), k(\mathbf{x}, \mathbf{x}))$ is always larger than 0 which implies $k(\mathbf{x}, \mathbf{x}) \geq k_M(\mathbf{x}, \mathbf{x})$ for all \mathbf{x} . Thus $\psi_M(k(\mathbf{x}, \mathbf{z}), k(\mathbf{x}, \mathbf{z}))$ is bounded from below by $\frac{1}{2}k(\mathbf{x}, \mathbf{z})^2 - \frac{1}{2}k_M(\mathbf{x}, \mathbf{z})^2$ from which we can conclude

$$\begin{aligned} \frac{(k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z}))^2}{\psi_M(k(\mathbf{x}, \mathbf{z}), k(\mathbf{x}, \mathbf{z}))} &\leq 2 \frac{(k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z}))^2}{k(\mathbf{x}, \mathbf{z})^2 - k_M(\mathbf{x}, \mathbf{z})^2} \\ &= 2 \frac{(k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z}))^2}{(k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z}))(k(\mathbf{x}, \mathbf{z}) + k_M(\mathbf{x}, \mathbf{z}))} \\ &= 2 \frac{|k(\mathbf{x}, \mathbf{z}) - k_M(\mathbf{x}, \mathbf{z})|}{k(\mathbf{x}, \mathbf{z}) + k_M(\mathbf{x}, \mathbf{z})} \\ &\leq 2. \end{aligned}$$

□

12.2 Related Work

In terms of using conjugate gradients for kernel machines there is related work by Filippone and Engler (2015). Their algorithm ULISSE is aimed at the estimation of the marginal likelihood $p(\boldsymbol{\theta} | \mathbf{y})$ where $\boldsymbol{\theta}$ are hyper-parameters of the kernel k . They use a randomized conjugate gradients to estimate gradients of the log-marginal likelihood (Eq. (3)) which in combination with Stochastic Gradient Langevin Dynamics (SGLD) (Welling and Teh 2011) allows to sample from $p(\boldsymbol{\theta} | \mathbf{y})$. Our work is complementary to ULISSE. While running CG the matrix multiplications the inference perspective in Chapter 11 can be used to build a low-rank approximation of the kernel matrix which can serve as preconditioner for the next SGLD step.

Using the Kronecker product for efficient inference has been explored before for example in the KISS-GP framework (Wilson and Nickisch 2015). The difference to this work is that Wilson and Nickisch (2015) factorize the kernel matrix \mathbf{K} into a Kronecker-product where here it is the covariance matrix of the prior $\psi(\mathbf{K}, \mathbf{K})$ over the kernel that has Kronecker structure (cf. Eq. (30)). A synergy between their and our approach is hard to imagine. However, the follow-up work by Pleiss, Gardner, Weinberger, and Wilson (2018) uses Lanczos iteration to build a low-rank approximation of a kernel matrix \mathbf{C} for the variance prediction. Presumably, one could use instead KMCG.

Empirical Comparison of CG and KMCG

This section elaborates the conceptual differences between CG and KMCG and then compares both algorithms with numerical experiments. Consider Eq. (1) restated below for convenience.

$$\bar{f}(\mathbf{x}_*) = \mathbf{k}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \quad (1)$$

CG computes an approximation to $(\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$ and uses the exact \mathbf{k}_* . In contrast, KMCG computes an approximation to k and substitutes \mathbf{k}_* as well. That the systematic replacement of the kernel can be of importance has been noted before by Rasmussen and Williams (2006, p. 177) when comparing SoR and the Nyström method (Williams and Seeger 2001). The SoR method approximates k with the kernel in Eq. (28). In contrast Nyström uses the exact \mathbf{k}_* such that the predictive variance (Eq. (2)) can become negative. They further observed that for large M , Nyström and SoR have a similar performance, yet for small M Nyström performs poorly. We will make the same observations for CG and KMCG in the following comparison.

Conjugate gradients is used to solve the equations $(\mathbf{K} + \sigma^2 \mathbf{I}) \boldsymbol{\alpha} = \mathbf{y}$. In contrast, since the goal of KMCG is to learn an approximation to the kernel, the algorithm runs conjugate gradients on $\mathbf{K} \boldsymbol{\alpha} = \mathbf{y}$, *i.e.* without noise term. Both methods were evaluated in terms of the average relative error

$$\epsilon_f := \frac{1}{n_*} \sum_{k=1}^{n_*} \left| \frac{\bar{f}(\mathbf{x}_{*,k}) - \hat{f}(\mathbf{x}_{*,k})}{\bar{f}(\mathbf{x}_{*,k})} \right|, \quad (42)$$

where $\mathbf{x}_{*,k}$ is a test-input not part of the training set.

The text-book version of conjugate gradients in Algorithm 2 is known to be numerically unstable which is demonstrated in Appendix E.3.3, and there exist different strategies to cope with this problem (Golub and Van Loan 2013, p. 635). To explore the potential of KMCG, we bypass this implementation issue using the slowest¹ yet most stable solution: complete reorthogonalization Golub and Van Loan 2013, p. 564 and the explicit projection-method formulation, Eq. (FOM), to compute $\boldsymbol{\alpha}$. Therefore the following comparison will be conceptually, *i.e.* over the number of conjugate gradient steps. For completeness, Appendix E.3.1 contains results how KMCG performs in wall-clock time. Often the baseline methods converge faster since block-matrix multiplication is faster than looped matrix-vector multiplication. Baseline methods are the Fully Independent Training Conditional (FITC) approximation

¹ Computing the exact solution is actually faster.

(Quiñonero-Candela and Rasmussen 2005) and the Variational Free Energy (VFE) method (Titsias 2009) with inducing inputs randomly selected from the dataset as recommended by Chalupka, Williams, and Murray (2013). The baseline runs were repeated 10 times and besides the average, each figure will show also the progressive minimum and maximum over all runs to take into account for more elaborate inducing-input selection schemes.

In all our experiments, we used two popular stationary kernel functions: automatic relevance determination (ARD) Squared Exponential (Eq. (43)) and ARD Matérn 5/2 (Rasmussen and Williams 2006, p. 83f, p. 106),

$$k_{SE}(d(\mathbf{x}, \mathbf{z}; \mathbf{\Lambda})) = \theta_f \exp\left(-\frac{1}{2}d^2\right) \quad (43)$$

$$k_{52}(d(\mathbf{x}, \mathbf{z}; \mathbf{\Lambda})) = \theta_f \left(1 + \sqrt{5}d + \frac{5}{3}d^2\right) \exp\left(-\sqrt{5}d\right) \quad (44)$$

where $d = d(\mathbf{x}, \mathbf{z}; \mathbf{\Lambda}) = \|\mathbf{x} - \mathbf{z}\|_{\mathbf{\Lambda}}$ and $\mathbf{\Lambda}$ is a diagonal matrix. All experiments were executed with Matlab R2019a on an Intel i7 CPU with 32 Gigabytes of RAM running Ubuntu 18.04.

13.1 Common Regression Datasets

The datasets chosen are small such that computation of the exact GP is still feasible. Origin and purpose the datasets can be found in Table 3 in Appendix A. Each dataset has been shuffled and split into two sets, using one for training and the other for testing. For each dataset we optimized the kernel parameters running Carl Rasmussen’s `minimize` function² for 100 optimization-steps, where initially all kernel hyper-parameters are set to 1.

Fig. 6 shows how the average relative error develops for the described setup. Since the Matérn kernel experiments look very similar, these figures are part of Appendix E.3.2. The upper x-axis displays the number of conjugate gradients steps, the lower x-axis, the number of inducing inputs. During early iterations the performance of CG is not as reliable as KMCG and the latter also improves more consistently. For the baselines, the number of inducing inputs M was set to $M = \sqrt{NP}$ such that \mathcal{O} -notation costs are equivalent to KMCG. (Since KMCG uses multiplications with \mathbf{K} for observations, the costs per CG-step are $\mathcal{O}(N^2)$.) In comparison to the baselines, KMCG often provides a worse approximation to start with but exhibits a faster convergence rate.

In contrast to plain conjugate gradients, KMCG naturally provides estimates for variance (Eq. (2)) and evidence (Eq. (3)). Define the average relative errors ϵ_{var} and ϵ_{ev} analogously to Eq. (42), respectively. Figs. 7 and 8 show the average relative error of these estimates in comparison to the baselines. For all datasets one can observe that the

²This method is part of the GPML toolbox (Rasmussen and Nickisch 2010), see <http://www.gaussianprocess.org/gpml/code/matlab/doc>.

approximation quality of KMCG for the evidence (Eq. (3)) is improving at first and then worsening. KMCG is better at approximating the quadratic form than the determinant. Therefore, the approximation often ‘overshoots’.

The baselines clearly outperform KMCG in these experiments. A possible explanation is that the baselines provide a better overall-approximation to the kernel matrix: After P CG-steps, the KMCG kernel is of rank P whereas using M inducing inputs, the VFE kernel is of rank M (so is the FITC kernel, putting the diagonal correction aside). Since $M = \sqrt{NP}$, the baselines can afford more inducing inputs M than KMCG can afford CG-steps P . Overall, when it comes to real-time, the baselines are preferable over KMCG. The picture changes when matrix-multiplication is less expensive than $\mathcal{O}(N^2)$ which is investigated in the next section.

13.2 Grid-structured Datasets

In the previous section the baselines are the preferable estimators over KMCG. This changes when matrix-multiplication costs less than $\mathcal{O}(N^2)$. For example when the kernel is a product kernel (such as squared exponential) and the dataset has grid-structure, the costs for matrix-multiplication are almost linearly in the number of data-points (Wilson, Gilboa, Nehorai, and Cunningham 2014) such that the number of CG-steps KMCG can take, matches the number of baseline inducing inputs.

13.2.1 Artificial Datasets

The datasets considered in the following are artificial multi-dimensional grids. For the training set³, along each axis, G points are equally spaced in $[-G/4, G/4]$ distorted by Gaussian noise $\mathcal{N}(0, 10^{-3})$. One-hundred test inputs are uniformly distributed over the $[-G/4, G/4]$ cube. Targets are drawn from a Gaussian process with squared exponential kernel (length scales and amplitude equal to 1). The number of inducing inputs had to be capped at 500 due to memory limitations.

Fig. 9 shows how the approximation error to mean, variance and likelihood term evolves, zoomed in on the first 100 steps. In Appendix E.3.1, Fig. 16 shows the same comparison over time for the whole 500 steps, stopping KMCG when it becomes slower than the baselines. For reference, we include a 10×10 dataset to give an idea how each method would evolve when investing more computational power would be feasible.

On these datasets KMCG dominates the baseline methods. After already one-hundred CG-steps, KMCG provides a useful approximation

³ Computing the exact solution is feasible exploiting the Kronecker structure of the kernel matrix which we use to evaluate the quality of the approximation methods. However, we may imagine datapoints missing, s.t. matrix-vector multiplication is fast but computing the exact solution is not.




Figure 6: progression of the relative error ϵ_f as a function of the number of iterations of CG and KMCG for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

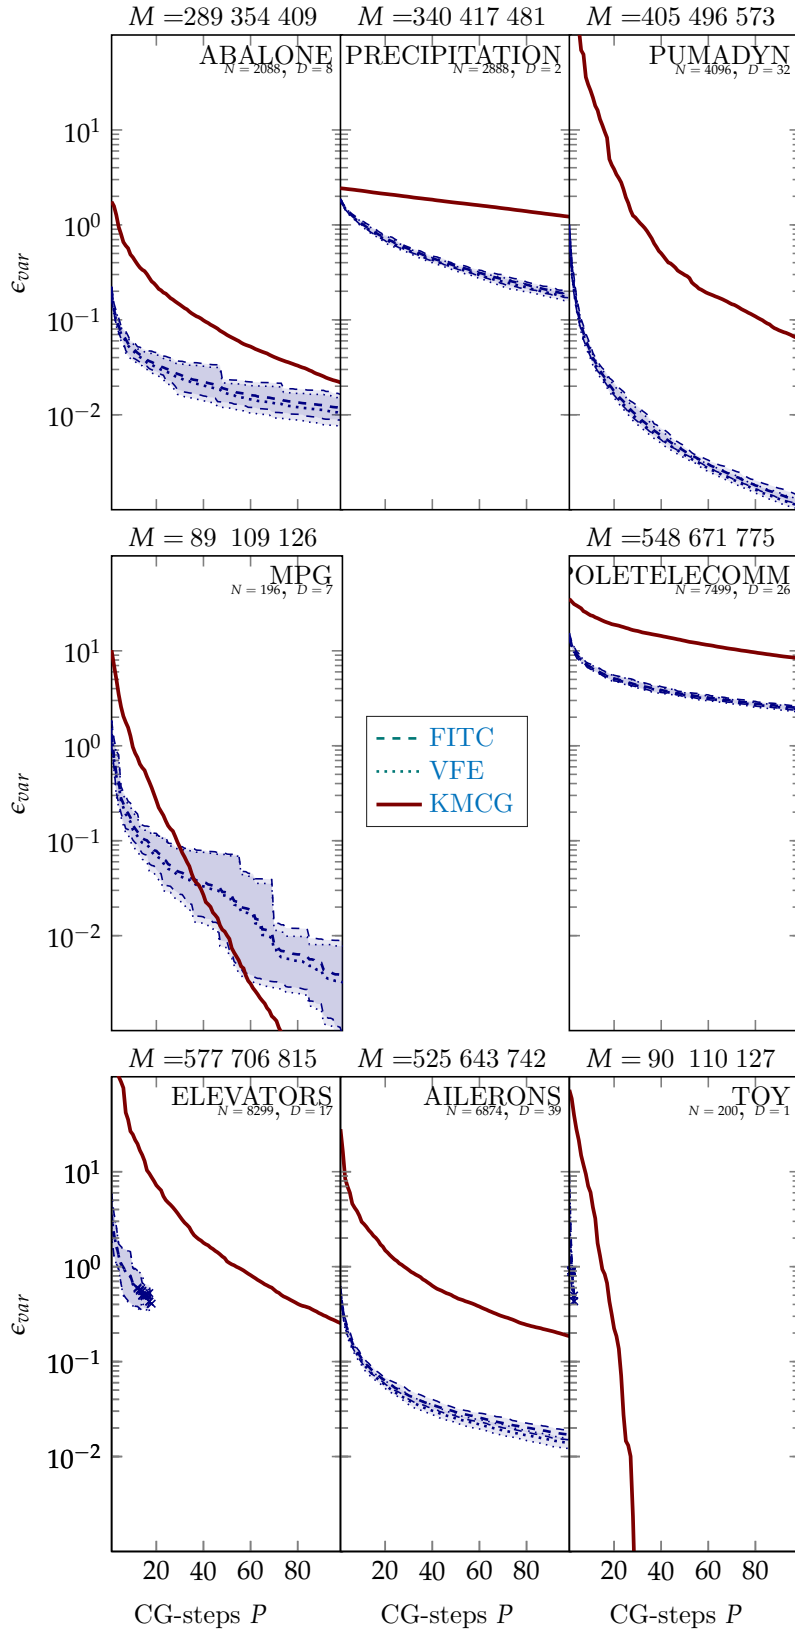


Figure 7: progression of the relative error of the variance ϵ_{var} as a function of the number of iterations of KMCG and baseline for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.




Figure 8: progression of the relative error of the evidence ϵ_{eV} as a function of the number of iterations of baseline and KMCG for different datasets using the squared-exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. The small spikes in the plots where KMCG appears to be close to the solution correspond to changes of the estimate from too small to too large.

to the posterior mean whereas the baselines hardly show any progress. For the variance, the same computational effort is not enough. Though the baselines find better solutions, all methods essentially fail to arrive at a satisfactory solution of a relative error below one. The issue is that all methods overestimate the posterior variance by two orders of magnitude. The picture is similar for the evidence, albeit the approximations are closer to the truth and KMCG performs slightly better on average.

13.2.2 Natural Sound Modeling

For a real-world example of a grid-structured dataset, we repeat the Natural Sound Modeling experiment considered by Turner (2010), Wilson and Nickisch (2015), and Dong, Eriksson, Nickisch, Bindel, and Wilson (2017). Given the intensity of a sound signal recorded over time, the objective is to recover the signal in missing regions. All inputs (*i.e.* including missing) are equidistant and hence the kernel matrix (over all inputs) is Toeplitz for stationary kernel. The kernel matrix over the given inputs is not Toeplitz, which forbids to use this structure for exact inference. Nevertheless matrix-vector-multiplication can be performed in linear time. We use the squared-exponential kernel with the hyper-parameters used by Dong, Eriksson, Nickisch, Bindel, and Wilson (2017). Since the exact posterior is infeasible to compute we report only the standardized mean squared-error:

$$SMSE := \frac{1}{\mathbb{V}[\mathbf{y}]} \sum_{j=1}^{N_*} (\mathbf{y}_{*,j} - \hat{f}(\mathbf{x}_{*,j}))^2.$$

To conform with the original experiment, we added for each baseline a run where the inducing inputs were chosen to be on a regular grid. The result of this run corresponds to the minimum. Fig. 10 confirms the observations of the previous section that KMCG arrives at satisfactory solutions faster than baseline, if matrix-vector multiplication is not an issue.

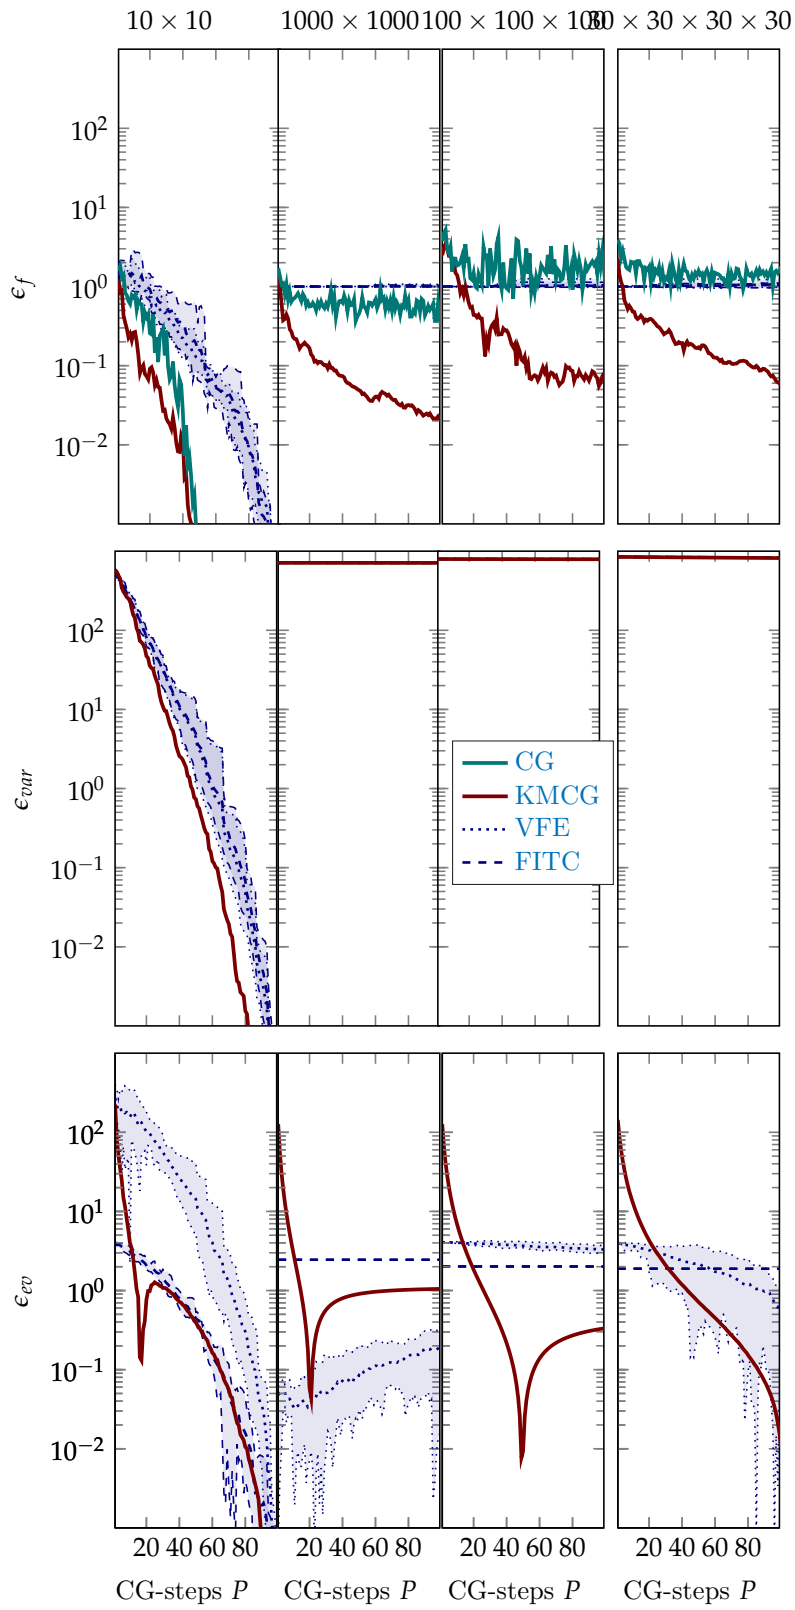


Figure 9: comparison of baseline and KMCG on grid-structured datasets using the squared exponential kernel (Eq. (43)). The shaded area visualizes minimum and maximum over all baseline runs.

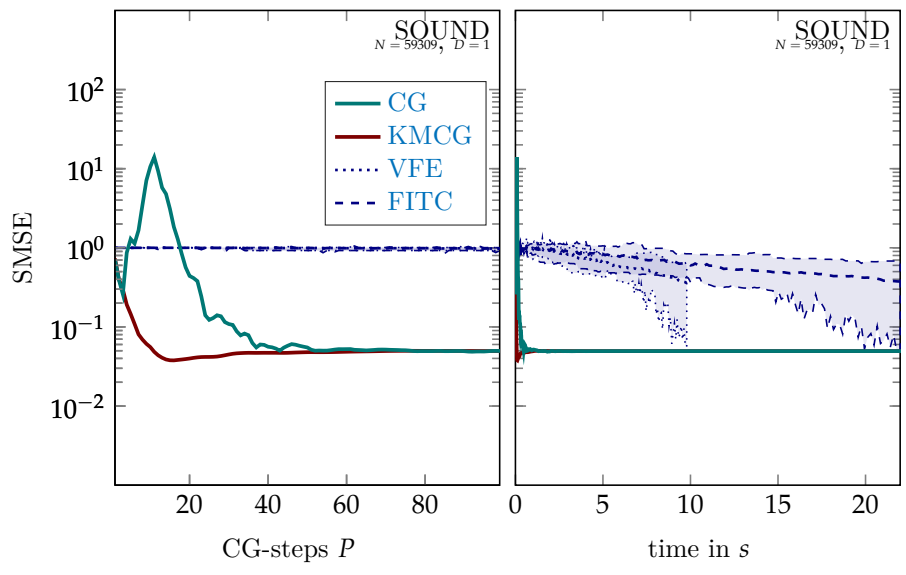


Figure 10: comparison of KMCG and CG on the SOUND dataset using the squared exponential kernel (Eq. (43)) with the hyper-parameters from Dong, Eriksson, Nickisch, Bindel, and Wilson (2017). The shaded area visualizes minimum and maximum over all baseline runs.

Discussion

14.1 Summary

This part presented a new approximate inference method for kernel machines that showed how linear solvers can be used in combination with low-rank kernel approximations. The approach is based on a probabilistic numerics viewpoint: the kernel k is treated as a latent quantity and conjugate gradients is used for collecting observations of k . By design, the resulting approximate kernel is of low rank and is plugged into the nonparametric least-squares problem. The approach is not restricted to least-squares problems but applicable in any scenario where the bottleneck is the inversion of a large kernel matrix.

Kernel machine conjugate gradients (KMCG) consistently outperforms plain conjugate gradients in numerical experiments. This does not change the fact that standard dense kernel least-squares problems are often more efficiently solved by inducing point methods. However, as demonstrated in Section 13.2, in the settings which allow fast multiplication with the kernel matrix, the new algorithm can improve upon the state of the art.

14.2 Future Directions

A hope associated with KMCG was that the probabilistic approach would allow to reason about the approximation error. However, this endeavor turned out to be more challenging than anticipated. For example, for an uncertainty over the approximate posterior mean, Eq. (1), this requires propagating the uncertainty over the kernel matrix through a matrix inverse operation. However, even if x is only univariate Gaussian, only the 0-th moment for $1/x$ exists. Hence, to derive a meaningful uncertainty, requires the additional step of formalizing another ground-truth. If there exist a lower bound or an upper bound on x which guarantee that x is necessarily larger or smaller than 0, one can use a truncated Gaussian. In that case higher moments of x^{-1} exist. Generalizing this approach to the multivariate case is more challenging.

Putting aside the issue of defining $p(\mathbf{K}^{-1})$, evaluating this expression will likely require approximation. Given a belief over \mathbf{K} , one direction

could be to use a consequence of the Cayley–Hamilton theorem: the inverse of an invertible $d \times d$ matrix \mathbf{A} can be written as

$$\mathbf{A}^{-1} = \sum_{j=0}^{d-1} c_j \mathbf{A}^j,$$

where the coefficients c_j also depend on \mathbf{A} . Using above expansion for \mathbf{K} and $\hat{\mathbf{K}}$, results in the following expression for the inverse:

$$\mathbf{K}^{-1} = \hat{\mathbf{K}}^{-1} + \sum_{j=0}^{N-1} (c_j \mathbf{K}^j - \hat{c}_j \hat{\mathbf{K}}).$$

Under the assumption that second order and higher polynomials contribute little, one can use the approximation

$$\mathbf{K}^{-1} \approx \hat{\mathbf{K}}^{-1} + (c_0 - \hat{c}_0) \mathbf{I} + c_1 \mathbf{K} - \hat{c}_1 \hat{\mathbf{K}}.$$

Treating c_0, c_1 as constant, above expression contains only \mathbf{K} as latent variable, and one obtains:

$$\begin{aligned} \mathbf{E}\mathbf{K}^{-1} &\approx \hat{\mathbf{K}}^{-1} + (c_0 - \hat{c}_0) \mathbf{I} + (c_1 - \hat{c}_1) \hat{\mathbf{K}} \text{ and} \\ \mathbf{V}\mathbf{K}^{-1} &\approx c_1^2 \mathbf{V}\mathbf{K}. \end{aligned}$$

To examine the usefulness of this idea, next steps would be to test the assumption if second order and above polynomials can be ignored, to find means of estimating $c_0 - \hat{c}_0$, $c_1 - \hat{c}_1$ and c_1^2 with reasonable effort, and to examine the relationship to a first-order Taylor expansion.

Part V

Probabilistic Kernel-Matrix Determinant Estimation

“This is just what you need, old pal. Me and you on a bender with a few beautiful ladies. I’m going over there.” “No.” “Oh, yes. I may be tiny, but I’ve got a certain *je ne sais quoi*.” “A certain what?” “I don’t know what,” admitted Zaphod, “But that’s never stopped me before.”

—Eoin Colfer, *And Another Thing...*

Preliminaries

15.1 Introduction

Whereas the previous parts have been concerned with the questions how new numerical approximation algorithms can be derived from a statistical perspective, this part deals with exploiting statistical properties of the numerical problem. The particular problem under consideration is the estimation of log-determinants of kernel matrices, a vital component of inference with models like determinantal point processes or Gaussian processes (*c.f.* Chapter 4). Given a set of $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{X}$ inputs and a kernel function k , the task is to compute

$$\ln |\mathbf{K}_N|$$

where

$$\mathbf{K}_N := \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \dots & k(\mathbf{x}_1, \mathbf{x}_N) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) & & \vdots \\ \vdots & & \ddots & \\ k(\mathbf{x}_N, \mathbf{x}_1) & \dots & & k(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix}.$$

Fig. 11 shows the progression of $\ln |\mathbf{K}_n + 0.001\mathbf{I}|$ with increasing n for different benchmark datasets and the squared-exponential kernel (Eq. (43)). The figure demonstrates that there exist cases in which a linear extrapolation from some $\ln |\mathbf{K}_n|$, $n < N$, is sufficient to estimate $\ln |\mathbf{K}_N|$. The success of such a strategy depends two factors: the kernel and the inputs $\mathbf{x}_1, \dots, \mathbf{x}_N$.

The goal of this part is to develop a stopping strategy that recognizes such “easy” cases while computing the log-determinant and terminates the calculation if possible. Early stopping would be helpful during hyper-parameter tuning for Gaussian processes. For example, consider the case of a kernel matrix generated from an ARD squared-exponential kernel (Eq. (43)) with lengthscales far too long for the dataset. In that case, a subset is sufficient to make accurate predictions and the progression of the determinant is easy to predict. An algorithm capable of early stopping, allows to try these kernel parameters “cheaply”. Contrarily, not early stopping can be a sign that the amount of data is not sufficient with respect to the flexibility of the model.

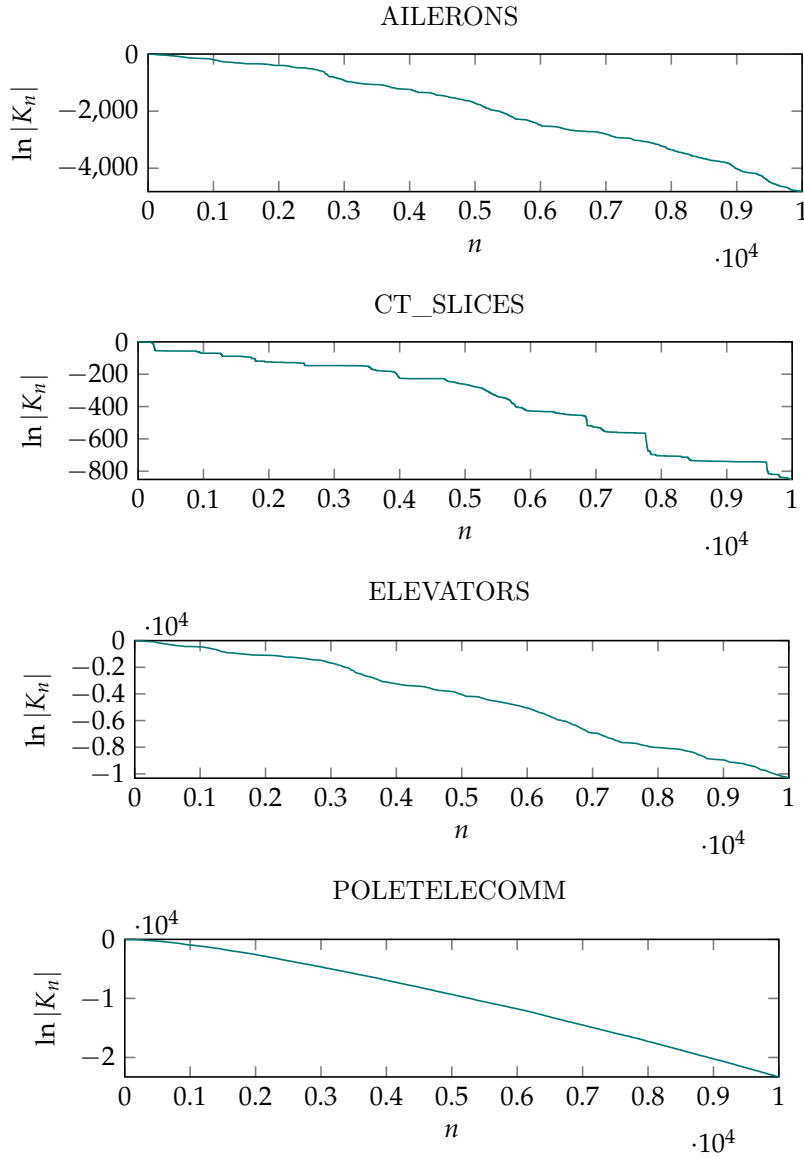


Figure 11: evolution of the log determinant $\ln|\mathbf{K}_n + 0.001\mathbf{I}|$ for common benchmark problems using the squared-exponential kernel Eq. (43). The datasets have been centered and standardized but are otherwise unprocessed. Note the different magnitudes of the y-axis. The figure demonstrates that depending on the kernel and the distribution of the dataset, there are cases in which a linear extrapolation from some $\ln|\mathbf{K}_n|$ is sufficient to estimate $\ln|\mathbf{K}_N|$ with high relative precision. More details about the datasets can be found in Appendix A.

15.2 The Cholesky Decomposition

To evaluate $\ln |\mathbf{K}_N|$ one usually computes the Cholesky decomposition \mathbf{C} which is a triangular matrix satisfying $\mathbf{C}\mathbf{C}^\top = \mathbf{K}_N$. Given the Cholesky, one obtains the log-determinant by using $\ln |\mathbf{K}_N| = 2 \sum_{j=1}^N \ln C_{jj}$ (Lemma 49, p. 130). In the following, we will focus on an implementation of the Cholesky decomposition that proceeds columnwise over the elements of \mathbf{K}_N , Algorithm 5. The motivation being, that a blocked version (Golub and Van Loan 2013, p. 170) of this algorithm is used in *OpenBLAS* (Wang, Zhang, Zhang, and Yi 2013). Presumably, closed source libraries, such as *Intel MKL*, use the same implementation. Observe that the costs per step j scale as $\mathcal{O}(j + (N - j) \cdot j)$, *i.e.* the costs increase until $j = N/2$ and then decrease again. Hence, the saved time when stopping “late” could be insignificant. However, in the algorithm as proposed by Cholesky (Benoit 1924), the effort per step scales as $\mathcal{O}(j^2)$, *s.t.* even when stopping only in one of the last steps, the saved effort is considerable.

Note that computing C_{jj} requires access only to the first $\mathbf{x}_1, \dots, \mathbf{x}_j$ datapoints. Under the assumption that the $\mathbf{x}_1, \dots, \mathbf{x}_N$ are independent and identically distributed, one can show that the C_{jj} decrease in expectation, *i.e.* $\mathbb{E}[C_{j+1,j+1}] \leq \mathbb{E}[C_{jj}]$ (see Lemma 32, p. 91). This allows to design a stopping condition that interrupts the calculation of the Cholesky decomposition and returns an estimate that is correct up to a relative error with high probability. In the following Chapter 16, we will consider a more general problem.

```

1  procedure CHOLESKY( $\mathbf{A}$ )
2      for  $j = 1, \dots, N$  do
3          for  $k = 1, \dots, j - 1$  do
4               $A_{jj} \leftarrow A_{jj} - A_{jk}A_{jk}$ 
5          end for
6           $A_{jj} \leftarrow \sqrt{A_{jj}}$ 
7          for  $\ell = j + 1, \dots, N$  do
8              for  $k = 1, \dots, j - 1$  do
9                   $A_{\ell j} \leftarrow A_{\ell j} - A_{ik}A_{jk}$ 
10             end for
11              $A_{\ell j} \leftarrow A_{\ell j} / A_{jj}$ 
12         end for
13     end for
14 end procedure

```

Algorithm 5: Cholesky decomposition according to Meister (2015, p. 48). The letters of the indices have been adapted to fit with the notation here. A blocked version of the algorithm can be found in Golub and Van Loan (2013, p. 170).

15.3 Martingales and Stopping Times

The following elaborations require some notation and terminology. For a more thorough introduction, these terms are discussed by Grimmett and Stirzaker (2001), and Davidson (1994).

For a monotonically increasing function $f : \mathbb{R} \mapsto \mathbb{R}$ and $\delta \in \mathbb{R}$, define $f^{-1}(\delta) := \arg \sup_{\epsilon \in \mathbb{R}} \{f(\epsilon) \leq \delta\}$. The following definitions can be more general, but for our purposes, discrete and positive indices in \mathbb{N} will be sufficient. A *filtration* is a sequence $(\mathcal{F}_j)_{j \in \mathbb{N}}$ of increasing σ -algebras, *i.e.* $\mathcal{F}_j \subseteq \mathcal{F}_{j+1}$ for all $j \in \mathbb{N}$. For random variables X_1, \dots, X_N denote by $\sigma(X_1, \dots, X_N)$ the generated σ -Algebra. A sequence of random variables $(X_j)_{j \in \mathbb{N}}$ is called *adapted* to a filtration, if X_j is \mathcal{F}_j -measurable for all $j \in \mathbb{N}$. Let $(X_j)_{j \in \mathbb{N}}$ be a sequence of random variables adapted to a filtration $\mathcal{F}_{j \in \mathbb{N}}$. A sequence $(X_j, \mathcal{F}_j)_{j \in \mathbb{N}}$ is called a *martingale* if

$$\mathbb{E}|X_j| < \infty, \text{ and}$$

$$\mathbb{E}[X_{j+1} \mid \mathcal{F}_j] = X_j \text{ almost surely,}$$

for all $j \in \mathbb{N}$. A sequence $(X_j, \mathcal{F}_j)_{j \in \mathbb{N}}$ is called a *martingale difference*, if

$$\mathbb{E}|X_j| < \infty, \text{ and}$$

$$\mathbb{E}[X_j \mid \mathcal{F}_{j-1}] = 0 \text{ almost surely,}$$

for all $j \in \mathbb{N}$. A random variable τ is called a *stopping time* (w.r.t. a filtration), if it takes values in \mathbb{N} and $\{\tau = j\} \in \mathcal{F}_j$ for all $j \in \mathbb{N}$.

filtration

adapted

martingale

martingale difference

stopping time

A Probably Approximately Correct Bound

16.1 Problem Definition

Let (Ω, \mathcal{F}, P) be a probability space and $(\mathcal{F}_j)_{j \in \{1, \dots, N\}}$ be a filtration. Furthermore, let $(f_j)_{j \in \mathbb{N}} \in [\mathbb{C}^-, \mathbb{C}^+]$ be a sequence of \mathcal{F}_j -measurable random variables that decrease conditionally in expectation,

$$\mathbb{E}[f_{j+1} \mid \mathcal{F}_j] \leq \mathbb{E}[f_j \mid \mathcal{F}_{j-1}] \quad (*)$$

for $j \in \{1, \dots, N-1\}$, where $\mathcal{F}_0 := \{\emptyset, \mathbb{R}\}$, and define

$$D_N := \sum_{j=1}^N f_j. \quad (45)$$

Given a desired precision $r \in (0, 1)$ and failure chance $\delta \in (0, 1)$, the problem is to devise a strategy that, being presented sequentially with the f_1, f_2, \dots , decides in each step whether to continue or to stop, and if stopping, provides an estimator \hat{D}_τ , s.t. its relative error is less than r with probability $1 - \delta$. Formally, the goal is to devise a stopping time τ and an estimator \hat{D}_τ , s.t.

$$P\left(\text{abs}\left(\frac{D_N - \hat{D}_\tau}{D_N}\right) > r\right) \leq \delta. \quad (46)$$

Remark 24. A trivial solution is to define $\tau := N$ and $\hat{D}_\tau := D_N$.

In absence of related work (c.f. Section 16.4), formal evaluation criteria for τ and \hat{D}_τ will not be developed. The goal will be to design a strategy that satisfies Eq. (46), and for at least one scenario: $\tau < N$.

16.2 Stopping Condition

The design of the stopping conditions are based on the following lemma.

Lemma 25 (Bound on Relative Error). *Let $D, \hat{D} \in [\mathcal{L}, \mathcal{U}]$, and assume $\text{sign}(\mathcal{L}) = \text{sign}(\mathcal{U}) \neq 0$. Then the relative error of the estimator \hat{D} can be bounded as*

$$\frac{\text{abs}(D - \hat{D})}{\text{abs}(D)} \leq \frac{\max(\mathcal{U} - \hat{D}, \hat{D} - \mathcal{L})}{\min(\text{abs}(\mathcal{L}), \text{abs}(\mathcal{U}))}.$$

Proof. The proof is part of Appendix F. □

Remark 26. *The bound is minimized for the choice $\hat{D} := 1/2(\mathcal{L} + \mathcal{U})$. In that case*

$$\frac{\text{abs}(D - \hat{D})}{\text{abs}(D)} \leq \frac{\mathcal{U} - \mathcal{L}}{2 \min(\text{abs}(\mathcal{L}), \text{abs}(\mathcal{U}))}.$$

In the following, we will devise lower bounds \mathcal{L}_n and upper bounds \mathcal{U}_n to D_N , respectively, for $n = 1, \dots, N$. The lower bounds \mathcal{L}_n will be deterministic, whereas $D_N \leq \mathcal{U}_n$ only with a certain probability. The stopping time τ will monitor these bounds and stop if they are large in magnitude (away from zero) and close enough that the relative error can not exceed the desired precision r .

Describing stopping strategy and estimator formally, will necessitate a number of definitions. The reader can quickly access all of them in Table 2 on p. 82. Define $C := C^+ - C^-$ and $l_j := f_j - C^-$. Observe that l_j is bounded from above by C and from below by 0. The latter property is the reason for introducing this definition.

Set

$$\mathcal{L}_n := NC^- + \sum_{j=1}^n l_j$$

and define an empirical mean of the “last” m elements, ending in n :

$$\hat{\mu}_n := \frac{1}{m} \sum_{j=n-m+1}^n l_j,$$

where $m \in \{1, \dots, N-1\}$ is a user defined nuisance parameter. Our estimator will be

$$\hat{D}_n := 1/2(\mathcal{L}_n + \mathcal{U}_n). \quad (47)$$

To construct \mathcal{U}_n , define the function

$$H_N(x) := H(x, N) = \mathbf{1}_{\{x \leq N\}} \sqrt{\left(\frac{N}{N+x}\right)^{N+x} \left(\frac{N}{N-x}\right)^{N-x}}$$

where $H(x, N)$ is defined in Theorem 29 (p. 84). The function $H_N(x)$ will be used to bound the probability that the upper bound \mathcal{U}_n fails. Fig. 12 visualizes this function for different N . Define the error tolerance for the mean $\hat{\mu}_n$ as $\epsilon_\mu := C/mH_m(\frac{\delta}{2|S|})$, where $S \subseteq \{m+1, \dots, N-1\}$ is another nuisance parameter. The set S defines possible stopping points. It is an artifact from the proof of Theorem 27 which may become obsolete (see Section 19.2). Define the error tolerance for the upper bound $\epsilon_n := (N-n)\epsilon_\mu + CH_N^{-1}(\delta/2)$, and the upper bound

$$\mathcal{U}_n := \mathcal{L}_n + \min((N-n)\hat{\mu}_n + \epsilon_n, (N-n)C).$$

$\mathcal{L}_n + (N - n)C$ is a deterministic upper bound to D_N and taking the minimum ensures that \mathcal{U}_n is never worse than that. At last, define the stopping time

$$\tau := \min \left\{ n \in S \mid \text{sign}(\mathcal{U}_n) = \text{sign}(\mathcal{L}_n) \neq 0, \text{ and} \right. \quad (\text{C1})$$

$$\left. \frac{\mathcal{U}_n - \mathcal{L}_n}{2 \min(\text{abs}(\mathcal{U}_n), \text{abs}(\mathcal{L}_n))} \leq r \right\} \cup \{N\}. \quad (\text{C2})$$

Note that the quantities in the stopping conditions are \mathcal{F}_n -measurable, *i.e.* τ defines indeed a stopping time.




Figure 12: the function $H_N(x)$ for different N . Increasing N requires a larger ϵ for $H_N(\epsilon)$ to fall below a certain threshold. However, in relation to N the increase is small. The thresholds $\delta/2$ and $\delta/2|S|$ become relevant for Fig. 14.

Theorem 27. Let (Ω, \mathcal{F}, P) be a probability space and $(\mathcal{F}_j)_{j \in \{1, \dots, N\}}$ be a filtration. Furthermore, let $(f_j)_{j \in \mathbb{N}} \in [C^-, C^+]$ be a sequence of \mathcal{F}_j -measurable random variables that decrease conditionally in expectation,

$$\mathbb{E}[f_{j+1} \mid \mathcal{F}_j] \leq \mathbb{E}[f_j \mid \mathcal{F}_{j-1}] \quad (*)$$

for $j \in \{1, \dots, N - 1\}$, where $\mathcal{F}_0 := \{\emptyset, \mathbb{R}\}$, and define

$$D_N := \sum_{j=1}^N f_j.$$

For $r, \delta \in (0, 1)$, the probability that the relative error of the estimator \hat{D}_τ defined by Eqs. (47), (C1) and (C2) is larger than r is less than δ .

$$P \left(\text{abs} \left(\frac{D_N - \hat{D}_\tau}{D_N} \right) > r \right) \leq \delta \quad (48)$$

The proof will be presented in Chapter 17.

Theorem 28. Assume $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{X}$ are independent and identically distributed. Denote with P the law of the $\mathbf{x}_1, \dots, \mathbf{x}_N$ and with \mathbf{C} the Cholesky decomposition of $\mathbf{K}_N + \sigma^2 \mathbf{I}$, where $\sigma^2 > 0$. Define

the probability space $(\mathbb{X}, \sigma(\mathbf{x}_1, \dots, \mathbf{x}_N), P)$ and the canonical filtration $\mathcal{F}_j := \sigma(\mathbf{x}_1, \dots, \mathbf{x}_j)$ for $j = 1, \dots, N$. Further, define

$$\begin{aligned} f_j &:= 2 \ln |\mathbf{C}_{jj}|, \text{ with} \\ \mathbf{C}^- &:= \ln \sigma^2 \text{ and} \\ \mathbf{C}^+ &:= \max_{j=1, \dots, N} \ln(k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2). \end{aligned}$$

Then, using the definitions of Theorem 27,

$$P \left(\frac{\text{abs}(|\mathbf{K}_N + \sigma^2 \mathbf{I}| - \hat{D}_\tau)}{\text{abs}(|\mathbf{K}_N + \sigma^2 \mathbf{I}|)} > r \right) \leq \delta.$$

The proof will be presented in Chapter 18.

16.3 Discussion

For Theorem 28 the quantities required for the stopping time are inexpensive to compute compared to the elements of the Cholesky. I modified the *OpenBLAS* (Wang, Zhang, Zhang, and Yi 2013) Cholesky decomposition to accommodate the stopping rule. Fig. 13 shows the relative overhead when computing the log-determinant of identity matrices of different size. The identity matrix is a special case, where early stopping is impossible—since $\ln |\mathbf{I}| = 0$, there exists no $\delta < 1$ for which a finite relative error could be guaranteed. One can see that with increasing dataset size the overhead becomes less than one percent on average. In that sense, Theorem 28 offers an “almost-free lunch”.




Figure 13: relative overhead of the OpenBLAS Cholesky decomposition with early stopping against without. Both algorithms compute the Cholesky decomposition of identity matrices. The nuisance parameters were set to $m := 256$ and $S := m \lfloor \{m, \dots, N - 1\} / m \rfloor$. The parameters r and δ are in this case irrelevant for the performance, since early stopping is impossible by design. With increasing matrix size, the measurement noise is decreasing and more runs remain below the one percent overhead mark. The experiments were executed on an Intel i7 CPU with 32 Gigabytes of RAM running Ubuntu 18.04.

Theorem 27 requires the user to define two nuisance parameters m and S . Fig. 14 demonstrates their influence on the bound \mathcal{U}_n and whether the stopping conditions can be met. The parameter m sets the number of elements l_j preceding the current step, used in the empirical mean $\hat{\mu}_n$ to build the upper bound \mathcal{U}_n . When setting the parameter too small, the target probability δ could be unattainable. On the other hand, for larger m , it is more likely that the upper bound \mathcal{U}_n is more conservative, hence delaying termination.

Furthermore, the user needs to select a subset $S \subseteq \{m, \dots, N - 1\}$ of possible stopping points. This is due to a union bound in the proof of

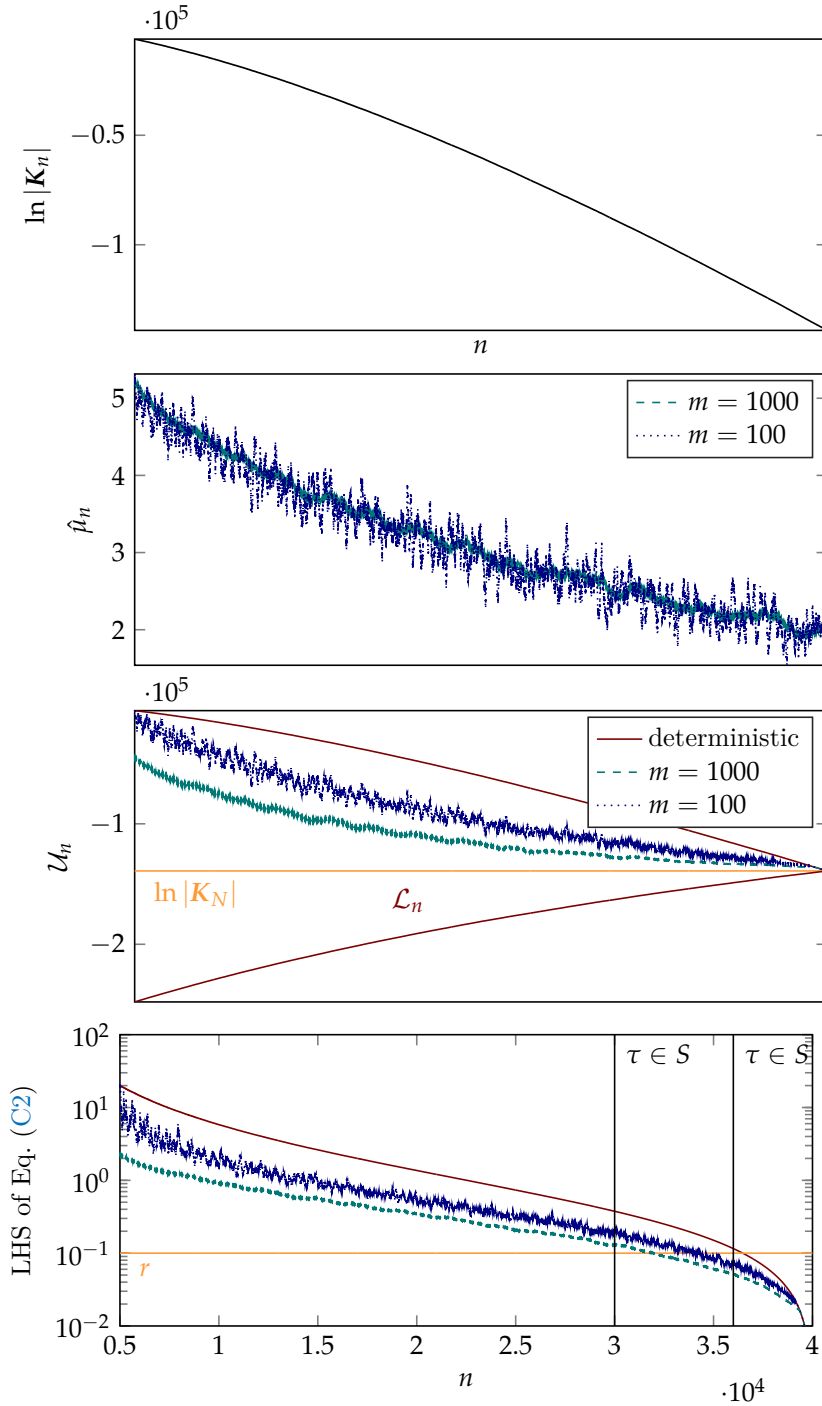


Figure 14: visualization of the quantities of Theorem 28 for one random shuffling of the SARCOS dataset ($N = 40000$), the squared-exponential kernel (Eq. (43)) and $\sigma^2 := 0.001$. The parameters have been set to $r := 0.1, \delta := 0.25, m := 100, 1000$ and $S := \{0.75N, 0.9N\}$. The deterministic upper bound is $\mathcal{U}_n := \mathcal{L}_n + (N - n)C^+$. It may seem surprising that the upper bound for $m = 100$ is higher than for $m = 1000$. However, recall that ϵ_μ has a term C/m . In the last picture, both stopping conditions are fulfilled for $\tau = 0.9N$, for both m . This figure demonstrates that Theorem 28 is not trivial. Arguably, it is desirable to achieve early stopping for less conservative parameters than $|S| = 2$ and $r = 0.1$. Section 19.2 discusses possibilities how to obtain more practical versions of Theorem 27.

Theorem 27 (Eq. (54)). If $|S|$ is too large, it may not be possible to stay below the desired probability δ . On the other hand, choosing fewer stopping points implies fewer opportunities to terminate. Section 19.2 discusses options how the union-bound could be avoided.

Fig. 14 demonstrates that Theorem 28 is not trivial: there exist computationally easy settings, it is possible to recognize these settings during runtime, and to guarantee a desired precision. However, the parameter configurations used in the experiments for Fig. 14 are pathological. The reader will agree that a relative precision of 0.1 and a failure chance of 25% can hardly be considered practically relevant values. In that sense, Theorem 27 should rather be taken as an encouraging feasibility study. Section 19.2 discusses possibilities how to obtain more practical versions.

16.4 Related Work

Most closely related is the work by Mnih, Szepesvári, and Audibert (2008) and references therein. They propose an algorithm called *EBStop* that returns an estimate of the mean of a sum of i.i.d. random variables. Similar to Theorem 27, they are able to guarantee a relative precision with high probability. However, Mnih, Szepesvári, and Audibert (2008) assume that the addends are independent and identically distributed, whereas Theorem 27 is more general and assumes only a (non-strict) decrease in expectation (Eq. (*)). Their approach is in a sense more sophisticated as they also monitor the empirical variance of the addends (*c.f.* remarks on future directions in Section 19.2). Bardenet, Doucet, and Holmes (2014) propose a related approach to stop Metropolis-Hastings sampling, but the algorithm is not applicable for the problem described in Section 16.1. Zhao, Zhou, Sabharwal, and Ermon (2016) present concentration inequalities for arbitrary stopping times, yet they also assume that the addends are independent and identically distributed.

The problem of estimating determinants has been studied extensively especially for symmetric and positive definite matrices (Skilling 1989; Dorn and Enßlin 2015; Fitzsimons, Granzio, Cutajar, Osborne, Filippone, and Roberts 2017; Fitzsimons, Cutajar, Osborne, Roberts, and Filippone 2017; Saibaba, Alexanderian, and Ipsen 2017; Boutsidis, Drineas, Kambadur, Kontopoulou, and Zouzias 2017). Yet, none of the aforementioned work considers kernel matrices in particular. In comparison, Theorem 28 offers an “almost-free lunch”. If early stopping is not possible, the algorithm returns the exact solution with negligible overhead. To my knowledge, Theorem 28 is the first distribution-free and deterministic approach to quantify uncertainty over approximate determinant estimation for kernel matrices.

I assume that the problem described in Section 16.1 has been approached by someone else before. However, I was unable to find literature that makes the same or less restrictive assumptions, as in the case described here. It appears as if the bandit and reinforcement learning community has not yet considered this scenario. A literature search on optimal stopping, empirical stopping, adaptive stopping, sequential analysis and racing, did not yield any further leads. An indicator why this problem may not have been considered yet is that the main tool for the proof of Theorem 27, Theorem 29 by Fan, Grama, and Liu (2012), has been published fairly recently.

Proof of Theorem 27

Theorem 27 from page 77 is restated here for the readers convenience. Recall that all definitions introduced in Section 16.2 are summarized in Table 2 on page 82.

Theorem 27. *Let (Ω, \mathcal{F}, P) be a probability space and $(\mathcal{F}_j)_{j \in \{1, \dots, N\}}$ be a filtration. Furthermore, let $(f_j)_{j \in \mathbb{N}} \in [C^-, C^+]$ be a sequence of \mathcal{F}_j -measurable random variables that decrease conditionally in expectation,*

$$\mathbb{E}[f_{j+1} \mid \mathcal{F}_j] \leq \mathbb{E}[f_j \mid \mathcal{F}_{j-1}] \quad (*)$$

for $j \in \{1, \dots, N-1\}$, where $\mathcal{F}_0 := \{\emptyset, \mathbb{R}\}$, and define

$$D_N := \sum_{j=1}^N f_j.$$

For $r, \delta \in (0, 1)$, the probability that the relative error of the estimator \hat{D}_τ defined by Eqs. (47), (C1) and (C2) is larger than r is less than δ .

$$P\left(\text{abs}\left(\frac{D_N - \hat{D}_\tau}{D_N}\right) > r\right) \leq \delta \quad (48)$$

PROOF OUTLINE The steps to bound the left hand side of Eq. (48) can be split into three parts. Section 17.1 contains the first part, where Eq. (48) is split into two cases. In the first case, D_N is smaller than the probabilistic upper bound \mathcal{U}_τ such that by the conditions of the stopping time τ , the probability of failure is zero. The second case is the main proof, which is to show that the probability $P(\mathcal{U}_\tau < D_N)$ is small. Part two, in Section 17.2, shows that $P(\mathcal{U}_\tau < D_N)$ is bound from above by two terms. The first term is essentially the probability that D_N is “much larger” than its expected value and we will show that this probability is below $\delta/2$. Section 17.3, the last part, is concerned with the second term which is the probability that the upper bound \mathcal{U}_τ is less than the expected value of D_N . There we will make use of Eq. (*) to show that this probability is also less than $\delta/2$.

Proof. As a preliminary observation note that

$$\begin{aligned} D_N &= \sum_{j=1}^N f_j \\ &\quad // \text{ by definition} \\ &= NC^- + \sum_{j=1}^N l_j \end{aligned}$$

Symbol	Definition	Intuition
l_j	$f_j - C^-$	the (positive) random elements
$\hat{\mu}_n$	$\frac{1}{m} \sum_{j=n-m+1}^n l_j$	mean estimate
\mathcal{L}_n	$NC^- + \sum_{j=1}^n l_j$	lower bound on D_N (deterministic)
C	$C^+ - C^-$	upper bound on l_j
$H_N(x)$	$\mathbf{1}_{\{x \leq N\}} \sqrt{\left(\frac{N}{N+x}\right)^{N+x} \left(\frac{N}{N-x}\right)^{N-x}}$	bounding function for failure probability (definition of $H(x, N)$ in Theorem 29, p. 84)
ϵ_μ	$C/mH_m^{-1} \left(\frac{\delta}{2 S }\right)$	error tolerance for the mean
ϵ_n	$(N-n)\epsilon_\mu + CH_N^{-1}(\delta/2)$	error tolerance for the upper bound
\mathcal{U}_n	$\mathcal{L}_n + \min((N-n)\hat{\mu}_n + \epsilon_n, (N-n)C)$	upper bound on D_N (probabilistic)
\hat{D}_n	$1/2(\mathcal{L}_n + \mathcal{U}_n)$	estimate for D_N
Eq. (C1)	$\text{sign}(\mathcal{U}_n) = \text{sign}(\mathcal{L}_n) \neq 0$	first condition for τ
Eq. (C2)	$\frac{\mathcal{U}_n - \mathcal{L}_n}{2 \min(\text{abs}(\mathcal{U}_n), \text{abs}(\mathcal{L}_n))} \leq r$	second condition for τ

Table 2: compact overview of all the definitions of Part v.

$$\begin{aligned}
& \text{// using the definition of } l_j \\
& = \mathcal{L}_n + \sum_{j=n+1}^N l_j \text{ for all } n = 0, \dots, N \tag{49} \\
& \text{// using the definition of } \mathcal{L}_n
\end{aligned}$$

and hence, for all $n = 0, \dots, N$, \mathcal{L}_n is a (deterministic) lower bound to D_N , since $l_j \geq 0$ by definition.

17.1 Using the Stopping Conditions

This section is the first part of the proof in which we split Eq. (48) into two cases and resolve the first case. We will make use of Lemma 25 (p. 75).

$$\begin{aligned}
& P\left(\text{abs}\left(\frac{D_N - \hat{D}_\tau}{D_N}\right) > r\right) \\
& = P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N \leq \mathcal{U}_\tau\right) \\
& \quad + P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N > \mathcal{U}_\tau\right) \\
& \text{// sum rule} \\
& \leq P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N \leq \mathcal{U}_\tau\right) + P(D_N > \mathcal{U}_\tau) \tag{50} \\
& \text{// upper-bounding joint by marginal}
\end{aligned}$$

Consider the first term. Recall that $\hat{D}_\tau \in [\mathcal{L}_\tau, \mathcal{U}_\tau]$.

$$\begin{aligned}
& P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N \leq \mathcal{U}_\tau\right) \\
& \leq P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N \leq \mathcal{U}_\tau, \text{sign}(\mathcal{L}_\tau) = \text{sign}(\mathcal{U}_\tau) \neq 0\right) \\
& \quad + P(\text{sign}(\mathcal{U}_\tau) \neq \text{sign}(\mathcal{L}_\tau)) \\
& \text{// sum rule and upper bounding joint by marginal} \\
& = P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, D_N \leq \mathcal{U}_\tau, \text{sign}(\mathcal{L}_\tau) = \text{sign}(\mathcal{U}_\tau) \neq 0\right) \\
& \text{// using the first stopping condition of } \tau, \text{ Eq. (C1)} \\
& = P\left(\frac{\text{abs}(D_N - \hat{D}_\tau)}{\text{abs}(D_N)} > r, \mathcal{L}_\tau \leq D_N \leq \mathcal{U}_\tau, \text{sign}(\mathcal{L}_\tau) = \text{sign}(\mathcal{U}_\tau) \neq 0\right) \\
& \text{// since } \mathcal{L}_\tau \text{ is a deterministic lower bound to } D_N, \text{ Eq. (49)} \\
& \leq P\left(\frac{\max(\mathcal{U}_\tau - \hat{D}_\tau, \hat{D}_\tau - \mathcal{L}_\tau)}{\min(\text{abs}(\mathcal{L}_\tau), \text{abs}(\mathcal{U}_\tau))} > r, \mathcal{L}_\tau \leq D_N \leq \mathcal{U}_\tau, \text{sign}(\mathcal{L}_\tau) = \text{sign}(\mathcal{U}_\tau) \neq 0\right) \\
& \text{// by Lemma 25}
\end{aligned}$$

$$\begin{aligned}
&= P \left(\frac{\mathcal{U}_\tau - \mathcal{L}_\tau}{2 \min(\text{abs}(\mathcal{L}_\tau), \text{abs}(\mathcal{U}_\tau))} > r, \mathcal{L}_\tau \leq D_N \leq \mathcal{U}_\tau, \text{sign}(\mathcal{L}_\tau) = \text{sign}(\mathcal{U}_\tau) \neq 0 \right) \\
&\quad // \text{ plugging in definition of } \hat{D}_\tau \\
&= 0 \\
&\quad // \text{ by the second condition of } \tau, \text{ Eq. (C2)}
\end{aligned}$$

17.2 D_N is probably close to its Expected Value

In this section we will consider the remaining term from Eq. (50). Again, this term will be split into two cases, where this section resolves the first and the second is taken care of in the next section. The following parts of the proof rely on a recent theorem presented by Fan, Grama, and Liu (2012).

Theorem 29 (Hoeffding's inequality for supermartingales (Fan, Grama, and Liu 2012)). *Assume that $(\xi_j, \mathcal{F}_j)_{j=1, \dots, N}$ are supermartingale differences satisfying $\xi_j \leq 1$. Then, for any $x \geq 0$ and $v > 0$,*

$$P \left(\sum_{j=1}^n \xi_j \geq x \text{ and } \sum_{j=1}^n \mathbb{V}[\xi_j | \mathcal{F}_{j-1}] \leq v \text{ for some } n \in [1, N] \right) \leq H_N(x, v),$$

where

$$H_N(x, v) := \mathbf{1}_{\{x \leq N\}} \left\{ \left(\frac{v}{v+x} \right)^{v+x} \left(\frac{N}{N-x} \right)^{N-x} \right\}^{\frac{N}{N+v}}.$$

To apply Theorem 29 define $Z'_j := l_j - \mathbb{E}[l_j | \mathcal{F}_{j-1}]$ and $Z_j := Z'_{\tau+j}$. Since $(Z'_j, \mathcal{F}_j)_{j \in \{1, N\}}$ is a martingale difference,

$$\left(Z_{\min(j, N)}, \mathcal{F}_{\min(\tau+j, N)} \right)_{j \in \mathbb{N}_0},$$

is a martingale difference as well (Corollary 48 in Appendix F).¹ The random variables Z_j/C are bounded from below by -1 and from above by 1 . Now consider the second, remaining term in Eq. (50) from page 83.

¹This follows from a result that is sometimes referred to as Doob's Optional Sampling Theorem (Theorem 46 in Appendix F).

$$\begin{aligned}
&P(D_N > \mathcal{U}_\tau) \\
&= P \left(L_\tau + \sum_{j=\tau+1}^N l_j > L_\tau + \min((N-\tau)\hat{\mu}_\tau + \epsilon_\tau, (N-\tau)C) \right) \\
&\quad // \text{ using Eq. (49) and definition of } \mathcal{U}_\tau \\
&= P \left(\sum_{j=\tau+1}^N l_j > \min((N-\tau)\hat{\mu}_\tau + \epsilon_\tau, (N-\tau)C) \right) \\
&\quad // \text{ simplifying} \\
&= P \left(\sum_{j=\tau+1}^N l_j > (N-\tau)\hat{\mu}_\tau + \epsilon_\tau \text{ or } \sum_{j=\tau+1}^N l_j > (N-\tau)C \right)
\end{aligned}$$

// exchanging min for logical or

$$= P \left(\sum_{j=\tau+1}^N l_j > (N - \tau) \hat{\mu}_\tau + \epsilon_\tau \right)$$

// since $l_j \leq C$

$$= P \left(\sum_{j=1}^{N-\tau} [Z_j + \mathbb{E}[l_{\tau+j} | \mathcal{F}_{\tau+j-1}]] > (N - \tau) \hat{\mu}_\tau + \epsilon_\tau \right)$$

// definition of Z_j

$$\leq P \left(\sum_{j=1}^{N-\tau} Z_j + \sum_{j=\tau+1}^N \mathbb{E}[l_j | \mathcal{F}_{j-1}] > (N - \tau) \hat{\mu}_\tau + \epsilon_\tau, \quad (51) \right.$$

$$\left. \sum_{j=\tau+1}^N \mathbb{E}[l_j | \mathcal{F}_{j-1}] \leq (N - \tau)(\hat{\mu}_\tau + \epsilon_\mu) \right)$$

$$+ P \left(\sum_{j=\tau+1}^N \mathbb{E}[l_j | \mathcal{F}_{j-1}] > (N - \tau)(\hat{\mu}_\tau + \epsilon_\mu) \right)$$

// sum rule and upper-bounding joint by marginal

Consider the first term in Eq. (51). We want to apply Theorem 29.

$$P \left(\sum_{j=1}^{N-\tau} Z_j + \sum_{j=\tau+1}^N \mathbb{E}[l_j | \mathcal{F}_{j-1}] > (N - \tau) \hat{\mu}_\tau + \epsilon_\tau, \right.$$

$$\left. \sum_{j=\tau+1}^N \mathbb{E}[l_j | \mathcal{F}_{j-1}] \leq (N - \tau)(\hat{\mu}_\tau + \epsilon_\mu) \right)$$

$$\leq P \left(\sum_{j=1}^{N-\tau} Z_j + (N - \tau)(\hat{\mu}_\tau + \epsilon_\mu) > (N - \tau) \hat{\mu}_\tau + \epsilon_\tau \right)$$

// combining the two events

$$= P \left(\sum_{j=1}^{N-\tau} Z_j > \epsilon_\tau - (N - \tau) \epsilon_\mu \right)$$

// simplifying

$$= P \left(\sum_{j=1}^{N-\tau} \frac{Z_j}{C} > H_N^{-1}(\delta/2) \right)$$

// definition of ϵ_τ and dividing by C

$$\leq P \left(\sum_{j=1}^n \frac{Z_j}{C} > H_N^{-1}(\delta/2) \text{ for some } n \in \{1, \dots, N\} \right) \quad (52)$$

// enlargening the event

$$= P \left(\sum_{j=1}^n \frac{Z_j}{C} > H_N^{-1}(\delta/2), \sum_{j=1}^n \mathbb{V}[Z_j/C | \mathcal{F}_{j-1}] \leq N \text{ for some } n \in \{1, \dots, N\} \right)$$

// by Popoviciu's inequality (Theorem 44 on p. 129): $\mathbb{V}[Z_j/C | \mathcal{F}_{j-1}] \leq 1$.

$$\leq H(H_N^{-1}(\delta/2), N) \quad (53)$$

// by Theorem 29, where H is defined in that theorem

$$= H_N(H_N^{-1}(\delta/2)) \leq \delta/2.$$

// definition of H_N

17.3 \mathcal{U}_τ is probably large enough

This part takes care of the second term in Eq. (51). In essence, we will show that \mathcal{U}_τ is large enough with high probability. Now we will make use of the assumption that the l_j decrease in expectation: Eq. (*). We will again apply Theorem 29.

$$\begin{aligned}
& P\left(\sum_{j=\tau+1}^N \mathbb{E}[l_j \mid \mathcal{F}_{j-1}] > (N - \tau)(\hat{\mu}_\tau + \epsilon_\mu)\right) \\
& \leq P\left(\mathbb{E}[l_{\tau+1} \mid \mathcal{F}_\tau] - \hat{\mu}_\tau > \epsilon_\mu\right) \\
& \quad \text{// using Eq. (*)} \\
& = P\left(\mathbb{E}[l_{\tau+1} \mid \mathcal{F}_\tau] - \frac{1}{m} \sum_{j=\tau-m+1}^{\tau} l_j > \epsilon_\mu\right) \\
& \quad \text{// definition of } \hat{\mu}_\tau \\
& = P\left(\sum_{j=\tau-m+1}^{\tau} (\mathbb{E}[l_{\tau+1} \mid \mathcal{F}_\tau] - l_j) > m\epsilon_\mu\right) \\
& \quad \text{// multiplication by } m \text{ and moving the conditional expectation into the sum} \\
& \leq P\left(\sum_{j=\tau-m+1}^{\tau} (\mathbb{E}[l_j \mid \mathcal{F}_{j-1}] - l_j) > m\epsilon_\mu\right) \\
& \quad \text{// using again Eq. (*)} \\
& = \sum_{s \in S} P\left(\sum_{j=\tau-m+1}^{\tau} (\mathbb{E}[l_j \mid \mathcal{F}_{j-1}] - l_j) > m\epsilon_\mu, \tau = s\right) \\
& \quad \text{// sum rule} \\
& = \sum_{s \in S} P\left(\sum_{j=s-m+1}^s (\mathbb{E}[l_j \mid \mathcal{F}_{j-1}] - l_j) > CH_m^{-1} \left(\frac{\delta}{2|S|}\right)\right) \quad (54) \\
& \quad \text{// definition of } \epsilon_\mu \\
& = \sum_{s \in S} P\left(\sum_{j=s-m+1}^s -\frac{Z'_j}{C} > H_m^{-1} \left(\frac{\delta}{2|S|}\right)\right) \\
& \quad \text{// definition of } Z'_j \\
& \leq \sum_{s \in S} P\left(\sum_{j=1}^{m'} -\frac{Z'_{s-m+j}}{C} > H_m^{-1} \left(\frac{\delta}{2|S|}\right) \text{ for some } m' \in \{1, \dots, m\}\right) \\
& \quad \text{// enlargening the event}
\end{aligned}$$

Changing the sign does not change the martingale difference property and hence, $(-Z'_j, \mathcal{F}_j)_{j \in \{s-m+1, \dots, s\}}$ is a martingale difference for all $s \in S$. We can apply the same argument as in Eq. (53).

$$\sum_{s \in S} P\left(\sum_{j=1}^{m'} -\frac{Z'_{s-m+j}}{C} > H_m^{-1} \left(\frac{\delta}{2|S|}\right) \text{ for some } m' \in \{1, \dots, m\}\right)$$

$$\leq \sum_{s \in S} H \left(H_m^{-1} \left(\frac{\delta}{2|S|} \right), m \right)$$

// using the same argument as in Eq. (53)

$$\leq \sum_{s \in S} \frac{\delta}{2|S|} = \frac{\delta}{2}$$

To see that this completes the proof, we recall below the proof outline from page 81. In Section 17.1 we showed that in order to proof

$$P \left(\text{abs} \left(\frac{D_N - \hat{D}_\tau}{D_N} \right) > r \right) \leq \delta,$$

using the definition of the stopping time τ , it is sufficient to show $P(D_N > \mathcal{U}_\tau) \leq \delta$. In Section 17.2, more specifically in Eq. (51), the term $P(D_N > \mathcal{U}_\tau)$ was split into a sum of two terms. The remainder of Section 17.2 showed that the first term is less than $\delta/2$. This section showed above that the second term is less than $\delta/2$. Hence, the proof is complete.

□

Application to Kernel Matrix Determinant Estimation

This chapter shows that the stopping time proposed in Theorem 27 can be used to terminate the computation of kernel matrix determinants. We will make the mild assumption that $\mathbf{x}_1, \dots, \mathbf{x}_n, \dots, \mathbf{x}_N \in \mathbb{X}$ are independent and identically distributed. This assumption is not always fulfilled, *e.g.* when the inputs are sorted. By shuffling the dataset, the assumption can be established.

In the following, we will consider the problem of evaluating $\ln |\mathbf{K}_N + \sigma^2 \mathbf{I}_N|$ for $\sigma^2 > 0$. This is a common expression required for example in Gaussian process regression where σ^2 represents the observational noise (see *e.g.* Rasmussen and Williams (2006, p. 13ff.)). In case $\sigma^2 = 0$ one needs another, strictly positive, lower bound on the smallest eigenvalue of \mathbf{K}_N .

Theorem 28. *Assume $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{X}$ are independent and identically distributed. Denote with P the law of the $\mathbf{x}_1, \dots, \mathbf{x}_N$ and with \mathbf{C} the Cholesky decomposition of $\mathbf{K}_N + \sigma^2 \mathbf{I}$, where $\sigma^2 > 0$. Define the probability space $(\mathbb{X}, \sigma(\mathbf{x}_1, \dots, \mathbf{x}_N), P)$ and the canonical filtration $\mathcal{F}_j := \sigma(\mathbf{x}_1, \dots, \mathbf{x}_j)$ for $j = 1, \dots, N$. Further, define*

$$\begin{aligned} f_j &:= 2 \ln |\mathbf{C}_{jj}|, \text{ with} \\ C^- &:= \ln \sigma^2 \text{ and} \\ C^+ &:= \max_{j=1, \dots, N} \ln(k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2). \end{aligned}$$

Then, using the definitions of Theorem 27,

$$P \left(\frac{\text{abs}(|\mathbf{K}_N + \sigma^2 \mathbf{I}| - \hat{D}_\tau)}{\text{abs}(|\mathbf{K}_N + \sigma^2 \mathbf{I}|)} > r \right) \leq \delta.$$

Proof. The proof follows from Theorem 27. To apply the theorem we need to show that for all $j = 1, \dots, N$, the \mathbf{C}_{jj} are functions of $\mathbf{x}_1, \dots, \mathbf{x}_j$ (Lemma 30, p. 89), that $f_j \in [C^-, C^+]$ (Lemma 31, p. 90), and that $\mathbb{E}[f_{j+1} | \mathcal{F}_j] \leq \mathbb{E}[f_j | \mathcal{F}_{j-1}]$ (Lemma 32, p. 91). \square

To proof the three lemmata referenced above, define

$$\begin{aligned} \mathbf{k}_j(\mathbf{x}) &:= [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_j)]^\top \in \mathbb{R}^j \text{ and} \\ \mathbf{k}_{j+1} &:= \mathbf{k}_j(\mathbf{x}_j) \in \mathbb{R}^j. \end{aligned}$$

Lemma 30. Denote with \mathbf{C}_N the Cholesky decomposition of $\mathbf{K}_N + \sigma^2 \mathbf{I}_N$. The n -th diagonal element of \mathbf{C}_N , squared, is the posterior variance of a GP conditioned on the previous data-points:

$$[\mathbf{C}_N]_{nn}^2 = k(\mathbf{x}_n, \mathbf{x}_n) + \sigma^2 - \mathbf{k}_n^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_n.$$

Proof. With abuse of notation, define $\mathbf{C}_1 := \sqrt{k(\mathbf{x}_1, \mathbf{x}_1)}$ and

$$\mathbf{C}_N := \begin{bmatrix} \mathbf{C}_{N-1} & \mathbf{0} \\ \mathbf{k}_N^\top \mathbf{C}_{N-1}^{-\top} & \sqrt{k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N} \end{bmatrix}.$$

We will show that the lower triangular matrix \mathbf{C}_N satisfies $\mathbf{C}_N \mathbf{C}_N^\top = \mathbf{K}_N + \sigma^2 \mathbf{I}_N$. Since the Cholesky decomposition is unique (Golub and Van Loan 2013, Theorem 4.2.7), \mathbf{C}_N must be the Cholesky decomposition of $\mathbf{K}_N + \sigma^2 \mathbf{I}_N$. Furthermore, by definition of \mathbf{C}_N , $[\mathbf{C}_N]_{NN}^2 = k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N$. The statement then follows by induction.

To remain within the text margins, define

$$x := \mathbf{k}_N^\top \mathbf{C}_{N-1}^{-\top} \mathbf{C}_{N-1}^{-1} \mathbf{k}_N + k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N.$$

We want to show that $\mathbf{C}_N \mathbf{C}_N^\top = \mathbf{K}_N + \sigma^2 \mathbf{I}_N$.

$$\begin{aligned} \mathbf{C}_N \mathbf{C}_N^\top &= \begin{bmatrix} \mathbf{C}_{N-1} & \mathbf{0} \\ \mathbf{k}_N^\top \mathbf{C}_{N-1}^{-\top} & \sqrt{k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N} \end{bmatrix} \\ &\quad \cdot \begin{bmatrix} \mathbf{C}_{N-1}^\top & \mathbf{C}_{N-1}^{-1} \mathbf{k}_N \\ \mathbf{0}^\top & \sqrt{k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{C}_{N-1} \mathbf{C}_{N-1}^\top & \mathbf{C}_{N-1} \mathbf{C}_{N-1}^{-1} \mathbf{k}_N \\ \mathbf{k}_N^\top \mathbf{C}_{N-1}^{-\top} \mathbf{C}_{N-1}^\top & x \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{K}_{N-1} + \sigma^2 \mathbf{I}_{N-1} & \mathbf{k}_N \\ \mathbf{k}_N^\top & x \end{bmatrix} \end{aligned}$$

Also x can be simplified further.

$$\begin{aligned} x &= \mathbf{k}_N^\top \mathbf{C}_{N-1}^{-\top} \mathbf{C}_{N-1}^{-1} \mathbf{k}_N + k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N \\ &= \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N + k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2 - \mathbf{k}_N^\top (\mathbf{K}_{n-1} + \sigma^2 \mathbf{I}_{n-1})^{-1} \mathbf{k}_N \\ &= k(\mathbf{x}_N, \mathbf{x}_N) + \sigma^2. \end{aligned}$$

□

Lemma 31. The $f_j := 2 \ln \mathbf{C}_{jj}$ are bounded between $\mathbf{C}^- := \ln \sigma^2$ and $\mathbf{C}^+ := \max_{j=1, \dots, N} \ln(k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2)$.

Proof. By Lemma 30, $\mathbf{C}_{jj}^2 = k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2 - \mathbf{k}_j^\top (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}_{j-1})^{-1} \mathbf{k}_j$ which is the posterior variance of a Gaussian process. The third term

in the posterior variance is always positive since $(\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}_{j-1})^{-1}$ is an s.p.d. matrix. Hence, $k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2$ is an upper bound to \mathbf{C}_{jj} . On the other hand, since k is a kernel, $k(\mathbf{x}_j, \mathbf{x}_j) - \mathbf{k}_j^\top (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}_{j-1})^{-1} \mathbf{k}_j$ can not be negative and σ^2 is therefore a lower bound to \mathbf{C}_{jj}^2 . Since both values are positive and the logarithm is an increasing function on the positive real axis, the proof is complete. \square

Lemma 32. *The $f_j := 2 \ln \mathbf{C}_{jj}$ decrease in expectation:*

$$\mathbb{E}[f_{j+1} \mid \sigma(\mathbf{x}_1, \dots, \mathbf{x}_j)] \leq \mathbb{E}[f_j \mid \sigma(\mathbf{x}_1, \dots, \mathbf{x}_{j-1})].$$

Proof. By Lemma 30, $\mathbf{C}_{jj}^2 = k(\mathbf{x}_j, \mathbf{x}_j) + \sigma^2 - \mathbf{k}_j^\top (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}_{j-1})^{-1} \mathbf{k}_j$ which is the posterior variance of a Gaussian process. Since a variance is always positive and the logarithm is an increasing function on the positive real axis it is sufficient to show that the \mathbf{C}_{jj}^2 decrease in expectation. Define the two shorthands $p_j := \mathbf{C}_{jj}^2$ and $q_j(\mathbf{x}) := \mathbf{k}_j(\mathbf{x})^\top (\mathbf{K}_j + \sigma^2 \mathbf{I}_j)^{-1} \mathbf{k}_j(\mathbf{x})$. We will show below, in Eq. (56), that $q_j(\mathbf{x}) = q_{j-1}(\mathbf{x}) + r_{j-1}(\mathbf{x})$ where $r_{j-1}(\mathbf{x}) \geq 0$.

$$\begin{aligned} & \mathbb{E}[p_{j+1} \mid \sigma(\mathbf{x}_1, \dots, \mathbf{x}_j)] \\ &= \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - \mathbf{k}_j(\mathbf{x})^\top (\mathbf{K}_j + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_j(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) \\ & \quad \text{// property of conditional expectation} \\ &= \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - q_j(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) \\ & \quad \text{// definition of } q_j(\mathbf{x}) \\ &= \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - q_{j-1}(\mathbf{x}) - r_{j-1}(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) \\ & \quad \text{// using Eq. (55)} \\ &= \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - q_{j-1}(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) - \int r_{j-1}(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) \\ & \quad \text{// splitting the integral} \\ &\leq \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - q_{j-1}(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_j) \\ & \quad \text{// using Eq. (56)} \\ &= \int k(\mathbf{x}, \mathbf{x}) + \sigma^2 - q_{j-1}(\mathbf{x}) P(\mathbf{d}\mathbf{x} \mid \mathbf{x}_1, \dots, \mathbf{x}_{j-1}) \\ & \quad \text{// with Fubini's theorem} \\ &= \mathbb{E}[p_j \mid \sigma(\mathbf{x}_1, \dots, \mathbf{x}_{j-1})] \\ & \quad \text{// property of conditional expectation} \end{aligned}$$

It remains to show $q_j(\mathbf{x}) = q_{j-1}(\mathbf{x}) + r_{j-1}(\mathbf{x})$ where $r_{j-1}(\mathbf{x}) \geq 0$. Define $\mathbf{v}_x := (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{j-1}(\mathbf{x})$. First note, that using block-matrix inversion we can write

$$(\mathbf{K}_j + \sigma^2 \mathbf{I})^{-1} = \begin{bmatrix} (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I})^{-1} + \mathbf{v}_x p_j^{-1} \mathbf{v}_x^\top & -\mathbf{v}_x p_j^{-1} \\ -\mathbf{v}_x^\top p_j^{-1} & p_j^{-1} \end{bmatrix}.$$

Using above observation, we can transform $q_j(\mathbf{x})$.

$$\begin{aligned} q_j(\mathbf{x}) &= \begin{bmatrix} \mathbf{k}_{j-1}(\mathbf{x})^\top & k(\mathbf{x}, \mathbf{x}_j) \end{bmatrix} \\ &\cdot \begin{bmatrix} (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I})^{-1} + \mathbf{v}_x p_j^{-1} \mathbf{v}_x^\top & -\mathbf{v}_x p_j^{-1} \\ -\mathbf{v}_x^\top p_j^{-1} & p_j^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{k}_{j-1}(\mathbf{x}) \\ k(\mathbf{x}, \mathbf{x}_j) \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{k}_{j-1}(\mathbf{x})^\top & k(\mathbf{x}, \mathbf{x}_j) \end{bmatrix} \\ &\cdot \begin{bmatrix} (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{j-1}(\mathbf{x}) + \mathbf{v}_x p_j^{-1} \mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}) - \mathbf{v}_x p_j^{-1} k(\mathbf{x}, \mathbf{x}_j) \\ -\mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}) p_j^{-1} + p_j^{-1} k(\mathbf{x}, \mathbf{x}_j) \end{bmatrix} \\ &= \mathbf{k}_{j-1}(\mathbf{x})^\top (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}) \mathbf{k}_{j-1}(\mathbf{x}) + p_j^{-1} (\mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}))^2 \\ &\quad - 2 \mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}) p_j^{-1} k(\mathbf{x}, \mathbf{x}_j) + p_j^{-1} k(\mathbf{x}, \mathbf{x}_j)^2 \\ &= \mathbf{k}_{j-1}(\mathbf{x})^\top (\mathbf{K}_{j-1} + \sigma^2 \mathbf{I}) \mathbf{k}_{j-1}(\mathbf{x}) + p_j^{-1} (k(\mathbf{x}, \mathbf{x}_j)^2 - \mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}))^2 \\ &= q_{j-1}(\mathbf{x}) + p_j^{-1} (k(\mathbf{x}, \mathbf{x}_j)^2 - \mathbf{v}_x^\top \mathbf{k}_{j-1}(\mathbf{x}))^2 \\ &= q_{j-1}(\mathbf{x}) + p_j^{-1} (k(\mathbf{x}, \mathbf{x}_j)^2 - q_{j-1}(\mathbf{x}))^2 \end{aligned}$$

This shows that

$$q_j(\mathbf{x}) = q_{j-1}(\mathbf{x}) + r_{j-1}(\mathbf{x}), \text{ where} \quad (55)$$

$$r_{j-1}(\mathbf{x}) := p_j^{-1} (k(\mathbf{x}, \mathbf{x}_j)^2 - q_{j-1}(\mathbf{x}))^2 \geq 0. \quad (56)$$

□

Discussion

19.1 Summary

This part presented a stopping strategy for the Cholesky decomposition that provides probably approximately correct kernel matrix determinants, under the assumptions that the dataset inputs are independent and identically distributed. Fig. 14 demonstrated that Theorem 27, though not trivial, needs further improvement. Nevertheless, Fig. 13 demonstrated that the practical realization of this stopping strategy is feasible and promising. And, though, behind expectations, using the stopping strategy is an “almost free lunch.”

19.2 Future Directions

19.2.1 Improving the Proof

The need to select a subset $S \subseteq \{m, \dots, N - 1\}$ of possible stopping points, is due to the union bound in Eq. (54). The proof of Theorem 29 by Fan, Grama, and Liu (2012) avoids a union bound with a technique they call the *conjugate probability measure*. It should be possible to apply the same technique.

The union bound may be unnecessary, altogether. Theorem 29 is a generalization of Hoeffding’s inequality (Fan, Grama, and Liu 2012). Schölkopf and Smola (2002, p. 181f.) proof with Hoeffding’s inequality that a subsample from a larger sum of i.i.d. random variables can be used to estimate the latter. Their proof does not require to measure separately how much estimate and remaining terms deviate from their expectation (*c.f.* Eq. (51)). Yet, in the case described here the summands are neither identically distributed nor independent, and there comes additional difficulty from the stopping condition.

Boucheron, Lugosi, and Massart (2013) describe a number of concentration inequalities for functions of random variables that are called self-bounding. It turns out, that the log-determinant of a kernel matrix falls into this category (Lemma 33, below). Self-bounding functions have the property that their variance is bounded by their expected value. Theorem 29 provides stronger guarantees when providing better bounds on the variance. Originally, my plan was to monitor the empirical variance as well, to obtain a Bernstein-like bound. For the examples from Figure 11 in Section 15.1 the empirical variance of the l_j/C is

lower than the bound by Popoviciu's inequality (c.f. Eq. (52)). One can expect a stopping rule that takes the variance into account to be more efficient than one which does not. Furthermore, the concentration inequalities for self-bounding functions often allow to reason about the probability of the function falling *below* its expectation. In the proof, the lower-bound for the determinant is deterministic. An estimated, probabilistic lower-bound could be less conservative.

Another direction to obtain sharper bounds is to stipulate further assumptions in Theorem 27. Properties of the kernel, such as Lipschitzness and differentiability could be bequeathed to the f_j .

Lemma 33. *Denote with \mathbf{C} the Cholesky decomposition of \mathbf{K}_N . Assume that there exist constants $C^-, C^+ \in \mathbb{R}$, s.t. $2 \ln \mathbf{C}_{jj} \in [C^-, C^+]$ for all $j \in \{1, \dots, N\}$ (cf. Lemma 31). Then the function*

$$f(\mathbf{x}_1, \dots, \mathbf{x}_N) := \frac{\ln |\mathbf{K}_N| - NC^-}{C^+ - C^-}$$

is self-bounding (definition in proof).

Proof. For the definition of self-bounding (Boucheron, Lugosi, and Massart 2013, p. 60) we have to show that there exist functions $f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N)$ s.t. for all $\mathbf{x}_1, \dots, \mathbf{x}_N$:

$$0 \leq f(\mathbf{x}_1, \dots, \mathbf{x}_N) - f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N) \leq 1, \quad (57)$$

and

$$\sum_{i=1}^N (f(\mathbf{x}_1, \dots, \mathbf{x}_N) - f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N)) \leq f(\mathbf{x}_1, \dots, \mathbf{x}_N). \quad (58)$$

Choose

$$f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N) := \frac{\ln |\mathbf{K}_{N \setminus i}| - (N-1)C^-}{C^+ - C^-}$$

where $\mathbf{K}_{N \setminus i}$ is obtained from \mathbf{K}_N by deleting the i -th row and column. Note that, the order of the inputs does not matter. Swapping two rows or columns changes the sign of the determinant. Swapping the order of the datapoints i and j corresponds to a swap of the corresponding columns and rows, and the sign change cancels. Denote with \mathbf{K}_N^i the resulting kernel matrix when moving the i -th datapoint to the end of the dataset, and denote with \mathbf{C}^i the Cholesky decomposition. This allows to write rewrite Eq. (57) as follows:

$$\begin{aligned} & f(\mathbf{x}_1, \dots, \mathbf{x}_N) - f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N) \\ &= \frac{\ln |\mathbf{K}_N| - \ln |\mathbf{K}_{N \setminus i}| - C^-}{C^+ - C^-} \\ & \quad // \text{definition of } f \text{ and } f_i \end{aligned}$$

$$\begin{aligned}
&= \frac{\ln |\mathbf{K}_N^i| - \ln |\mathbf{K}_{N-1}^i| - C^-}{C^+ - C^-} \\
&\quad // \text{ using the determinant's invariance} \\
&= \frac{2 \ln \mathbf{C}_{NN}^i - C^-}{C^+ - C^-} \\
&\quad // \text{ by Lemma 49 (p. 130)}
\end{aligned}$$

By assumption on \mathbf{C} , which carries over to \mathbf{C}^i , the last expression is bounded from below by 0 and from above by 1. This shows Eq. (57). Extending above argument, we obtain for Eq. (58):

$$\begin{aligned}
&\sum_{i=1}^N (f(\mathbf{x}_1, \dots, \mathbf{x}_N) - f_i(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N)) \\
&= \frac{1}{C^+ - C^-} \sum_{i=1}^N (2 \ln \mathbf{C}_{NN}^i - C^-) \\
&\leq \frac{1}{C^+ - C^-} \sum_{i=1}^N (2 \ln \mathbf{C}_{ii} - C^-) \\
&\quad // \text{ using Lemma 30, see below} \\
&= \frac{\ln |\mathbf{K}_N| - N C^-}{C^+ - C^-} \\
&\quad // \text{ by Lemma 49 (p. 130)} \\
&= f(\mathbf{x}_1, \dots, \mathbf{x}_N) \\
&\quad // \text{ definition of } f
\end{aligned}$$

Recall that Lemma 30 relates the diagonal elements of a Cholesky decomposition, to the posterior variance of a Gaussian process: $(\mathbf{C}_{NN}^i)^2$ is the posterior variance in \mathbf{x}_i given $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N$, whereas \mathbf{C}_{ii}^2 is the posterior variance in \mathbf{x}_i given only $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$. The posterior variance of a GP can never increase with more datapoints and therefore $\mathbf{C}_{NN}^i \leq \mathbf{C}_{ii}$. \square

19.2.2 Beyond Determinant Estimation

Often (e.g. for GP regression) it is necessary to evaluate a quadratic form $\mathbf{y}^\top (\mathbf{K}_N + \sigma^2 \mathbf{I}) \mathbf{y}$ and the determinant $|\mathbf{K}_N|$ at the same time. The analysis for the quadratic form is similar and it would be possible to compute both terms in parallel, stopping prematurely if one or both can be sufficiently approximated. In a conversation, Carl Rasmussen pointed out that a combination of above ideas with variational approximation methods for GPs would be interesting. For example to decide the number of necessary inducing inputs.

Part VI

Epilogue

Ich klage mich an, bekenne mich ohne Einschränkung schuldig, und vielleicht senken sich jetzt die schon zu verzweifelter Ringen erhobenen Hände, glätten sich die gerunzelten Stirnen wieder und wischt sich der eine oder andere den Schaum aus den Mundwinkeln. Ich verspreche hiermit feierlich, daß ich am Schluß dieses Erzählwerks ein umfassendes Geständnis ablegen, eine fix und fertige Moral liefern werde, auch eine Interpretation, die allen Interpreten vom Obertertianer bis zum Meisterinterpreten im Oberseminar Seufzen und Nachdenken ersparen wird. Sie wird so abgefaßt sein, daß auch der einfache, der unbefangene Leser sie "mit nach Hause nehmen kann", weit weniger kompliziert als die Anleitung zur Ausfüllung des Antrags auf Lohnsteuerjahresausgleich. Geduld, Geduld, wir sind noch nicht am Ende.

—Heinrich Böll, *Entfernung von der Truppe*

Conclusions

This thesis contains three main parts, each concerned with a different aspect of probabilistic numerics. Part [iii](#) examined the connection between probabilistic numerical methods and classic linear solvers. Part [iv](#) asked the question how probabilistic numerics can be used to develop novel methods. Part [v](#) showed how structure induced by probability can be exploited in classic numerical algorithms. This chapter summarizes the discussions and future directions that have been presented in Chapters [9](#), [14](#) and [19](#) for each part, respectively.

20.1 Discussion

Part [iii](#) showed that solution-based inference (SBI) proposed by Cockayne, Oates, Ipsen, and Girolami ([2018](#)) always constitutes a projection method ([Proposition 6](#)) and that for a given projection method there exist different SBI interpretations. Furthermore, SBI was shown to be subsumed in the matrix-based inference perspective, motivating a classification of probabilistic perspectives into left-multiplied and right-multiplied information. Hence, when reasoning about linear solvers, the unification allows to switch between perspectives as appropriate or convenient. The probabilistic perspective on preconditioning, though not unique, is the first and justifiably intuitive: when using left-multiplied information, right preconditioning corresponds to a change in prior, whereas left preconditioning corresponds to a change in information. When using right-multiplied information, the converse holds. The probabilistic interpretations of GMRES and FOM are currently a rather theoretical contribution as posterior variances are not available in practice.

Equipped with the insights from Part [iii](#), Part [iv](#) presented a probabilistic numerical method to solve kernel least-squares problems. In essence, the approach is based on a finite-rank inducing Gaussian process prior over kernel functions and using conjugate gradients matrix-vector products with the kernel matrix to collect observations. The *kernel machine conjugate gradients* (KMCG) approximation consistently outperforms plain conjugate gradients in numerical experiments. This improvement does not change the fact, that standard dense kernel least-squares problems are often more efficiently solved by inducing point methods. Nevertheless, one contribution is a principled procedure to use linear solvers in the context of kernel least-squares problems.

Part v presented a new approach to the approximation of kernel matrix determinants. More generally, Part v posed the problem of approximating the sum of a sequence of N elements that decrease in expectation from a subsequence of eligible length, s.t. a certain precision is reached with a desired probability. Section 16.2 presented a non-trivial solution and the proof of Theorem 27 showed that the combination of proposed stopping strategy and estimator provide a prediction that has indeed the relative precision with the desired probability. Theorem 28 showed that Theorem 27 can be used to terminate the Cholesky decomposition to approximate log-determinants of kernel matrices. When the stopping condition is not triggered, the implementation delivers the exact solution with negligible overhead. In that sense, the stopping strategy is an “almost free lunch.” Arguably, the stopping conditions are more conservative than necessary, s.t. the scenarios in which a termination could occur, are rare in practice. Nevertheless, this is the first contribution regarding a distribution-free and deterministic approach to approximate determinants of kernel matrices. To the best of my knowledge, there is no literature that focuses on the approximation of *specifically* kernel-matrix determinants—usually the more general case of symmetric and positive definite matrices is considered.


20.2 Future Work

This thesis left the question how to interpret and evaluate posterior uncertainty provided by probabilistic linear solvers unanswered. The χ^2 statistic proposed by Cockayne, Oates, Ipsen, and Girolami (2018) is a first step but is restricted to Gaussian linear solvers and not applicable for the probabilistic interpretation of Jacobi-iteration from Section 9.2.

A hope associated with KMCG from Part iv was that the probabilistic approach would allow to reason about the approximation error. Propagating probability measures through inverse operations remains an open challenge. First, this requires defining rigorously a ground-truth that can be aimed for. Then the question remains how to approximate the inversion operation. The Cayley–Hamilton theorem might provide an answer.

The work presented in Part v can be extended in different directions. Besides improving the bound in Theorem 27 with other mathematical tools, another option is to change the assumptions. Often Lipschitzness and differentiability of the kernel are known properties and using these could give less conservative bounds. The analysis could be extended to other operations associated with the Cholesky decomposition, *e.g.* estimation of the quadratic form $\mathbf{y}^\top \mathbf{K}^{-1} \mathbf{y}$. In a more distant future, the stopping strategy could be a starting point for a modified Cholesky decomposition that constructs a low-rank approximations to \mathbf{K}_N with

probabilistic guarantees. When treating inducing input variational inference methods (*c.f.* Titsias (2009)) as such low-rank approximations, the considerations in Part v might be useful to decide the number of necessary inducing inputs, as Carl Rasmussen pointed out in a personal conversation.



I DON'T KNOW HOW TO PROPAGATE
ERROR CORRECTLY, SO I JUST PUT
ERROR BARS ON ALL MY ERROR BARS.

R. Munroe (2019). *Error Bars*. URL: <https://xkcd.com/2110/>.
License: Creative Commons 2.5 BY-NC-SA

Part VII

Appendix

“You know,” said Arthur, “it’s at times like this, when I’m trapped in a Vogon airlock with a man from Betelgeuse, and about to die of asphyxiation in deep space that I really wish I’d listened to what my mother told me when I was young.” “Why, what did she tell you?” “I don’t know, I didn’t listen.” “Oh.” Ford carried on humming.

—Douglas Adams, *The Hitchhiker’s Guide to the Galaxy*



Benchmark Datasets

Table 3 describes purpose and origin of standard benchmark datasets used for Gaussian process regression. More information on PRECIPITATION can be found at <http://www.image.ucar.edu/Data/US.monthly.met/>. It appears that the datasets AILERONS, ELEVATORS and POLETELECOMM are no longer available under the link <https://www.dcc.fc.up.pt/~ltorgo/Regression/DataSets.html>.

name	reference	url	description
ABALONE	Nash, Sellers, Talbot, Cawthorn, and Ford (1994), Waugh and Dua and Graff (2019)	https://archive.ics.uci.edu/ml/datasets/Abalone	age prediction of abalone from physical measurements
AILERONS	Camachol (1998)	n/a	control action prediction on the ailerons of an F16 aircraft
CT_SLICES	Graf, Kriegel, Schubert, Pölslerl, and Cavallaro (2011) and Dua and Graff (2019)	https://archive.ics.uci.edu/ml/datasets/Relative+location+of+CT+slices+on+axial+axis	prediction of relative location of a computer tomographic image in the human body in polar coordinates
ELEVATORS	Camachol (1998)	n/a	control action prediction on the elevators of an F16 aircraft
MPG	Quinlan (1993) and Dua and Graff (2019)	https://archive.ics.uci.edu/ml/datasets/auto+mpg	fuel consumption prediction in miles per gallon for different attributes of cars
POLET-ELECOMM	Weiss and Indurkha (1995)	n/a	commercial telecommunication application—no further information
PRECIPITATION	Vanhatalo and Vehtari (2008)	github.com/gpstuff-dev/gpstuff/blob/master/gp/demo_regression_ppcs.m	US annual precipitation prediction for the year 1995
PUMADYN	Snelson and Ghahramani (2006)	ftp://ftp.cs.toronto.edu/pub/neuron/delve/data/tarfiles/pumadyn-family/pumadyn-32nm.tar.gz	acceleration prediction one of the arm links given angles, positions and velocities of other links of a <i>Puma560</i> robot
SARCOS	Vijayakumar and Schaal (2000)	http://www.gaussianprocess.org/gpml/data/	torque prediction for the seven degrees-of-freedom SARCOS anthropomorphic robot arm
SOUND	Turner (2010) and Wilson and Nickisch (2015)	https://github.com/kd383/GPML_SLD/blob/master/demo/sound/audio_data.mat	sound intensity prediction of a signal recorded over time for missing regions
TOY	Bartels and Henig (2019)	n/a	targets are a draw from a zero-mean Gaussian process with squared exponential kernel (Eq. (43) with $\mathbf{\Lambda} = 0.25$ and $\theta_f = 2$), inputs stem in equal parts from a Gaussian mixture ($\mathcal{N}(0, 1) + \mathcal{N}(1, 0.1) + \mathcal{N}(-0.5, 0.05)$) and the uniform distribution over $[0, 1]$

Table 3: descriptions and sources for all datasets considered in this work.

Gaussian Processes

The density of an l -dimensional standard normal random vector \mathbf{z} is defined as

$$p(\mathbf{z}) := \frac{1}{|2\pi|^l} \exp\left(-\frac{1}{2}\mathbf{z}^\top\mathbf{z}\right).$$

A real d -dimensional random vector \mathbf{x} is said to be distributed Gaussian $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ iff there exist $l \in \mathbb{N}$ and $\mathbf{A} \in \mathbb{R}^{d \times l}$ s.t. $\mathbf{x} = \mathbf{A}\mathbf{z} + \boldsymbol{\mu}$, where \mathbf{z} is distributed standard normal (Gut 2009, p. 454). These parameters also define mean and covariance of \mathbf{x} with $\mathbb{E}\mathbf{x} = \boldsymbol{\mu}$ and $\text{cov}\mathbf{x} = \boldsymbol{\Sigma}$. Gaussian random variables exhibit convenient properties as described in the two lemmata below.

Lemma 34. *Let $\mathbf{x} \in \mathbb{R}^d$ be Gaussian distributed with density $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma})$ for $\mathbf{x}_0 \in \mathbb{R}^d$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ a positive semi-definite matrix. Let $\mathbf{M} \in \mathbb{R}^{n \times d}$ and $\mathbf{z} \in \mathbb{R}^n$. Then $\mathbf{v} = \mathbf{M}\mathbf{x} + \mathbf{z}$ is also Gaussian, with*

$$p(\mathbf{v}) = \mathcal{N}(\mathbf{v}; \mathbf{M}\mathbf{x}_0 + \mathbf{z}, \mathbf{M}\boldsymbol{\Sigma}\mathbf{M}^\top).$$

Lemma 35. *Let $\mathbf{x} \in \mathbb{R}^d$ be distributed as in Lemma 34, and let observations $\mathbf{y} \in \mathbb{R}^n$ be generated from the conditional density*

$$p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y}; \mathbf{M}\mathbf{x} + \mathbf{z}, \boldsymbol{\Lambda})$$

with $\mathbf{M} \in \mathbb{R}^{n \times d}$, $\mathbf{z} \in \mathbb{R}^n$, and $\boldsymbol{\Lambda} \in \mathbb{R}^{n \times n}$ again positive-semidefinite. Then the associated conditional distribution on \mathbf{x} after observing \mathbf{y} is again Gaussian, with

$$\begin{aligned} p(\mathbf{x} | \mathbf{y}) &= \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \bar{\boldsymbol{\Sigma}}) \quad \text{where} \\ \bar{\mathbf{x}} &= \mathbf{x}_0 + \boldsymbol{\Sigma}\mathbf{M}^\top(\mathbf{M}\boldsymbol{\Sigma}\mathbf{M}^\top + \boldsymbol{\Lambda})^{-1}(\mathbf{y} - \mathbf{M}\mathbf{x}_0 - \mathbf{z}) \\ \bar{\boldsymbol{\Sigma}} &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{M}^\top(\mathbf{M}\boldsymbol{\Sigma}\mathbf{M}^\top + \boldsymbol{\Lambda})^{-1}\mathbf{M}\boldsymbol{\Sigma}. \end{aligned}$$

This formula also applies if $\boldsymbol{\Lambda} = \mathbf{0}$, i.e. observations are made without noise, with the caveat that if $\mathbf{M}\boldsymbol{\Sigma}\mathbf{M}^\top$ is singular, the inverse should be interpreted as a pseudo-inverse.

The following definition of a Gaussian process follows Rasmussen and Williams (2006, p. 13). A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution. It is completely specified by a mean function $\mu : \mathbb{R}^D \rightarrow \mathbb{R}$ and a covariance function $k : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$. Hence, mean vector and covariance matrix remain implicit. The indices of the variables change from natural numbers to input locations of the function.

Additional Material for Part ii

Proposition 2. Let $V, W \in \mathbb{R}^{N \times N}$ be square matrices and $A^\top, B \in \mathbb{R}^{N \times M}$ be rectangular.

$$W \otimes W = \Gamma_N(W \otimes W) \quad (\text{SK1})$$

$$\Gamma_M(A \otimes A) = (A \otimes A)\Gamma_N \quad (\text{SK2})$$

$$V \otimes W = W \otimes V \quad (\text{SK3})$$

$$(A \otimes A)(W \otimes W)(B \otimes B) = (AWB) \otimes (AWB) \quad (\text{SK4})$$

$$W \otimes W - V \otimes V = (W + V) \otimes (W - V) \quad (\text{SK5})$$

$$(W \otimes W)^{-1} = (W^{-1} \otimes W^{-1}). \quad (\text{SK6})$$

The interpretation of Eq. (SK6) requires some care: symmetric Kronecker product matrices are rank deficient. Eq. (SK6) is to be read in the sense that for symmetric $Y \in \mathbb{R}^{N \times N}$, i.e. $Y = Y^\top$, $X := \text{mat}\left((W^{-1} \otimes W^{-1}) \text{vec}(Y)\right)$ satisfies $\text{vec}(Y) = (W \otimes W) \text{vec}(X)$ and X is the unique symmetric solution.

Proof. The proofs for Eqs. (SK1) and (SK2) can be found in Magnus and Neudecker 1999[p. 46-50]. In the notation of Magnus and Neudecker 1999 $\Gamma = N_n = D_n D_n^+$ and $K = 2\Gamma - 2I$. Eq. (SK1) is Theorem 13 (a). Eq. (SK2) follows from Theorem 9 (a).

To show $(W \otimes V) = (V \otimes W)$, let $X \in \mathbb{R}^{N \times N}$ be an arbitrary matrix.

$$\begin{aligned} (V \otimes W) \text{vec}(X) &= \Gamma(V \otimes W) \Gamma \text{vec}(X) \\ &= \frac{1}{2} \Gamma(V \otimes W) \text{vec}(X + X^\top) \\ &= \frac{1}{2} \Gamma \text{vec}(V(X + X^\top)W^\top) \\ &= \frac{1}{4} \text{vec}(V(X + X^\top)W^\top + W(X + X^\top)V^\top) \\ &= \frac{1}{2} \Gamma \text{vec}(W(X + X^\top)V^\top) \\ &= \frac{1}{2} \Gamma(W \otimes V) \text{vec}(X + X^\top) \\ &= \Gamma(W \otimes V) \Gamma \text{vec}(X) \\ &= (W \otimes V) \text{vec}(X) \end{aligned}$$

To show Eq. (SK4), use (SK2).

$$\begin{aligned} (A \otimes A)(W \otimes W)(B \otimes B) &= (A \otimes A)\Gamma_N(W \otimes W)\Gamma_N(B \otimes B) \\ &= \Gamma_M(A \otimes A)(W \otimes W)(B \otimes B)\Gamma_M \\ &= \Gamma_M(AWB \otimes AWB)\Gamma_M \\ &= AWB \otimes AWB \end{aligned}$$

The proof of Eq. (SK5) uses (SK3).

$$\begin{aligned}
(\mathbf{A} + \mathbf{B}) \otimes (\mathbf{A} - \mathbf{B}) &= \mathbf{\Gamma}(\mathbf{A} + \mathbf{B}) \otimes (\mathbf{A} - \mathbf{B})\mathbf{\Gamma} \\
&= \mathbf{\Gamma}(\mathbf{A} \otimes \mathbf{A} - \mathbf{A} \otimes \mathbf{B} + \mathbf{B} \otimes \mathbf{A} - \mathbf{B} \otimes \mathbf{B})\mathbf{\Gamma} \\
&= \mathbf{A} \otimes \mathbf{A} - \mathbf{A} \otimes \mathbf{B} + \mathbf{B} \otimes \mathbf{A} - \mathbf{B} \otimes \mathbf{B} \\
&= \mathbf{A} \otimes \mathbf{A} - \mathbf{B} \otimes \mathbf{A} + \mathbf{B} \otimes \mathbf{A} - \mathbf{B} \otimes \mathbf{B} \\
&= \mathbf{A} \otimes \mathbf{A} - \mathbf{B} \otimes \mathbf{B}
\end{aligned}$$

It remains to prove Eq. (SK6). Assume \mathbf{Z} satisfies $(\mathbf{W} \otimes \mathbf{W}) \text{vec}(\mathbf{Z}) = \text{vec}(\mathbf{Y})$ and $\mathbf{Z} = \mathbf{Z}^\top$. Then,

$$\begin{aligned}
\text{vec}(\mathbf{Y}) &= (\mathbf{W} \otimes \mathbf{W}) \text{vec}(\mathbf{Z}) \\
&= (\mathbf{W} \otimes \mathbf{W})\mathbf{\Gamma}_N \text{vec}(\mathbf{Z}) \\
&\quad // \text{ using Eq. (SK1) and Eq. (SK2)} \\
&= (\mathbf{W} \otimes \mathbf{W}) \text{vec}(\mathbf{Z}) \\
&\quad // \text{ since } \mathbf{Z} = \mathbf{Z}^\top
\end{aligned}$$

and hence, $\text{vec}(\mathbf{Z}) = (\mathbf{W} \otimes \mathbf{W})^{-1} \text{vec}(\mathbf{Y})$. Using Eq. (K3) and again Eq. (SK1),

$$\begin{aligned}
\mathbf{Z} &= (\mathbf{W}^{-1} \otimes \mathbf{W}^{-1}) \text{vec}(\mathbf{Y}) \\
&\quad // \text{ by Eq. (K3)} \\
&= (\mathbf{W}^{-1} \otimes \mathbf{W}^{-1})\mathbf{\Gamma}_N \text{vec}(\mathbf{Y}) \\
&\quad // \text{ since } \mathbf{Y} = \mathbf{Y}^\top \\
&= (\mathbf{W}^{-1} \otimes \mathbf{W}^{-1}) \text{vec}(\mathbf{Y}) \\
&\quad // \text{ by Eq. (SK2) and Eq. (SK1)}
\end{aligned}$$

which is the definition of \mathbf{X} . □

Additional Material for Part [iii](#)D.1 Proof of Proposition [5](#)

Proposition 5. *Consider a Gaussian MBI prior*

$$p(\mathbf{A}^{-1}) = \mathcal{N}(\mathbf{A}^{-1}; \text{vec}(\mathbf{A}_0^{-1}), \boldsymbol{\Sigma}_0 \otimes \mathbf{W}_0),$$

conditioned on the left-multiplied information of Eq. [\(L\)](#). The associated marginal on \mathbf{x} (Eq. [\(MBI\)](#)) is identical to the SBI posterior on \mathbf{x} arising in Lemma [3](#) from $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \boldsymbol{\Sigma}_0)$, under the conditions

$$\mathbf{A}_0^{-1} \mathbf{b} = \mathbf{x}_0 \quad \text{and} \quad \mathbf{b}^\top \mathbf{W}_0 \mathbf{b} = 1.$$

Proof of Proposition 5. The proof is analogous to the proof of Lemma 2.1 in Hennig [\(2015\)](#). Let $\mathbf{H} := \mathbf{A}^{-1}$ and let $\mathbf{H}_0 := \mathbf{A}_0^{-1}$. First note that by right-multiplying the information in Eq. [\(L\)](#) by \mathbf{H} :

$$\begin{aligned} \mathbf{Y}_m^\top \mathbf{H} &= \mathbf{S}_m^\top \\ &\quad \text{// Eq. [\(L\)](#)} \\ \implies \text{vec}(\mathbf{Y}_m^\top \mathbf{H}) &= \text{vec}(\mathbf{S}_m^\top) \\ &\quad \text{// applying } \text{vec}(\cdot) \text{ on both sides} \\ \implies (\mathbf{Y}_m \otimes \mathbf{I}) \text{vec}(\mathbf{H}) &= \text{vec}(\mathbf{S}_m^\top) \\ &\quad \text{// from Eq. [\(K1\)](#)} \end{aligned}$$

Let $\boldsymbol{\Omega}_0 = \boldsymbol{\Sigma}_0 \otimes \mathbf{W}_0$ and let $\mathbf{P} = \mathbf{Y}_m^\top \otimes \mathbf{I}$. Now the implied posterior on $\text{vec}(\mathbf{H})$ can be computed using the standard laws of Gaussian conditioning (Lemma [35](#)): $p(\mathbf{H} \mid \mathbf{S}_m, \mathbf{Y}_m) = \mathcal{N}(\text{vec}(\mathbf{H}), \text{vec}(\mathbf{H}_m), \boldsymbol{\Omega}_m)$ where

$$\begin{aligned} \text{vec}(\mathbf{H}_m) &= \text{vec}(\mathbf{H}_0) + [\mathbf{P}\boldsymbol{\Omega}_0]^\top [\mathbf{P}\boldsymbol{\Omega}_0\mathbf{P}^\top]^{-1} (\text{vec}(\mathbf{S}_m^\top) - \text{vec}(\mathbf{Y}_m^\top \mathbf{H}_0)) \\ \boldsymbol{\Omega}_m &= \boldsymbol{\Omega}_0 - [\mathbf{P}\boldsymbol{\Omega}_0]^\top [\mathbf{P}\boldsymbol{\Omega}_0\mathbf{P}^\top]^{-1} (\mathbf{P}\boldsymbol{\Omega}_0) \end{aligned}$$

Using Eq. [\(K2\)](#), $\mathbf{P}\boldsymbol{\Omega}_0 = (\mathbf{Y}_m^\top \otimes \mathbf{I})(\boldsymbol{\Sigma}_0 \otimes \mathbf{W}) = (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \otimes \mathbf{W}$ which implies $(\mathbf{P}\boldsymbol{\Omega}_0)^\top = (\boldsymbol{\Sigma}_0 \mathbf{Y}_m) \otimes \mathbf{W}$ by Eq. [\(K4\)](#). Thus

$$\begin{aligned} \mathbf{P}\boldsymbol{\Omega}_0\mathbf{P}^\top &= (\mathbf{Y}_m^\top \otimes \mathbf{I})(\boldsymbol{\Sigma}_0 \otimes \mathbf{W}_0)(\mathbf{Y}_m^\top \otimes \mathbf{I})^\top \\ &\quad \text{// definitions of } \mathbf{P} \text{ and } \boldsymbol{\Omega}_0 \\ &= (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m) \otimes \mathbf{W}_0 \\ &\quad \text{// using Eqs. [\(K2\)](#) and [\(K4\)](#)} \\ \implies (\mathbf{P}\boldsymbol{\Omega}_0\mathbf{P}^\top)^{-1} &= (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \otimes \mathbf{W}_0^{-1} \\ &\quad \text{// using Eq. [\(K3\)](#)} \end{aligned}$$

which allows us to conclude that

$$\begin{aligned}
& (\mathbf{P}\boldsymbol{\Omega}_0)^\top (\mathbf{P}\boldsymbol{\Omega}_0 \mathbf{P}^\top)^{-1} \\
&= [(\boldsymbol{\Sigma}_0 \mathbf{Y}_m) \otimes \mathbf{W}] [(\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \otimes \mathbf{W}^{-1}] \\
&= (\boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1}) \otimes \mathbf{I} \\
\implies & (\mathbf{P}\boldsymbol{\Omega}_0)^\top (\mathbf{P}\boldsymbol{\Omega}_0 \mathbf{P}^\top)^{-1} (\mathbf{P}\boldsymbol{\Omega}_0) \\
&= (\boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \otimes \mathbf{W}_0.
\end{aligned}$$

From these expressions one can simplify the expressions for $\text{vec}(\mathbf{H}_m)$:

$$\begin{aligned}
\text{vec}(\mathbf{H}_m) &= \text{vec}(\mathbf{H}_0) + (\boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \otimes \mathbf{I}) (\text{vec}(\mathbf{S}_m^\top) - \text{vec}(\mathbf{Y}_m^\top \mathbf{H}_0)) \\
&= \text{vec}\left(\mathbf{H}_0 + \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} (\mathbf{S}_m^\top - \mathbf{Y}_m^\top \mathbf{H}_0)\right) \\
&\quad // \text{ by Eq. (K1)}.
\end{aligned}$$

Analogously, for $\boldsymbol{\Omega}_m$:

$$\begin{aligned}
\boldsymbol{\Omega}_m &= \boldsymbol{\Sigma}_0 \otimes \mathbf{W} - (\boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \otimes \mathbf{W}_0 \\
&= (\boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \otimes \mathbf{W}_0 \\
&\quad // \text{ by Eq. (K5)}
\end{aligned}$$

It remains to project the posterior into \mathbb{R}^d by performing the matrix-vector product $\mathbf{H}\mathbf{b} = \mathbf{x} = (\mathbf{I} \otimes \mathbf{b}^\top) \mathbf{H}$ using Eq. (K1). Thus, the implied posterior is $\mathbf{x} \sim \mathcal{N}(\bar{\mathbf{x}}_m, \bar{\boldsymbol{\Sigma}}_m)$, with

$$\begin{aligned}
\bar{\mathbf{x}}_m &= (\mathbf{I} \otimes \mathbf{b}^\top) \text{vec}\left(\mathbf{H}_0 + \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} (\mathbf{S}_m^\top - \mathbf{Y}_m^\top \mathbf{H}_0)\right) \\
&= \text{vec}\left(\mathbf{H}_0 \mathbf{b} + \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} (\mathbf{S}_m^\top \mathbf{b} - \mathbf{Y}_m^\top \mathbf{H}_0 \mathbf{b})\right) \\
&= \mathbf{x}_0 + \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m (\mathbf{S}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m)^{-1} \mathbf{S}_m^\top (\mathbf{b} - \mathbf{A} \mathbf{x}_0). \\
&\quad // \text{ since } \mathbf{H}_0 \mathbf{b} = \mathbf{x}_0 \text{ and } \mathbf{Y}_m = \mathbf{A}^\top \mathbf{S}_m
\end{aligned}$$

Furthermore

$$\begin{aligned}
\bar{\boldsymbol{\Sigma}}_m &= (\mathbf{I} \otimes \mathbf{b}^\top) \cdot \left[(\boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \otimes \mathbf{W}_0 \right] \cdot (\mathbf{I} \otimes \mathbf{b}^\top)^\top \\
&= (\boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0 \mathbf{Y}_m (\mathbf{Y}_m^\top \boldsymbol{\Sigma}_0 \mathbf{Y}_m)^{-1} \mathbf{Y}_m^\top \boldsymbol{\Sigma}_0) \times \mathbf{b}^\top \mathbf{W}_0 \mathbf{b} \\
&\quad // \text{ using Eq. (K2) and that } \mathbf{b}^\top \mathbf{W}_0 \mathbf{b} \text{ is scalar} \\
&= \boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m (\mathbf{S}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \mathbf{A}^\top \mathbf{S}_m)^{-1} \mathbf{S}_m^\top \mathbf{A} \boldsymbol{\Sigma}_0 \\
&\quad // \text{ since } \mathbf{b}^\top \mathbf{W}_0 \mathbf{b} = 1 \text{ and } \mathbf{Y}_m = \mathbf{A}^\top \mathbf{S}_m
\end{aligned}$$

Now note that $\mathbf{x}_m = \bar{\mathbf{x}}_m$ and $\boldsymbol{\Sigma}_m = \bar{\boldsymbol{\Sigma}}_m$, as in Lemma 3. Thus, the proof is complete. \square

D.2 Proof of Theorem 15

Theorem 36 (Theorem 2.3 in Hennig (2015)). *Let $\mathbf{W} \in \mathbb{R}^{d \times d}$ be symmetric and positive definite. Assume a Gaussian prior of symmetric mean \mathbf{A}_0^{-1} and covariance $\mathbf{W} \otimes \mathbf{W}$ on the elements of a symmetric matrix \mathbf{A}^{-1} . After m linearly independent noise-free observations*

of the form $\mathbf{S} = \mathbf{A}^{-1}\mathbf{Y}$, $\mathbf{Y} \in \mathbb{R}^{d \times m}$, $\text{rk}(\mathbf{Y}) = m$, the posterior belief over \mathbf{A}^{-1} is a Gaussian with mean

$$\begin{aligned} \mathbf{A}_m^{-1} &= \mathbf{A}_0^{-1} + (\mathbf{S} - \mathbf{A}_0^{-1}\mathbf{Y})\mathbf{G}\mathbf{Y}^\top\mathbf{W} \\ &\quad + \mathbf{W}\mathbf{Y}\mathbf{G}(\mathbf{S} - \mathbf{A}_0^{-1}\mathbf{Y})^\top \\ &\quad + \mathbf{W}\mathbf{Y}\mathbf{G}\mathbf{Y}^\top(\mathbf{S} - \mathbf{A}_0^{-1}\mathbf{Y})\mathbf{G}\mathbf{Y}^\top\mathbf{W} \end{aligned} \quad (59)$$

and posterior covariance

$$\mathbf{V}_m = (\mathbf{W} - \mathbf{W}\mathbf{Y}\mathbf{G}\mathbf{Y}^\top\mathbf{W}) \otimes (\mathbf{W} - \mathbf{W}\mathbf{Y}\mathbf{G}\mathbf{Y}^\top\mathbf{W})$$

where $\mathbf{G} := (\mathbf{Y}^\top\mathbf{W}\mathbf{Y})^{-1}$.

The following proofs are due to Philipp Hennig. My contribution are simplifications that reduce the length by approximately one page.

Proof of Theorem 15 by Philipp Hennig. Denote by \mathbf{x}_i^{CG} the conjugate gradient estimate in iteration i and with \mathbf{p}_i the search direction in that iteration. From one iteration to the next, the update to the solution can be written as (Nocedal and Wright 1999, p. 108)

$$\mathbf{x}_{i+1}^{\text{CG}} = \mathbf{x}_i^{\text{CG}} + \frac{\mathbf{r}_i^\top \mathbf{p}_i}{\mathbf{p}_i^\top \mathbf{A} \mathbf{p}_i} \mathbf{p}_i.$$

Comparing this update to lines 7 to 10 in Algorithm 3 it is sufficient to show that $\mathbf{d}_i \propto \mathbf{p}_i$ which follows from Lemma 37. \square

Lemma 37. *Assume that CG does not terminate before d iterations. Using the prior of Theorem 15 in Algorithm 3, the directions \mathbf{d}_i are scaled conjugate gradients search directions, i.e.*

$$\mathbf{d}_i = \gamma_i \mathbf{p}_i^{\text{CG}}$$

where \mathbf{p}_i^{CG} is the CG search direction in iteration i and $\gamma_i \in \mathbb{R} \setminus \{0\}$.

Proof. The proof proceeds by induction. Throughout we will suppress the superscript CG on the CG search directions, i.e. $\mathbf{p}_i^{\text{CG}} = \mathbf{p}_i$. For $i = 1$, $\mathbf{A}_{i-1}^{-1} = \alpha \mathbf{I}$ by assumption and therefore $\mathbf{d}_i = \alpha \mathbf{r}_0$ which is the first CG search direction scaled by $\gamma_1 = \alpha \neq 0$.

For the inductive step, suppose that the search directions $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}$ are scaled CG directions and that the vectors $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ are the same as the first $i-1$ solution estimates produced by CG. We will prove that \mathbf{s}_i is the i^{th} CG search direction, and that \mathbf{x}_i is the i^{th} solution estimate from CG. Lemma 39 states that \mathbf{d}_i can be written as

$$\mathbf{d}_i = \mathbf{A}_{i-1}^{-1} \mathbf{r}_{i-1} = \sum_{j < i} v_j \mathbf{s}_j + v_i \mathbf{r}_{i-1}.$$

where $v_j \in \mathbb{R}$, $j = 1, \dots, i$. Under the prior, the posterior mean \mathbf{A}_i^{-1} is always symmetric. This allows application of Lemma 38, so that $\{\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{d}_i\}$ is an \mathbf{A} -conjugate set. Thus we have, for $\ell < i$:

$$0 = \mathbf{s}_\ell^\top \mathbf{A} \mathbf{d}_i = v_\ell \mathbf{s}_\ell^\top \mathbf{A} \mathbf{s}_\ell + v_i \mathbf{s}_\ell^\top \mathbf{A} \mathbf{r}_{i-1}$$

$$= v_\ell \mathbf{s}_\ell^\top \mathbf{A} \mathbf{s}_\ell + v_i \mathbf{y}_\ell^\top \mathbf{r}_{i-1}. \quad (60)$$

Now note that

$$\mathbf{y}_\ell^\top \mathbf{r}_{i-1} = (\mathbf{r}_\ell - \mathbf{r}_{\ell-1})^\top \mathbf{r}_{i-1}.$$

This follows from Line 10 of Algorithm 3, from which it is clear that $\mathbf{y}_\ell = \mathbf{r}_\ell - \mathbf{r}_{\ell-1}$. Recall that the CG residuals \mathbf{r}_j are orthogonal Nocedal and Wright 1999, p. 109, and that from the inductive assumption, Algorithm 3 is equivalent to CG up to iteration $i-1$). Thus, for $\ell < i-1$ we have that

$$\begin{aligned} \mathbf{y}_\ell^\top \mathbf{r}_{i-1} &= 0 \\ \implies \mathbf{s}_\ell^\top \mathbf{A} \mathbf{d}_i &= v_\ell \mathbf{s}_\ell^\top \mathbf{A} \mathbf{s}_\ell = 0 \quad \forall \ell < i-1 \end{aligned}$$

where the second line is from application of the first line in Eq. (60). However, \mathbf{A} is positive definite and by assumption the algorithm has not converged, so $\mathbf{d}_\ell \neq \mathbf{0}$. Furthermore clearly $\mathbf{s}_\ell^\top \mathbf{A} \mathbf{s}_\ell \neq 0$. Hence we must have that

$$v_\ell = 0 \quad \forall j < i-1.$$

Equation (D.2) thus simplifies to

$$\mathbf{d}_i = v_{i-1} \mathbf{s}_{i-1} + v_i \mathbf{r}_{i-1} = v_{i-1} \alpha_{i-1} \mathbf{d}_{i-1} + v_i \mathbf{r}_{i-1}.$$

Now, again by Lemma 38, \mathbf{d}_i must be conjugate to \mathbf{s}_{i-1} which implies $v_i \neq 0$. Pre-multiplying Eq. (D.2) by $\mathbf{s}_{i-1}^\top \mathbf{A}$ gives

$$\begin{aligned} 0 &= v_{i-1} \alpha_{i-1} \mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{d}_{i-1} + v_i \mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{r}_{i-1} \\ \implies v_{i-1} \alpha_{i-1} &= -v_i \frac{\mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{r}_{i-1}}{\mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{d}_{i-1}}. \end{aligned}$$

Thus, \mathbf{d}_i can be written as

$$\begin{aligned} \mathbf{d}_i &= v_i \left(\mathbf{r}_{i-1} - \frac{\mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{r}_{i-1}}{\mathbf{s}_{i-1}^\top \mathbf{A} \mathbf{d}_{i-1}} \mathbf{d}_{i-1} \right) \\ &= v_i \left(\mathbf{r}_{i-1} - \frac{\mathbf{p}_{i-1}^\top \mathbf{A} \mathbf{r}_{i-1}}{\mathbf{p}_{i-1}^\top \mathbf{A} \mathbf{p}_{i-1}} \mathbf{p}_{i-1} \right) \end{aligned} \quad (61)$$

where the second line again applies the inductive assumption, that \mathbf{d}_{i-1} and \mathbf{s}_{i-1} are proportional to the CG search direction \mathbf{p}_{i-1} , noting that the proportionality constants on numerator and denominator cancel. The term inside the brackets is precisely the i^{th} CG search direction. This completes the result. \square

Lemma 38. *If the belief over \mathbf{A}_m^{-1} is symmetric for all $m = 0, \dots, d$ and \mathbf{A} is symmetric and positive definite, then Algorithm 3 produces \mathbf{A} -conjugate directions.*

Proof. The proof is by induction. Note that the case $i = 1$ is irrelevant since a set consisting of one element is trivially \mathbf{A} -conjugate. On many occasions the proof relies on the consistency

of the MBI belief, *i.e.* $\mathbf{A}_i^{-1}\mathbf{z}_k = \mathbf{d}_k$ for $k \leq i$ and by symmetry $\mathbf{z}_k^\top \mathbf{A}_i^{-1} = \mathbf{d}_k^\top$. Thus, for the base case $i = 2$ we have:

$$\begin{aligned} \mathbf{d}_1^\top \mathbf{A} \mathbf{d}_2 &= -\mathbf{d}_1^\top \mathbf{A} (\mathbf{A}_1^{-1} \mathbf{r}_1) \\ &= -\mathbf{d}_1^\top \mathbf{A} (\mathbf{A}_1^{-1} (\mathbf{y}_1 + \mathbf{r}_0)) \\ &= -\mathbf{d}_1^\top \mathbf{A} (\mathbf{s}_1 + \mathbf{A}_1^{-1} \mathbf{r}_0) \end{aligned}$$

where the second line is by Line 10 of Algorithm 3. Now recall that $\alpha_1 = -\mathbf{d}_1^\top \mathbf{r}_0 / \mathbf{d}_1^\top \mathbf{A} \mathbf{d}_1$ to give:

$$\begin{aligned} \mathbf{d}_1^\top \mathbf{A} \mathbf{d}_2 &= -\alpha_1 \mathbf{d}_1^\top \mathbf{A} \mathbf{d}_1 - \mathbf{d}_1^\top \mathbf{A} \mathbf{A}_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{d}_1^\top \mathbf{A} \mathbf{A}_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{z}_1^\top \mathbf{A}_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{d}_1^\top \mathbf{r}_0 \\ &= 0. \end{aligned} \tag{62}$$

Here, the symmetry of the estimator \mathbf{A}_i^{-1} is used in Eq. (62). For the inductive step, assume $\{\mathbf{d}_0, \dots, \mathbf{d}_{i-1}\}$ are pairwise \mathbf{A} -conjugate. For any $k < i$ we have:

$$\begin{aligned} \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_i &= -\mathbf{d}_k^\top \mathbf{A} (\mathbf{A}_i^{-1} \mathbf{r}_i) \\ &= -\mathbf{d}_k^\top \mathbf{A} \mathbf{A}_i^{-1} \left(\sum_{j \leq i} \mathbf{y}_j + \mathbf{r}_0 \right) \end{aligned}$$

where the second line follows from the fact that $\mathbf{r}_i = \mathbf{r}_{i-1} + \mathbf{y}_i$. Thus, we have:

$$\begin{aligned} \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_i &= -\mathbf{d}_k^\top \mathbf{A} \left(\sum_{j \leq i} \mathbf{s}_j + \mathbf{A}_i^{-1} \mathbf{r}_0 \right) \\ &= -\mathbf{d}_k^\top \mathbf{A} \left(\sum_{j \leq i} \alpha_j \mathbf{d}_j + \mathbf{A}_i^{-1} \mathbf{r}_0 \right). \end{aligned}$$

Now, applying the conjugacy from the inductive assumption:

$$\begin{aligned} \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_i &= -\alpha_k \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_k - \mathbf{d}_k^\top \mathbf{A} (\mathbf{A}_i^{-1} \mathbf{r}_0) \\ &= \mathbf{d}_k^\top \mathbf{r}_{k-1} - \mathbf{d}_k^\top \mathbf{r}_0 \\ &= \mathbf{d}_k^\top \left(\sum_{j < k} \mathbf{y}_j + \mathbf{r}_0 \right) - \mathbf{d}_k^\top \mathbf{r}_0 = 0 \\ &= \sum_{j < k} \alpha_j \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_j = 0. \end{aligned}$$

where the second line rearranges line 6 of the algorithm to obtain $\alpha_i \mathbf{d}_i^\top \mathbf{z}_i = -\mathbf{d}_i^\top \mathbf{r}_{i-1}$. The third line again uses that $\mathbf{r}_i = \mathbf{r}_{i-1} + \mathbf{y}_i$, while the fourth line is from the assumed conjugacy. \square

Lemma 39. *Under the prior in Theorem 15 and given scaled CG search directions $\mathbf{p}_1, \dots, \mathbf{p}_i$, it holds that $\mathbf{A}_i^{-1} \mathbf{r}_i \in \text{span}\{\mathbf{p}_1, \dots, \mathbf{p}_i, \mathbf{r}_i\}$.*

Proof. Recall first that under the prior in Theorem 15, $A_0^{-1} = \alpha I$. Then by inspection of Eq. (59) we have $A_i^{-1}r_i \in \mathcal{S}$ where

$$\mathcal{S} = \text{span}\{r_i, p_1, \dots, p_i, y_1, \dots, y_i, Wy_1, \dots, Wy_i\}$$

By choice of $W = \beta I + \gamma A^{-1}$, $\mathcal{S} = \text{span}\{r_i, p_1, \dots, p_i, y_1, \dots, y_i\}$. From line 10 of Algorithm 3 $y_i = r_i - r_{i-1}$ and therefore $\mathcal{S} = \text{span}\{r_1, \dots, r_i, p_1, \dots, p_i\}$. By Theorem 5.3 in Nocedal and Wright 1999, p. 109 the span of the conjugate gradients residuals and search directions are equivalent. Therefore $\mathcal{S} \subseteq \{r_i, p_1, \dots, p_i\}$. \square

Additional Material for Part iv

E.1 Sampling from a Gaussian with Symmetric Kronecker Covariance

To sample matrices from the KMCG posterior (Eq. 34) the following proposition will be useful.

Proposition 40. *Let $\mathbf{W}, \mathbf{W}_M \in \mathbb{R}^{N \times N}$ be symmetric and positive semi-definite matrices s.t. $\mathbf{W} - \mathbf{W}_M$ is symmetric positive-semidefinite as well. Further let $\text{vec}(\mathbf{Y}) \sim \mathcal{N}(\mathbf{0}, \mathbf{W} \otimes \mathbf{W} - \mathbf{W}_M \otimes \mathbf{W}_M)$, denote with \mathbf{L}_+ the Cholesky of $\mathbf{W} + \mathbf{W}_M$, with \mathbf{L}_- the Cholesky of $\mathbf{W} - \mathbf{W}_M$ and let $\text{vec}(\mathbf{X}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N^2})$, then $\Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-) \text{vec}(\mathbf{X})$ and $\text{vec}(\mathbf{Y})$ have the same distribution and $\mathbf{Y} = \mathbf{Y}^\top$.*

Proof. As $\text{vec}(\mathbf{X})$ is standard normal, $\Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-) \text{vec}(\mathbf{X})$ is distributed Gaussian with mean $\mathbf{0}$ and covariance matrix $\Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-) (\Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-))^\top$.

$$\begin{aligned} \Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-) [\Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-)]^\top &= \Gamma(\mathbf{L}_+ \otimes \mathbf{L}_-) (\mathbf{L}_+^\top \otimes \mathbf{L}_-^\top) \Gamma \\ &= (\mathbf{L}_+ \mathbf{L}_+^\top) \otimes (\mathbf{L}_- \mathbf{L}_-^\top) \\ &= (\mathbf{W} + \mathbf{W}_M) \otimes (\mathbf{W} - \mathbf{W}_M) \end{aligned}$$

According to Equation (SK5): $(\mathbf{W} + \mathbf{W}_M) \otimes (\mathbf{W} - \mathbf{W}_M) = \mathbf{W} \otimes \mathbf{W} - \mathbf{W}_M \otimes \mathbf{W}_M$. \mathbf{Y} is symmetric due to the application of the Γ -operator. \square

E.2 Proofs

This section contains the proof of Proposition 20:

Proposition 20 (Subset of Regressors). *Consider the prior of Eq. (30) with $k_0 := 0$ and $w := k$ and the likelihood defined in Eq. (33) with $s_{ij} = \delta_{ij}$. Then the posterior mean k_M is equivalent to that of SoR:*

$$k_M(\mathbf{x}, \mathbf{z}) = k_{\text{SoR}} = k(\mathbf{x}, \mathbf{X}_U) k(\mathbf{X}_U, \mathbf{X}_U)^{-1} k(\mathbf{X}_U, \mathbf{z})$$

where \mathbf{X}_U are inducing inputs, not necessarily part of \mathbf{X} .

The mentioned prior is $p(k) := \mathcal{N}(0, \psi)$ where $\psi(k(\mathbf{a}, \mathbf{b}), k(\mathbf{c}, \mathbf{d})) := \frac{1}{2} k(\mathbf{a}, \mathbf{c}) k(\mathbf{b}, \mathbf{d}) + \frac{1}{2} k(\mathbf{a}, \mathbf{d}) k(\mathbf{b}, \mathbf{c})$. Denote with \mathbf{Y} the noise-free observations of k , $\mathbf{Y} = \text{mat}(\mathbf{T}_p k)$ where

$$\mathbf{T}_p : k \mapsto \text{vec} \left(\left[\iint k(\mathbf{x}, \mathbf{z}) p_i(\mathbf{x}) p_j(\mathbf{z}) \, d\mathbf{x} \, d\mathbf{z} \right]_{ij} \right) \quad (32)$$

$$p_i(\mathbf{x}) = \sum_{j=1}^M s_{ij} \delta(\mathbf{x} - \mathbf{x}_{u_j}) \quad (33)$$

which implies the likelihood $p(\mathbf{Y} \mid \mathbf{T}_p k) = \delta(\mathbf{Y} - \mathbf{T}_p k)$. Proposition 20 follows from the more general Proposition 41.

Proposition 41. *Consider the prior of Eq. (30) (without the restriction $w = k$) and the likelihood defined in Eq. (33). The posterior over k is $p(k \mid \mathbf{Y} = \mathbf{T}_p k) = \mathcal{N}(k; k_M, \psi_M)$ with posterior mean*

$$k_M(\mathbf{a}, \mathbf{b}) = k_0(\mathbf{a}, \mathbf{b}) + w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{K}_M \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b}) \\ - w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top k_0(\mathbf{X}_U, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b}) \quad (63)$$

and posterior variance

$$\psi_M(k(\mathbf{a}, \mathbf{b}), k(\mathbf{c}, \mathbf{d})) = \frac{1}{2} w(\mathbf{a}, \mathbf{c}) w(\mathbf{b}, \mathbf{d}) + \frac{1}{2} w(\mathbf{a}, \mathbf{d}) w(\mathbf{b}, \mathbf{c}) \\ - \frac{1}{2} w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{c}) w(\mathbf{b}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{d}) \\ - \frac{1}{2} w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{d}) w(\mathbf{b}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{c}) \quad (64)$$

where $\mathbf{W}_M = w(\mathbf{X}_U, \mathbf{X}_U)$.

Proof. Proving the proposition is tedious linear Algebra which is what follows now. If prior and likelihood are Gaussian, so is the posterior with mean and variance:

$$k_M(\mathbf{a}, \mathbf{b}) = k_0(\mathbf{a}, \mathbf{b}) - (\mathbf{T}_p \psi(k(\mathbf{a}, \mathbf{b}), \cdot))^\top (\mathbf{T}_p (\mathbf{T}_p w)^\top)^{-1} \text{vec}(\mathbf{Y} - \mathbf{S}^\top k_0(\mathbf{X}_U, \mathbf{X}_U) \mathbf{S}), \\ \psi_M(k(\mathbf{a}, \mathbf{b}), k(\mathbf{c}, \mathbf{d})) = \psi((\mathbf{a}, \mathbf{b}), (\mathbf{c}, \mathbf{d})) - (\mathbf{T}_p \psi((\mathbf{a}, \mathbf{b}), (\cdot, \cdot)))^\top (\mathbf{T}_p (\mathbf{T}_p \psi)^\top)^{-1} \mathbf{T}_p \psi(\mathbf{c}, \mathbf{d}), (\cdot, \cdot)$$

Lemma 42 allows to write

$$(\mathbf{T}_p \psi(k(\mathbf{a}, \mathbf{b}), \cdot))^\top (\mathbf{T}_p (\mathbf{T}_p w)^\top)^{-1} \\ = \frac{1}{2} \text{vec}(\mathbf{S}^\top w(\mathbf{X}_U, \mathbf{a}) w(\mathbf{b}, \mathbf{X}_U) + w(\mathbf{X}_U, \mathbf{b}) w(\mathbf{a}, \mathbf{X}_U)) \mathbf{S}^\top \left((\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \otimes (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \right) \\ = \frac{1}{2} \text{vec} \left((\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{a}) w(\mathbf{b}, \mathbf{X}_U) + w(\mathbf{X}_U, \mathbf{b}) w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \right)^\top$$

and thus for Eq. (63)

$$k_M(\mathbf{a}, \mathbf{b}) = k_0(\mathbf{a}, \mathbf{b}) \\ + \frac{1}{2} \text{tr}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{a}) w(\mathbf{b}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} (\mathbf{Y} - k_0(\mathbf{X}_U, \mathbf{X}_U)) \\ + \frac{1}{2} \text{tr}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b}) w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} (\mathbf{Y} - k_0(\mathbf{X}_U, \mathbf{X}_U)) \\ = k_0(\mathbf{a}, \mathbf{b}) + w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{K}_M \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b}) \\ - w(\mathbf{a}, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top k_0(\mathbf{X}_U, \mathbf{X}_U) \mathbf{S} (\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b})$$

The derivation for Eq. (64) follows analogously. \square

Lemma 42. *Let \mathbf{T}_p be as in Eq. (36).*

$$\mathbf{T}_p w(k(\mathbf{a}, \mathbf{b}), \cdot) = \frac{1}{2} \text{vec}(\mathbf{S}^\top (w(\mathbf{X}_U, \mathbf{a}) w(\mathbf{b}, \mathbf{X}_U) + w(\mathbf{X}_U, \mathbf{b}) w(\mathbf{a}, \mathbf{X}_U)) \mathbf{S}) \quad (65)$$

$$\mathbf{T}_p (\mathbf{T}_p w(\cdot, \cdot))^\top = (\mathbf{S}^\top \mathbf{W}_M \mathbf{S}) \otimes (\mathbf{S}^\top \mathbf{W}_M \mathbf{S}) \quad (66)$$

Proof. Denote with $\text{mat}(\cdot)$ the complement of the vectorization operator, i.e. $\text{mat}(\text{vec}(\mathbf{A})) = \mathbf{A}$. Define the matrix $\mathbf{S} \in \mathbb{R}^{N \times M}$ as $S_{ij} = s_{ij}$ and denote with \mathbf{S}_l the l -th column of \mathbf{S} . Also recall that by Eq. (31) $\psi(k(\mathbf{a}, \mathbf{b}), k(\mathbf{x}, \mathbf{z})) = \frac{1}{2}(w(\mathbf{a}, \mathbf{x})w(\mathbf{b}, \mathbf{z}) + w(\mathbf{a}, \mathbf{z})w(\mathbf{b}, \mathbf{x}))$.

$$\begin{aligned}
& [\text{mat}(\mathbf{T}_p[\psi(k(\mathbf{a}, \mathbf{b}), k(\cdot, \cdot))])]_{ij} \\
&= \iint \psi(k(\mathbf{a}, \mathbf{b}), k(\mathbf{x}, \mathbf{z})) \left(\sum_{l=1}^M s_{il} \delta(\mathbf{x} - \mathbf{u}_l) \right) \left(\sum_{l=1}^M s_{jl} \delta(\mathbf{z} - \mathbf{u}_l) \right) \mathrm{d}\mathbf{x} \mathrm{d}\mathbf{z} \\
&= \iint \frac{1}{2} (w(\mathbf{a}, \mathbf{x})w(\mathbf{b}, \mathbf{z}) + w(\mathbf{a}, \mathbf{z})w(\mathbf{b}, \mathbf{x})) \left(\sum_{l=1}^M s_{il} \delta(\mathbf{x} - \mathbf{u}_l) \right) \left(\sum_{l=1}^M s_{jl} \delta(\mathbf{z} - \mathbf{u}_l) \right) \mathrm{d}\mathbf{x} \mathrm{d}\mathbf{z} \\
&= \frac{1}{2} \sum_{m=1}^M \sum_{l=1}^M S_{im} S_{jl} (w(\mathbf{a}, \mathbf{u}_m)w(\mathbf{b}, \mathbf{u}_l) + w(\mathbf{a}, \mathbf{u}_l)w(\mathbf{b}, \mathbf{u}_m)) \\
&= \frac{1}{2} [\mathbf{S}^\top w(\mathbf{X}_U, \mathbf{a})w(\mathbf{b}, \mathbf{X}_U)\mathbf{S} + \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{b})w(\mathbf{a}, \mathbf{X}_U)\mathbf{S}]_{ij} \\
&= \frac{1}{2} [\mathbf{S}^\top (w(\mathbf{X}_U, \mathbf{a})w(\mathbf{b}, \mathbf{X}_U) + w(\mathbf{X}_U, \mathbf{b})w(\mathbf{a}, \mathbf{X}_U))\mathbf{S}]_{ij}
\end{aligned}$$

which shows Eq. (65)

$$\begin{aligned}
&= [\text{mat}((\mathbf{S}^\top \otimes \mathbf{S}^\top)\mathbf{\Gamma} \text{vec}(w(\mathbf{X}_U, \mathbf{a})w(\mathbf{b}, \mathbf{X}_U)))]_{ij} \\
&= [\text{mat}((\mathbf{S}^\top \otimes \mathbf{S}^\top)\mathbf{\Gamma}(w(\mathbf{X}_U, \mathbf{a}) \otimes w(\mathbf{X}_U, \mathbf{b})))]_{ij}
\end{aligned}$$

Repeating above derivations shows the second statement:

$$\begin{aligned}
\mathbf{T}_p(\mathbf{T}_p\psi)^\top &= (\mathbf{S}^\top \otimes \mathbf{S}^\top)\mathbf{\Gamma}(w(\mathbf{X}_U, \mathbf{X}_U) \otimes w(\mathbf{X}_U, \mathbf{X}_U))\mathbf{\Gamma}(\mathbf{S} \otimes \mathbf{S}) \\
&= (\mathbf{S} \otimes \mathbf{S})^\top (w(\mathbf{X}_U, \mathbf{X}_U) \otimes w(\mathbf{X}_U, \mathbf{X}_U))(\mathbf{S} \otimes \mathbf{S}) \\
&= (\mathbf{S}^\top \mathbf{W}_M \mathbf{S}) \otimes (\mathbf{S}^\top \mathbf{W}_M \mathbf{S}) \\
&\quad // \text{Eq. (SK4)}
\end{aligned}$$

□

E.2.1 Positive Semi-definiteness of the Approximate Kernel

This section shows that the KMCG kernel (Eq. 34) is always positive semi-definite.

Proposition 43. *If $k_0 = 0$, \mathbf{S} has rank M and k and w are positive definite kernel functions then the posterior mean in Equation (63) is symmetric and positive semi-definite.*

Proof. With $k_0 = 0$ the expression for k_M simplifies to

$$k_M(\mathbf{x}, \mathbf{z}) = w(\mathbf{x}, \mathbf{X}_U)\mathbf{S}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{K}_M \mathbf{S}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top w(\mathbf{X}_U, \mathbf{z})$$

The function k_M is symmetric since k is symmetric. The bivariate function k_M is said to be positive (semi-)definite iff for all $n \in \mathbb{N}$ and for all $\mathbf{Z} \in \mathbb{X}$, $k_M(\mathbf{Z}, \mathbf{Z})$ is a positive (semi-)definite matrix. Since $k(\mathbf{X}_U, \mathbf{X}_U)$ is an s.p.d. matrix, so is $\mathbf{S}^\top k(\mathbf{X}_U, \mathbf{X}_U)\mathbf{S}$ for arbitrary \mathbf{S} . The same argument holds for $\mathbf{S}^\top \mathbf{W}_M \mathbf{S}$. Since \mathbf{S} is rank M , $(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1}$ exists and the inverse of an s.p.d. matrix is s.p.d. as well. Therefore $\mathbf{S}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{K}_M \mathbf{S}(\mathbf{S}^\top \mathbf{W}_M \mathbf{S})^{-1} \mathbf{S}$ is symmetric and positive semi-definite. This completes the proof. □

E.3 Additional Experiments and Results

This section consists of figures showing the results of the experiments-section (Section 13) for the Matérn kernel, real-time experiments and experiments with the textbook version of conjugate gradients. All figures have been trimmed to the slowest baseline method.

E.3.1 Real-time Results

This section shows the same results as in Section 13 but over training-time instead of CG-steps. Figure 15 shows how the relative error ϵ_f develops over time for the squared exponential kernel and Figure 16 shows the same for experiments over grid-structured datasets from Section 13.2. For the x-axis values we took the median of all measurements and fitted a quadratic function to these.

E.3.2 Matérn Kernel Results

The figures in this section show the results for the Matérn $5/2$ kernel (Eq. (44)) for the experiment setup described Section 13. Figure 17 shows the results for the relative error ϵ_f , Figure 18 and Figure 19 the results for ϵ_{var} and ϵ_{ev} , respectively. Figure 20 displays the relative error over time.

E.3.3 Instability of Textbook Conjugate Gradients

The experiments in Section 13, where carried out by running conjugate gradients with full reorthogonalization. Figure 21 demonstrates that for the problems under consideration, the textbook version of conjugate gradients is not sufficiently numerically stable. With vanilla conjugate gradients in the background, KMCG can run only for a couple of steps before the necessary Cholesky decompositions fail to be computable.

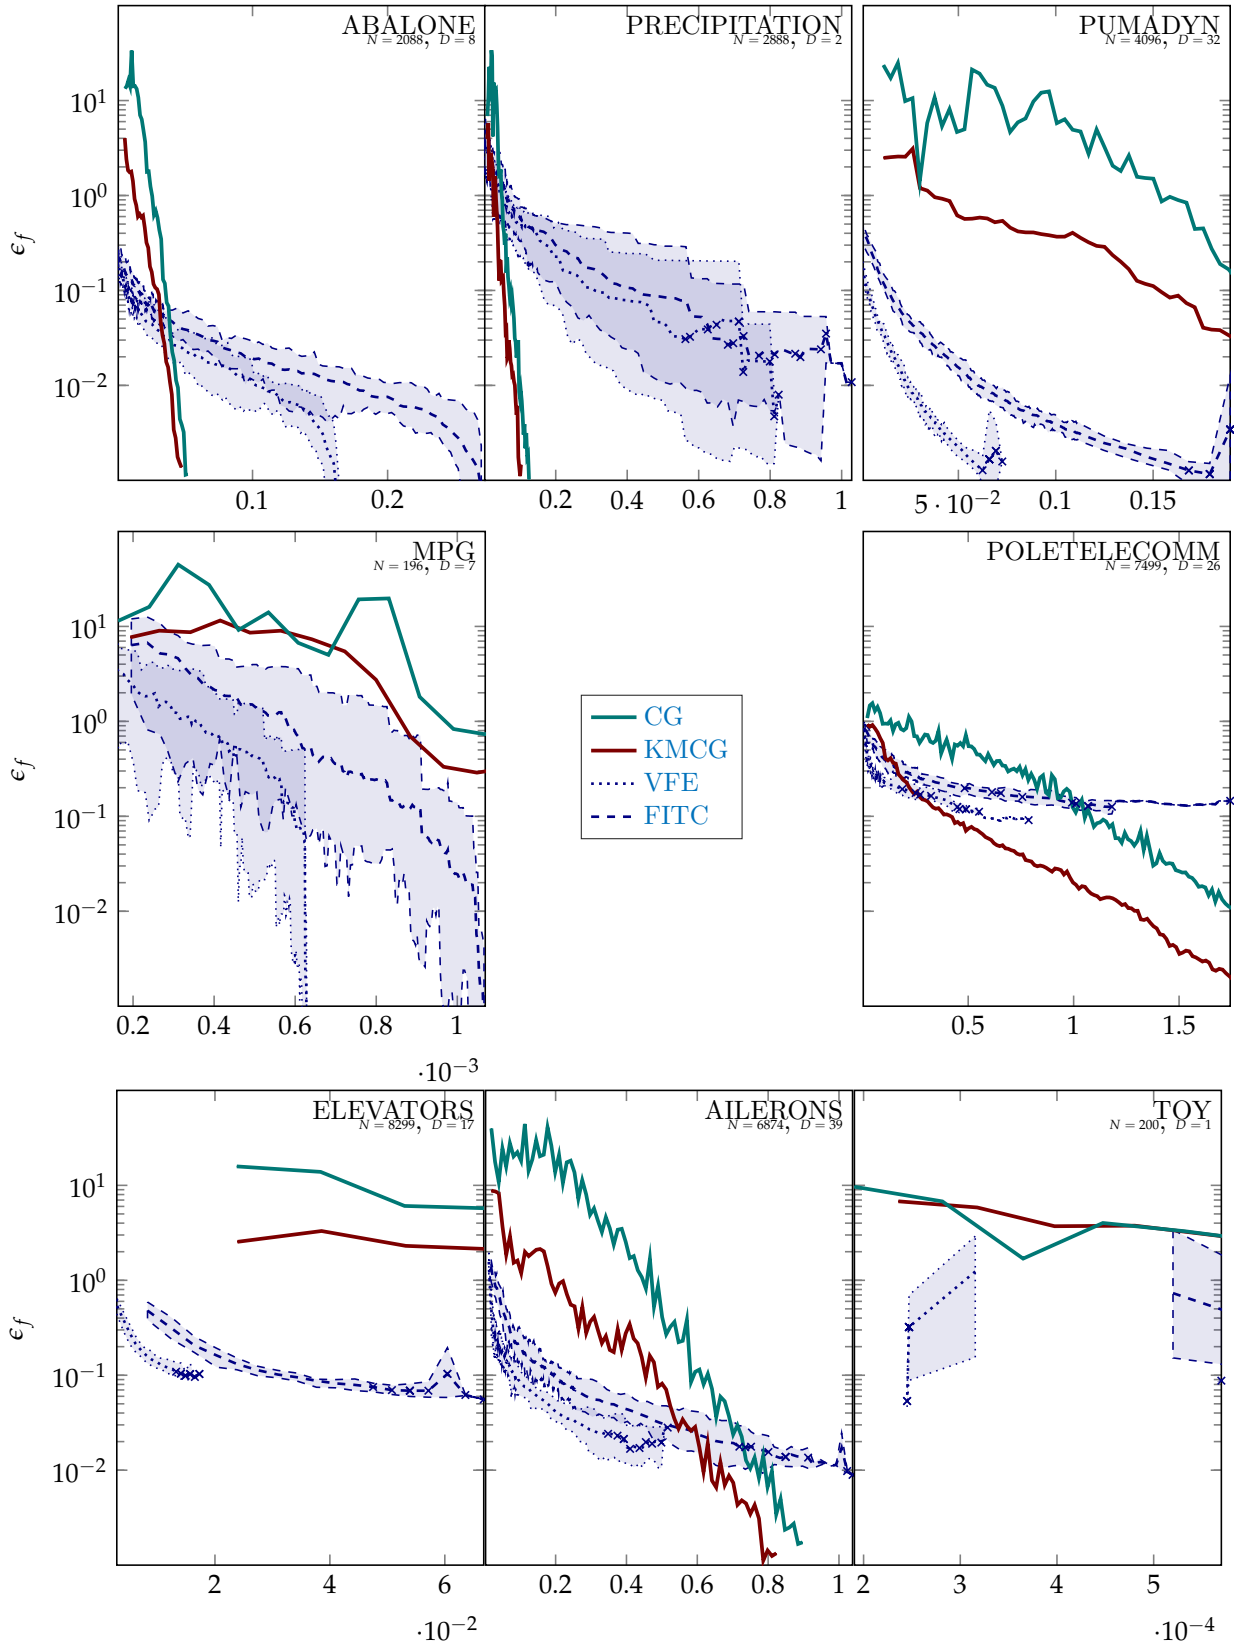


Figure 15: progression of the relative error ϵ_f over training time for different datasets using the squared-exponential kernel (Eq. 43). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.




Figure 16: comparison of baseline and KMCG on grid-structured datasets using the squared exponential kernel (Eq. 43). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run. It may seem surprising that the runs on the $100 \times 100 \times 100$ dataset take more than twice as long. By chance the dataset contains more extreme values in the kernel matrix, *i.e.* smaller than $1e^{-50}$. Multiplication with these elements takes more time.

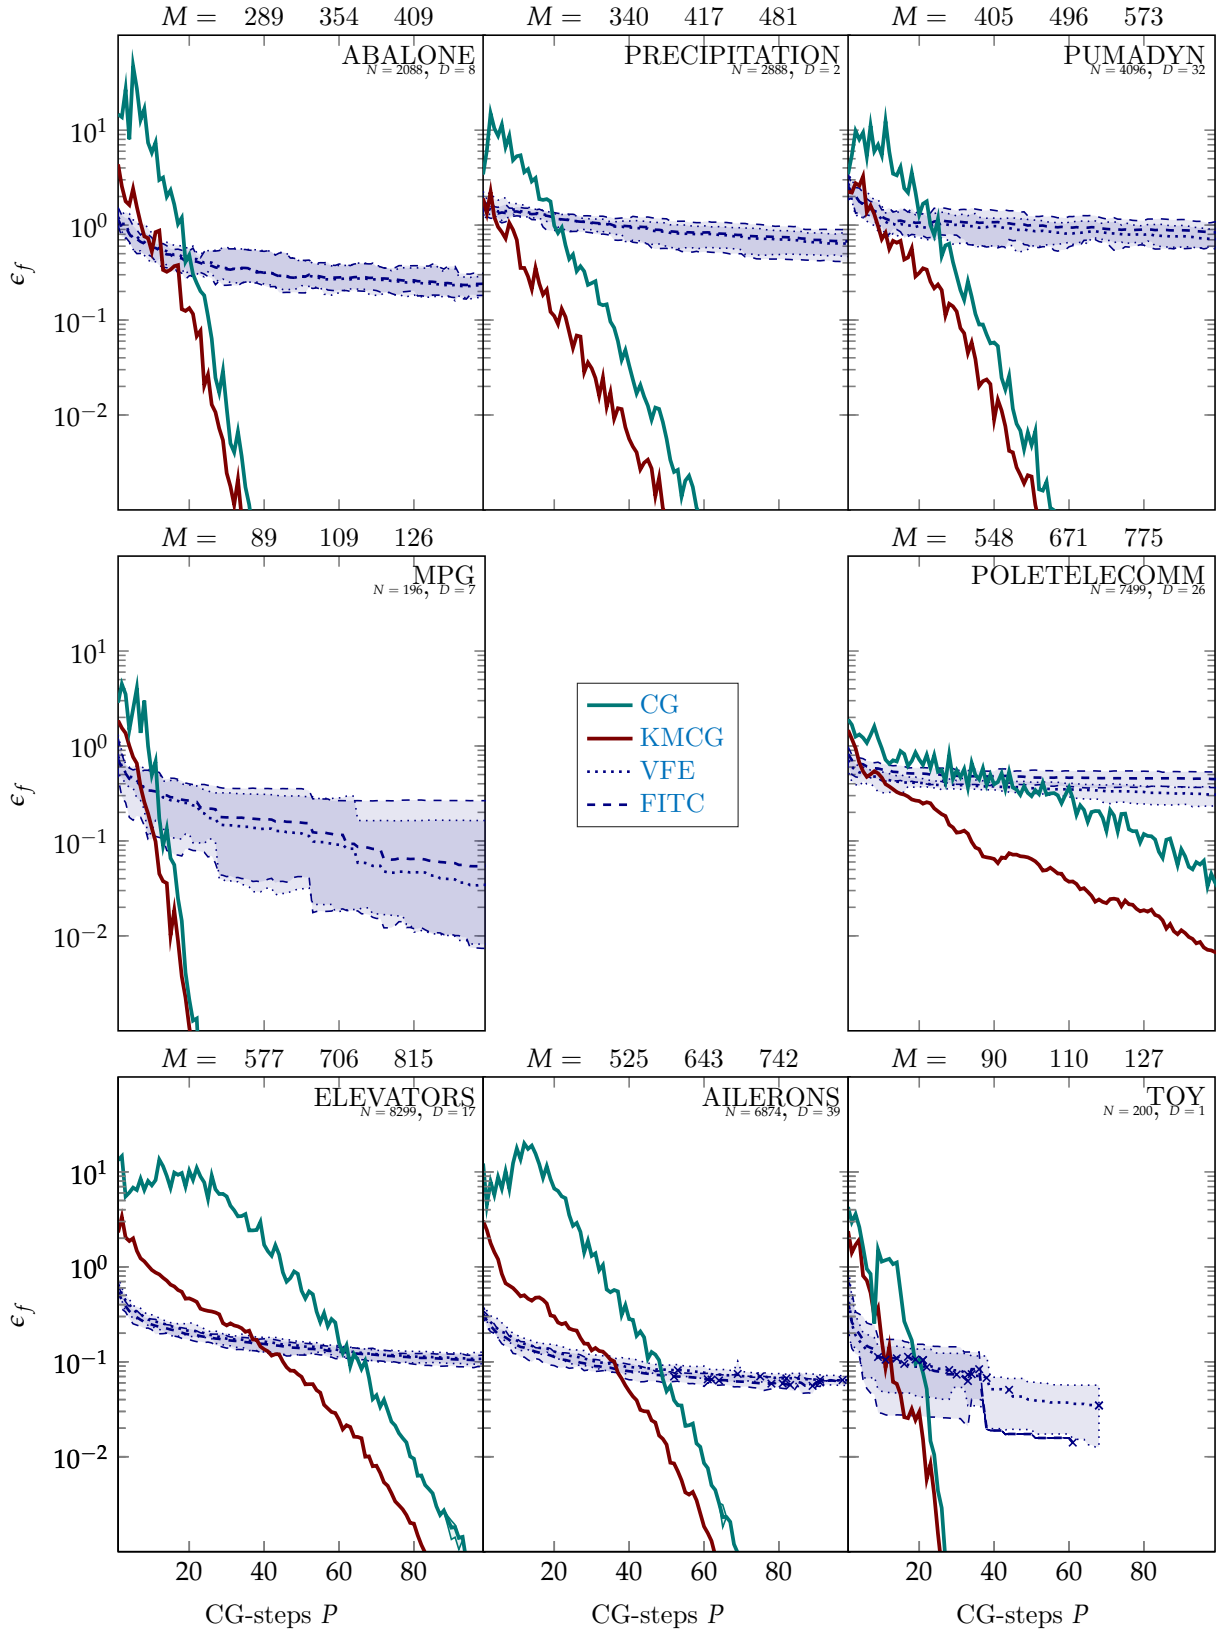


Figure 17: progression of the relative error ϵ_f as a function of the number of iterations of CG and KMCG for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.




Figure 18: progression of the relative error of the variance as a function of the number of iterations of KMCG and baseline for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

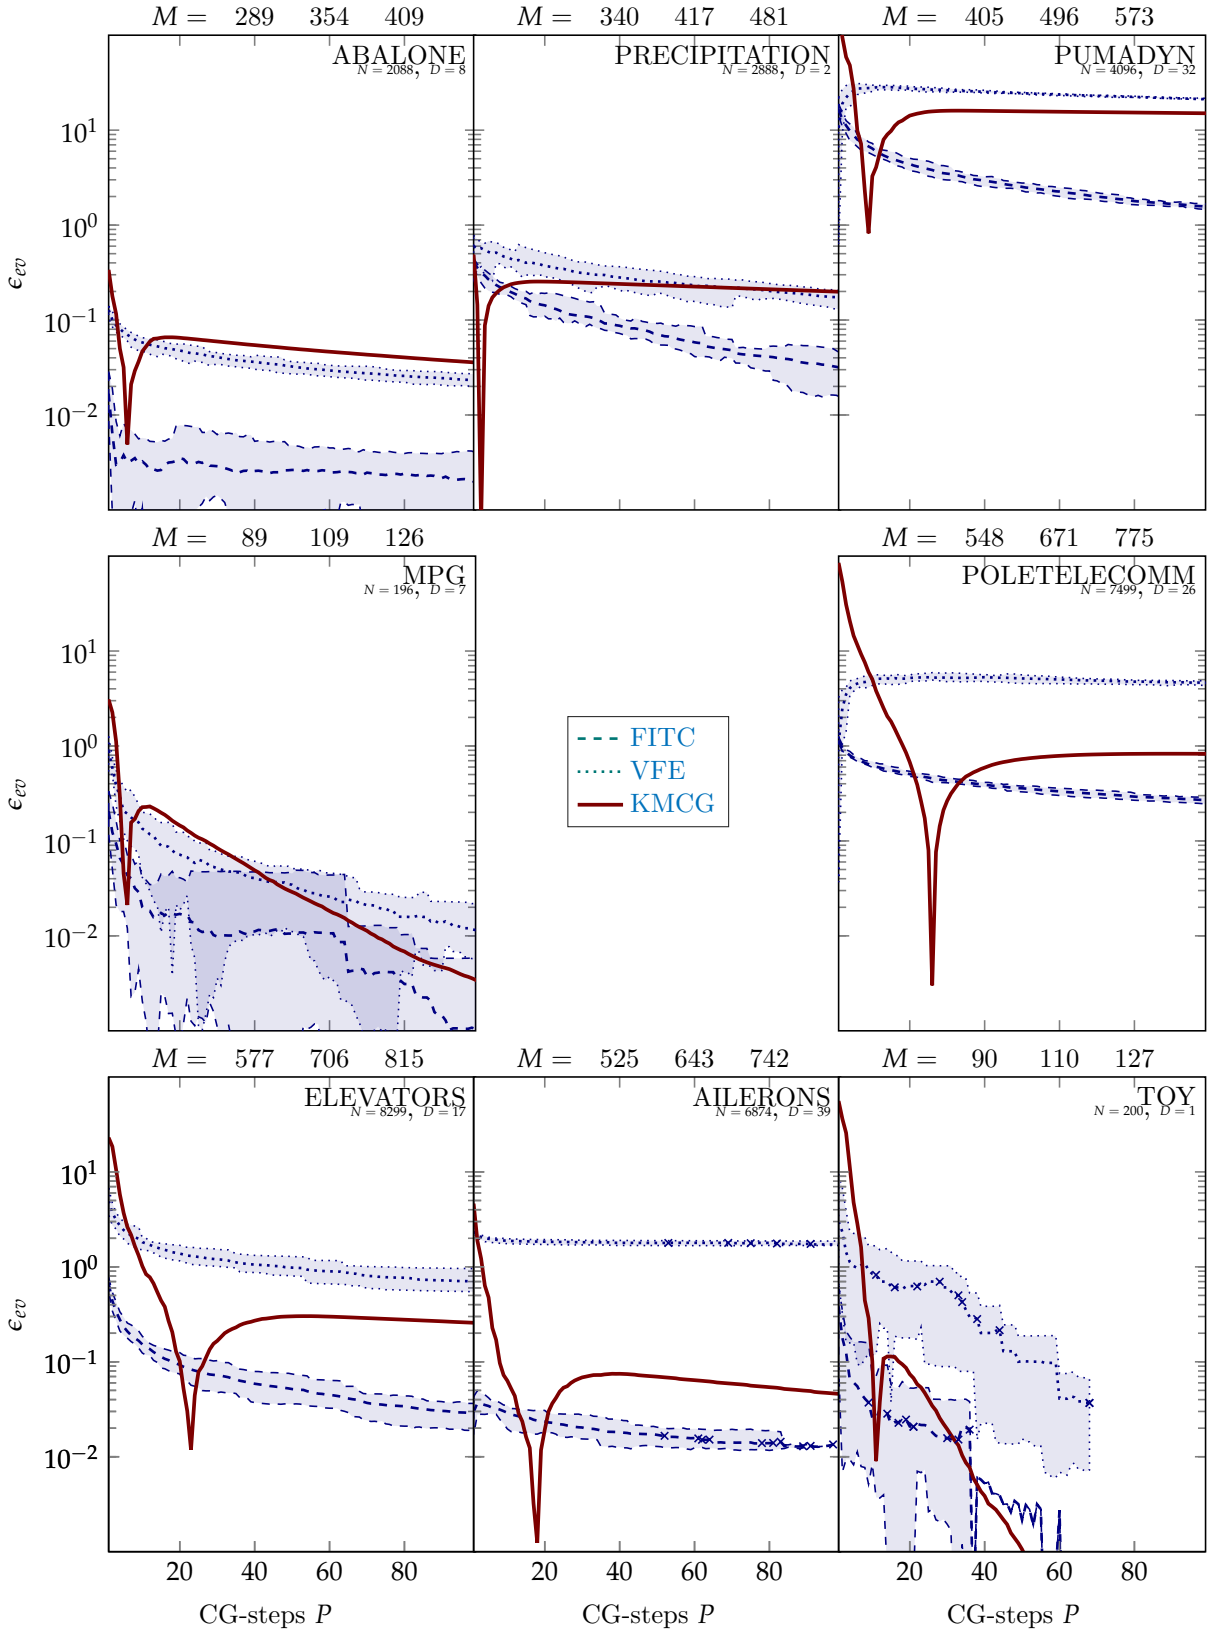


Figure 19: progression of the relative error of the evidence as a function of the number of iterations of baseline and KMCG for different datasets using the Matérn kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

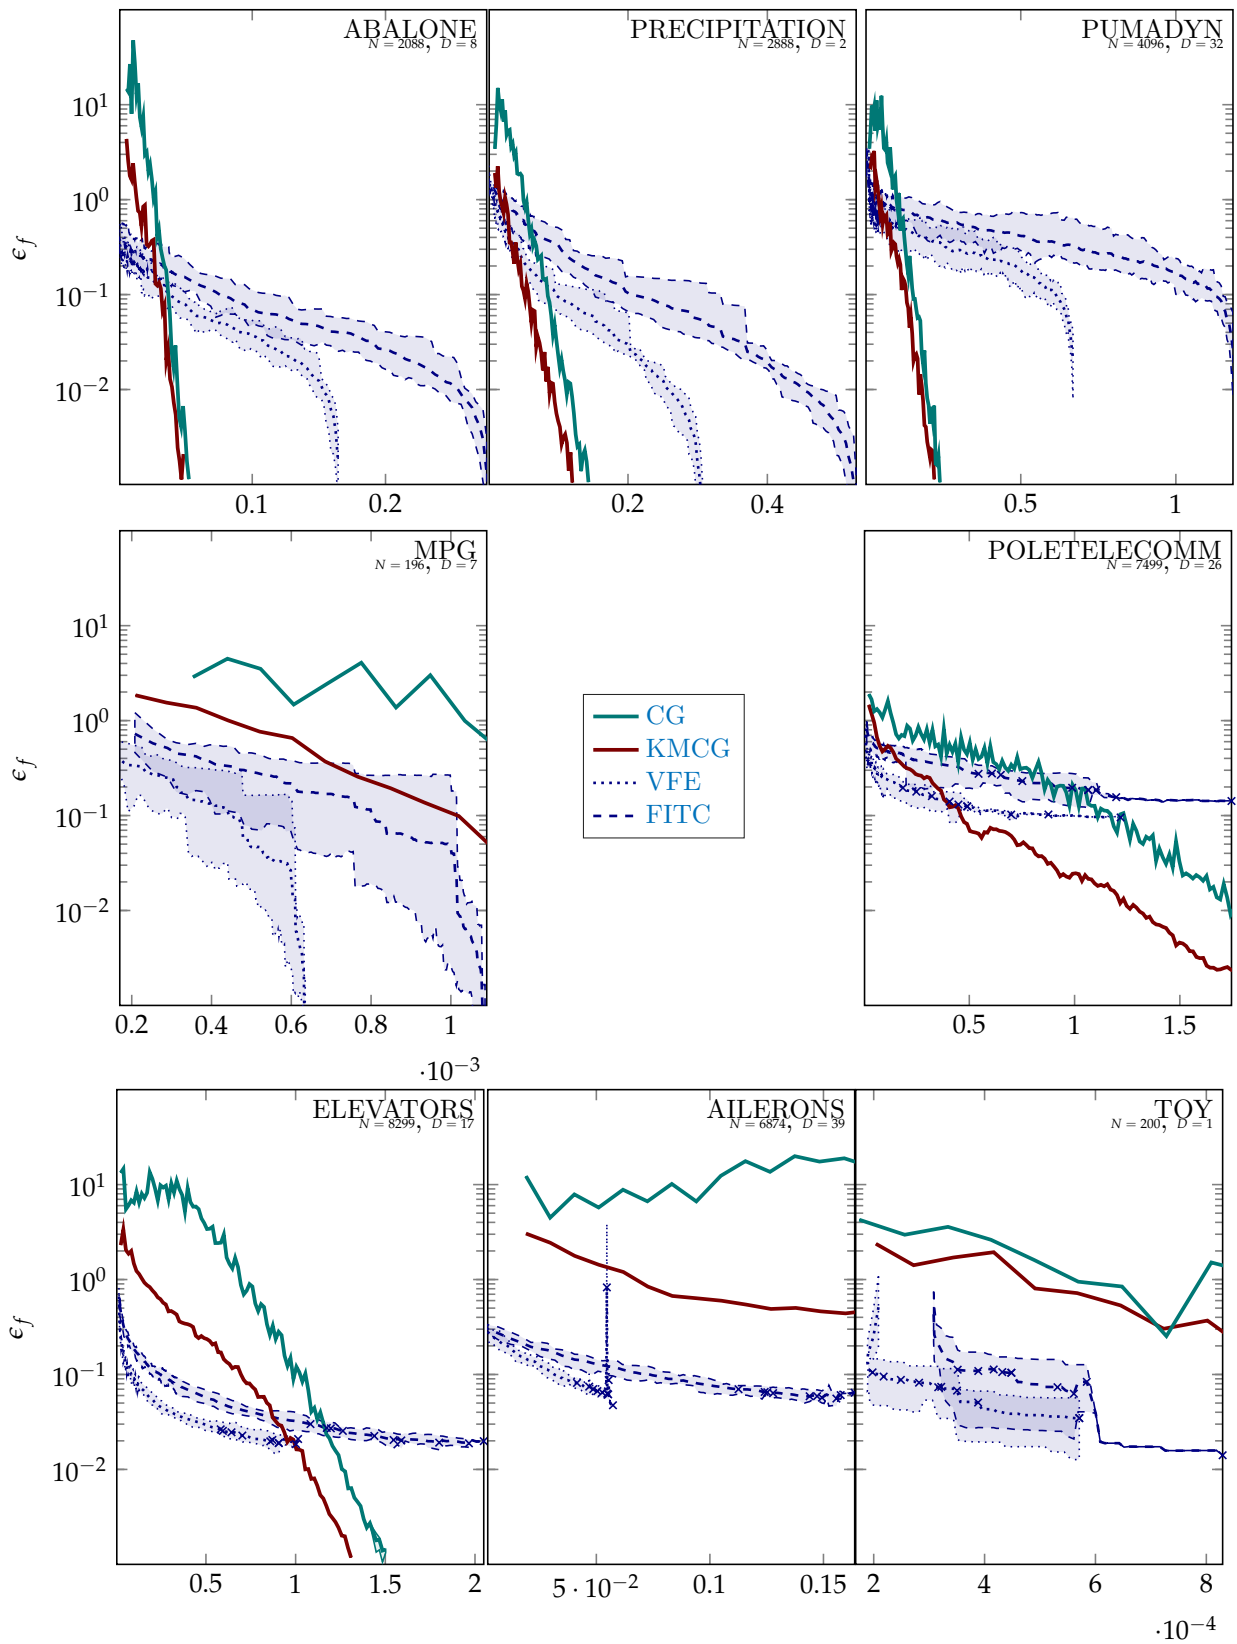


Figure 20: progression of the relative error ϵ_f over training time for different datasets using the Matérn $5/2$ kernel (Eq. 44). The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

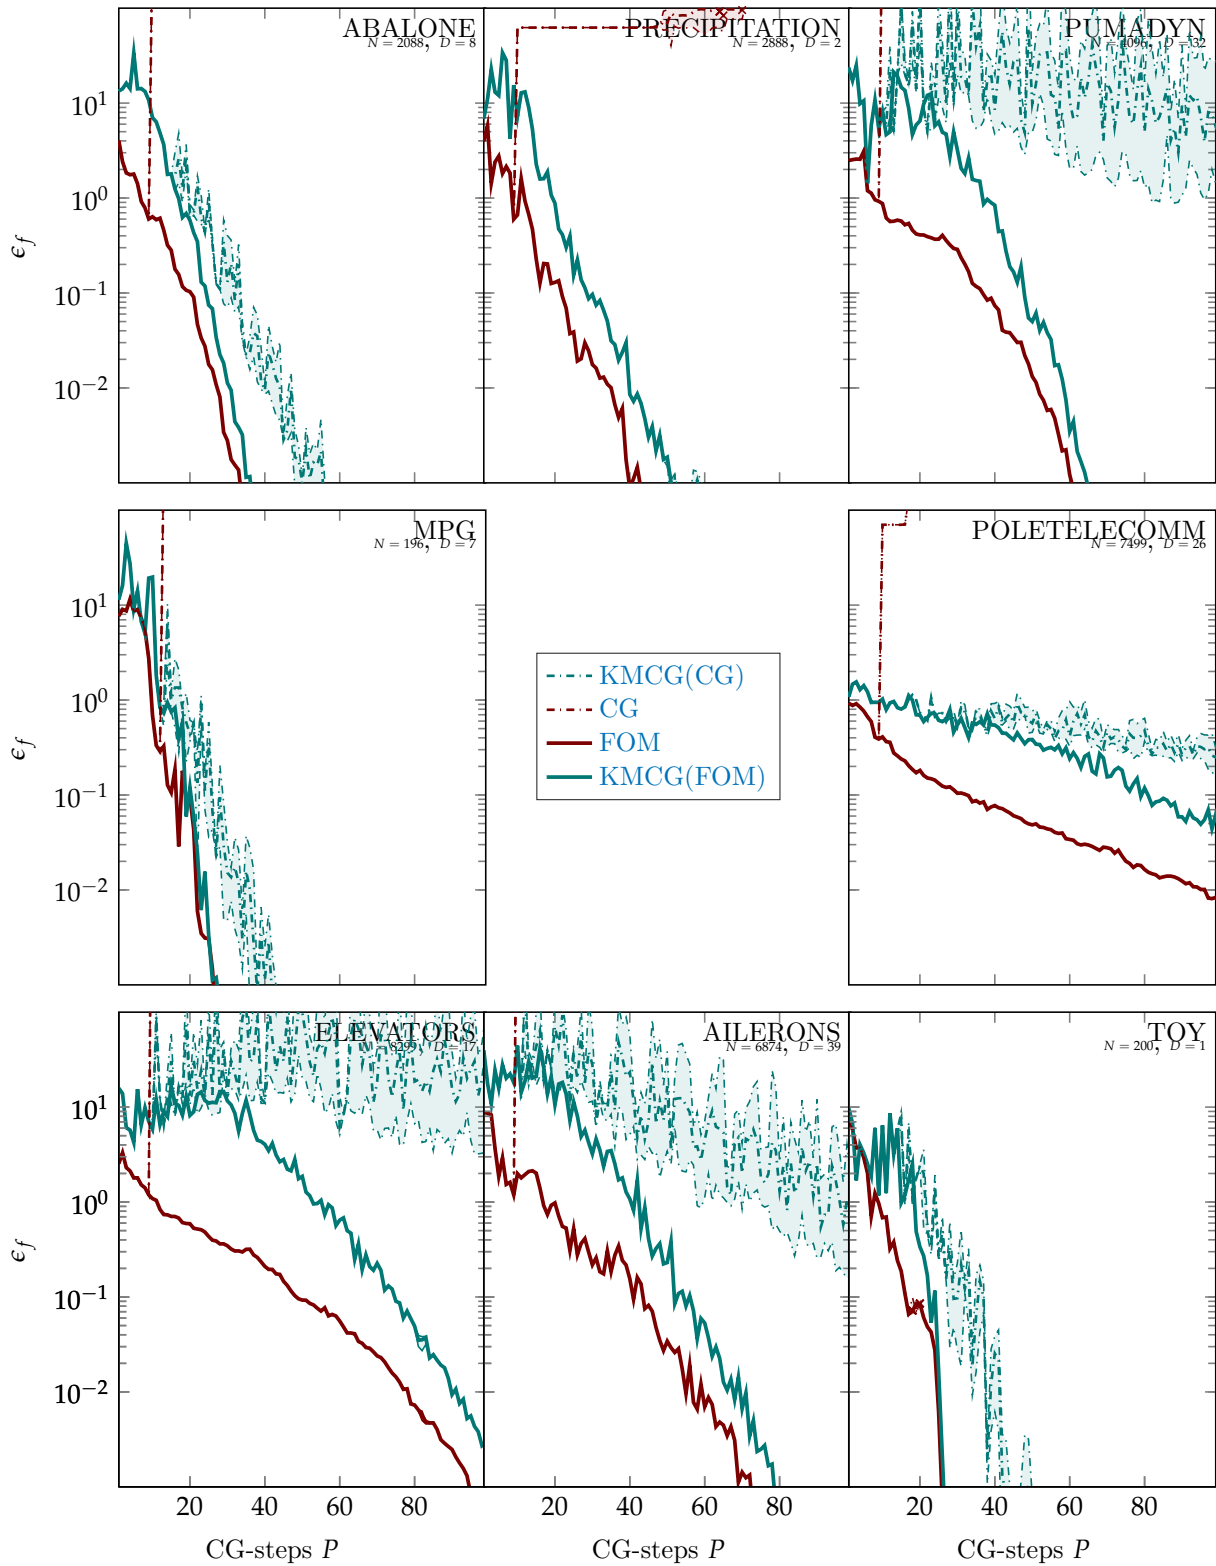


Figure 21: progression of the relative error ϵ_f over 100 CG-steps for different datasets using the squared exponential kernel (Eq. 43), comparing CG and FOM. The shaded area visualizes minimum and maximum over all baseline runs. A cross denotes the end of a crashed run.

Additional Material for Part v

Theorem 44 (Popoviciu (1935) and Sharma, Gupta, and Kapoor (2010)). *For a sequence of real numbers $x_1, \dots, x_n \in [m, M]$, define $\mu := \frac{1}{n} \sum_{j=1}^n x_j$ and $\sigma^2 := \frac{1}{n} \sum_{j=1}^n (x_j - \mu)^2$, then*

$$\sigma^2 \leq 1/4(M - m)^2.$$

Remark 45. *Above theorem can be used for a bound on the (conditional) variance as well. Let $x_1, \dots, x_n \sim P(\cdot \mid \mathcal{F})$ be independent. Then,*

$$\begin{aligned} \mathbb{V}[X \mid \mathcal{F}] &= \mathbb{E}[(X - \mathbb{E}[X \mid \mathcal{F}])^2 \mid \mathcal{F}] \\ &= \frac{n}{n-1} \mathbb{E}[\sigma^2 \mid \mathcal{F}] \\ &\quad // \text{ using Bessel's correction} \\ &\leq \frac{n}{4(n-1)} (M - m)^2 \end{aligned}$$

which holds for all $n \in \mathbb{N}$. Hence, $\mathbb{V}[X \mid \mathcal{F}] \leq 1/4(M - m)^2$.

Theorem 46 (Doob's Optional Sampling Theorem (Grimmett and Stirzaker 2001, p. 489)). *Let $(X_j, \mathcal{F}_j)_{j \in \mathbb{N}}$ be a submartingale and $\tau_1 \leq \tau_2 \leq \dots$ be a sequence of stopping times s.t. $P(\tau_j \leq n_j) = 1$ for some deterministic real sequence n_j , then the stopped process $(X_{\tau_j}, \mathcal{F}_{\tau_j})_{j \in \mathbb{N}}$ is also a submartingale.*

Remark 47. *By exchanging X_j for $-X_j$ the theorem can be shown to hold for supermartingales as well.*

Corollary 48. *Let $(\xi_j, \mathcal{F}_j)_{j \in \mathbb{N}}$ be a submartingale-difference and let τ be a stopping time, then the stopped process $(\xi_{\min(j, \tau)}, \mathcal{F}_{\min(j, \tau)})_{j \in \mathbb{N}}$ is also a submartingale-difference.*

Proof. Define $X_l := \sum_{j=1}^l \xi_j$ and observe that this defines a submartingale. By Theorem 46 $(X_{\min(j, \tau)}, \mathcal{F}_{\min(j, \tau)})_{j \in \mathbb{N}}$ is a submartingale. Then $X_{\min(j, \tau)} - X_{\min(j, \tau)-1} = \xi_{\min(j, \tau)}$ is again a submartingale-difference. \square

The following lemma appears as well in Mnih (2008). It was relieving to see that other researchers have the same ideas and I took this as a sign to be on the right track.

Lemma 25 (Bound on Relative Error). *Let $D, \hat{D} \in [\mathcal{L}, \mathcal{U}]$, and assume $\text{sign}(\mathcal{L}) = \text{sign}(\mathcal{U}) \neq 0$. Then the relative error of the estimator \hat{D} can be bounded as*

$$\frac{\text{abs}(D - \hat{D})}{\text{abs}(D)} \leq \frac{\max(\mathcal{U} - \hat{D}, \hat{D} - \mathcal{L})}{\min(\text{abs}(\mathcal{L}), \text{abs}(\mathcal{U}))}.$$

Proof. First observe that if $D_N > \hat{D}$ then $\text{abs}(D_N - \hat{D}) = D_N - \hat{D} \leq \mathcal{U} - \hat{D}$. If not, then $\text{abs}(D_N - \hat{D}) = \hat{D} - D_N \leq \hat{D} - \mathcal{L}$. Hence,

$$\text{abs}(D_N - \hat{D}) \leq \max(\mathcal{U} - \hat{D}, \hat{D} - \mathcal{L}).$$

Case $\mathcal{L} > 0$: This implies $\text{abs}(D_N) = D_N$.

$$\frac{\text{abs}(D_N - \hat{D})}{\text{abs}(D_N)} \leq \frac{\text{abs}(D_N - \hat{D})}{\mathcal{L}}$$

Case $\mathcal{U} < 0$: In that case $\text{abs}(\mathcal{L}) \geq \text{abs}(D_N) \geq \text{abs}(\mathcal{U})$.

$$\frac{\text{abs}(D_N - \hat{D})}{\text{abs}(D_N)} \leq \frac{\text{abs}(D_N - \hat{D})}{\text{abs}(\mathcal{U})}$$

Since we assumed $\text{sign}(\mathcal{L}) = \text{sign}(\mathcal{U})$ these were all cases that required consideration. Note that, in the first case $1/\mathcal{L} \geq 1/\text{abs}(\mathcal{U})$, and in the second case, the $1/\mathcal{L} \leq 1/\text{abs}(\mathcal{U})$, since $\mathcal{L} < 0$. Combining all observations gives

$$\frac{\text{abs}(D_N - \hat{D})}{\text{abs}(D_N)} \leq \max(\mathcal{U} - \hat{D}, \hat{D} - \mathcal{L}) \max\left(\frac{1}{\text{abs}(\mathcal{U})}, \frac{1}{\mathcal{L}}\right)$$

If $\mathcal{L} > 0$, $\max\left(\frac{1}{\text{abs}(\mathcal{U})}, \frac{1}{\mathcal{L}}\right) = \min(\text{abs}(\mathcal{U}), \mathcal{L}) = \min(\text{abs}(\mathcal{U}), \text{abs}(\mathcal{L}))$. In the other case, $\mathcal{U} < 0$, we can write

$$\max\left(\frac{1}{\text{abs}(\mathcal{U})}, \frac{1}{\mathcal{L}}\right) = \max\left(\frac{1}{\text{abs}(\mathcal{U})}, \frac{1}{\text{abs}(\mathcal{L})}\right) = \min(\text{abs}(\mathcal{U}), \text{abs}(\mathcal{L})).$$

□

Lemma 49. Denote with \mathbf{C} the Cholesky decomposition of a symmetric and positive definite matrix \mathbf{K} . The log-determinant of \mathbf{K} equals two times the sum over the logarithm of the diagonal elements of the Cholesky decomposition \mathbf{C} :

$$\ln |\mathbf{K}| = 2 \sum_{j=1}^N \ln C_{jj}.$$

Proof.

$$\begin{aligned} \ln |\mathbf{K}| &= \ln |\mathbf{C}\mathbf{C}^\top| \\ &\quad // \text{ using } \mathbf{K} = \mathbf{C}\mathbf{C}^\top \\ &= \ln(|\mathbf{C}| \cdot |\mathbf{C}^\top|) \\ &\quad // \text{ property of the determinant} \\ &= \ln(|\mathbf{C}|^2) \\ &\quad // \text{ transposition does not affect the determinant} \\ &= \ln\left(\prod_{j=1}^N C_{jj}\right)^2 \\ &\quad // \text{ property of triangular matrices} \\ &= 2 \sum_{j=1}^N \ln C_{jj} \\ &\quad // \text{ property of the logarithm} \end{aligned}$$

□

Bibliography

- Arras, P., J. Knollmüller, H. Junklewitz, and T. A. Enßlin (2018). “Radio Imaging with Information Field Theory.” In: *2018 26th European Signal Processing Conference (EUSIPCO)*, pp. 2683–2687.
- Bardenet, R., A. Doucet, and C. Holmes (2014). “Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach.” In: *Proceedings of the 31st International Conference on Machine Learning*. Ed. by E. P. Xing and T. Jebara. Vol. 32. Proceedings of Machine Learning Research 1, pp. 405–413.
- Bartels, S., J. Cockayne, I. C. F. Ipsen, and P. Hennig (2019). “Probabilistic Linear Solvers: A Unifying View.” In: *Statistics and Computing* 29.6, pp. 1249–1263.
- Bartels, S. and P. Hennig (2016). “Probabilistic Approximate Least-Squares.” In: *Proceedings of Artificial Intelligence and Statistics (AISTATS)*.
- (2019). “Conjugate Gradients for Kernel Machines.” In: *ArXiv e-prints* 1911.06048. [arXiv:1911.06048](https://arxiv.org/abs/1911.06048).
- Benoit (1924). “Note sûre une méthode de résolution des équations normales provenant de l’application de la méthode des moindres carrés a un système d’équations linéaires en nombre inférieure a celui des inconnues. Application de la méthode a la résolution d’un système défini d’équations linéaires. (Procédé du Commandant Cholesky).” In: *Bulletin Geodesique* 7.1, pp. 67–77.
- Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Boucheron, S., G. Lugosi, and P. Massart (2013). *Concentration Inequalities: A Nonasymptotic Theory of Independence*. 1st. Oxford University Press.
- Boutsidis, C., P. Drineas, P. Kambadur, E.-M. Kontopoulou, and A. Zouzias (2017). “A randomized algorithm for approximating the log determinant of a symmetric positive definite matrix.” In: *Linear Algebra and its Applications* 533, pp. 95–117.
- Camachol, R. (1998). “Inducing models of human control skills.” In: *Machine Learning: ECML-98*. Ed. by C. Nédellec and C. Rouveirol, pp. 107–118.
- Chalupka, K., C. K. I. Williams, and I. Murray (2013). “A Framework for Evaluating Approximation Methods for Gaussian Process Regression.” In: *Journal of Machine Learning Research* 14.1, pp. 333–350.
- Cockayne, J., C. Oates, I. C. F. Ipsen, and M. Girolami (2018). “A Bayesian Conjugate Gradient Method.” In: *ArXiv e-prints* 1801.05242. [arXiv:1801.05242](https://arxiv.org/abs/1801.05242).
- Cockayne, J., C. Oates, T. J. Sullivan, and M. Girolami (2016). “Probabilistic Numerical Methods for Partial Differential Equations and Bayesian Inverse Problems.” In: *ArXiv e-prints* 1605.07811. [arXiv:1605.07811](https://arxiv.org/abs/1605.07811).
- (2017). “Bayesian Probabilistic Numerical Methods.” In: *ArXiv e-prints* 1702.03673. [arXiv:1702.03673](https://arxiv.org/abs/1702.03673).
- Csató, L. and M. Opper (2002). “Sparse On-line Gaussian Processes.” In: *Neural Computation* 14.3, pp. 641–668.
- Davidson, J. (1994). *Stochastic Limit Theory: An Introduction for Econometricians*. Oxford University Press.

- Davies, A. (2015). “Effective implementation of Gaussian process regression for machine learning.” PhD thesis. University of Cambridge.
- DeGroot, M. H. and M. J. Schervish (2012). *Probability and Statistics*. 4th. Pearson.
- Diaconis, P. and M. Shahshahani (1987). “The Subgroup Algorithm for Generating Uniform Random Variables.” In: *Probability in the Engineering and Informational Sciences* 1.01, p. 15. DOI: [10.1017/s0269964800000255](https://doi.org/10.1017/s0269964800000255).
- Dong, K., D. Eriksson, H. Nickisch, D. Bindel, and A. G. Wilson (2017). “Scalable Log Determinants for Gaussian Process Kernel Learning.” In: *Advances in Neural Information Processing Systems*, pp. 6330–6340.
- Dorn, S. and T. A. Enßlin (2015). “Stochastic determination of matrix determinants.” In: *Physical Review E* 92 (1), p. 013302.
- Dua, D. and C. Graff (2019). *UCI Machine Learning Repository*. URL: <http://archive.ics.uci.edu/ml>.
- Fan, X., I. Grama, and Q. Liu (2012). “Hoeffding’s inequality for supermartingales.” In: *Stochastic Processes and their Applications* 122.10, pp. 3545–3559.
- Filippone, M. and R. Engler (2015). “Enabling scalable stochastic gradient-based inference for Gaussian processes by employing the Unbiased Linear System SolvEr (ULISSE).” In: *Proceedings of the 32nd International Conference on Machine Learning*. Lille, France, pp. 1015–1024.
- Fitzsimons, J., K. Cutajar, M. Osborne, S. Roberts, and M. Filippone (2017). “Bayesian Inference of Log Determinants.” In: *Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, August 11-15, 2017, Sydney, Australia*. Ed. by G. Elidan, K. Kersting, and A. T. Ihler.
- Fitzsimons, J., D. Granziol, K. Cutajar, M. Osborne, M. Filippone, and S. Roberts (2017). “Entropic Trace Estimates for Log Determinants.” In: *Machine Learning and Knowledge Discovery in Databases*. Ed. by M. Ceci, J. Hollmén, L. Todorovski, C. Vens, and S. Džeroski, pp. 323–338.
- Geman, S. and D. Geman (1984). “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images.” In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6.6, pp. 721–741.
- Golub, G. and C. Van Loan (2013). *Matrix computations*. 4th ed. Johns Hopkins Univ Pr.
- Graf, F., H.-P. Kriegel, M. Schubert, S. Pölsterl, and A. Cavallaro (2011). “2D Image Registration in CT Images Using Radial Image Descriptors.” In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011*. Ed. by G. Fichtinger, A. Martel, and T. Peters, pp. 607–614.
- Grimmett, G. and D. Stirzaker (2001). *Probability and Random Processes*. 3rd. Oxford University Press.
- Gut, A. (2009). *An Intermediate Course in Probability*. 2nd. Springer.
- Hennig, P. (2015). “Probabilistic Interpretation of Linear Solvers.” In: *SIAM J on Optimization* 25.1, pp. 210–233.
- Hennig, P. and M. Kiefel (2012). “Quasi-Newton methods – a new direction.” In: *International Conference on Machine Learning (ICML)*.
- (2013). “Quasi-Newton Methods – a new direction.” In: *Journal of Machine Learning Research* 14, pp. 834–865.

- Hennig, P., M. Osborne, and M. Girolami (2015). “Probabilistic numerics and uncertainty in computations.” In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 471.2179.
- Hensman, J., N. Durrande, and A. Solin (2018). “Variational Fourier Features for Gaussian Processes.” In: *Journal of Machine Learning Research* 18.151, pp. 1–52.
- Hestenes, M. and E. Stiefel (1952). “Methods of conjugate gradients for solving linear systems.” In: *Journal of Research of the National Bureau of Standards* 49.6, pp. 409–436.
- Hoerl, A. E. and R. W. Kennard (1970). “Ridge regression: Biased estimation for nonorthogonal problems.” In: *Technometrics* 12.1, pp. 55–67.
- Karvonen, T. and S. Sarkka (2017). “Classical quadrature rules via Gaussian processes.” In: *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE.
- Kersting, H., T. J. Sullivan, and P. Hennig (2018). “Convergence Rates of Gaussian ODE Filters.” In: *ArXiv e-prints* 1807.09737. [arXiv:1807.09737](https://arxiv.org/abs/1807.09737).
- Kimeldorf, G. S. and G. Wahba (1970). “A correspondence between Bayesian estimation on stochastic processes and smoothing by splines.” In: *The Annals of Mathematical Statistics*, pp. 495–502.
- Lázaro-Gredilla, M., J. Quiñero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal (2010). “Sparse Spectrum Gaussian Process Regression.” In: *Journal of Machine Learning Research* 11, pp. 1865–1881.
- Le, Q., T. Sarlos, and A. Smola (2013). “Fastfood - Computing Hilbert Space Expansions in loglinear time.” In: *Proceedings of the 28th International Conference on Machine Learning*, pp. 244–252.
- Loan, C. F. V. (2000). “The ubiquitous Kronecker product.” In: *Journal of Computational and Applied Mathematics* 123.1. Numerical Analysis 2000. Vol. III: Linear Algebra, pp. 85–100.
- Magnus, J. R. and H. Neudecker (1980). “The Elimination Matrix: Some Lemmas and Applications.” In: *SIAM Journal on Algebraic Discrete Methods* 1.4, pp. 422–449. DOI: [10.1137/0601049](https://doi.org/10.1137/0601049).
- Magnus, J. R. and H. Neudecker (1999). *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Second. John Wiley.
- Matheron, G. (1973). “The intrinsic random functions and their applications.” In: *Advances in applied probability*, pp. 439–468.
- Meister, A. (2015). *Numerik Linearer Gleichungssysteme*. 5th. Springer Fachmedien Wiesbaden.
- Mnih, V. (2008). “Efficient stopping rules.” MA thesis. University of Alberta, Canada.
- Mnih, V., C. Szepesvári, and J. Audibert (2008). “Empirical Bernstein stopping.” In: ed. by A. McCallum and S. Roweis, pp. 672–679.
- Munroe, R. (2019). *Error Bars*. URL: <https://xkcd.com/2110/>. License: Creative Commons 2.5 BY-NC-SA.
- Nash, W. J., T. L. Sellers, S. R. Talbot, A. J. Cawthorn, and W. B. Ford (1994). *The Population biology of abalone (Haliotis species) in Tasmania. 1, Blacklip abalone (H. rubra) from the north coast and the islands of Bass Strait*. Tech. rep. 48. Sea Fisheries Division, Marine Research Laboratories - Tarooma, Department of Primary Industry and Fisheries, Tasmania. URL: <https://trove.nla.gov.au/work/11326142>.

- Nille, D., U. von Toussaint, B. Sieglin, and M. Faitsch (2018). “Probabilistic Inference of Surface Heat Flux Densities from Infrared Thermography.” In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. Ed. by A. Polpo, J. Stern, F. Louzada, R. Izbicki, and H. Takada, pp. 55–64.
- Nocedal, J. and S. Wright (1999). *Numerical Optimization*. Springer Verlag.
- Pleiss, G., J. Gardner, K. Weinberger, and A. G. Wilson (2018). “Constant-Time Predictive Distributions for Gaussian Processes.” In: *Proceedings of the 35th International Conference on Machine Learning*, pp. 4114–4123.
- Popoviciu, T. (1935). “Sur les equations algebriques ayant toutes leurs racines reelles.” In: *Mathematica (Cluj)* 9, pp. 129–145.
- Quinlan, J. R. (1993). “Combining Instance-based and Model-based Learning.” In: *Proceedings of the Tenth International Conference on International Conference on Machine Learning*. ICML’93, pp. 236–243.
- Quiñero-Candela, J. and C. E. Rasmussen (2005). “A unifying view of sparse approximate Gaussian process regression.” In: *Journal of Machine Learning Research* 6, pp. 1939–1959.
- Rahimi, A. and B. Recht (2008). “Random Features for Large-Scale Kernel Machines.” In: *Advances in Neural Information Processing Systems 20*. Ed. by J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, pp. 1177–1184.
- (2009). “Weighted Sums of Random Kitchen Sinks: Replacing minimization with randomization in learning.” In: *Advances in Neural Information Processing Systems 23*, pp. 1313–1320.
- Rasmussen, C. E. and H. Nickisch (2010). “Gaussian Processes for Machine Learning (GPML) Toolbox.” In: *Journal of Machine Learning Research* 11, pp. 3011–3015.
- Rasmussen, C. and C. Williams (2006). *Gaussian Processes for Machine Learning*. MIT.
- Roeckner, E., G. Bäuml, L. Bonaventura, R. Brokopf, M. Esch, and M. Giorgetta (2003). *The atmospheric general circulation model ECHAM 5. part I: Model description*. Tech. rep. MPI für Meteorologie.
- Saad, Y. and M. H. Schultz (1986). “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems.” In: *SIAM Journal on Scientific and Statistical Computing* 7.3, pp. 856–869.
- Saad, Y. (2003). *Iterative Methods for Sparse Linear Systems*. 2nd. Society for Industrial and Applied Mathematics.
- Saibaba, A. K., A. Alexanderian, and I. C. F. Ipsen (2017). “Randomized matrix-free trace and log-determinant estimators.” In: *Numerische Mathematik* 137.2, pp. 353–395.
- Schober, M., D. Duvenaud, and P. Hennig (2014). “Probabilistic ODE Solvers with Runge-Kutta Means.” In: *Advances in Neural Information Processing Systems 27*, pp. 739–747.
- Schober, M., S. Särkkä, and P. Hennig (2018). “A probabilistic model for the numerical solution of initial value problems.” In: *Statistics and Computing*.
- Schölkopf, B. and A. Smola (2002). *Learning with Kernels*. MIT Press.
- Sharma, R., M. Gupta, and G. Kapoor (2010). “Some better bounds on the variance with applications.” In: *Journal of Mathematical Inequalities* 4, pp. 355–363.
- Skilling, J. (1989). “The Eigenvalues of Mega-dimensional Matrices.” In: *Maximum Entropy and Bayesian Methods: Cambridge, England, 1988*. Ed. by J. Skilling, pp. 455–466.

- (1993). “Bayesian Numerical Analysis.” In: *Physics and Probability: Essays in Honor of Edwin T. Jaynes*. Ed. by J. W. T. Grandy and P. W. Milonni, pp. 207–222.
- Snelson, E. and Z. Ghahramani (2006). “Sparse Gaussian Processes using Pseudo-inputs.” In: *Advances in Neural Information Processing Systems 18*. Ed. by Y. Weiss, B. Schölkopf, and J. C. Platt, pp. 1257–1264.
- (2007). “Local and global sparse Gaussian process approximations.” In: *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, pp. 524–531.
- Solin, A. and S. Särkkä (2014). “Hilbert Space Methods for Reduced-Rank Gaussian Process Regression.” In: *ArXiv e-prints* 1401.5508. [arXiv:1401.5508](https://arxiv.org/abs/1401.5508).
- Soodhalter, K. M., D. B. Szyld, and F. Xue (2014). “Krylov subspace recycling for sequences of shifted linear systems.” In: *Applied Numerical Mathematics* 81, pp. 105–118. DOI: [10.1016/j.apnum.2014.02.006](https://doi.org/10.1016/j.apnum.2014.02.006). URL: <https://doi.org/10.1016/j.apnum.2014.02.006>.
- Stark, P. B. and D. A. Freedman (2003). “What is the chance of an earthquake?” In: *Earthquake Science and Seismic Risk Reduction*, pp. 201–216.
- Titsias, M. (2009). “Variational Learning of Inducing Variables in Sparse Gaussian Processes.” In: *Artificial Intelligence and Statistics (AISTATS), JMLR W&CP 5*.
- Trecate, G. F., C. K. I. Williams, and M. Opper (1999). “Finite-dimensional Approximation of Gaussian Processes.” In: *Advances in Neural Information Processing Systems 2*, pp. 218–224.
- Tronarp, F., H. Kersting, S. Särkkä, and P. Hennig (2019). “Probabilistic solutions to ordinary differential equations as nonlinear Bayesian filtering: a new perspective.” In: *Statistics and Computing* 29.6, pp. 1297–1315.
- Turner, R. E. (2010). “Statistical Models for Natural Sounds.” PhD thesis. University College London.
- Vanhatalo, J. and A. Vehtari (2008). “Modelling local and global phenomena with sparse Gaussian processes.” In: *UAI 2008, Twenty-Fourth Conference on Uncertainty in Artificial Intelligence, Helsinki, Finland, July 9-12, 2008*. Ed. by D. McAllester and P. Myllymäki.
- Vijayakumar, S. and S. Schaal (2000). “Locally Weighted Projection Regression: An $\mathcal{O}(n)$ Algorithm for Incremental Real Time Learning in High Dimensional Space.” In: *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, pp. 1079–1086.
- Wahba, G. (1990). *Spline models for observational data*. CBMS-NSF Regional Conferences series in applied mathematics 59. SIAM.
- Walder, C., K. I. Kim, and B. Schölkopf (2008). “Sparse Multiscale Gaussian Process Regression.” In: *Proceedings of the 25th International Conference on Machine Learning*, pp. 1112–1119.
- Wang, Q., X. Zhang, Y. Zhang, and Q. Yi (2013). “AUGEM: Automatically Generate High Performance Dense Linear Algebra Kernels on x86 CPUs.” In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. SC ’13, 25:1–25:12.
- Waugh, S. (1995). “Extending and benchmarking Cascade-Correlation.” PhD thesis. University of Tasmania.
- Weiss, S. M. and N. Indurkha (1995). “Rule-based Machine Learning Methods for Functional Prediction.” In: *Journal of Artificial Intelligence Research* 3.1, pp. 383–403.

- Welling, M. and Y. W. Teh (2011). “Bayesian Learning via Stochastic Gradient Langevin Dynamics.” In: *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pp. 681–688.
- Williams, C. and M. Seeger (2001). “Using the Nyström Method to Speed Up Kernel Machines.” In: *Advances in Neural Information Processing Systems 13*.
- Wilson, A. and H. Nickisch (2015). “Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP).” In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by F. Bach and D. Blei. Vol. 37. Proceedings of Machine Learning Research, pp. 1775–1784.
- Wilson, A. G., E. Gilboa, A. Nehorai, and J. P. Cunningham (2014). “Fast Kernel Learning for Multidimensional Pattern Extrapolation.” In: *Advances in Neural Information Processing Systems 27*, pp. 3626–3634.
- Yan, F. and Y. Qi (2010). “Sparse Gaussian Process Regression via l1 Penalization.” In: *Proceedings of the 27th International Conference on Machine Learning*, pp. 1183–1190.
- Zhao, S., E. Zhou, A. Sabharwal, and S. Ermon (2016). “Adaptive Concentration Inequalities for Sequential Decision Problems.” In: *Advances in Neural Information Processing Systems 29*. Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, pp. 1343–1351.
- Zhu, H., C. K. I. Williams, R. J. Rohwer, and M. Morciniec (1998). “Gaussian regression and optimal finite dimensional linear models.” In: *Neural Networks and Machine Learning*.