The Generalised Gaussian Process Convolution Model



Wessel Pieter Bruinsma

Department of Engineering

University of Cambridge

This dissertation is submitted for the degree of

Master of Philosophy

Sidney Sussex College

12 August 2016

Declaration

I, Wessel Pieter Bruinsma of Sidney Sussex College, being a candidate for the M.Phil in Machine Learning, Speech and Language Technology, hereby declare that this thesis and the work described in it are my own work, unaided except as may be specified below, and that the thesis does not contain material that has already been used to any substantial extent for a comparable purpose. This thesis, excluding appendices, contains approximately 7300 words.

> Wessel Pieter Bruinsma 12 August 2016

LABRUINSAGE

Acknowledgements

First and foremost, I thank my supervisor Dr. Richard Turner. His guidance and support throughout the project have been indispensable.

I also thank Dr. Felipe Tobar, whose thoughts have been essential in the development of the project.

Abstract

This thesis formulates the Generalised Gaussian Process Convolution Model (GGPCM), which is a generalisation of the Gaussian Process Convolution Model presented by Tobar et al. [2015b]. The GGPCM provides a theoretical framework for nonparametric kernel models of multi-dimensional signals defined on multidimensional input spaces. We show that the GGPCM generalises and connects existing work; most notably, we derive a dual formulation of the cross-spectral mixture kernel presented by Ulrich et al. [2015]. Finally, we use the GGPCM to develop the Deep Kernel Model, which presents a new network structure for unsupervised learning.

Contents

Li	st of	Figures	ciii
\mathbf{Li}	st of	Tables	xv
Li	st of	Models x	vii
N	otati	n 2	xix
1	Intr	oduction	1
	1.1	Motivation	1
	1.2	Contribution of Thesis and Outline	3
2	Mo	lelling Dynamical Signals	5
	2.1	Introduction	5
	2.2	The Linear State-Space Model	5
	2.3	A General-Purpose Model of Dynamical Signals	6
	2.4	Conclusion	7
3	$Th\epsilon$	Generalised Gaussian Process Convolution Model	9
	3.1	Introduction	9
	3.2	Gaussian Process Regression	9
	3.3	The Generalised Gaussian Process Convolution Model	11
		3.3.1 Interpretation and Choice of the Kernel in the Nonparametric Ker-	
		nel Model	14
		3.3.2 Illustrative Samples of the Nonparametric Kernel Model	16
		3.3.3 Expressivity of White Noise Excitation	19
	3.4	The Approximate Kernel Model	21
		3.4.1 Interpretation of the Kernel Approximation	23

Contents

		3.4.2 The Case of the Diagonal Multi-Output Decaying Exponentiated-	
		Quadratic Kernel	23
	3.5	Related Work	25
	3.6	Inference in the Nonparametric Kernel Model	26
	3.7	Conclusion	30
	3.8	Discussion	30
4	Mu	lti-Task Learning	31
	4.1	Introduction	31
	4.2	Mixing Models	31
	4.3	The Mixing Model Hierarchy	33
	4.4	Conclusion	35
5	The	e Deep Kernel Model	37
	5.1	Introduction	37
	5.2	The Deep Kernel Model	37
		5.2.1 Network Interpretation	41
	5.3	Illustrative Samples	43
	5.4	Conclusion	44
	5.5	Discussion	44
\mathbf{A}	Solı	ntion of the Linear State-Space Model	47
	A.1	Time-Variant Solution	47
	A.2	Time-Invariant Solution	48
в	Pro	perties of the Multivariate Gaussian Distribution	51
	B.1	Marginal and Conditional Distribution	51
	B.2	Kullback-Leibler Divergence	51
С	Mu	ltivariate Matrix-Valued Gaussian Processes	53
D	Gau	assian Processes in Practice	55
	D.1	Implementation of Gaussian Process Models	55
	D.2	Nearest Symmetric Positive-Semidefinite Matrix	56
\mathbf{E}	Circ	culant Approximation of Stationary Multi-Output Kernel Matrices	61
	E.1	Introduction	61

${\rm Contents}$

	E.2	Circulant Approximation of Toeplitz Matrices	62
	E.3	Circulant Approximation of Stationary Multi-Output Kernel Matrices	63
	E.4	Approximating Determinants	65
	E.5	Approximating Products Involving an Inverse	67
	E.6	Conclusion	69
\mathbf{F}	\mathbf{Exp}	onentiated Quadratic Forms	71
	F.1	Introduction	71
	F.2	General Form	71
	F.3	Kronecker-Structured Form	75
	F.4	Conclusion	78
G	Roo	ts of Kernels	79
н	App	proximate Kernel Model	81
I	Vari	ational Free Energy of the Nonparametric Kernel Model	83
	I.1	Introduction	83
	I.2	Predictive Mean	84
	I.3	Predictive Autocovariance	85
	I.4	Integrals $I_{\cdot}^{(\cdot)}$	87
		I.4.1 Kernels	88
		I.4.2 Integral $I_{;;}^{(L,;,)}$	89
		I.4.3 Integral $I^{(Q_1,\cdot,\cdot)}$	90
		I.4.4 Integral $I^{(Q_2,\cdot,\cdot)}_{\cdot,\cdot}$	91
		I.4.5 Integral $I^{(Q_3,\cdot,\cdot,\cdot)}_{\cdot,\cdot}$	92
		I.4.6 Integral $I^{(Q_4,\cdot,\cdot)}_{\cdot,\cdot}$	94
	I.5	Conclusion	94
R			

Index

101

List of Figures

1.1	Posterior distribution over an unknown function as evidence is accumulated	2
1.2	Posterior distribution over an unknown function for different kernels	2
3.1	Generation of the outputs of in Model 5	13
3.2	Generative process of Model 5	16
3.3	One-dimensional samples from Model 5	17
3.4	Two-dimensional samples from Model 5	17
3.5	Interpolation between two kernels sampled from Model 5 \ldots .	18
3.6	Relationship of Model 5 and Model 6 to current literature	26
4.1	Organisation of multi-output models from current literature \ldots	34
5.1	Graphical model of Model 13	42
5.2	Generative process of Model 13 in the case that $N = 10 \dots \dots \dots \dots$	45
5.3	Observations from Model 13 in the cases that $N = 1$ and $N = 20$	46
E.1	Circulant approximation of a stationary kernel	62

List of Tables

4.1 Identification of multi-output models from current literature $\ . \ . \ . \ . \ . \ 33$

List of Models

1	Time-Variant Model	6
2	Time-Invariant Model	6
3	Multidimensional Time-Invariant Model	9
4	Generalised Gaussian Process Convolution Model (GGPCM) $\ . \ . \ . \ .$	11
5	Nonparametric Kernel Model (NKM)	11
6	Approximate Kernel Model (AKM)	22
7	Basis Function Model	29
8	Instantaneous Mixing Model (IMM)	32
9	Convolutional Mixing Model (CMM)	32
10	Gaussian Process Convolution Model [Tobar et al., 2015b]	37
11	Gaussian Process Convolution Model (Explicit Decay)	38
12	Deep Gaussian Process Convolution Model of Order N	39
13	Deep Kernel Model of Order N	40
14	Approximate Deep Kernel Model of Order N	41

Notation

General

We make use of Lagrange's notation for differentation.

 $(1,1),\ldots,(N,M)$ Shorthand for

 $(1,1),\ldots,(N,1),(1,2),\ldots,(N,2),\ldots,(1,M),\ldots,(N,M)$

i	Imaginary unit or, depending on the context, index
•*	Complex conjugate
·	Modulus or, depending on the context, determinant
$\mathbb{1}_X$	Indicator function
1	Shorthand for $\mathbb{1}_{\{0\}}$
δ	Dirac delta function
R	Reversal function; for any function $f_1: X \to Y$ the reversal function
	$R: Y^X \to Y^X$ yields the function $f_2 = R(f_1)$ such that
	$f_2(x) = f_1(-x)$ for all $x \in X$
\mathcal{O}	Bachmann-Landau notation for asymptotic behaviour

Probability Theory

p	Probability measure
$\stackrel{d}{=}$	Equality in distribution
\mathbb{E}	Expected value
$\mathcal{N}(\cdot;oldsymbol{\mu},oldsymbol{\Sigma})$	Multivariate Gaussian distribution with mean μ and covariance matrix Σ
$D_{KL}(\cdot \ \cdot)$	Kullback-Leibler divergence

Linear Algebra

Unbolded symbols x denote scalars, bolded lower-case symbols x vectors and bolded upper-case symbols X matrices or arrays—by array we mean the multi-dimensional generalisation of the two-dimensional matrix. Vectors are column vectors unless specified otherwise.

Ι	Identity matrix
$oldsymbol{I}_N$	$N \times N$ identity matrix
. <i>T</i>	Transpose
. ^H	Conjugate transpose
·	Determinant or, depending on the context, modulus
$\ \cdot\ _p$	<i>p</i> -norm
$\ \cdot\ _F$	Frobenius norm
\otimes	Kronecker product
$\operatorname{diag}(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_N)$	Block diagonal matrix; yields

$oxed{X}_1$	0	•••	0	0
0	$oldsymbol{X}_2$	•••	0	0
÷	÷	·	÷	÷
0	0	•••	$oldsymbol{X}_{N-1}$	0
0	0	•••	0	$oldsymbol{X}_N$

 $\operatorname{circ}(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_N)$ Block circulant matrix; yields

$$\begin{bmatrix} \boldsymbol{X}_1 & \boldsymbol{X}_N & \cdots & \boldsymbol{X}_3 & \boldsymbol{X}_2 \\ \boldsymbol{X}_2 & \boldsymbol{X}_1 & \cdots & \boldsymbol{X}_4 & \boldsymbol{X}_3 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \boldsymbol{X}_{N-1} & \boldsymbol{X}_{N-2} & \cdots & \boldsymbol{X}_1 & \boldsymbol{X}_N \\ \boldsymbol{X}_N & \boldsymbol{X}_{N-1} & \cdots & \boldsymbol{X}_2 & \boldsymbol{X}_1 \end{bmatrix}$$

vec

$\begin{bmatrix} X_1 \end{bmatrix}$	$oldsymbol{X}_2$		$oldsymbol{X}_{N-1}$	$oldsymbol{X}_N$
X_2	$oldsymbol{X}_1$	•••	$oldsymbol{X}_{N-2}$	$oldsymbol{X}_{N-1}$
:	:	·	÷	÷
$oldsymbol{X}_{N-1}$	$oldsymbol{X}_{N-2}$	•••	$oldsymbol{X}_1$	$oldsymbol{X}_2$
X_N	$oldsymbol{X}_{N-1}$	•••	$oldsymbol{X}_2$	$oldsymbol{X}_1$

	$\operatorname{vec} \boldsymbol{X} = \begin{bmatrix} X_{1,1} & \cdots & X_{N,M} \end{bmatrix}^T$
X_{i_1,\ldots,i_N}	Vector, matrix and array indexing; yields element (i_1, \ldots, i_N) of the array \boldsymbol{X}
$oldsymbol{X}_{i,:}$	Fixation of first dimension of a matrix and ranging over the second; yields vector containing the i 'th row of \boldsymbol{X}
$oldsymbol{X}_{:,i}$	Fixation of second dimension of a matrix and ranging over the first; yields vector containing the i 'th column of X
$X_{\underbrace{i_1,,i_n,:,,:}_{N ext{ indices}}}$	Fixation of an arbitrary subset of dimensions and ranging over the remaining; yields the array whose (i_{n+1}, \ldots, i_N) 'th element is given by X_{i_1,\ldots,i_N}
$f(oldsymbol{X})$	Element-wise function application; if \boldsymbol{X} is M -dimensional— $\boldsymbol{X} \in \mathbb{R}^{N_1 \times \cdots \times N_M}$ —but f takes a $(M-m)$ -dimensional object— $f : \mathbb{R}^{N_{m+1} \times \cdots \times N_M} \to C$ —then $f(\boldsymbol{X})$ is the m -dimensional array whose (i_1, \ldots, i_m) 'th element
$f(oldsymbol{X},oldsymbol{Y})$	is given by $f(\mathbf{X}_{i_1,\dots,i_m,i,\dots,i})$ Element-wise function application; if $\mathbf{X} \in \mathbb{R}^{N_1 \times \dots \times N_M}$ and $\mathbf{X} \in \mathbb{R}^{N_1 \times \dots \times N_K}$, but $f : \mathbb{R}^{N_{m+1} \times \dots \times N_M} \times \mathbb{R}^{N_{k+1} \times \dots \times N_K} \to C$, then $f(\mathbf{X}, \mathbf{Y})$ is the array whose $(i_1, \dots, i_m, j_1, \dots, j_k)$ 'th element is given by $f(\mathbf{X}_{i_1,\dots,i_m,i,\dots,i}, \mathbf{Y}_{i_1,\dots,i_k,i,\dots,i})$

Matrix vectorisation function; if $\boldsymbol{X} \in \mathbb{R}^{N \times M}$, then

Miscellaneous

Convolution; for matrix-valued functions \boldsymbol{F}_1 and \boldsymbol{F}_2 defined as

$$(\boldsymbol{F}_1 * \boldsymbol{F}_2)(\boldsymbol{x}) = \int \boldsymbol{F}_1(\boldsymbol{x} - \boldsymbol{y}) \boldsymbol{F}_2(\boldsymbol{y}) \,\mathrm{d}\boldsymbol{y}$$

${\cal F}_N$	$N \times N$ unitary discrete Fourier transform matrix
DFT	Discrete Fourier transform
$\mathcal{F}_{m{y}}\{f\}(m{x})$	Continuous Fourier transform of f ; defined as

$$\mathcal{F}_{\boldsymbol{y}}{f}(\boldsymbol{x}) = \int f(\boldsymbol{y}) \exp(-2\pi i \boldsymbol{x}^T \boldsymbol{y}) \,\mathrm{d}\boldsymbol{y}$$

 $\begin{array}{ll} \mathcal{F}_{f}(\boldsymbol{x}) & \text{Shorthand for } \mathcal{F}_{\boldsymbol{y}}\{f\}(\boldsymbol{x}) \\ \mathcal{S}_{\boldsymbol{y}}\{f,g\}(\boldsymbol{x}) & \text{Cross-spectral density between two wide-sense stationary processes } f \text{ and} \\ g; \text{ equal to} \end{array}$

$$\mathcal{S}_{\boldsymbol{y}}\{f,g\}(\boldsymbol{x}) = \mathcal{F}_{\boldsymbol{z}}\{\mathbb{E}[f(\boldsymbol{x}+\boldsymbol{z})g^{*}(\boldsymbol{x})]\}(\boldsymbol{y})$$

 $\begin{aligned} \mathcal{S}_{f,g}(\boldsymbol{y}) & \text{Shorthand for } \mathcal{S}_{\boldsymbol{y}}\{f,g\}(\boldsymbol{x}) \\ (C,\boldsymbol{A},\boldsymbol{b},c) & \text{See Appendix F.2; shorthand for} \end{aligned}$

$$C \exp\left(-\frac{1}{2}\boldsymbol{x}^{T}\boldsymbol{A}\boldsymbol{x} + \boldsymbol{x}^{T}\boldsymbol{b} + c\right)$$

 (C, \mathbf{A}) See Appendix F.3; shorthand for $(C, \mathbf{A} \otimes \mathbf{I}, \mathbf{0}, 0)$

Acronyms

AKM	Approximate Kernel Model (Model 6)
CGPM	Collaborative Gaussian processes model [Nguyen and Bonilla, 2014]
CMOGPM	Convolved multi-output Gaussian process model [Álvarez and Lawrence,
	2011]
CSMK	Cross-spectral mixture kernel [Ulrich et al., 2015]
GGPCM	Generalised Gaussian Process Convolution Model (Model 4)
GPCM	Gaussian Process Convolution Model [Tobar et al., 2015b]
GPRN	Gaussian process regression network [Wilson et al., 2012]

*

ICM	Intrinsic coregionalisation model [Goovaerts, 1997]
LCM	Linear coregionalisation model [Goovaerts, 1997]
LFM	Latent force model [Álvarez et al., 2009]
MTGPM	Multi-task Gaussian process [Bonilla et al., 2008]
NKM	Nonparametric Kernel Model (Model 5)
PSD	Power spectral density
SLFM	Semiparametric latent factor model [Teh and Seeger, $2005]$
SMK	Spectral mixture kernel [Wilson and Adams, 2013]

1 | Introduction

1.1 Motivation

Gaussian processes elegantly provide means to model an unknown function. They give rise to Bayesian regression models in which one maintains a posterior distribution over the unknown function as evidence is accumulated; Figure 1.1 illustrates this process. Gaussian processes have been succesfully applied in a wide variety of contexts—Rasmussen and Williams [2006] provide an excellent overview.

Gaussian processes are *nonparametric* models. As opposed to *parametric* models, there is no finite number of parameters that parametrises a Gaussian processes. Instead, the number of parameters grows with the amount of evidence that is accumulated. This property allows Gaussian processes to learn complex functions if plenty of evidence is available. Conversely, this property makes them robust against overfitting if only little evidence is available.

Their expressiveness and robustness however, come at a cost: Gaussian process models are often computationally expensive and one is forced to choose a *kernel*. The kernel of a Gaussian process reflects one's assumption on how the unknown function autocovaries.¹ The choice of the kernel is crucial; Figure 1.2 illustrates that the posterior distribution can wildly vary for different kernels. Unfortunately, determining which kernel to use is hard; the kernel thus poses a difficult design problem. We specifically refer to this problem as the *kernel design problem*.

A number of recent works address the kernel design problem: Duvenaud [2014] presents a way to search over a space of kernels through composition of existing kernels, and Wilson and Adams [2013] present a flexible kernel by modellings its power spectral

¹The autocovariance of an unknown function f specifies the covariance between any two function values $f(t_1)$ and $f(t_2)$.



Figure 1.1: Posterior distribution over an unknown function as evidence is accumulated. The evidence is represented by blue dots. The red line represents the mean of the posterior distribution and the gradient indicates the marginal variance up to two standard deviations.



Figure 1.2: Posterior distribution over an unknown function for different kernels. The evidence is represented by the blue dots and is the same as in Figure 1.1. The red line represents the mean of the posterior distribution and the gradient indicates the marginal variance up to two standard deviations.

1.2 Contribution of Thesis and Outline

density with a sum of Gaussians. As noted by Tobar et al. [2015b], these approaches do not completely resolve the kernel design problem: computational expenses restrict the space of kernels that can be searched, and the large number of parameters introduced by the sum of Gaussians reintroduces the problem of overfitting. Instead, Tobar et al. [2015b] note that the kernel constitutes just another unknown function and thus propose to model it by another Gaussian; the results are encouraging. They call their model the Gaussian Process Convolution Model (GPCM).

The GPCM has only been formulated for one-dimensional signals defined on one-dimensional input spaces. Extending this model to multi-dimensional signals defined on multi-dimensional input spaces would enable numerous applications in econometrics, geostatistics and signal processing, among others.

1.2 Contribution of Thesis and Outline

The main contribution of this thesis is the formulation of the Generalised Gaussian Process Convolution Model (GGPCM). The GGPCM provides a theoretical framework for nonparametric kernel models of multi-dimensional signals defined on multidimensional input spaces.

Chapter 2 develops a general-purpose model of dynamical signals. This model will suggest a generalisation of the GPCM: the GGPCM. Thus, in hindsight, Chapter 2 motivates why this particular generalisation is appropriate.

Chapter 3 develops the GGPCM. We use the GGPCM to address the kernel design problem in multi-output Gaussian processes on multidimensional input spaces. Furthermore, we develop a dual formulation of the cross-spectral mixture kernel [Ulrich et al., 2015] and we show how the GGPCM connects to existing work.

Chapter 4 examines the GGPCM's connection to existing multi-output Gaussian process models. We provide an overview of the current literature.

Chapter 5 develops the Deep Kernel Model, which is a new network structure for unsupervised learning. We sample from the Deep Kernel Model and investigate its properties.

2 | Modelling Dynamical Signals

2.1 Introduction

For many years scientists have used the concept of a system to study dynamic processes of diverse nature. A system is a mathematical abstraction that is used to describe properties of the process we intend to study. Dynamical systems describe processes that evolve in time; the system then usually describes a transformation $T : (\mathbb{R}^M)^{\mathbb{R}} \to (\mathbb{R}^N)^{\mathbb{R}}$ from an input signal \boldsymbol{x} to an output signal $\boldsymbol{f} = T(\boldsymbol{x})$ —the latter is also called the system response.

This chapter develops a general-purpose model of dynamical signals. In doing so we take a systems-modelling perspective by postulating that every signal is the response of a particular system. Thus, equivalently, we develop a model of system responses. We therefore utilise a general description of dynamical *systems* to construct the desired model of dynamical *signals*.

2.2 The Linear State-Space Model

A widely used dynamical system model is one that describes mechanical, electrical, hydraulic, and thermal systems, among others. These systems are all compositions of resistive, inductive, capactive and memristive elements that are connected through lossless transfer of energy; hence, they allow for a unified description. This description embodies the concept of *system state*, whose manifestation is a vector s that at any time determines future output given future input, thus rendering past input irrelevant. More precisely, the description is a so-called *state-space model* whose general *linear* form is

given by

$$\boldsymbol{s}'(t) = \boldsymbol{A}(t)\boldsymbol{s}(t) + \boldsymbol{B}(t)\boldsymbol{x}(t), \qquad (2.1)$$

$$\boldsymbol{f}(t) = \boldsymbol{C}(t)\boldsymbol{s}(t) + \boldsymbol{D}(t)\boldsymbol{x}(t). \tag{2.2}$$

Examples of systems that are succesfully described by linear state-space models are ubiquitous: linear state-space models describe electronics that are essential in everyday life; systems that control your car, airplanes and even rockets; and many phenomena in nature.

2.3 A General-Purpose Model of Dynamical Signals

Suppose that f is a dynamical signal. By the postulate that every signal is the response of some system, we can assume that some system generated f. Now, Section 2.2 suggests that we can safely assume this system to be of the form of Equations (2.1) and (2.2). In that case Appendix A.1 shows that f admits the following parametrisation:

Model 1 (Time-Variant Model).

$$\boldsymbol{f}(t) = \int_{\mathbb{R}} \boldsymbol{H}(t,\tau) \boldsymbol{x}(\tau) \,\mathrm{d}\boldsymbol{\tau}.$$

Observe that f is parametrised in terms of some matrix-valued function H—the *impulse* response¹—and the postulated input x.

In many applications of state-space models $\mathbf{A}(t)$, $\mathbf{B}(t)$, $\mathbf{C}(t)$ and $\mathbf{D}(t)$ are approximately constant.² In that case Appendix A.2 shows that $\mathbf{H}(t,\tau) = \mathbf{H}(t-\tau)$. Model 1 then reduces to the following model:

¹Letting $\boldsymbol{x}(t) = \delta(t - t_0)\boldsymbol{x}_0, t \in \mathbb{R}$ yields the response $\boldsymbol{f}(t) = \boldsymbol{H}(t, t_0), t \in \mathbb{R}$. That is, \boldsymbol{H} is the system response of an *impulse excitation*; in this sense \boldsymbol{H} is called the *impulse response*.

²For example, consider an electric circuit. In that case, approximating A(t), B(t), C(t) and D(t) as time invariant roughly corresponds to approximating resistances, capacitances and inductances as time invariant. This approximation is often reasonable.

2.4 Conclusion

Model 2 (Time-Invariant Model).

$$\boldsymbol{f}(t) = \int_{\mathbb{R}} \boldsymbol{H}(t-\tau)\boldsymbol{x}(\tau) \,\mathrm{d}\tau = (\boldsymbol{H} \ast \boldsymbol{x})(t).$$

Observe that Model 2 attains the form of a convolution.

2.4 Conclusion

We have developed a general-purpose model of dynamical signals: Model 2. Unfortunately, Model 2 leaves the impulse reponse H and postulated input x unspecified. This issue will be addressed in Chapter 3.

3 | The Generalised Gaussian Process Convolution Model

3.1 Introduction

Chapter 2 developed a general-purpose model of dynamical signals: the Time-Invariant Model (Model 2). By extending Model 2's input space to \mathbb{R}^{K} we derive the Multidimensional Time-Invariant Model:

Model 3 (Multidimensional Time-Invariant Model).

$$oldsymbol{f}(oldsymbol{t}) = \int_{\mathbb{R}^K} oldsymbol{H}(oldsymbol{t}-oldsymbol{ au}) x(oldsymbol{ au}) \, \mathrm{d}oldsymbol{ au} = (oldsymbol{H} stoldsymbol{x})(oldsymbol{t})$$

Model 3 is a deterministic model of multidimensional signals defined on multidimensional input spaces. This chapter shows that a stochastic version of Model 3 can be used to address the kernel design problem in multi-output Gaussian processes on multidimensional input spaces.

Recall that $\boldsymbol{f}: \mathbb{R}^K \to \mathbb{R}^N, \, \boldsymbol{H}: \mathbb{R}^K \to \mathbb{R}^{M \times N} \text{ and } \boldsymbol{x}: \mathbb{R}^K \to \mathbb{R}^M.$

3.2 Gaussian Process Regression

A Gaussian process defines a distribution over functions. More precisely, a stochastic process $f(t), t \in \mathbb{R}$ is Gaussian if and only if for every $\mathbf{t} = \begin{bmatrix} t_1 & \cdots & t_T \end{bmatrix}^T$ it holds that $f(\mathbf{t}) = \begin{bmatrix} f(t_1) & \cdots & f(t_T) \end{bmatrix}^T$ is multivariate Gaussian distributed. We denote $f \sim \mathcal{GP}(m_f, \mathcal{K}_f)$ where m_f and \mathcal{K}_f denote respectively the mean function and kernel of f. This implies that $f(\mathbf{t}) \sim \mathcal{N}[m_f(\mathbf{t}), \mathcal{K}_f(\mathbf{t}, \mathbf{t})]$. Furthermore, we usually let $m_f = 0$ without

loss of generality [Rasmussen and Williams, 2006] and denote $K_{f_1,f_2} = \mathcal{K}_f(t_1, t_2)$ where $f_1 = f(t_1)$ and $f_2 = f(t_2)$.

Suppose that we observe data \boldsymbol{y} at \boldsymbol{t}_{y} generated by the process $y(t) = f(t) + \varepsilon(t), t \in \mathbb{R}$ where $\varepsilon(t) \sim \mathcal{N}(0, \sigma^{2}), t \in \mathbb{R}$. For unobserved function values \boldsymbol{f}^{*} at $\boldsymbol{t}_{f^{*}}$ it then holds that

$$egin{bmatrix} egin{aligned} egi$$

where $\boldsymbol{\theta}$ denotes σ and the parameters of \mathcal{K}_f and $m_f - \boldsymbol{\theta}$ are also called the *hyperparameters*. We usually omit explicit conditioning \boldsymbol{t}_y and \boldsymbol{t}_{f^*} . According to Appendix B.1 we can perform prediction via

$$f^* | y, \theta \sim \mathcal{N}[K_{f^*,y}(K_{y,y} + \sigma^2 I)^{-1}y, K_{f^*,f^*} - K_{f^*,y}(K_{y,y} + \sigma^2 I)^{-1}K_{y,f^*}].$$

This is the Gaussian process regression framework.

The values of the hyperparameters should be chosen such that the model best fits the true data-generating process. In this sense, as argued by Rasmussen and Williams [2006], their values can be determined by maximising the *marginal likelihood*—also called the *evidence*—given by

$$p(\boldsymbol{y} \mid \boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \propto \underbrace{|\boldsymbol{\Sigma}|^{-1/2}}_{\text{complexity penalty}} \underbrace{\exp\left[-\frac{1}{2}\|\boldsymbol{\Sigma}^{-1/2}(\boldsymbol{y}-\boldsymbol{\mu})\|_{2}^{2}\right]}_{\text{compatability with data}}$$

for some $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Observe that the marginal likelihood is a trade-off between the model's compatibility with the data and $|\boldsymbol{\Sigma}|^{-1/2}$ being small; that is, the evidence is large if the model is compatible with the data, but only if the model is not too *complex* in the sense that it is compatible with *any* data—if $\boldsymbol{\Sigma}^{-1/2}$ is such that $\|\boldsymbol{\Sigma}^{-1/2}(\boldsymbol{y}-\boldsymbol{\mu})\|_2^2$ is small for any \boldsymbol{y} , then $|\boldsymbol{\Sigma}^{-1/2}| = |\boldsymbol{\Sigma}|^{-1/2}$ is small and thereby the evidence must be small. Thus maximising the evidence tends to explain the data in a *simple* way. This tendency is commonly recognised as a manifestation of *Occam's razor* [MacKay, 2002].

Finally, Appendix C shows how Gaussian processes can be used to define distributions over multivariate matrix-valued functions.

3.3 The Generalised Gaussian Process Convolution Model

Model 3 leaves the input signal \boldsymbol{x} and impulse response \boldsymbol{H} unspecified. Since they constitute unknown functions, it is sensible to model them by two Gaussian processes; we then obtain the Generalised Gaussian Process Convolution Model (GGPCM):

Model 4 (Generalised Gaussian Process Convolution Model (GGPCM)). Draw

 $egin{aligned} m{H} &\sim \mathcal{GP}(m{0},m{\mathcal{K}}_{m{H}}), \ m{x} &\sim \mathcal{GP}(m{0},m{\mathcal{K}}_{x}), \ m{arepsilon} & & arepsilon & areps$

independently for some kernel \mathcal{K}_x , some kernel \mathcal{K}_H and some diagonal matrix Λ . Then observations are generated by $\mathbf{y} = \mathbf{f} + \boldsymbol{\varepsilon} = \mathbf{H} * \mathbf{x} + \boldsymbol{\varepsilon}$.

In the remainder of this chapter we study the case that $\mathcal{K}_{x}(t_{1}, t_{2}) = \delta(t_{1} - t_{2})I$ —we let x be white noise. The general case will be studied in Chapter 4.

Observe that $\boldsymbol{f} \mid \boldsymbol{H}$ is a linear combination of Gaussian processes. Hence $\boldsymbol{f} \mid \boldsymbol{H}$ is another Gaussian process, which thus can be identified by its mean function and kernel:

$$\mathbb{E}[\boldsymbol{f}(\boldsymbol{t}) \mid \boldsymbol{H}] = \int_{\mathbb{R}^{K}} \boldsymbol{H}(\boldsymbol{t} - \boldsymbol{\tau}) \mathbb{E}[\boldsymbol{x}(\boldsymbol{\tau})] \, \mathrm{d}\boldsymbol{\tau} = \boldsymbol{0},$$

$$\mathcal{K}_{\boldsymbol{f} \mid \boldsymbol{H}}(\boldsymbol{t}_{1}, \boldsymbol{t}_{2}) = \mathbb{E}[\boldsymbol{f}(\boldsymbol{t}_{1}) \boldsymbol{f}^{T}(\boldsymbol{t}_{2}) \mid \boldsymbol{H}]$$

$$= \int_{\mathbb{R}^{K}} \int_{\mathbb{R}^{K}} \boldsymbol{H}(\boldsymbol{t}_{1} - \boldsymbol{\tau}_{1}) \underbrace{\mathbb{E}[\boldsymbol{x}(\boldsymbol{\tau}_{1}) \boldsymbol{x}^{T}(\boldsymbol{\tau}_{2})]}_{\delta(\boldsymbol{\tau}_{1} - \boldsymbol{\tau}_{2})I} \boldsymbol{H}^{T}(\boldsymbol{t}_{2} - \boldsymbol{\tau}_{2}) \, \mathrm{d}\boldsymbol{\tau}_{1} \, \mathrm{d}\boldsymbol{\tau}_{2}$$

$$= \int_{\mathbb{R}^{K}} \boldsymbol{H}[\boldsymbol{\tau} - (\boldsymbol{t}_{2} - \boldsymbol{t}_{1})] \boldsymbol{H}^{T}(\boldsymbol{\tau}) \, \mathrm{d}\boldsymbol{\tau}$$

$$= [R(\boldsymbol{H}) * \boldsymbol{H}^{T}](\boldsymbol{t}_{2} - \boldsymbol{t}_{1}) \qquad (3.1)$$

$$= \mathcal{K}_{\boldsymbol{f} \mid \boldsymbol{H}}(\boldsymbol{t}_{1} - \boldsymbol{t}_{2})$$

where R is the reversal function. We have established the following equivalent model:

Model 5 (Nonparametric Kernel Model (NKM)). Draw

$$oldsymbol{H} \sim \mathcal{GP}(oldsymbol{0}, oldsymbol{\mathcal{K}_H}), \ oldsymbol{arepsilon} \sim \mathcal{GP}[oldsymbol{0}, \delta(oldsymbol{t}_1 - oldsymbol{t}_2) oldsymbol{\Lambda}^2]$$

independently for some kernel \mathcal{K}_{H} and some diagonal matrix Λ . Afterwards draw

$$\boldsymbol{f} \mid \boldsymbol{H} \sim \mathcal{GP} \{ \boldsymbol{0}, [R(\boldsymbol{H}) * \boldsymbol{H}^T] (\boldsymbol{t}_2 - \boldsymbol{t}_1) \}$$

Then observations are generated by $\boldsymbol{y} = \boldsymbol{f} + \boldsymbol{\varepsilon}$.

This equivalent formulation reveals that white noise excitation in Model 4 yields ordinary Gaussian process regression in which the kernel is modelled nonparametrically. We cannot model the kernel of \boldsymbol{f} simply with a sample from a Gaussian process—a kernel has to be positive semidefinite, which in general a sample from a Gaussian process is not. Thus Model 5 shows a way to nonparametrically model a positive-semidefinite matrixvalued function. We verify that $\mathcal{K}_{\boldsymbol{f}|\boldsymbol{H}} = R(\boldsymbol{H}) * \boldsymbol{H}^T$ is indeed positive semidefinite: Let $\boldsymbol{g} : \mathbb{R}^K \to \mathbb{R}^N$ be square integrable. Then

$$\begin{split} \int_{\mathbb{R}^{2K}} \boldsymbol{g}^{T}(\boldsymbol{t}_{1}) \boldsymbol{\mathcal{K}}(\boldsymbol{t}_{1},\boldsymbol{t}_{2}) \boldsymbol{g}(\boldsymbol{t}_{2}) \, \mathrm{d}\boldsymbol{t}_{1} \, \mathrm{d}\boldsymbol{t}_{2} &= \int_{\mathbb{R}^{3K}} \boldsymbol{g}^{T}(\boldsymbol{t}_{1}) \boldsymbol{H}[\boldsymbol{\tau} - (\boldsymbol{t}_{2} - \boldsymbol{t}_{1})] \boldsymbol{H}^{T}(\boldsymbol{\tau}) \boldsymbol{g}(\boldsymbol{t}_{2}) \, \mathrm{d}\boldsymbol{t}_{1} \, \mathrm{d}\boldsymbol{t}_{2} \, \mathrm{d}\boldsymbol{\tau} \\ &= \int_{\mathbb{R}^{3K}} \boldsymbol{g}^{T}(\boldsymbol{t}_{1}) \boldsymbol{H}(\boldsymbol{t}_{1} - \boldsymbol{\tau}) \, \mathrm{d}\boldsymbol{t}_{1} \boldsymbol{H}^{T}(\boldsymbol{t}_{2} - \boldsymbol{\tau}) \boldsymbol{g}(\boldsymbol{t}_{2}) \, \mathrm{d}\boldsymbol{t}_{2} \, \mathrm{d}\boldsymbol{\tau} \\ &= \int_{\mathbb{R}^{K}} \left\| \int_{\mathbb{R}^{K}} \boldsymbol{H}^{T}(\boldsymbol{t} - \boldsymbol{\tau}) \boldsymbol{g}(\boldsymbol{t}) \, \mathrm{d}\boldsymbol{t} \right\|^{2} \, \mathrm{d}\boldsymbol{\tau} \\ &\geq 0. \end{split}$$

Alternatively, Model 5 can be interpreted in the frequency domain. It is clear that $f \mid H$


Figure 3.1: Generation of the outputs of in Model 5 and relationships between their cross-spectral densities

is wide-sense stationary; hence,

$$S_{f_{i},f_{j}\mid\boldsymbol{H}}(\boldsymbol{\xi}) = \mathcal{F}_{\boldsymbol{\tau}}\{\mathcal{K}_{f_{i},f_{j}\mid\boldsymbol{H}}(\boldsymbol{t}+\boldsymbol{\tau},\boldsymbol{t})\}(\boldsymbol{\xi})$$

$$= \sum_{m=1}^{M} \int_{\mathbb{R}^{2K}} H_{i,m}(\boldsymbol{\tau}'+\boldsymbol{\tau})H_{j,m}(\boldsymbol{\tau}')\exp(-2\pi i\boldsymbol{\xi}^{T}\boldsymbol{\tau})\,\mathrm{d}\boldsymbol{\tau}'\,\mathrm{d}\boldsymbol{\tau}$$

$$= \sum_{m=1}^{M} \int_{\mathbb{R}^{2K}} H_{i,m}(\boldsymbol{\tau})\exp(-2\pi i\boldsymbol{\xi}^{T}\boldsymbol{\tau})\,\mathrm{d}\boldsymbol{\tau}[H_{j,m}(\boldsymbol{\tau}')\exp(-2\pi i\boldsymbol{\xi}^{T}\boldsymbol{\tau}')]^{*}\,\mathrm{d}\boldsymbol{\tau}'$$

$$= \sum_{m=1}^{M} \mathcal{F}_{H_{i,m}}(\boldsymbol{\xi})\mathcal{F}_{H_{j,m}}^{*}(\boldsymbol{\xi}) \qquad (3.2)$$

$$= \mathcal{F}_{H_{i,:}}^{T}(\boldsymbol{\xi})\mathcal{F}_{H_{j,:}}^{*}(\boldsymbol{\xi}). \qquad (3.3)$$

Since each $H_{i,j}$ is modelled nonparametrically, each $\mathcal{F}_{H_{i,j}}$ is modelled nonparametrically. Model 5 therefore models the cross-spectral densities \mathcal{S}_{f_i,f_j} nonparametrically by complex inner products between the stacked spectra $\mathcal{F}_{H_{i,:}}$ and $\mathcal{F}_{H_{j,:}}$ associated to f_i and f_j . This provides insight in how correlations between the outputs are induced; Figure 3.1 depicts in more detail how Model 5 generates outputs and how their cross-spectral densities relate.

Finally, Model 5 can be interpreted as a continuous-time moving average model. Equivalently, consider white noise excitation in Model 4. Then discretising the input space results in a model of the form

$$oldsymbol{f}_t = \sum_{ au \in \mathbb{R}^K} oldsymbol{H}_{t- au} oldsymbol{x}_{ au}$$

where \boldsymbol{x} is white noise. This is exactly a multi-output moving average model on a multidimensional input space. Hence Model 5 is a continuous-time Bayesian multi-output moving average model on a multidimensional input space.

3.3.1 Interpretation and Choice of the Kernel in the Nonparametric Kernel Model

In this section we resolve an apparent modelling issue concerning Model 5: it is not clear how to interpret and choose \mathcal{K}_H and thereby M—recall that $\mathcal{K}_H : \mathbb{R}^K \times \mathbb{R}^K \to \mathbb{R}^{NM \times NM}$.

First, we examine M. We already determined that $f \mid H$ is a Gaussian process with zero mean function. Thus Model 5 attains its full expressive power if it can generate any $\mathcal{K}_{f \mid H}$, or equivalently any $\mathcal{F}{\mathcal{K}_{f \mid H}} = S_{f \mid H}$. To begin with,

$$oldsymbol{S}_{oldsymbol{f}\midoldsymbol{H}}^{H}=\mathcal{F}_{ au}^{*}\{\mathbb{E}^{T}[oldsymbol{f}(oldsymbol{t}+ au)oldsymbol{f}^{T}(oldsymbol{t})]\}=\mathcal{F}_{ au}\{\mathbb{E}[oldsymbol{f}(oldsymbol{t})oldsymbol{f}^{T}(oldsymbol{t}- au)]\}=oldsymbol{S}_{oldsymbol{f}\midoldsymbol{H}}$$

shows that $S_{f|H}$ is Hermitian. Therefore, by the Complex Spectral Theorem, generating any $S_{f|H}$ is equivalent to generating any spectral decomposition. Now, utilising Equation (3.2) to express

$$oldsymbol{S_{f|H}}(oldsymbol{\xi}) = \sum_{m=1}^M \mathcal{F}_{oldsymbol{H}_{:,m}}(oldsymbol{\xi}) \mathcal{F}_{oldsymbol{H}_{:,m}}^H(oldsymbol{\xi})$$

shows that any spectral decomposition can be generated only if $M \ge N$. Thus M = N is the minimal M for which Model 5 attains its full expressive power. Conversely, Model 5 has limited expressive power if M < N. More precisely, Equation (3.3) then shows that at each frequency $\boldsymbol{\xi}$ at most M outputs can be independent; f_i and f_j are independent if and only if $S_{f_i,f_j|H}(\boldsymbol{\xi}) = \mathcal{F}_{H_{j,i}}^H, \mathcal{F}_{H_{i,i}} = 0$, and at most M vectors $\mathcal{F}_{H_{i,i}}$ can simultaneously be orthogonal as they are elements of \mathbb{C}^M . We let M = N in further development of Model 5.

Second, we examine \mathcal{K}_H . It turns out that we must be careful in choosing \mathcal{K}_H . Sup-

pose the simplest case; that is, let the components of H be independent and share an exponentiated-quadratic kernel:

$$\mathcal{K}_{\boldsymbol{H}}(\boldsymbol{t}_1, \boldsymbol{t}_2) = \sigma_h^2 \exp(-\gamma \|\boldsymbol{t}_1 - \boldsymbol{t}_2\|^2) \boldsymbol{I}.$$

Having already observed that Model 5 induces correlations between its outputs justifies modelling the components of H independently. We call \mathcal{K}_H the diagonal multi-output exponentiated-quadratic kernel. Now, by equivalently considering white noise excitation in Model 4,

$$\begin{split} \mathbb{E}[f_i^2(\boldsymbol{t})] &= \int_{\mathbb{R}^{2K}} \mathbb{E}[\boldsymbol{H}_{i,:}^T(\boldsymbol{t} - \boldsymbol{\tau}_1) \boldsymbol{x}(\boldsymbol{\tau}_1) \boldsymbol{x}(\boldsymbol{\tau}_2)^T \boldsymbol{H}_{i,:}(\boldsymbol{t} - \boldsymbol{\tau}_2)] \, \mathrm{d}\boldsymbol{\tau}_1 \, \mathrm{d}\boldsymbol{\tau}_2 \\ &= \int_{\mathbb{R}^{2K}} \mathrm{tr} \, \mathbb{E}[\underbrace{\boldsymbol{H}_{i,:}(\boldsymbol{t} - \boldsymbol{\tau}_1) \boldsymbol{H}_{i,:}^T(\boldsymbol{t} - \boldsymbol{\tau}_2)}_{\mathcal{K}_{\boldsymbol{H}_{i,:}}[(\boldsymbol{t} - \boldsymbol{\tau}_1) - (\boldsymbol{t} - \boldsymbol{\tau}_2)]} \underbrace{\boldsymbol{x}(\boldsymbol{\tau}_1) \boldsymbol{x}^T(\boldsymbol{\tau}_2)}_{\delta(\boldsymbol{\tau}_1 - \boldsymbol{\tau}_2)\boldsymbol{I}}] \, \mathrm{d}\boldsymbol{\tau}_1 \, \mathrm{d}\boldsymbol{\tau}_2 \\ &= \int_{\mathbb{R}^K} \mathrm{tr}(\sigma_h^2 \boldsymbol{I}) \, \mathrm{d}\boldsymbol{\tau} \\ &= \infty. \end{split}$$

Intuitively, this means that the model prior assigns each f_i infinitely large error bars, or equivalently assigns each f_i infinite signal power. This is unlike real-world signals, which have finite power. We therefore let the components of \boldsymbol{H} have a decaying exponentiatedquadratic kernel [Tobar et al., 2015b] instead:

$$\mathcal{K}_{\boldsymbol{H}}(\boldsymbol{t}_1, \boldsymbol{t}_2) = \sigma_h^2 \exp(-\alpha \|\boldsymbol{t}_1\|^2 - \alpha \|\boldsymbol{t}_2\|^2 - \gamma \|\boldsymbol{t}_1 - \boldsymbol{t}_2\|^2) \boldsymbol{I}.$$

We call \mathcal{K}_{H} the diagonal multi-output decaying exponentiated-quadratic kernel. Now, using the notation and identities from Appendix F with composite vector $\begin{bmatrix} t^{T} & \tau^{T} \end{bmatrix}^{T}$,

$$\mathbb{E}[f_i^2(\boldsymbol{t})] = \int_{\mathbb{R}^K} \operatorname{tr}[\sigma_h^2 \exp(-2\alpha \|\boldsymbol{t} - \boldsymbol{\tau}\|^2) \boldsymbol{I}] \, \mathrm{d}\boldsymbol{\tau}$$
$$= I_{\boldsymbol{\tau}}[(K\sigma_h^2, 4 \begin{bmatrix} \alpha & -\alpha \\ -\alpha & \alpha \end{bmatrix})]$$
$$= K\sigma_h^2 \frac{\pi^{K/2}}{(2\alpha)^{K/2}}$$

and so our modelling issue is resolved. We decide to use the diagonal multi-output decaying exponentiated-quadratic kernel in further development of Model 5.



Figure 3.2: Generative process of Model 5 in the case that N = M = K = 1. Illustrates the interpretation that Model 5 corresponds ordinary Gaussian process regression in which the kernel is modelled nonparametrically.

Note that the components of $\mathcal{K}_{H}(t_{1}, t_{2})$ are small for $||t_{1}||$ or $||t_{2}||$ much larger than $1/\sqrt{\alpha}$. Therefore for such t_{1} or t_{2} the posteriors of the components of H are likely to be near zero; that is, $1/\sqrt{\alpha}$ is the effective extent of the components of H. Similarly, the components of $\mathcal{K}_{H}(t_{1}, t_{2})$ are small for $||t_{1} - t_{2}||$ much larger than $1/\sqrt{\gamma}$; thus $1/\sqrt{\gamma}$ determines the effective length scale on which the components of H vary.

We can let the effective extent and effective length scale be different for different dimensions of the input space. We then obtain

$$\mathcal{K}_{H}(t_{1}, t_{2}) = \sigma_{h}^{2} \exp[-t_{1}^{T} A t_{1} - t_{2}^{T} A t_{2} - (t_{1} - t_{2})^{T} \Gamma(t_{1} - t_{2})] I$$

for positive-semidefinite matrices A and Γ . Since A and Γ reflect the relevance of each input variable, we call \mathcal{K}_H the diagonal multi-output decaying automatic relevance determination kernel [MacKay, 2002].

3.3.2 Illustrative Samples of the Nonparametric Kernel Model

Figure 3.2 shows the generative process of Model 5 in the simplest case that M = N = 1and K = 1; that is, a sample is defined on a one-dimensional space—for example, time—and constitutes a single output. This figure illustrates the interpretation that Model 5 corresponds ordinary Gaussian process regression in which the kernel is modelled nonparametrically: First, $H_{1,1}$ is generated. Then, a kernel for $f_1 | H_{1,1}$ is constructed by computing $R(H_{1,1}) * H_{1,1}$. Finally, $f_1 | H_{1,1}$ is generated by drawing from a Gaussian process with $f_1 | H_{1,1}$'s generated kernel.



Figure 3.3: Samples from Model 5 in the case that N = M = 2 and K = 1. Shows the sampled kernel and four observations.



Figure 3.4: Samples from Model 5 in the case that N = M = 2 and K = 2. Shows the sampled kernel and two observations.



Figure 3.5: Interpolation between two kernels sampled from Model 5 in the case that N = M = 2 and K = 1. Shows the way an observation changes as its kernel is interpolated between the two kernels.

Furthermore, Figure 3.3 shows samples from Model 5 in the case that M = N = 2 and K = 1—a sample now constitutes two outputs. First, observe that the sampled kernel is indeed symmetric: $\mathcal{K}_{H_{1,1}}$ and $\mathcal{K}_{H_{2,2}}$ are symmetric and $\mathcal{K}_{H_{1,2}} = R(\mathcal{K}_{H_{2,1}})$. Second, observe that $\mathcal{K}_{H_{1,2}}$ and $\mathcal{K}_{H_{1,2}}$ attain reasonably peaked maxima at a nonzero lag. Hence we expect the outputs of a sample to be positively correlated at a nonzero lag; indeed, Figure 3.3 shows that the red lines essentially lead the blue lines. Thus Model 5 shows an certain ability to model multi-output time series. In the same way Model 5 is also able to model multi-output signals defined on multidimensional spaces; Figure 3.4 illustrates that the case K = 2 generates correlated surfaces.

Finally, Figure 3.5 illustrates that different samples from Model 5 yield different correlation structures between the outputs; in the case that M = N = 2 and K = 1, Figure 3.5 shows the way an observation changes as its kernel is interpolated between two kernels sampled from Model 5. Observe that kernel $\mathcal{K}_{H}^{(1)}$ corresponds to positively correlated outputs, whereas kernel $\mathcal{K}_{H}^{(2)}$ corresponds to negatively corresponds outputs; indeed, as the kernel of the observation is interpolated, the outputs change from being positively correlated to being negatively correlated.

Appendix D was utilised in generating Figures 3.2 to 3.5.

3.3.3 Expressivity of White Noise Excitation

Thus far is not clear how Model 5's expressivity compares to that of Model 4. In this section we shed some light on this matter.

Model 4 shows that f is generated by filtering x with H. This means that frequencies in x are attenuated according to the spectrum of H. In other words, f consists only of frequencies that are also present in x. Thus, if x's spectrum has effectively limited support, then f is limited to only those frequencies. This makes choosing a white noise kernel for x a great choice; a white noise process has a constant power spectrum and therefore contains all frequencies, meaning that f can also contain all frequencies.

More concretely, we show that an instance $(\mathcal{K}_{x^{(1)}}^{(1)}, \mathcal{K}_{H^{(1)}}^{(1)}, \Lambda^{(1)})$ of Model 4 can be expressed as an instance $(\mathcal{K}_{H^{(2)}}^{(2)}, \Lambda^{(2)})$ of Model 5 if $\mathcal{K}_{x^{(1)}}^{(1)}(t_1, t_2) = [R(\mathbf{R}) * \mathbf{R}^T](t_2 - t_1)$ for some \mathbf{R} —we call \mathbf{R} the root of $\mathcal{K}_{x^{(1)}}^{(1)}$. This is equivalent to showing that $(\mathcal{K}_{x^{(1)}}^{(1)}, \mathcal{K}_{H^{(1)}}^{(1)}, \Lambda^{(1)})$ can be expressed as another instance $(\mathcal{K}_{x^{(2)}}^{(2)}, \mathcal{K}_{H^{(2)}}^{(2)}, \Lambda^{(2)})$ of Model 4 where $\mathbf{x}^{(2)}$ is white noise.

Observe that $\mathbf{R} * \mathbf{x}^{(2)}$ and $\mathbf{H}^{(1)} * \mathbf{R}$ are linear combinations of Gaussian processes. Hence $\mathbf{R} * \mathbf{x}^{(2)}$ and $\mathbf{H}^{(1)} * \mathbf{R}$ are also Gaussian processes, which thus can be identified by their mean functions and kernels. It is readily verified that $\mathbf{R} * \mathbf{x}^{(2)}$ and $\mathbf{H}^{(1)} * \mathbf{R}$ have zero mean function. The kernel of $\mathbf{R} * \mathbf{x}^{(2)}$ is derived in a similar fashion as Equation (3.1), yielding $\mathcal{K}_{\mathbf{R}*\mathbf{x}^{(2)}}(\mathbf{t}_1, \mathbf{t}_2) = [\mathbf{R}(\mathbf{R}) * \mathbf{R}^T](\mathbf{t}_2 - \mathbf{t}_1)$. Thus $\mathbf{x}^{(1)} \stackrel{d}{=} \mathbf{R} * \mathbf{x}^{(2)}$. Now let $\mathcal{K}_{\mathbf{H}^{(2)}} = \mathcal{K}_{\mathbf{H}^{(1)}*\mathbf{R}}$ and $\mathbf{\Lambda}^{(2)} = \mathbf{\Lambda}^{(1)}$. Then also $\mathbf{H}^{(1)} * \mathbf{R} \stackrel{d}{=} \mathbf{H}^{(2)}$ and $\boldsymbol{\varepsilon}^{(1)} \stackrel{d}{=} \boldsymbol{\varepsilon}^{(2)}$. Consequently, by associativity of the convolution operator,

$$f^{(1)} = H^{(1)} * x^{(1)} + \varepsilon^{(1)}$$

= $H^{(1)} * (R * x^{(2)}) + \varepsilon^{(2)}$
= $(H^{(1)} * R) * x^{(2)} + \varepsilon^{(2)}$
= $H^{(2)} * x^{(2)} + \varepsilon^{(2)}$
= $f^{(2)}$

so that $\boldsymbol{f}^{(1)} \stackrel{d}{=} \boldsymbol{f}^{(2)}$.

We have established that an instance of Model 4 can be expressed of an instance of Model 5 if \boldsymbol{x} 's kernel has a root and is thereby stationary. Appendix G shows that any kernel of exponentiated quadratic form has a root, that the diagonal multi-output exponentiated-quadratic kernel has a root and that the white noise kernel is its own root. This sheds some light on the class of kernels of \boldsymbol{x} for which Model 4 and Model 5 have equal expressivity.

We can more generally identify which kernels have roots in the case that M = N = 1. Suppose that \mathcal{K}_x is stationary and Schwartz—that is, rapidly decreasing. Then, by Bochner's Theorem, its Fourier transform $\mathcal{F}_{\mathcal{K}_x}$ exists and is real. Therefore $\mathcal{F}_{\mathcal{K}_x}^{1/2}$ is well defined. As \mathcal{K}_x is Schwartz, $\mathcal{F}_{\mathcal{K}_x}$ is also Schwartz and thereby absolutely integrable. Thus, by

$$\int_{\mathbb{R}^K} |\mathcal{F}_{\mathcal{K}_x}(\boldsymbol{\xi})| \, \mathrm{d}\boldsymbol{\xi} = \int_{\mathbb{R}^K} |\mathcal{F}_{\mathcal{K}_x}^{1/2}(\boldsymbol{\xi})|^2 \, \mathrm{d}\boldsymbol{\xi} < \infty,$$

3.4 The Approximate Kernel Model

 $\mathcal{F}_{\mathcal{K}_x}^{1/2}$ is square integrable and so $\mathcal{F}^{-1}\{\mathcal{F}_{\mathcal{K}_x}^{1/2}\}$ exists. Let $r = \mathcal{F}^{-1}\{\mathcal{F}_{\mathcal{K}_x}^{1/2}\}$. Then

$$\begin{split} [R(r)*r](\boldsymbol{\tau}) &= \int_{\mathbb{R}^{3K}} \mathcal{F}_{\mathcal{K}_x}^{1/2}(\boldsymbol{\xi}_1) \exp[2\pi i \boldsymbol{\xi}_1^T(\boldsymbol{\tau}'-\boldsymbol{\tau})] \mathcal{F}_{\mathcal{K}_x}^{1/2}(\boldsymbol{\xi}_2) \exp(2\pi i \boldsymbol{\xi}_2^T \boldsymbol{\tau}') \, \mathrm{d} \boldsymbol{\xi}_1 \, \mathrm{d} \boldsymbol{\xi}_2 \, \mathrm{d} \boldsymbol{\tau}' \\ &= \int_{\mathbb{R}^{3K}} \mathcal{F}_{\mathcal{K}_x}^{1/2}(\boldsymbol{\xi}_1) \mathcal{F}_{\mathcal{K}_x}^{1/2}(\boldsymbol{\xi}_2) \exp(2\pi i \boldsymbol{\xi}_2^T \boldsymbol{\tau}) \underbrace{\exp[2\pi i (\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2)^T \boldsymbol{\tau}']}_{\delta(-\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2)} \, \mathrm{d} \boldsymbol{\xi}_1 \, \mathrm{d} \boldsymbol{\xi}_2 \, \mathrm{d} \boldsymbol{\tau}' \\ &= \int_{\mathbb{R}^K} \mathcal{F}_{\mathcal{K}_x}(\boldsymbol{\xi}) \exp(2\pi i \boldsymbol{\xi}^T \boldsymbol{\tau}) \, \mathrm{d} \boldsymbol{\xi} \\ &= \mathcal{K}_x(\boldsymbol{\tau}) \end{split}$$

and so \mathcal{K}_x has a root.

In summary, white noise excitation in Model 4 yields great expressivity because the spectrum of white noise has infinite support. Concretely, an instance of Model 4 can be expressed as an instance of Model 5 if \boldsymbol{x} 's kernel has a root. In the case that M = N = 1, a kernel has a root if it is stationary and rapidly decreasing.

3.4 The Approximate Kernel Model

Model 5 is troublesome from a computational perspective. Namely, $\mathcal{K}_{f|H} = R(H) * H^T$ is a convolution between two matrix-valued stochastic processes and it is not clear how such a quantity can be computed. We therefore develop a model that is approximate to Model 5, but whose computation is tractable.

Recall that \boldsymbol{H} is smooth and that its support is effectively limited due to use of the decaying exponentiated-quadratic kernel. In that case, as demonstrated by Minka [2000], we can numerically approximate the integral $R(\boldsymbol{H}) * \boldsymbol{H}^T$ by its expectation conditioned on some finite number of observations $\boldsymbol{H}(\boldsymbol{T})$ where $\boldsymbol{T} \in \mathbb{R}^{T \times K}$ —we call these observations *inducing points*. Symbolically,

$$\mathcal{K}_{f|H} = R(H) * H^T \approx \mathbb{E}[R(H) * H^T | H(T)] = \tilde{\mathcal{K}}_{f|H(T)}$$

Note that $\tilde{\mathcal{K}}_{f|H(T)}$ numerically approximates $\mathcal{K}_{f|H}$ and at the same time maintains knowledge of its own uncertainty; in this sense $\tilde{\mathcal{K}}_{f|H(T)}$ is called a *Bayesian numerical approximation* of $\mathcal{K}_{f|H}$ [Tobar et al., 2015b]. Appendix H shows that the analytical form of $\tilde{\mathcal{K}}_{f \mid H(T)}$ is given by

$$\tilde{\mathcal{K}}_{f_i,f_j \mid \boldsymbol{H}(\boldsymbol{T})} = \sum_{k=1}^{M} \left\{ \underbrace{\mathbb{1}(i-j)I^{(1,H_{i,k})}}_{\text{prior}} - \underbrace{\mathbb{1}(i-j)\operatorname{tr}(\boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1}\boldsymbol{I}^{(2,H_{i,k},H_{i,k})})}_{\text{conditioning on inducing points}} + \underbrace{H_{i,k}^{T}(\boldsymbol{T})\boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1}\boldsymbol{I}^{(2,H_{i,k},H_{j,k})}\boldsymbol{K}_{H_{j,k}(\boldsymbol{T})}^{-1}H_{j,k}(\boldsymbol{T})}_{\text{learned through inducing points}} \right\}.$$

We clearly distinguish the prior covariance structure, the covariance structure due to conditioning on the inducing points and the covariance structure learned through the inducing points.

We have derived the following approximation of Model 5:

Model 6 (Approximate Kernel Model (AKM)). Draw

independently for some diagonal kernel \mathcal{K}_H and some diagonal matrix Λ . Afterwards let $\tilde{\mathcal{K}}_{f|H(T)}$ be such that

$$\mathcal{K}_{f_{i},f_{j}|\boldsymbol{H}(\boldsymbol{T})} = \sum_{k=1}^{M} \Big\{ \mathbb{1}(i-j) [I^{(1,H_{i,k})} - \operatorname{tr}(\boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1} \boldsymbol{I}^{(2,H_{i,k},H_{i,k})})] \\ + H_{i,k}^{T}(\boldsymbol{T}) \boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1} \boldsymbol{I}^{(2,H_{i,k},H_{j,k})} \boldsymbol{K}_{H_{j,k}(\boldsymbol{T})}^{-1} H_{j,k}(\boldsymbol{T}) \Big\}.$$

Finally draw

$$\boldsymbol{f} \mid \boldsymbol{H}(\boldsymbol{T}) \sim \mathcal{GP}(\boldsymbol{0}, \tilde{\mathcal{K}}_{\boldsymbol{f} \mid \boldsymbol{H}(\boldsymbol{T})}).$$

Then observations are generated by $y = f + \varepsilon$.

Observe that Model 6 is ordinary Gaussian process regression with a specific kernel and a prior over the hyperparameters. Thus Model 6 can take advantage of existing literature on ordinary Gaussian process regression.

Furthermore, the fact that Model 6 approximates Model 5 implies that the prior over the hyperparameters specified by Model 6 is one that is sensible to use. This is an advantage

of Model 6 over most other multi-output kernels; priors over hyperparameters are usually chosen without principled motivation.

3.4.1 Interpretation of the Kernel Approximation

We investigate the kernel approximation in the case that M = N = 1. Denote $f = f_1$, $h = H_{1,1}$ and $h(\mathbf{T}) = \mathbf{h}$. Let $\mathcal{K}_{\tilde{h}|\mathbf{A}}$ be the posterior kernel of the sparse approximation of h according to [Titsias, 2009] where \mathbf{h} are the approximation's inducing points and \mathbf{A} is their posterior covariance matrix. We compare $\tilde{\mathcal{K}}_{f|\mathbf{h}}$ to $\mathcal{K}_{\tilde{h}|\mathbf{A}}$:

$$\begin{split} \tilde{\mathcal{K}}_{f|h}(t_1, t_2) &= I^{(1,h)}(t_1, t_2) - \operatorname{tr}[\boldsymbol{K}_h^{-1} I^{(2,h,h)}(t_1, t_2)] \\ &+ \operatorname{tr}[\boldsymbol{K}_h^{-1} I^{(2,h,h)}(t_1, t_2) \boldsymbol{K}_h^{-1} h h^T], \\ \mathcal{K}_{\tilde{h}|\boldsymbol{A}}(t_1, t_2) &= \mathcal{K}_h(t_1, t_2) - \operatorname{tr}[\boldsymbol{K}_h^{-1} \mathcal{K}_h(\boldsymbol{T}, t_2) \mathcal{K}_h(t_1, \boldsymbol{T})] \\ &+ \operatorname{tr}[\boldsymbol{K}_h^{-1} \mathcal{K}_h(\boldsymbol{T}, t_2) \mathcal{K}_h(t_1, \boldsymbol{T}) \boldsymbol{K}_h^{-1} \boldsymbol{A}] \end{split}$$

Thus, by Equations (H.1) and (H.2), we establish that

$$\tilde{\mathcal{K}}_{f|h}(\boldsymbol{t}_1, \boldsymbol{t}_2) = \int_{\mathbb{R}^K} \mathcal{K}_{\tilde{h}|\boldsymbol{A}=\boldsymbol{h}\boldsymbol{h}^T}(\boldsymbol{t}_1 - \boldsymbol{\tau}, \boldsymbol{t}_2 - \boldsymbol{\tau}) \,\mathrm{d}\boldsymbol{\tau}.$$

This means that the approximate integration performed by $\tilde{\mathcal{K}}_{f|h}$ corresponds to first sparsely approximating h and then exactly integrating its posterior kernel.

3.4.2 The Case of the Diagonal Multi-Output Decaying Exponentiated-Quadratic Kernel

Let \mathcal{K}_{H} be a diagonal multi-output decaying exponentiated-quadratic kernel. It then turns out that $I^{(1,H_{i,k})}$, $I^{(2,H_{i,k},H_{j,k})}$ and their Fourier transforms admit simple forms. We now compute these simple forms, which in Section 3.5 will reveal Model 6's connection to existing work. To begin with, Appendices I.4.3 and I.4.5 show that

$$\begin{split} I^{(1,H_{i,k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2}) &= \underbrace{\sigma_{h}^{2} \frac{\pi^{K/2}}{(2\alpha)^{K/2}}}_{C^{(1)}} \exp\left[-\frac{1}{2} \underbrace{(\alpha+2\gamma)}_{l^{-(1)}} \|\boldsymbol{t}_{1}-\boldsymbol{t}_{2}\|^{2}\right], \\ &= C^{(1)} \exp\left(-\frac{1}{2}l^{-(1)} \|\boldsymbol{t}_{1}-\boldsymbol{t}_{2}\|^{2}\right) \\ &= I^{(1)}(\boldsymbol{t}_{1}-\boldsymbol{t}_{2}), \\ I^{(2,H_{i,k},H_{j,k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2}) &= \sigma_{h}^{4} \frac{\pi^{K/2}}{(2\alpha+2\gamma)^{K/2}} \exp\left[-\frac{1}{2}\frac{(\alpha+\gamma)^{2}-\gamma^{2}}{\alpha+\gamma} (\|\boldsymbol{T}_{m,:}-\boldsymbol{T}_{n,:}\|^{2}\right) \\ & \underbrace{-2\|\boldsymbol{T}_{m,:}+\boldsymbol{T}_{n,:}\|^{2}\right] \\ & \underbrace{\exp\left[-\frac{1}{2}\underbrace{(\alpha+\gamma)}_{l^{-(2)}} \|(\boldsymbol{t}_{1}-\boldsymbol{t}_{2}) - \underbrace{\frac{\gamma}{\alpha+\gamma}(\boldsymbol{T}_{m,:}-\boldsymbol{T}_{n,:})}_{\mu_{m,n}^{(2)}}\right]^{2}\right] \\ &= C^{(2)}_{m,n} \exp\left[-\frac{1}{2}l^{-(2)}\|(\boldsymbol{t}_{1}-\boldsymbol{t}_{2}) - \mu_{m,n}^{(2)}\|^{2}\right] \\ &= I^{(2)}_{m,n}(\boldsymbol{t}_{1}-\boldsymbol{t}_{2}). \end{split}$$

Therefore

$$\mathcal{F}_{t_1-t_2}\{I^{(1)}\}(\boldsymbol{f}) = C^{(1)}\sqrt{2\pi l^{(1)}}\exp(-2\pi^2 l^{(1)}\|\boldsymbol{f}\|^2), \qquad (3.4)$$

$$\mathcal{F}_{t_1-t_2}\{I_{m,n}^{(2)}\}(\boldsymbol{f}) = C_{m,n}^{(2)}\sqrt{2\pi l^{(2)}}\exp(-2\pi^2 l^{(2)}\|\boldsymbol{f}\|^2 + 2\pi i\boldsymbol{\mu}_{m,n}^T\boldsymbol{f}).$$
(3.5)

Now, consider the case that i = j. First, all $\mathcal{K}_{H_{i,j}}$ are equal and so all $\mathbf{K}_{H_{i,j}(T)} = \mathbf{K}$; thus the kernel approximation simplifies to

$$\mathcal{K}_{f_{i},f_{i} \mid \boldsymbol{H}(\boldsymbol{T})} = M[I^{(1)} - \operatorname{tr}(\boldsymbol{K}^{-1}\boldsymbol{I}^{(2)})] + \sum_{k=1}^{M} H_{i,k}^{T}(\boldsymbol{T})\boldsymbol{K}^{-1}\boldsymbol{I}^{(2)}\boldsymbol{K}^{-1}H_{i,k}(\boldsymbol{T}), \qquad (3.6)$$

which shows that we can equivalently let $\mathbf{I}^{(2)} \leftarrow (\mathbf{I}^{(2)} + \mathbf{I}^{(2)T})/2$. Second, note that $C_{m,n}^{(2)} = C_{n,m}^{(2)}$ and $-\boldsymbol{\mu}_{n,m}^{(2)} = \boldsymbol{\mu}_{m,n}^{(2)}$; hence $\mathcal{F}\{I_{m,n}^{(2)}\} = \mathcal{F}^*\{I_{n,m}^{(2)}\}$. Therefore, if i = j, then

3.5 Related Work

equivalently

$$\mathcal{F}_{t_1-t_2}\{I_{m,n}^{(2)}\}(\boldsymbol{f}) = \frac{1}{2}[\mathcal{F}_{t_1-t_2}\{I_{m,n}^{(2)}\}(\boldsymbol{f}) + \mathcal{F}_{t_1-t_2}\{I_{n,m}^{(2)}\}(\boldsymbol{f})]$$

$$= \frac{1}{2}[\mathcal{F}_{t_1-t_2}\{I_{m,n}^{(2)}\}(\boldsymbol{f}) + \mathcal{F}_{t_1-t_2}^*\{I_{m,n}^{(2)}\}(\boldsymbol{f})]$$

$$= C_{m,n}^{(2)}\sqrt{2\pi l^{(2)}}\exp(-2\pi^2 l^{(2)}\|\boldsymbol{f}\|^2)\cos(2\pi i\boldsymbol{\mu}_{m,n}^T\boldsymbol{f}).$$
(3.7)

3.5 Related Work

Model 5 and its approximation Model 6 are closely related to recent work.

First, Model 5 is a generalisation of the Gaussian Process Convolution Model (GPCM) [Tobar et al., 2015b]; the case N = M = K = 1 recovers their model exactly.

Second, Equations (3.4), (3.6) and (3.7) show that the power spectral densities of the kernel in Model 6 take the form of spectral mixture kernels (SMKs) [Wilson and Adams, 2013]. In other words, the diagonal entries of the kernel in Model 6 form *Fourier pairs* with spectral mixture kernels. Therefore the kernel in Model 6 is a *dual* multi-output generalisation of the spectral mixture kernel. Furthermore, recall that Model 6 approximates Model 5 more accurately as the number of inducing points increases. This reveals the kernel in Model 5 as a multi-output generalisation of the spectral mixture kernel with an infinite number of components.

Third, Equations (3.4) to (3.6) show that the kernel in Model 6 is a dual formulation of the cross-spectral mixture kernel (CSMK) [Ulrich et al., 2015]. Similar to the spectral mixture kernel, this reveals the kernel in Model 5 as a cross-spectral mixture kernel with an infinite number of components. Furthermore, the connection between Model 6 and the cross-spectral mixture kernel suggests that results in [Ulrich et al., 2015] might carry over to Model 6; indeed, inspired by [Ulrich et al., 2015], Appendix E presents an approximation of stationary multi-output kernel matrices that can be used to efficiently perform inference in Model 6 if H(T) are instead treated as hyperparameters.

Figure 3.6 summarises how Model 5 and Model 6 relate to current literature on flexible kernel models.



Figure 3.6: Relationship of Model 5 and Model 6 to current literature on flexible kernel models. Single-headed arrows indicate generalisation and double-headed arrows indicate duality. "Model limit" refers to taking the model's number of components to infinity.

3.6 Inference in the Nonparametric Kernel Model

Performing inference in Model 5 directly is troublesome, because \boldsymbol{H} parametrises the kernel of Model 5's likelihood $p(\boldsymbol{y} | \boldsymbol{H})$ which thus is intractable. Instead, equivalently consider white noise excitation in Model 4. Model 4's likelihood $p(\boldsymbol{y} | \boldsymbol{H}, \boldsymbol{x})$ is Gaussian with mean function $\boldsymbol{f} = \boldsymbol{H} * \boldsymbol{x}$. As we show shortly, this likelihood is manageable.

We perform inference in Model 5 by inferring H and x in Model 4 where x is white noise. This immediately raises a concern: f = H * x requires knowledge of the whole processes H and x, but x cannot wholly be learned because $x(t_1)$ and $x(t_2)$ are independent for $t_1 \neq t_2$; that is, we would have to learn x(t) for every t separately, but a computer can only store x(t) for finitely many t.

To resolve this issue, consider another instance of Model 4 whose excitation is denoted by \tilde{x} . Let \tilde{x} have a diagonal multi-output exponentiated-quadratic kernel. In that case \tilde{x} is smooth and can therefore be learned by conditioning on $\tilde{x}(t)$ for finitely many t[Titsias, 2009]. Thus we should be able to perform inference in this new instance. Now, Appendix G shows that \tilde{x} 's kernel has a root R. Therefore, by Section 3.3.3, Model 5 can be expressed in terms of this new instance where $\tilde{x} = R * x$. Crucially, the fact that we can learn this new instance then implies that we can learn Model 5 by learning $\boldsymbol{R} * \boldsymbol{x}$ instead of \boldsymbol{x} .

Learning $\mathbf{R} * \mathbf{x}$ instead of \mathbf{x} has a number of equivalent interpretations. First, $\mathbf{R} * \mathbf{x}$ represents a smoothed version of \mathbf{x} where each point is a linear combination of the points of \mathbf{x} . Thus learning one point on $\mathbf{R} * \mathbf{x}$ induces knowledge about the whole process \mathbf{x} . Second, \mathbf{R} acts as a filter, which means that $\mathbf{R} * \mathbf{x}$ represents a band-limited version of \mathbf{x} . Hence we restrict ourselves to only learning frequencies in the band of \mathbf{R} ; this is an easier problem than learning all frequencies.

The latter interpretation raises another concern; if \boldsymbol{H} admits frequencies not in the band of \boldsymbol{R} , then those frequencies cannot be learned. Thus \boldsymbol{R} should be chosen carefully; in the case that M = N = K = 1, Tobar et al. [2015a] consider how \boldsymbol{R} can be designed to promote training performance. We decide to simply use

$$\boldsymbol{R}(\boldsymbol{t}) = \exp(-\omega \|\boldsymbol{t}\|^2) \boldsymbol{I}.$$

This choice should enable us to learn \boldsymbol{x} ; namely, Appendix G shows that this \boldsymbol{R} is a root of a diagonal multi-output exponentiated-quadratic kernel, and we previously determined that a process with a diagonal multi-output exponentiated-quadratic kernel can be learned.

Concretely, we perform approximate inference in Model 5 via a structured mean field approximation in which we sparsely approximate \boldsymbol{H} and $\boldsymbol{R} * \boldsymbol{x}$ through inducing points [Titsias, 2009]. Let \boldsymbol{Y} denote Y observations. Consider the family \mathcal{Q} of distributions of the form

$$q(\boldsymbol{H}, \boldsymbol{x}, \underbrace{\boldsymbol{u}_{H_{1,1}}, \ldots, \boldsymbol{u}_{H_{N,M}}}_{\boldsymbol{U}_{\boldsymbol{H}}}, \underbrace{\boldsymbol{u}_{\tilde{x}_{1}}, \ldots, \boldsymbol{u}_{\tilde{x}_{M}}}_{\boldsymbol{U}_{\tilde{\boldsymbol{x}}}}) = p(\boldsymbol{H}, \boldsymbol{x} \mid \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \prod_{i=1, j=1}^{N, M} q(\boldsymbol{u}_{H_{i,j}}) \prod_{j=1}^{M} q(\boldsymbol{u}_{\tilde{x}_{j}})$$

where $\boldsymbol{u}_{H_{i,j}} = H_{i,j}(\boldsymbol{T}_{H_{i,j}})$ and $\boldsymbol{u}_{\tilde{x}_j} = \tilde{x}_j(\boldsymbol{T}_{\tilde{x}_j})$ are inducing points for respectively the processes $H_{i,j}$ and $\tilde{x}_j = \exp(-\omega \|\boldsymbol{t}\|^2) * x_j$. Let

$$egin{aligned} q(oldsymbol{u}_{H_{i,j}}) &= \mathcal{N}(oldsymbol{u}_{H_{i,j}};oldsymbol{\mu}_{H_{i,j}},oldsymbol{\Sigma}_{H_{i,j}}), \ q(oldsymbol{u}_{ ilde{x}_j}) &= \mathcal{N}(oldsymbol{u}_{ ilde{x}_j};oldsymbol{\mu}_{ ilde{x}_j},oldsymbol{\Sigma}_{ ilde{x}_j}). \end{aligned}$$

We then perform approximate inference via

$$p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}} | \boldsymbol{Y})$$

$$= \underset{q}{\operatorname{argmin}} D_{KL}[q(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \| p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}} | \boldsymbol{Y})]$$

$$\approx \underset{q \in \mathcal{Q}}{\operatorname{argmin}} D_{KL}[q(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \| p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}} | \boldsymbol{Y})]$$

$$= \underset{q \in \mathcal{Q}}{\operatorname{argmax}} \{ \log p(\boldsymbol{Y}) - D_{KL}[q(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \| p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}} | \boldsymbol{Y})] \}$$

$$= \underset{q \in \mathcal{Q}}{\operatorname{argmax}} \mathcal{F}(q).$$

Since $D_{KL}(\cdot \| \cdot) \geq 0$, $\mathcal{F}(q)$ —also called the *variational free energy*, or *free energy* in short—is a lower bound on the marginal likelihood; hence performing inference is equivalent to maximising this lower bound. Now, recall that H's diagonal multi-output decaying exponentiated-quadratic kernel implies independence between different $H_{i,j}$'s; as a result,

$$p(\boldsymbol{H}, \boldsymbol{x} \mid \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) = p(\boldsymbol{H} \mid \boldsymbol{U}_{\boldsymbol{H}}) p(\boldsymbol{x} \mid \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) = \prod_{i=1,j=1}^{N,M} p(H_{i,j} \mid \boldsymbol{u}_{H_{i,j}}) \prod_{j=1}^{M} p(x_j \mid \boldsymbol{u}_{\tilde{x}_j}).$$

Then

$$\begin{aligned} \mathcal{F}(q) &= \mathbb{E}_{q}[\log p(\mathbf{Y})] + \mathbb{E}_{q}\left[\log \frac{p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}} \mid \mathbf{Y})}{q(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}})}\right] \\ &= \mathbb{E}_{q}\left[\log \frac{p(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}}, \mathbf{Y})}{q(\boldsymbol{H}, \boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{H}}, \boldsymbol{U}_{\tilde{\boldsymbol{x}}})}\right] \\ &= \mathbb{E}_{q}\left[\log \frac{p(\mathbf{Y} \mid \boldsymbol{H}, \boldsymbol{x}) p(\boldsymbol{H} \mid \boldsymbol{U}_{\boldsymbol{H}}) p(\boldsymbol{x} \mid \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \prod_{i=1,j=1}^{N,M} p(\boldsymbol{u}_{H_{i,j}}) \prod_{j=1}^{M} p(\boldsymbol{u}_{\tilde{x}_{j}})}{p(\boldsymbol{H} \mid \boldsymbol{U}_{\boldsymbol{H}}) p(\boldsymbol{x} \mid \boldsymbol{U}_{\tilde{\boldsymbol{x}}}) \prod_{i=1,j=1}^{N,M} q(\boldsymbol{u}_{H_{i,j}}) \prod_{j=1}^{M} q(\boldsymbol{u}_{\tilde{x}_{j}})}\right] \\ &= \mathbb{E}_{q}[\log p(\mathbf{Y} \mid \boldsymbol{H}, \boldsymbol{x})] - \sum_{i=1,j=1}^{N,M} D_{KL}[q(\boldsymbol{u}_{H_{i,j}}) \mid p(\boldsymbol{u}_{H_{i,j}})] - \sum_{j=1}^{M} D_{KL}[q(\boldsymbol{u}_{\tilde{x}_{j}}) \mid p(\boldsymbol{u}_{\tilde{x}_{j}})] \\ &= -\frac{1}{2}\log[(2\pi)^{N}|\mathbf{A}|^{2}] - \underbrace{\frac{1}{2}\sum_{i=1}^{Y} \mathbb{E}_{q}\{\|\mathbf{A}^{-1}[\mathbf{Y}_{i,:} - \boldsymbol{y}(\mathbf{T}_{i,:})]\|^{2}\}}_{\text{data reconstruction cost}} \\ &- \underbrace{\sum_{i=1,j=1}^{N,M} D_{KL}[q(\boldsymbol{u}_{H_{i,j}}) \mid p(\boldsymbol{u}_{H_{i,j}})] - \underbrace{\sum_{j=1}^{M} D_{KL}[q(\boldsymbol{u}_{\tilde{x}_{j}}) \mid p(\boldsymbol{u}_{\tilde{x}_{j}})]}_{\text{divergence from model}}.\end{aligned}$$

We observe that the variational free energy can interpreted as negative a cost of reconstructing the data regularised by a divergence from the model prior; thus, maximisation of the free energy seeks an explanation of the data that is compatible with the model prior.

Appendix I derives an analytical expression for the variational free energy. Appendix D discusses some useful techniques concerning implementation of the free energy.

Furthermore, we consider the asymptotic time complexity of computing the free energy. Denote the number of inducing points for the processes $H_{i,j}$ and \tilde{x}_j by $T_{H_{i,j}}$ and $T_{\tilde{x}_j}$ respectively. Let $T_H = \max_{i,j} T_{H_{i,j}}$ and $T_{\tilde{x}} = \max_j T_{\tilde{x}_j}$. Then Appendix I.5 shows that the time complexity of computing the free energy is given by

$$\mathcal{O}[NMT_{H}^{3} + MT_{\tilde{x}}^{3} + YNM(T_{H}^{2}T_{\tilde{x}} + KT_{H}^{2} + KT_{H}T_{\tilde{x}} + KT_{\tilde{x}}^{2} + T_{H}T_{\tilde{x}}^{2}) + YNM^{2}].$$

Recall that \boldsymbol{H} has effectively limited support due to use of the decaying exponentiatedquadratic kernel; thus likely $T_{\tilde{x}} \gg T_H$. Also, likely $T_{\tilde{x}} \gg M$. In that case the asymptotic time complexity of computing the free energy simplies to $\mathcal{O}[MT_{\tilde{x}}^3 + YNM(K + T_H)T_{\tilde{x}}^2]$.

Finally, we discuss initialisation of the parameters of the variational free energy—that is, initialisation of $\mu_{H_{i,j}}$, $\Sigma_{H_{i,j}}$, $\mu_{\tilde{x}_j}$ and $\Sigma_{\tilde{x}_j}$. Optimisation of the free energy is usually performed via gradient-based methods. This means that the optimiser is likely to become trapped in a local minimum if the parameters of the free energy are initialised far from their optimal values. Hence proper initialisation of the parameters is of crucial importance.

We propose the following initialisation: To begin with, fit the predictive mean to the data using weighted least squares [Tobar and Turner, 2016]. Appendix I.2 shows that this is equivalent to performing inference in the following model:

Model 7 (Basis Function Model). Let f be such that

$$f_i(t) = \sum_{j=1}^{M} \boldsymbol{\mu}_{H_{i,j}}^T \boldsymbol{K}_{\boldsymbol{u}_{H_{i,j}}}^{-1} \boldsymbol{I}^{(L,H_{i,j},x_j)}(t) \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_j}}^{-1} \boldsymbol{\mu}_{\tilde{x}_j}$$

for some kernel \mathcal{K}_x , some kernel \mathcal{K}_H and some vectors $\boldsymbol{\mu}_{H_{i,j}}$ and $\boldsymbol{\mu}_{\tilde{x}_j}$. Afterwards draw $\boldsymbol{\varepsilon} \sim \mathcal{GP}(\mathbf{0}, \Lambda^2)$ for some diagonal matrix Λ . Then observations are generated by $\boldsymbol{y} = \boldsymbol{f} + \boldsymbol{\varepsilon}$.

Learning $\mu_{H_{i,j}}$ and $\mu_{\tilde{x}_j}$ through Model 7 yields arguably sensible initial values. Finally, we initialise $\Sigma_{H_{i,j}}$ and $\Sigma_{\tilde{x}_j}$ to fractions of their priors $K_{u_{H_{i,j}}}$ and $K_{u_{\tilde{x}_j}}$.

Observe that Model 7 is a multi-output basis function model that is structured in its parametrisation. Model 7 is a generalisation of the basis function model presented by Tobar and Turner [2016]; the case N = M = K = 1 recovers their model exactly.

3.7 Conclusion

The Generalised Gaussian Process Convolution Model enabled us to formulate Model 5. Model 5 addresses the kernel design problem in multi-output Gaussian processes on multidimensional input spaces by parameterising the kernel with another Gaussian process. By modelling the kernel nonparametrically we avoid choosing a kernel of parametric form; instead, we infer its form from the data.

Finally, Model 5 is intimately connected to other models (Section 3.5). Most notably, Model 5 enabled us to formulate Model 6, which forms a dual formulation of the cross-spectral mixture kernel and reveals Model 5 as a cross-spectral mixture kernel with an infinite number of components.

3.8 Discussion

If $T_{\tilde{x}} \gg T_H$ and $T_{\tilde{x}} \gg M$, then Section 3.6 shows that the variational free energy of Model 5 can be computed in $\mathcal{O}[MT_{\tilde{x}}^3 + YNM(K + T_H)T_{\tilde{x}}^2]$ time. This complexity scales linearly in Y. However, in many settings—for example, time series—the size of the data is proportional to the size of the space the data occupies. Now, the inducing points for \boldsymbol{x} must cover the space the data occupies. Therefore, in such settings, $T_{\tilde{x}}$ effectively scales with Y, which means that cost of computing the free energy effectively scales with Y^3 .

4 | Multi-Task Learning

4.1 Introduction

Many inference problems involve dealing with multiple signals. These problems are best solved not only by learning signals individually, but also by simultaneously exploiting dependencies between signals. The signals can be interpreted as the output of multiple *tasks*; a problem involving multiple signals is therefore commonly referred to as a *multi-task* problem. We will use the terms "multi-output" and "multi-task" interchangeably.

A prominent example of a multi-task problem is the prediction of concentration levels of pollutants in geostatistics [Álvarez and Lawrence, 2011; Álvarez et al., 2009; Wilson et al., 2012]. Pollutants are often expensive to sample, which hinders accurately predicting their concentration levels. Fortunately, pollutants often strongly correlate with other substances, which can be cheap to sample. The multi-task predictor aims to improve predicting concentration levels of scarcely sampled pollutants by exploiting their correlation with densely sampled cheap substances.

Numerous multi-task models have been developed, many of which are Gaussian process models. Now, the Generalised Gaussian Process Convolution Model (Model 4) is also a multi-output model. This chapter provides an overview of multi-output models from the geostatistics and machine learning literature and shows how Model 4 fits in.

4.2 Mixing Models

A model for a multi-output signal f can be constructed by assuming that at any t the output f(t) is explained by the value x(t) of some latent process x with independent components. If the relationship between f(t) and x(t) is linear, then we arrive at the

Instantaneous Mixing Model (IMM):

Model 8 (Instantaneous Mixing Model (IMM)). Let H be a matrix. Draw

$$oldsymbol{x} \sim \mathcal{GP}(oldsymbol{0}, oldsymbol{\mathcal{K}}_{oldsymbol{x}}), \ oldsymbol{arepsilon} \sim \mathcal{GP}(oldsymbol{0}, oldsymbol{\Lambda}^2)$$

independently for some diagonal kernel \mathcal{K}_x and some diagonal matrix Λ . Then observations are generated by $\mathbf{y} = \mathbf{f} + \boldsymbol{\varepsilon} = \mathbf{H}\mathbf{x} + \boldsymbol{\varepsilon}$.

Observe that Model 8 corresponds to ordinary Gaussian process regression with a kernel of the form $\mathcal{K}_f = H \mathcal{K}_x H^T$ for some matrix H and some diagonal kernel \mathcal{K}_x . Alternatively, Model 8 can be interpreted as a factor analysis model at all t.

Model 8 assumes that f(t) depends only on x(t), meaning that f(t) is independent of x(t') for $t' \neq t$. This assumption limits the expressivity of f significantly; for example, f cannot be a smoothed version of x. We therefore relax Model 8 by instead letting f(t) depend on x(t) for all t. If their relationship is again linear, then we arrive at the Convolutional Mixing Model (CMM):

Model 9 (Convolutional Mixing Model (CMM)). Let H be a matrix-valued function. Draw

$$oldsymbol{x} \sim \mathcal{GP}(oldsymbol{0}, oldsymbol{\mathcal{K}}_{oldsymbol{x}}), \ oldsymbol{arepsilon} \sim \mathcal{GP}(oldsymbol{0}, oldsymbol{\Lambda}^2)$$

independently for some diagonal kernel \mathcal{K}_x and some diagonal matrix Λ . Then observations are generated by $\mathbf{y} = \mathbf{f} + \boldsymbol{\varepsilon} = \mathbf{H} * \mathbf{x} + \boldsymbol{\varepsilon}$.

Observe that Model 9 is a generalisation of Model 8, since letting $H = \delta H'$ for some constant matrix H' yields that

$$\boldsymbol{f}(\boldsymbol{t}) = (\delta \boldsymbol{H}' * \boldsymbol{x})(\boldsymbol{t}) = \int_{\mathbb{R}^K} \delta(\boldsymbol{t} - \boldsymbol{\tau}) \boldsymbol{H}' \boldsymbol{x}(\boldsymbol{\tau}) \, \mathrm{d} \boldsymbol{\tau} = \boldsymbol{H}' \boldsymbol{x}(\boldsymbol{t}).$$

Model	Form of \boldsymbol{H}	Form of \mathcal{K}_x	Mixing type
ICM [Goovaerts, 1997]	H	$(\sum_{q=1}^Q k_x^{(q)}) oldsymbol{I}$	Instantaneous
LCM [Goovaerts, 1997]	$egin{bmatrix} m{H}_1 & \cdots & m{H}_Q \end{bmatrix}$	$\operatorname{diag}(k_x^{(1)}\boldsymbol{I},\ldots,k_x^{(Q)}\boldsymbol{I})$	Instantaneous
SLFM [Teh and Seeger, 2005]	H	\mathcal{K}_x	Instantaneous
MTGPM [Bonilla et al., 2008]	H	$k_x I$	Instantaneous
LFM [Álvarez et al., 2009]	Green's function	\mathcal{K}_x	Convolutional
CMOGPM [Álvarez and Lawrence, 2011]	$egin{bmatrix} m{H}_1 & \cdots & m{H}_Q \end{bmatrix}$	$\operatorname{diag}(k_x^{(1)}\boldsymbol{I},\ldots,k_x^{(Q)}\boldsymbol{I})$	Convolutional
CGPM [Nguyen and Bonilla, 2014]	$\begin{bmatrix} H & I \end{bmatrix}$	\mathcal{K}_x	Instantaneous

Table 4.1: Identification of multi-output models from the geostatistics and machine learning literature as specialisations of Models 8 and 9

4.3 The Mixing Model Hierarchy

Many multi-output Gaussian process models from the geostatistics and machine learning literature can be identified as specialisations of Models 8 and 9; Table 4.1 shows the identification of the intrinsic coregionalisation model (ICM) [Goovaerts, 1997], the linear coregionalisation model (LCM) [Goovaerts, 1997], the semiparametric latent factor model (SLFM) [Teh and Seeger, 2005], the multi-task Gaussian process model (MT-GPM) [Bonilla et al., 2008], the latent force model (LFM) [Álvarez et al., 2009], the convolved multi-output Gaussian process model (CMOGPM) [Álvarez and Lawrence, 2011] and the collaborative Gaussian processes model (CGPM) [Nguyen and Bonilla, 2014].

Model 9 is also closely related to Model 4. Specifically, Model 4 is a generalisation of Model 9 where H is modelled stochastically. This shows that Models 4, 8 and 9 form a hierarchy—the *mixing model hierarchy*—in which Model 4 generalises Model 9 and Model 9 generalises Model 8.

Figure 4.1 organises the main models presented in this thesis, the cross-spectral mixture kernel (CSMK) [Ulrich et al., 2015], and the models in Table 4.1 according to the mixing model hierarchy. This yields an overview of many multi-output Gaussian process models from the geostatistics and machine learning literature that emphasises their distinctive modelling assumptions. An immediate result is that the cross-spectral mixture kernel can now be connected to other multi-output models.

An important model that does not fit in the mixing model hierarchy is the Gaussian



Figure 4.1: Organisation of multi-output models from the geostatistics and machine learning literature according to the mixing model hierarchy. Single-header arrows indicate generalisation and double-headed arrows indicate equivalence. "Model limit" refers to taking the model's number of components to infinity.

process regression network (GPRN) [Wilson et al., 2012].

4.4 Conclusion

We have presented the mixing model model hierarchy, which organises many multioutput Gaussian process from the geostatistics and machine learning literature according to their distinctive modelling assumptions.

Model 4 fits in as a generalisation of Model 9 and Model 8. This shows that many multi-output Gaussian process models from the current literature can be derived from Model 4, with the exception being the Gaussian process regression network.

5 | The Deep Kernel Model

5.1 Introduction

Chapter 3 developed the idea that the kernel design problem can be addressed by modelling the kernel nonparametrically. In retrospect, this approach actually replaces the kernel design problem with another one: we now have to design the *kernel of the kernel*. Fortunately, we know already how to solve this kernel design problem: we can parametrise the *kernel of the kernel* with another Gaussian process. But then we have to design the *kernel of the kernel of the kernel*...

This chapter investigates a model that not only models the kernel nonparametrically, but also the kernel of the kernel, the kernel of the kernel of the kernel, and further "deeper" kernels. We confine the presentation to one-dimensional signals f on one-dimensional input spaces.

5.2 The Deep Kernel Model

Consider the Nonparametric Kernel Model (Model 5). In the case of one-dimensional signals on one-dimensional input spaces Model 5 recovers the Gaussian Process Convolution Model (GPCM) [Tobar et al., 2015b] (Section 3.5):

Model 10 (Gaussian Process Convolution Model [Tobar et al., 2015b]). Draw

$$h \sim \mathcal{GP}(0, \mathcal{K}_h),$$

$$x \sim \mathcal{GP}[0, \delta(t_1 - t_2)],$$

$$\varepsilon \sim \mathcal{GP}[0, \sigma^2 \delta(t_1 - t_2)]$$

independently for some kernel \mathcal{K}_h and some constant σ . Then observations are generated

 $by \ y = f + \varepsilon = h * x + \varepsilon.$

Recall that Model 10 corresponds to ordinary Gaussian process regression in which the kernel is modelled nonparametrically (Section 3.3). We immediately recognise that the kernel design problem is replaced with another one: we now have to choose \mathcal{K}_h —the kernel of the kernel. We investigate whether \mathcal{K}_h can be modelled nonparametrically.

Modelling \mathcal{K}_h nonparametrically immediately presents a problem. In Section 3.3.1 we ensured that f has finite variance by letting \mathcal{K}_h be a *decaying* exponentiated-quadratic kernel. However, if we model \mathcal{K}_h nonparametrically, then it is not clear how we must restrict the prior on \mathcal{K}_h to ensure that f has finite variance. We investigate by examining the case that \mathcal{K}_h is a *non-decaying* exponentiated-quadratic kernel. In that case f has infinite variance (Section 3.3.1). Now, let $w(t) = \exp(-\alpha t^2), t \in \mathbb{R}$ and consider the process wh. It holds that

$$\begin{aligned} \mathcal{K}_{wh}(t_1, t_2) &= \mathbb{E}[w(t_1)h(t_1)w(t_2)h(t_2)] \\ &= w(t_1)w(t_2)\mathbb{E}[h(t_1)h(t_2)] \\ &= \exp(-\alpha t_1^2)\exp(-\alpha t_2^2)\sigma_h^2\exp[-\gamma(t_1-t_2)^2] \\ &= \sigma_h^2\exp[-\alpha t_1^2 - \alpha t_2^2 - \gamma(t_1-t_2)^2], \end{aligned}$$

which is a *decaying* exponentiated-quadratic kernel. In other words, despite h having a kernel for which f has infinite variance, multiplication by w yields a kernel for which f has finite variance. This suggests the following, more robust formulation of the GPCM:

Model 11 (Gaussian Process Convolution Model (Explicit Decay)). Draw

$$h \sim \mathcal{GP}(0, \mathcal{K}_h),$$

$$x \sim \mathcal{GP}[0, \delta(t_1 - t_2)],$$

$$\varepsilon \sim \mathcal{GP}[0, \sigma^2 \delta(t_1 - t_2)]$$

independently for some kernel \mathcal{K}_h and some constant σ . Let $w(t) = \exp(-\alpha t^2)$, $w \in \mathbb{R}$. Then observations are generated by $y = f + \varepsilon = wh * x + \varepsilon$. It now holds that

$$\mathbb{E}[f^2(t)] = \int_{\mathbb{R}^2} w(t-\tau_1)w(t-\tau_2)\mathbb{E}[h(t-\tau_1)h(t-\tau_2)]\underbrace{\mathbb{E}[x(\tau_1)x(\tau_2)]}_{\delta(\tau_1-\tau_2)} d\tau_1 d\tau_2$$
$$= \int_{\mathbb{R}} w^2(\tau)\mathcal{K}_h(\tau,\tau) d\tau,$$

which clearly is finite if \mathcal{K}_h is bounded. In other words, if h has finite variance, then f has finite variance. Therefore, we can safely model \mathcal{K}_h nonparametrically if we ensure that h has finite variance.

To model \mathcal{K}_h nonparametrically, consider the following extension of Model 11:

Model 12 (Deep Gaussian Process Convolution Model of Order N). Draw

$$h_{1} \sim \mathcal{GP}(0, \mathcal{K}_{h}),$$

$$x_{1} \sim \mathcal{GP}[0, \delta(t_{1} - t_{2})],$$

$$\vdots$$

$$x_{N} \sim \mathcal{GP}[0, \delta(t_{1} - t_{2})],$$

$$\varepsilon \sim \mathcal{GP}[0, \sigma^{2}\delta(t_{1} - t_{2})]$$

independently for some kernel \mathcal{K}_h and some constant σ . Let $w(t) = \exp(-\alpha t^2), t \in \mathbb{R}$. Then observations are generated by y where

$$h_{1} = h,$$

$$h_{2} = wh_{1} * x_{1},$$

$$\vdots$$

$$h_{N+1} = wh_{N} * x_{N},$$

$$f = h_{N+1},$$

$$y = f + \varepsilon.$$

Observe that each $h_{i+1} | h_i$ is a linear combination of Gaussian processes. Hence $h_{i+1} | h_i$ is another Gaussian process, which thus can be identified by its mean function and

kernel:

$$\begin{split} \mathbb{E}[h_{i+1}(t) \mid h_i] &= \int_{\mathbb{R}} w(t-\tau)h_i(t-\tau)\mathbb{E}[x_i(\tau)] \,\mathrm{d}\tau = 0, \\ \mathcal{K}_{h_{i+1} \mid h_i}(t_1, t_2) &= \mathbb{E}[h_{i+1}(t_1)h_{i+1}(t_2) \mid h_i] \\ &= \int_{\mathbb{R}^2} w(t_1 - \tau_1)h_i(t_1 - \tau_1)w(t_2 - \tau_2)h_i(t_2 - \tau_2)\underbrace{\mathbb{E}[x_i(\tau_1)x_i(\tau_2)]}_{\delta(\tau_1 - \tau_2)} \,\mathrm{d}\tau_1 \,\mathrm{d}\tau_2 \\ &= \int_{\mathbb{R}} w[\tau - (t_2 - t_1)]h_i[\tau - (t_2 - t_1)]w(\tau)h_i(\tau) \,\mathrm{d}\tau \\ &= [R(wh_i) * wh_i](t_2 - t_1) \\ &= \mathcal{K}_{h_{i+1}}(t_1 - t_2). \end{split}$$

We have established the following equivalent model:

Model 13 (Deep Kernel Model of Order N). Draw

$$h \sim \mathcal{GP}(0, \mathcal{K}_h),$$

$$\varepsilon \sim \mathcal{GP}[0, \sigma^2 \delta(t_1 - t_2)]$$

independently for some kernel \mathcal{K}_h and some constant σ . Let $h = h_1$ and $w(t) = \exp(-\alpha t^2)$, $t \in \mathbb{R}$. Afterwards draw

$$h_{i+1} | h_i \sim \mathcal{GP}\{0, [R(wh_i) * wh_i](t_2 - t_1)\}$$

in the order i = 1, ..., N. Then observations are generated by $y = f + \varepsilon = h_{N+1} + \varepsilon$.

Model 13 shows that the kernel of $f = h_{N+1}$ is now parametrised by h_N , whose kernel is parametrised by h_{N-1} , whose kernel is parametrised by h_{N-2} , et cetera. That is, we have successfully formulated a nonparametric model of the kernel, the *kernel of the kernel*, the *kernel of the kernel*, and further "deeper" kernels. Furthermore, it holds that

$$\mathbb{E}[f^{2}(t)] = \mathbb{E}[(w\{\cdots [w(wh_{1} * x_{1}) * x_{2}] \cdots \} * x_{N})^{2}(t)]$$

$$= \int \left[\prod_{i=1}^{N} w\left(t - \sum_{j=i}^{N} \tau_{j}^{(1)}\right) w\left(t - \sum_{j=i}^{N} \tau_{j}^{(2)}\right)\right]$$

$$\mathbb{E}\left[h_{1}\left(t - \sum_{i=1}^{N} \tau_{i}^{(1)}\right) h_{1}\left(t - \sum_{i=1}^{N} \tau_{i}^{(2)}\right)\right] \prod_{i=1}^{N} \underbrace{\mathbb{E}[x_{i}(\tau_{i}^{(1)}) x_{i}(\tau_{i}^{(2)})]}_{\delta(\tau_{i}^{(1)} - \tau_{i}^{(2)})} d\tau_{i}^{(1)} d\tau_{i}^{(2)}$$

$$= \int \left[\prod_{i=1}^{N} w^{2}\left(t - \sum_{j=i}^{N} \tau_{j}\right)\right] \mathcal{K}_{h_{1}}\left(t - \sum_{i=1}^{N} \tau_{i}, t - \sum_{i=1}^{N} \tau_{i}\right) \prod_{i=1}^{N} d\tau_{i}$$

$$= \int_{\mathbb{R}} w^{2}(\tau_{1}) \int_{\mathbb{R}} w^{2}(\tau_{1} + \tau_{2}) \cdots \int_{\mathbb{R}} w^{2}\left(\sum_{i=1}^{N} \tau_{i}\right) \mathcal{K}_{h_{1}}\left(\sum_{i=1}^{N} \tau_{i}, \sum_{i=1}^{N} \tau_{i}\right) d\tau_{N} \cdots d\tau_{2} d\tau_{1}.$$

Observe that the most inner integral is finite if \mathcal{K}_{h_1} is bounded, in which case all integrals are finite. In other words, if h_1 has finite variance, then f has finite variance. Therefore Models 12 and 13 are well defined.

5.2.1 Network Interpretation

The left side of Figure 5.1 depicts the graphical model of Model 13. We recognise the "deep" kernel structure by the chain of latent h_i 's.

The right side of Figure 5.1 depicts the graphical model of Model 13 if the latent processes h_i are expanded into their function values; that is, any h_i is instead represented by $h_i(t)$ for all *t*—uncountably many *t*. Now, since the kernel of h_{i+1} is a convolution of h_i , it holds that every $h_{i+1}(t)$ depends on $h_i(t')$ for all *t'*. Hence the resulting graphical model is a fully connected network with layers of infinite size. This is very reminiscent of neural networks. The key difference between a neural network and Model 13 is that in a neural network layer *i* forms the input of layer i + 1, whereas in Model 13 layer *i parametrises* layer i + 1.

We can truncate these layers of infinite size to finite size by approximating the kernels in Model 13 by their expectations conditioned on some finite number of observations. We then obtain the Approximate Deep Kernel Model:



Figure 5.1: Graphical model of Model 13. Shows the resulting network structure if the latent h_i are expanded into their function values.

Model 14 (Approximate Deep Kernel Model of Order N). Draw

$$h \sim \mathcal{GP}(0, \mathcal{K}_h),$$

$$\varepsilon \sim \mathcal{GP}[0, \sigma^2 \delta(t_1 - t_2)]$$

independently for some kernel \mathcal{K}_h and some constant σ . Let $h = h_1$ and $w(t) = \exp(-\alpha t^2)$, $w \in \mathbb{R}$. Afterwards draw

$$\mathcal{K}_{h_{i+1}|h_i}(t_1, t_2) = \mathbb{E}\{[R(wh_i) * wh_i](t_2 - t_1) | h_i(t)\},\$$

$$h_{i+1}|h_i \sim \mathcal{GP}(0, \mathcal{K}_{h_{i+1}|h_i})$$

in the order i = 1, ..., N. Then observations are generated by $y = f + \varepsilon = h_{N+1} + \varepsilon$.

Note that Model 14 approximates Model 13 in same way that Model 6 approximates Model 5.

5.3 Illustrative Samples

In the following we let h_1 have an exponentiated-quadratic kernel.

Figure 5.2 illustrates the generative process of Model 13 in the case that N = 10. Observe that the produced $h_i | h_{i-1}$'s and $\mathcal{K}_{h_i | h_{i-1}}$'s change in complexity as the generation progresses; that is, if by chance $h_i | h_{i-1}$ is slightly more complicated, then $\mathcal{K}_{h_{i+1} | h_i}$ and thereby $h_{i+1} | h_i$ tend to be more complicated as well. Thus, in each generation step Model 13 can either simplify or complicate the kernel, which means that $\mathcal{K}_{f | h_{10}}$ can be of a vastly different form than the kernel produced by h_1 ; Model 13 therefore seems to be less restricted by h_1 's kernel than Model 10 is by h's kernel.

Figure 5.3 shows observations from Model 13 in the cases that N = 1 and N = 20—a shallow and deep model respectively. Consider the shallow model. Observe that its observations exhibit barely any variation in smoothness and degree of periodicity; that is, their kernels are all of a similar form and their power spectral densities all show a similar low-pass structure. This is due h_1 's exponentiated-quadratic kernel, which imposes smoothness on h_1 and thereby directly on the observation's kernel. On the other hand, consider the deep model. Observe that its observations are much richer: they can be periodic, aperiodic, smoothly varying, or anywhere inbetween. This is reflected in the generated kernels, which now show great diversity. Furthermore, instead

of the low-pass structure observed for the shallow model, the power spectral densities now show low-pass structures, high-pass structures, band-pass structures, and structures inbetween.

5.4 Conclusion

We have presented Model 13, which not only models the kernel nonparametrically, but also the kernel of the kernel, the kernel of the kernel of the kernel, and further "deeper" kernels. Experiments showed that Model 13 exhibits greatly increased expressivity compared to its shallow counterpart Model 10, or equivalently Model 5 where M = N = K = 1.

5.5 Discussion

We can perform approximate inference in Model 13 along the lines of Section 3.6. Although the free energy allows to be solved for analytically, its resulting complexity is exponential in the number of layers. Therefore, to perform inference in Model 13, additional approximations are probably necessary.



Figure 5.2: Generative process of Model 13 in the case that N = 10

Observation $(N = 1)$	Kernel	PSD O	bservation $(N = 2$	0) Kernel	PSD
Whenter			WWW//ww//		MA
MWWWWWW	MM	- MM	$\bigwedge \bigwedge$		M
MMMMM		- MMM	Murham		MMM
MMmmm		-Manny-	mmMr		M
NMMMM		- Marth	MmM		MM
mmmmm		-Martin	h <mark>hl</mark> innt/myth	www.	MMM
Muhummy		- Marine	mmmm		AMA
WMMMMMMM		- Army	/w		MA
myMmm		- Martin -	MMM MM	~~~~~	MM
MMMMM					AMA

Figure 5.3: Observations from Model 13 in the cases that N = 1 and N = 20. Also shows the kernels and power spectral densities (PSDs) of the observations.

A | Solution of the Linear State-Space Model

A.1 Time-Variant Solution

 $\operatorname{Consider}$

$$s'(t) = A(t)s(t) + B(t)x(t), \qquad (A.1)$$
$$f(t) = C(t)s(t) + D(t)x(t).$$

Let $\Psi(t)$ be a fundamental matrix of the homogeneous system—that is, $\Psi'(t) = \mathbf{A}(t)\Psi(t)$ where $\Psi(t)$'s columns are independent. We proceed to find the general solution via variation of parameters. Assume a solution of the form $\mathbf{s}(t) = \Psi(t)\mathbf{c}(t)$. Then application of the product rule yields that

$$\boldsymbol{s}'(t) = \boldsymbol{\Psi}'(t)\boldsymbol{c}(t) + \boldsymbol{\Psi}(t)\boldsymbol{c}'(t) = \boldsymbol{A}(t)\boldsymbol{\Psi}(t)\boldsymbol{c}(t) + \boldsymbol{\Psi}(t)\boldsymbol{c}'(t)$$

while substitution into Equation (A.1) yields that

$$\boldsymbol{s}'(t) = \boldsymbol{A}(t)\boldsymbol{\Psi}(t)\boldsymbol{c}(t) + \boldsymbol{B}(t)\boldsymbol{x}(t).$$

Hence $\Psi(t)c'(t) = B(t)x(t)$. Since $\Psi(t)$ is invertible we have that $c'(t) = \Psi^{-1}(t)B(t)x(t)$. Therefore

$$\begin{aligned} \boldsymbol{f}(t) &= \boldsymbol{C}(t)\boldsymbol{\Psi}(t)\boldsymbol{c}(t) + \boldsymbol{D}(t)\boldsymbol{x}(t) \\ &= \boldsymbol{C}(t)\boldsymbol{\Psi}(t)\int_{-\infty}^{t}\boldsymbol{\Psi}^{-1}(\tau)\boldsymbol{x}(\tau)\boldsymbol{B}(t)\,\mathrm{d}\tau + \boldsymbol{D}(t)\boldsymbol{x}(t) \\ &= \int_{-\infty}^{t}\boldsymbol{C}(t)\boldsymbol{\Psi}(t)\boldsymbol{\Psi}^{-1}(\tau)\boldsymbol{B}(t)\boldsymbol{x}(\tau)\,\mathrm{d}\tau + \boldsymbol{D}(t)\boldsymbol{x}(t). \end{aligned}$$

Now

$$\begin{aligned} \boldsymbol{D}(t)\boldsymbol{x}(t) &= \int_{(-\infty,\infty)} \delta(t-\tau)\boldsymbol{D}(t)\boldsymbol{x}(\tau) \,\mathrm{d}\tau \\ &= \int_{(-\infty,t]} \delta(t-\tau)\boldsymbol{D}(t)\boldsymbol{x}(\tau) \,\mathrm{d}\tau + \int_{(t,-\infty)} \delta(t-\tau)\boldsymbol{D}(t)\boldsymbol{x}(\tau) \,\mathrm{d}\tau \\ &= \int_{(-\infty,t]} \delta(t-\tau)\boldsymbol{D}(t)\boldsymbol{x}(\tau) \,\mathrm{d}\tau \end{aligned}$$

so that finally

$$\begin{aligned} \boldsymbol{f}(t) &= \int_{\mathbb{R}} \underbrace{[\boldsymbol{C}(t)\boldsymbol{\Psi}(t)\boldsymbol{\Psi}^{-1}(\tau)\boldsymbol{B}(t) + \delta(t-\tau)\boldsymbol{D}(t)]\boldsymbol{\mathbb{1}}(\tau \leq t)}_{\boldsymbol{H}(t,\tau)} \boldsymbol{x}(\tau) \,\mathrm{d}\tau \\ &= \int_{\mathbb{R}} \boldsymbol{H}(t,\tau)\boldsymbol{x}(\tau) \,\mathrm{d}\tau. \end{aligned}$$

A.2 Time-Invariant Solution

Let A(t), B(t), C(t) and D(t) be independent of t; that is, let A(t) = A, B(t) = B, C(t) = C and D(t) = D. The fundamental matrix is then given by the matrix exponential

$$\Psi(t) = \sum_{i=0}^{\infty} \frac{t^i A^i}{i!},$$
denoted as $\Psi(t) = \exp t \mathbf{A}$. To show this, note that by the binomial theorem and a change of variables

$$\exp[(s+t)\mathbf{A}] = \sum_{i=0}^{\infty} \frac{(s+t)^i \mathbf{A}^i}{i!}$$
$$= \sum_{i=0,j=0}^{\infty,\infty} \frac{s^{i-j} t^j \mathbf{A}^i}{(i-j)! j!}$$
$$= \sum_{i=0,j=0}^{\infty,\infty} \frac{s^i t^j \mathbf{A}^{i+j}}{i! j!}$$
$$= (\exp s\mathbf{A})(\exp t\mathbf{A})$$

so that

$$\Psi'(t) = \lim_{h \to 0} \frac{\exp[(s+h)\mathbf{A}] - \exp s\mathbf{A}}{h}$$
$$= \left(\lim_{h \to 0} \frac{\exp h\mathbf{A} - \mathbf{I}}{h}\right) \exp s\mathbf{A}$$
$$= \mathbf{A}\Psi(t).$$

Now

$$\boldsymbol{I} = \exp \boldsymbol{0} = \exp(t\boldsymbol{A} - t\boldsymbol{A}) = (\exp t\boldsymbol{A})[\exp(-t\boldsymbol{A})]$$

implies that $(\exp t\mathbf{A})^{-1} = \exp(-t\mathbf{A})$. Consequently

$$H(t,\tau) = C\Psi(t)\Psi^{-1}(\tau)B + \delta(t-\tau)D$$

= $C(\exp tA)(\exp \tau A)^{-1}B + \delta(t-\tau)D$
= $C(\exp tA)[\exp(-\tau A)]B + \delta(t-\tau)D$
= $C\exp[(t-\tau)A]B + \delta(t-\tau)D$
= $H(t-\tau).$

B | Properties of the Multivariate Gaussian Distribution

B.1 Marginal and Conditional Distribution

Let

$$egin{bmatrix} oldsymbol{x}_1\ oldsymbol{x}_2 \end{bmatrix} \sim \mathcal{N}\left(egin{bmatrix} oldsymbol{\mu}_1\ oldsymbol{\mu}_2 \end{bmatrix}, egin{bmatrix} oldsymbol{\Sigma}_{1,1} & oldsymbol{\Sigma}_{1,2}\ oldsymbol{\Sigma}_{2,1} & oldsymbol{\Sigma}_{2,2} \end{bmatrix}
ight).$$

Then Murphy [2012] shows that $\boldsymbol{x}_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_{1,1})$ and

$$m{x}_1 \,|\, m{x}_2 \sim \mathcal{N}[m{\mu}_1 + m{\Sigma}_{1,2} m{\Sigma}_{2,2}^{-1} (m{x}_2 - m{\mu}_2), m{\Sigma}_{1,1} - m{\Sigma}_{1,2} m{\Sigma}_{2,2}^{-1} m{\Sigma}_{2,1}].$$

B.2 Kullback-Leibler Divergence

Let $\boldsymbol{x}_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $\boldsymbol{x}_2 \sim \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ where \boldsymbol{x}_1 and \boldsymbol{x}_2 attain values in \mathbb{R}^N . Then Murphy [2012] shows that

$$D_{KL}[p(\boldsymbol{x}_1) \| p(\boldsymbol{x}_2)] = \frac{1}{2} \left[\operatorname{tr}(\boldsymbol{\Sigma}_2^{-1}\boldsymbol{\Sigma}_1) + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_2^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) + \log \frac{|\boldsymbol{\Sigma}_2|}{|\boldsymbol{\Sigma}_1|} - N \right].$$

C | Multivariate Matrix-Valued Gaussian Processes

A stochastic process $F(t) \in \mathbb{R}^{N \times M}$, $t \in \mathbb{R}^{K}$ is Gaussian if and only if for every $T \in \mathbb{R}^{T \times K}$ it holds that (vec F)(T) is multivariate Gaussian distributed where

$$(\operatorname{vec} \boldsymbol{F})(\boldsymbol{T}) = \begin{bmatrix} \operatorname{vec} \boldsymbol{F}(\boldsymbol{T}_{1,:}) \\ \vdots \\ \operatorname{vec} \boldsymbol{F}(\boldsymbol{T}_{T,:}) \end{bmatrix}$$
$$= \begin{bmatrix} F_{1,1}(\boldsymbol{T}_{1,:}) & \cdots & F_{N,M}(\boldsymbol{T}_{1,:}) & \cdots & F_{1,1}(\boldsymbol{T}_{T,:}) & \cdots & F_{N,M}(\boldsymbol{T}_{T,:}) \end{bmatrix}^{T}.$$

The mean function is then defined by

$$\boldsymbol{m}_{\boldsymbol{F}}(\boldsymbol{T}) = \begin{bmatrix} \boldsymbol{m}_{\boldsymbol{F}}^{T}(\boldsymbol{T}_{1,:}) & \cdots & \boldsymbol{m}_{\boldsymbol{F}}^{T}(\boldsymbol{T}_{T,:}) \end{bmatrix}^{T} \\ = \begin{bmatrix} m_{F_{1,1}}(\boldsymbol{T}_{1,:}) & \cdots & m_{F_{N,M}}(\boldsymbol{T}_{1,:}) & \cdots & m_{F_{1,1}}(\boldsymbol{T}_{T,:}) & \cdots & m_{F_{N,M}}(\boldsymbol{T}_{T,:}) \end{bmatrix}^{T}$$

and the kernel by

$$\mathcal{K}_{F}(T^{(1)},T^{(2)}) = egin{bmatrix} \mathcal{K}_{F}(T^{(1)}_{1,:},T^{(2)}_{1,:}) & \cdots & \mathcal{K}_{F}(T^{(1)}_{1,:},T^{(2)}_{T,:}) \ dots & \ddots & dots \ \mathcal{K}_{F}(T^{(1)}_{T,:},T^{(2)}_{1,:}) & \cdots & \mathcal{K}_{F}(T^{(1)}_{T,:},T^{(2)}_{T,:}) \end{bmatrix}$$

where

$$\boldsymbol{\mathcal{K}_{F}}(\boldsymbol{T}_{i,:}^{(1)},\boldsymbol{T}_{j,:}^{(2)}) = \begin{bmatrix} \mathcal{K}_{F_{1,1},F_{1,1}}(\boldsymbol{T}_{i,:}^{(1)},\boldsymbol{T}_{j,:}^{(2)}) & \cdots & \mathcal{K}_{F_{1,1},F_{N,M}}(\boldsymbol{T}_{i,:}^{(1)},\boldsymbol{T}_{j,:}^{(2)}) \\ \vdots & \ddots & \vdots \\ \mathcal{K}_{F_{N,M},F_{1,1}}(\boldsymbol{T}_{i,:}^{(1)},\boldsymbol{T}_{j,:}^{(2)}) & \cdots & \mathcal{K}_{F_{N,M},F_{N,M}}(\boldsymbol{T}_{i,:}^{(1)},\boldsymbol{T}_{j,:}^{(2)}) \end{bmatrix}.$$

D | Gaussian Processes in Practice

D.1 Implementation of Gaussian Process Models

We briefly discuss some useful techniques concerning the implementation of Gaussian process models. Assume that all matrices have shape $N \times N$.

First, carefully considering the memory layout of large matrices can be beneficial. Namely, contiguous memory is accessed more cheaply than non-contiguous memory. This means that the cost of operations on large matrices greatly depends on the memory layout of the matrices.

Second, a covariance matrix Σ can be made positive semidefinite by adding a small diagonal. In theory, Σ is always positive semidefinite. However, when it is evaluated numerically, it can be indefinite or negative definite due round-off errors introduced by the floating-point representation of numbers. In that case Σ 's positive semidefiniteness can be ensured by adding a diagonal comparable to $\varepsilon \|\Sigma\|_{\infty}$ where ε is the machine epsilon. Usually the noise introduced by the diagonal is negligible.

In the case that Σ is indefinite or negative definite due to noise other than round-off errors, the diagonal required to make Σ positive semidefinite might be so large that the introduced noise becomes significant. In that case Theorem 1 can be used to compute the symmetric positive semidefinite matrix nearest to Σ in Frobenius norm.

Third, quantities of the form $A\Sigma^{-1}B$ where Σ is a covariance matrix should not be computed by first computing $\Sigma^{-1} = X$ explicity. Namely, the error in AXB due to round-off errors can be large, especially when Σ is ill-conditioned [Trefethen and Bau, 1997]. A numerically more stable approach is to first compute the Cholesky decomposition $\Sigma = \boldsymbol{L} \boldsymbol{L}^T$ where \boldsymbol{L} is lower triangular. Then

$$A\Sigma^{-1}B = A(LL^{T})^{-1}B = (AL^{-T})(L^{-1}B)$$

where the systems AL^{-T} and $L^{-1}B$ can be solved using respectively back substitution and forward substitution. The Cholesky decomposition has time complexity $\mathcal{O}(N^3)$ and back substitution and forward substitution have time complexity $\mathcal{O}(N^2)$ [Trefethen and Bau, 1997]. Thus $A\Sigma^{-1}B$ can be computed in $\mathcal{O}(N^3)$ time.

Fourth, the determinant Σ of a covariance matrix can efficiently be computed via its Cholesky decomposition $\Sigma = LL^{T}$. Specifically, as L is lower triangular,

$$|\boldsymbol{\Sigma}| = |\boldsymbol{L}|^2 = \prod_{i=1}^N L_{i,i}^2.$$

Thus the determinant can be computed in $\mathcal{O}(N^3)$ time.

D.2 Nearest Symmetric Positive-Semidefinite Matrix

Lemma 1 (Polar Decomposition of Symmetric Matrix). Let A be symmetrix and let $B = Z\Lambda Z^T$ be its spectral decomposition. Then there exists an orthogonal matrix U such that

$$oldsymbol{A} = oldsymbol{U}oldsymbol{Z}oldsymbol{S}oldsymbol{Z}^T$$

where $\mathbf{S} = \operatorname{diag}(|\Lambda_{1,1}|, \ldots, |\Lambda_{N,N}|).$

Proof. Let \boldsymbol{L} be diagonal such that $L_{i,i} = \operatorname{sign} \Lambda_{i,i}$. Then $\boldsymbol{L}\boldsymbol{L} = \boldsymbol{I}$. Therefore

$$oldsymbol{A} = oldsymbol{Z} oldsymbol{\Lambda} oldsymbol{Z}^T = oldsymbol{Z} oldsymbol{L} oldsymbol{\Lambda} oldsymbol{Z}^T = oldsymbol{Z} oldsymbol{L} oldsymbol{X}^T oldsymbol{Z} oldsymbol{L} oldsymbol{\Lambda} oldsymbol{Z}^T = oldsymbol{Z} oldsymbol{L} oldsymbol{\Lambda} oldsymbol{Z}^T oldsymbol{Z} oldsymbol{L} oldsymbol{\Lambda} oldsymbol{Z}^T$$

where $\boldsymbol{L}\boldsymbol{\Lambda} = \operatorname{diag}(|\Lambda_{1,1}|, \ldots, |\Lambda_{N,N}|)$ and

$$UU^T = ZLZ^TZLZ^T = ZLLZ^T = ZZ^T = I.$$

Lemma 2. Let A be symmetric and let B be antisymmetric. Then $||A + B||_F^2 = ||A||_F^2 + ||B||_F^2$.

Proof. To begin with, $\mathbf{A}^T \mathbf{B} = -\mathbf{A}\mathbf{B}^T$; thus $\operatorname{tr}(\mathbf{A}^T \mathbf{B}) + \operatorname{tr}(\mathbf{A}\mathbf{B}^T) = 0$. Therefore

$$\begin{split} \|\boldsymbol{A} + \boldsymbol{B}\|_{F}^{2} &= \operatorname{tr}[(\boldsymbol{A} + \boldsymbol{B})^{T}(\boldsymbol{A} + \boldsymbol{B})] \\ &= \operatorname{tr}(\boldsymbol{A}^{T}\boldsymbol{A}) + \operatorname{tr}(\boldsymbol{A}^{T}\boldsymbol{B}) + \operatorname{tr}(\boldsymbol{B}^{T}\boldsymbol{A}) + \operatorname{tr}(\boldsymbol{B}^{T}\boldsymbol{B}) \\ &= \operatorname{tr}(\boldsymbol{A}^{T}\boldsymbol{A}) + \operatorname{tr}(\boldsymbol{B}^{T}\boldsymbol{B}) \\ &= \|\boldsymbol{A}\|_{F}^{2} + \|\boldsymbol{B}\|_{F}^{2}. \end{split}$$

Lemma 3. Let A be a matrix and let U be orthogonal. Then $\|UAU^T\|_F^2 = \|A\|_F^2$.

Proof. Direct computation yields that

$$\begin{split} \|\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^{T}\|_{F}^{2} &= \operatorname{tr}(\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^{T}\boldsymbol{U}\boldsymbol{A}^{T}\boldsymbol{U}^{T}) \\ &= \operatorname{tr}(\boldsymbol{U}\boldsymbol{A}\boldsymbol{A}^{T}\boldsymbol{U}^{T}) \\ &= \operatorname{tr}(\boldsymbol{U}^{T}\boldsymbol{U}\boldsymbol{A}\boldsymbol{A}^{T}) \\ &= \operatorname{tr}(\boldsymbol{A}\boldsymbol{A}^{T}) \\ &= \|\boldsymbol{A}\|_{F}^{2}. \end{split}$$

 Γ.	-	-1	

The following proof is based on the proof by Higham [1988].

Theorem 1 (Nearest Symmetric Positive-Semidefinite Matrix). Let A be a matrix and $B = (A + A^T)/2$. Furthermore, let $B = Z\Lambda Z^T$ be B's spectral decomposition and let $B = UZSZ^T$ be B's polar decomposition (Lemma 1). Finally, let $A_F = (B + ZSZ^T)/2$. Then

- 1. A_F is symmetric positive semidefinite and
- 2. for any other symmetric positive semidefinite matrix \mathbf{X} it holds that $\|\mathbf{A} \mathbf{X}\|_F^2 \ge \|\mathbf{A} \mathbf{A}_F\|_F^2$.

Proof. First, we verify that A_F is symmetric positive semidefinite. Trivially, A_F is symmetric. To see that A_F is positive semidefinite, consider

$$\boldsymbol{A}_{F} = \frac{1}{2} (\boldsymbol{B} + \boldsymbol{Z} \boldsymbol{S} \boldsymbol{Z}^{T}) = \frac{1}{2} (\boldsymbol{Z} \boldsymbol{\Lambda} \boldsymbol{Z}^{T} + \boldsymbol{Z} \boldsymbol{S} \boldsymbol{Z}^{T}) = \boldsymbol{Z} \frac{\boldsymbol{\Lambda} + \boldsymbol{S}}{2} \boldsymbol{Z}^{T}.$$
 (D.1)

By Lemma 1,

$$\frac{\mathbf{\Lambda} + \mathbf{S}}{2} = \operatorname{diag}\left(\frac{\Lambda_{1,1} + |\Lambda_{1,1}|}{2}, \dots, \frac{\Lambda_{N,N} + |\Lambda_{N,N}|}{2}\right), \quad (D.2)$$

whose diagonal elements are all nonnegative. Therefore, for any \boldsymbol{x} , it holds that

$$\boldsymbol{x}^{T}\boldsymbol{A}_{F}\boldsymbol{x} = \boldsymbol{x}^{T}\boldsymbol{Z}\left(\frac{\boldsymbol{\Lambda}+\boldsymbol{S}}{2}\right)^{1/2}\left(\frac{\boldsymbol{\Lambda}+\boldsymbol{S}}{2}\right)^{1/2}\boldsymbol{Z}\boldsymbol{x}^{T} = \left\|\left(\frac{\boldsymbol{\Lambda}+\boldsymbol{S}}{2}\right)^{1/2}\boldsymbol{Z}^{T}\boldsymbol{x}\right\|_{2}^{2} \ge 0.$$

Second, let X be symmetric positive semidefinite. We then show that $||A - X||_F^2 \ge ||A - A_F||_F^2$. To begin with, let $C = (A - A^T)/2$. Then A = B + C where B is symmetric and C is antisymmetric. Thus, by Lemma 2,

$$\|\boldsymbol{A} - \boldsymbol{X}\|_{F}^{2} = \|(\boldsymbol{B} - \boldsymbol{X}) + \boldsymbol{C}\|_{F}^{2} = \|\boldsymbol{B} - \boldsymbol{X}\|_{F}^{2} + \|\boldsymbol{C}\|_{F}^{2} \ge \|\boldsymbol{B} - \boldsymbol{X}\|_{F}^{2}$$

Let $\boldsymbol{Y} = \boldsymbol{Z}^T \boldsymbol{X} \boldsymbol{Z}$. As \boldsymbol{X} is positive semidefinite, \boldsymbol{Y} is also positive semidefinite and so $Y_{i,i} \geq 0$. Therefore, by Lemma 3,

$$egin{aligned} \|oldsymbol{A}-oldsymbol{X}\|_F^2 &\geq \|oldsymbol{B}-oldsymbol{X}\|_F^2 \ &= \|oldsymbol{Z} (oldsymbol{\Lambda}-oldsymbol{Z}^Toldsymbol{X} oldsymbol{Z}^T -oldsymbol{X}\|_F^2 \ &= \|oldsymbol{\Lambda}-oldsymbol{Y}\|_F^2 \ &\geq \sum_{i:\Lambda_{i,i} < 0} (\Lambda_{i,i} - Y_{i,i})^2 \ &\geq \sum_{i:\Lambda_{i,i} < 0} \Lambda_{i,i}^2. \end{aligned}$$

This lower bound holds for any positive semidefinite X. Therefore, if A_F achieves this lower bound, then the result follows.

D.2 Nearest Symmetric Positive-Semidefinite Matrix

To this end, let $X = A_F$. As A_F is symmetric, C = 0 and so

$$\|oldsymbol{A}-oldsymbol{A}_F\|_F^2 = \|oldsymbol{\Lambda}-oldsymbol{Y}\|_F^2$$

where $\boldsymbol{Y} = \boldsymbol{Z} \boldsymbol{A}_F \boldsymbol{Z}^T$. Now, Equation (D.1) and Equation (D.2) show that

$$\boldsymbol{Y} = \operatorname{diag}\left(\frac{\Lambda_{1,1} + |\Lambda_{1,1}|}{2}, \dots, \frac{\Lambda_{N,N} + |\Lambda_{N,N}|}{2}\right).$$

Hence

$$\begin{split} \|\boldsymbol{A} - \boldsymbol{A}_F\|_F^2 &= \|\boldsymbol{\Lambda} - \boldsymbol{Y}\|_F^2 \\ &= \sum_{i=1}^N \left(\Lambda_{i,i} - \frac{\Lambda_{i,i} + |\Lambda_{i,i}|}{2} \right)^2 \\ &= \sum_{i:\Lambda_{i,i} < 0} \left(\Lambda_{i,i} - \frac{\Lambda_{i,i} + |\Lambda_{i,i}|}{2} \right)^2 \\ &= \sum_{i:\Lambda_{i,i} < 0} \Lambda_{i,i}^2. \end{split}$$

- L		J

E | Circulant Approximation of Stationary Multi-Output Kernel Matrices

E.1 Introduction

Let $\mathbf{K} : \mathbb{R}^K \times \mathbb{R}^K \to \mathbb{R}^{N \times N}$ be a stationary multi-output kernel and let \mathbf{K}_F be \mathbf{K} evaluated for points that lie on an evenly-spaced grid $\{t_1^{(1)}, \ldots, t_{Y_1}^{(1)}\} \times \cdots \times \{t_1^{(K)}, \ldots, t_{Y_K}^{(K)}\}$. The total number of points is then given by $Y_1 \cdots Y_K = Y$. Assume that all Y_i are even.

Consider computation of $|\mathbf{K}_{F}|$ and a product of the form $\mathbf{P} = \mathbf{U}^{T}\mathbf{K}_{F}^{-1}\mathbf{V}$ for some $NY \times P$ matrix \mathbf{U} and $NY \times Q$ matrix \mathbf{V} . Computing $|\mathbf{K}_{F}|$ directly costs $\mathcal{O}(N^{3}Y^{3})$ time (Appendix D). This does not scale for large Y. Furthermore, inverting \mathbf{K}_{F} costs $\mathcal{O}(N^{3}Y^{3})$ time and then computing $\mathbf{U}^{T}\mathbf{K}_{F}^{-1}\mathbf{V}$ costs $\mathcal{O}(PQN^{2}Y^{2})$ time, resulting in time complexity $\mathcal{O}(PQN^{2}Y^{2} + N^{3}Y^{3})$. This also does not scale for large Y.

This chapter develops an approximation of K_F based on the facts that K_F is stationary and is evaluated for points on a grid [Ulrich et al., 2015]. We show that this approximation can be leveraged to compute $|K_F|$ and $U^T K_F^{-1} V$ with complexities that scale favourably in Y.



Figure E.1: Circulant approximation of a stationary kernel. Shows how the space "wraps around" at the edge of the grid. Also shows that the approximation is excellent for points not in the vicinity of the edge.

E.2 Circulant Approximation of Toeplitz Matrices

Let $k : \mathbb{R} \to \mathbb{R}$ be a stationary kernel and let K_f be k evaluated for points that lie on an evenly-spaced grid $\{t_1, \ldots, t_Y\}$. Then, by k's stationary and the fact that the points lie on a grid,

$$\boldsymbol{K}_{\boldsymbol{f}} = \operatorname{stoep}[k(t_1 - t_1), \dots, k(t_1 - t_Y)].$$

A banded Toeplitz matrix can be approximated by a particular circulant matrix where the approximation becomes more accurate as the matrix becomes large [Gray, 2006]. More specifically, if k has effectively limited support, then we can approximate

$$\begin{aligned} \boldsymbol{K}_{\boldsymbol{f}} &= \operatorname{stoep}[k(t_1 - t_1), \dots, k(t_1 - t_{Y/2}), k(t_1 - t_{Y/2+1}), k(t_1 - t_{Y/2+2}), \dots, k(t_1 - t_Y)] \\ &\approx \operatorname{stoep}[k(t_1 - t_1), \dots, k(t_1 - t_{Y/2}), k(t_1 - t_{Y/2+1}), k(t_1 - t_{Y/2}), \dots, k(t_1 - t_2)] \\ &= \operatorname{circ}[k(t_1 - t_1), \dots, k(t_1 - t_{Y/2}), k(t_1 - t_{Y/2+1}), k(t_1 - t_{Y/2}), \dots, k(t_1 - t_2)]. \end{aligned}$$
(E.1)

We call this approximation the *circulant approximation* of K_f . Essentially, the approximation lets the space in which distance is measured now "wrap around" at the edge of the grid. As a consequence, points near the edge of the grid correlate. However, if the kernel is local, then correlations between points not in the vicinity of the edge are preserved; in that case K_f 's circulant approximation is an excellent approximation for most points. Figure E.1 illustrates the circulant approximation of a stationary kernel.

Finally, the $Y \times Y$ unitary discrete Fourier transform matrix \mathcal{F}_Y diagonalises any $Y \times Y$ circulant matrix [Gray, 2006]. Thus, by $\mathbf{K}_f = \mathbf{K}_f^T$, it holds that $\mathbf{K}_f \approx \mathcal{F}_Y \Lambda \mathcal{F}_Y^H = \mathcal{F}_Y^H \Lambda \mathcal{F}_Y$ where Λ is diagonal.

E.3 Circulant Approximation of Stationary Multi-Output Kernel Matrices

In the same way that a Toeplitz matrix has a circulant approximation, a block Toeplitz matrix has a block circulant approximation. We use this observation to approximate K_F .

We construct K_F in a particular way. Specifically, let $K^{(0)}(t,t') = K(t,t')$ and let

$$\boldsymbol{t}^{(>i)} = \begin{bmatrix} t_{i+1} & \cdots & t_K \end{bmatrix}^T, \\ \boldsymbol{t}^{(>i)'} = \begin{bmatrix} t'_{i+1} & \cdots & t'_K \end{bmatrix}^T, \\ \boldsymbol{K}^{(i)}(\boldsymbol{t}^{(>i)}, \boldsymbol{t}^{(>i)'}) = \begin{bmatrix} \boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_1^{(i)}, \boldsymbol{t}^{(>i)'}) & \cdots & \boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, \boldsymbol{t}^{(>i)'}) \end{bmatrix} \\ \vdots & \ddots & \vdots \\ \boldsymbol{K}^{(i-1)}(t_{Y_i}^{(i)}, \boldsymbol{t}^{(>i)}, t_1^{(>i)}, \boldsymbol{t}^{(>i)'}) & \cdots & \boldsymbol{K}^{(i-1)}(t_{Y_i}^{(i)}, \boldsymbol{t}^{(>i)}, t_{Y_i}^{(i)}, \boldsymbol{t}^{(>i)'}) \end{bmatrix}$$

for i = 1, ..., K. Then $\mathbf{K}_{\mathbf{F}} = \mathbf{K}^{(K)}$. Now, \mathbf{K} is stationary and the points for which we evaluate \mathbf{K} lie on a grid; thus we can use Equation (E.1) to approximate

$$\begin{split} \boldsymbol{K}^{(i)}(\boldsymbol{t}^{(>i)}, \boldsymbol{t}^{(>i)'}) \\ &= \operatorname{stoep}[\boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_1^{(i)}, \boldsymbol{t}^{(>i)'}), \dots, \boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_{Y_i}^{(i)}, \boldsymbol{t}^{(>i)'})] \\ &\approx \operatorname{circ}[\boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_{w^{(i)}(1)}^{(i)}, \boldsymbol{t}^{(>i)'}), \dots, \boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_{w^{(i)}(Y_i)}^{(i)}, \boldsymbol{t}^{(>i)'})] \end{split}$$

where $w^{(i)}(j_i) = \min\{j, Y_i - j_i + 2\}$. Then, by decomposing each $\boldsymbol{K}^{(i)}(\boldsymbol{t}^{(>i)}, \boldsymbol{t}^{(>i)'})$ as

$$\boldsymbol{K}^{(i)}(\boldsymbol{t}^{(>i)}, \boldsymbol{t}^{(>i)'}) \approx \sum_{j_i=1}^{Y_i} \overbrace{\operatorname{circ}(\underbrace{0, \dots, 0}_{j_i - 1 \text{ times}}, \underbrace{1, \underbrace{0, \dots, 0}_{Y_i - j_i \text{ times}}}^{\boldsymbol{C}^{(i,j_i)}}) \otimes \boldsymbol{K}^{(i-1)}(t_1^{(i)}, \boldsymbol{t}^{(>i)}, t_{w^{(i)}(j_i)}^{(i)}, \boldsymbol{t}^{(>i)'})$$

we obtain that

$$\begin{split} \boldsymbol{K}^{(K)} &\approx \sum_{j_{K}=1}^{Y_{K}} \boldsymbol{C}^{(K,j_{K})} \otimes \boldsymbol{K}^{(K-1)}(t_{1}^{(K)}, t_{w^{(K)}(j_{K})}^{(K)}) \\ &\approx \sum_{j_{K}=1, j_{K-1}=1}^{Y_{K}, Y_{K-1}} \boldsymbol{C}^{(K,j_{K})} \otimes \boldsymbol{C}^{(K-1, j_{K-1})} \otimes \boldsymbol{K}^{(K-1)}(t_{1}^{(K-1)}, t_{1}^{(K)}, t_{w^{(K-1)}(j_{K-1})}^{(K-1)}, t_{w^{(K)}(j_{K})}^{(K)}) \\ &\vdots \\ &\approx \sum_{j_{K}=1, \dots, j_{1}=1}^{Y_{K}, \dots, Y_{1}} \boldsymbol{C}^{(K, j_{K})} \otimes \dots \otimes \boldsymbol{C}^{(1, j_{1})} \otimes \boldsymbol{K}^{(0)}(\underbrace{t_{1}^{(1)}, \dots, t_{1}^{(K)}}_{t_{1}}, \underbrace{t_{w^{(1)}(j_{1})}, \dots, t_{w^{(K)}(j_{K})}^{(K)}}_{t_{w^{(j)}}}) \\ &= \sum_{j_{K}, \dots, Y_{1}}^{Y_{K}, \dots, Y_{1}} \boldsymbol{C}^{(K, j_{K})} \otimes \dots \otimes \boldsymbol{C}^{(1, j_{1})} \otimes \boldsymbol{K}(\boldsymbol{t}_{1}, \boldsymbol{t}_{w^{(j)}}) \end{split}$$

$$=\sum_{j_K=1,...,j_1=1}oldsymbol{C}^{(K,j_K)}\otimes\cdots\otimesoldsymbol{C}^{(1,j_1)}\otimesoldsymbol{K}(oldsymbol{t}_1,oldsymbol{t}_{w(j)}).$$

Let $\boldsymbol{\mathcal{F}}^{(i)} = \boldsymbol{\mathcal{F}}_{Y_i} \otimes \cdots \otimes \boldsymbol{\mathcal{F}}_{Y_1} \otimes \boldsymbol{I}$. Then

$$\mathcal{F}^{(K)} \mathbf{K}^{(K)} \mathcal{F}^{(K)H} \\ \approx \sum_{j_K=1,\dots,j_1=1}^{Y_K,\dots,Y_1} (\underbrace{\mathcal{F}_{Y_K} \mathbf{C}^{(K,j_K)} \mathcal{F}_{Y_K}^H}_{\mathbf{\Lambda}^{(K,j_k)}}) \otimes \cdots \otimes (\underbrace{\mathcal{F}_{Y_1} \mathbf{C}^{(1,j_1)} \mathcal{F}_{Y_1}^H}_{\mathbf{\Lambda}^{(1,j_k)}}) \otimes \mathbf{K}(\mathbf{t}_1, \mathbf{t}_{w(j)})$$
(E.2)

where all $\Lambda^{(i,j)}$ are diagonal because all $C^{(i,j)}$ are circulant by definition. Thus $\mathcal{F}^{(K)} \mathcal{K}^{(K)} \mathcal{F}^{(K)H}$ is approximately block diagonal. Let

$$\boldsymbol{B}_{m_1,\dots,m_K} = \sum_{j_K=1,\dots,j_1=1}^{Y_K,\dots,Y_1} \Lambda_{m_K,m_K}^{(K,j_k)} \cdots \Lambda_{m_1,m_1}^{(1,j_1)} \boldsymbol{K}(\boldsymbol{t}_1, \boldsymbol{t}_{w(j)}).$$
(E.3)

Then Equation (E.2) tells us that $\boldsymbol{B}_{m_1,\dots,m_K}$ is the $d(m_1,\dots,m_K)$ 'th block on the diagonal of $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{K}^{(K)}\boldsymbol{\mathcal{F}}^{(K)H}$ where

$$d(m_1,\ldots,m_K) = (m_K - 1)(Y_{K-1}\cdots Y_1) + \ldots + (m_2 - 1)Y_1 + (m_1 - 1) + 1.$$

Finally, let $m^{(i)}(d)$ be such that $m^{(i)}[d(m_1, \ldots, m_K)] = m_i$.¹ Then K_F is approximated by

$$\boldsymbol{K}_{\boldsymbol{F}} \approx \boldsymbol{\mathcal{F}}^{(K)H} \operatorname{diag}(\boldsymbol{B}_{m^{(1)}(1),\dots,m^{(K)}(1)},\dots,\boldsymbol{B}_{m^{(1)}(Y),\dots,m^{(K)}(Y)})\boldsymbol{\mathcal{F}}^{(K)}.$$
 (E.4)

E.4 Approximating Determinants

We can use Equation (E.4) to efficiently approximate $|\mathbf{K}_F|$. We first derive an approximation of $|\mathbf{K}_F|$ and then show that this approximation can be computed efficiently.

Note that $\mathcal{F}^{(K)H}\mathcal{F}^{(K)} = I$. Thus, by Equation (E.4),

$$|\boldsymbol{K}_{\boldsymbol{F}}| \approx |\boldsymbol{\mathcal{F}}^{(K)H}||\operatorname{diag}(\boldsymbol{B}_{m^{(1)}(1),\dots,m^{(K)}(1)},\dots,\boldsymbol{B}_{m^{(1)}(Y),\dots,m^{(K)}(Y)})||\boldsymbol{\mathcal{F}}^{(K)}|$$

= $\prod_{d=1}^{Y} |\boldsymbol{B}_{m^{(1)}(d),\dots,m^{(K)}(d)}|.$ (E.5)

We now show that Equation (E.5) can be computed efficiently. To begin with, we show that all $\boldsymbol{B}_{m_1,\ldots,m_K}$ can be computed efficiently. Let $\boldsymbol{S}^{(0)}(\boldsymbol{t}) = \boldsymbol{K}(\boldsymbol{t}_1, \boldsymbol{t})$ and let

$$\boldsymbol{t}_{w(\boldsymbol{j})}^{(>i)} = \begin{bmatrix} t_{w^{(i+1)}(j_{i+1})}^{(i+1)} & \cdots & t_{w^{(K)}(j_K)}^{(K)} \end{bmatrix}^T, \\ \boldsymbol{S}_{m_1,\dots,m_i}^{(i)}(\boldsymbol{t}_{w(\boldsymbol{j})}^{(>i)}) = \sum_{j_i=1}^{Y_i} \Lambda_{m_i,m_i}^{(i,j_i)} \boldsymbol{S}_{m_1,\dots,m_{i-1}}^{(i-1)}(t_{w^{(i)}(j_i)}^{(i)}, \boldsymbol{t}_{w(\boldsymbol{j})}^{(>i)})$$
(E.6)

for $i = 1, \ldots, K$. Then one verifies that $\boldsymbol{S}_{m_1, \ldots, m_K}^{(K)} = \boldsymbol{B}_{m_1, \ldots, m_K}$. Now, fix i and fix

¹That is, let
$$m^{(i)}(d) = 1 + [\lfloor (d-1)/(Y_{i-1}\cdots Y_i) \rfloor \mod Y_i]$$
. Then
 $m^{(i)}(d)\Big|_{d=d(m_1,\dots,m_K)} = 1 + \left(\left\lfloor \frac{d-1}{Y_{i-1}\cdots Y_1} \right\rfloor \mod Y_i \right) \Big|_{d=d(m_1,\dots,m_K)}$
 $= 1 + \left(\left\lfloor (m_K - 1)(Y_K \cdots Y_i) + \dots + (m_{i+1} - 1)Y_i + (m_i - 1) + \frac{m_{i-1}-Y}{Y_{i-1}} + \dots + \frac{m_1-1}{Y_{i-1}} \right\rfloor \mod Y_i \right)$
 $= 1 + \left[(m_K - 1)(Y_K \cdots Y_i) + \dots + (m_{i+1} - 1)Y_i + m_i - 1 \right] \mod Y_i$
 $= m_i.$

 m_1, \ldots, m_{i-1} and j_{i+1}, \ldots, j_K . It holds that [Gray, 2006]

$$\Lambda_{m_{i},m_{i}}^{(i,j_{i})} = (\mathcal{F}_{Y_{i}} \mathcal{C}^{(i,j_{i})} \mathcal{F}_{Y_{i}}^{H})_{m_{i},m_{i}}$$

= $\sum_{n=1}^{Y_{i}} \mathbb{1}(j_{i}-n) \exp\left[-2\pi\sqrt{-1}\frac{(m_{i}-1)(n-1)}{Y_{i}}\right]$
= $\exp\left[-2\pi\sqrt{-1}\frac{(m_{i}-1)(j_{i}-1)}{Y_{i}}\right].$

Hence, by Equation (E.6),

$$\underbrace{\mathbf{S}_{m_1,\dots,m_i}^{(i)}(\mathbf{t}_{w(j)}^{(>i)})}_{\mathbf{A}_{m_i}} = \sum_{j_i=1}^{Y_i} \underbrace{\mathbf{S}_{m_1,\dots,m_{i-1}}^{(i-1)}(t_{w^{(i)}(j_i)}^{(i)}, \mathbf{t}_{w(j)}^{(>i)})}_{\mathbf{B}_{j_i}} \exp\left[-2\pi\sqrt{-1}\frac{(m_i-1)(j_i-1)}{Y_i}\right]$$

Note that all \mathbf{A}_{m_i} and \mathbf{B}_{j_i} are $N \times N$ matrices. Importantly, observe that $\mathbf{A}_{:,n,n'} =$ DFT $\mathbf{B}_{:,n,n'}$. Thus, the fast Fourier transform algorithm can be used to compute a single $\mathbf{A}_{:,n,n'}$ in $\mathcal{O}(Y_i \log Y_i)$ time. By doing this for every (n, n') we can compute $\mathbf{S}_{m_1,\ldots,m_i}^{(i)}(\mathbf{t}_{w(j)}^{(>i)})$ for all m_i in $\mathcal{O}(N^2Y_i \log Y_i)$ time. Hence, we can compute $\mathbf{S}_{m_1,\ldots,m_i}^{(i)}(\mathbf{t}_{w(j)}^{(>i)})$ for all m_1,\ldots,m_i and j_{i+1},\ldots,j_K in

$$\mathcal{O}(Y_1 \cdots Y_{i-1} Y_{i+1} \cdots Y_K N^2 Y_i \log Y_i) = \mathcal{O}(N^2 Y \log Y_i)$$

time, which shows that recursively applying Equation (E.6) to compute $S_{m_1,\ldots,m_K}^{(K)} = B_{m_1,\ldots,m_K}$ for all m_1,\ldots,m_K costs

$$\mathcal{O}[N^2 Y(\log Y_1 + \ldots + \log Y_K)]$$

time. This time complexity is bounded by $\mathcal{O}(N^2 Y \log Y)$.

Finally, Equation (E.5) shows that after all $\boldsymbol{B}_{m_1,\dots,m_K}$ have been computed, $|\boldsymbol{K}_F|$ can be computed in $\mathcal{O}(N^3Y)$ time (Appendix D). Therefore, the resulting complexity of computing $|\boldsymbol{K}_F|$ is $\mathcal{O}(N^2Y \log Y + N^3Y)$. The complete procedure to compute $|\boldsymbol{K}_F|$ is outlined in Algorithm 1.

Algorithm 1 Efficient approximation of the determinant of a multi-output stationary

kernel matrix. Runs in $\mathcal{O}(N^2 Y \log Y + N^3 Y)$ time. 1: function DETERMINANT $(\mathbf{K}, \{t_1^{(1)}, \ldots, t_{Y_1}^{(1)}\} \times \cdots \times \{t_1^{(K)}, \ldots, t_{Y_K}^{(K)}\})$ 2: for $i = 1, \ldots, K$ do for $m_1 = 1, \ldots, Y_1$ to $m_{i-1} = 1, \ldots, Y_{i-1}$ and 3: $j_{i+1} = 1, \ldots, Y_{i+1}$ to $j_K = 1, \ldots, Y_K$ do for n = 1, ..., N and n' = 1, ..., N do 4: $A_{:.n.n'} \leftarrow \mathrm{DFT} \, B_{:.n.n'}$ 5: \triangleright Via fast Fourier transform algorithm $\triangleright \mathbf{S}^{(0)}(t) = \mathbf{K}(t_1, t), \ \mathbf{A}_{m_i} = \mathbf{S}^{(i)}_{m_1, \dots, m_i}(t_{w(j)}^{(>i)}) \ \text{and} \ \mathbf{B}_{j_i} = \mathbf{S}^{(i-1)}_{m_1, \dots, m_{i-1}}(t_{w(i)(j_i)}^{(i)}, t_{w(j)}^{(>i)}) \ \text{where} \ w^{(i)}(j_i) = \min\{j, Y_i - j_i + 2\}$ for $m_1 = 1, \ldots, Y_1$ to $m_K = 1, \ldots, Y_K$ do 6: Compute $|\mathbf{S}_{m_1,...,m_K}^{(K)}|$ return $\prod_{d=1}^{Y} |\mathbf{S}_{m^{(1)}(d),...,m^{(K)}(d)}^{(K)}|
ightarrow m^{(i)}(d) = 1 + [\lfloor (d-1)/(Y_{i-1}\cdots Y_i) \rfloor \mod Y_i]$ 7: 8:

Approximating Products Involving an E.5Inverse

We can also use Equation (E.4) to efficiently approximate a product of the form P = $\boldsymbol{U}^T \boldsymbol{K}_{\boldsymbol{F}}^{-1} \boldsymbol{V}$ for some $NY \times P$ matrix \boldsymbol{U} and $NY \times Q$ matrix \boldsymbol{V} . We first derive an approximation of P and then show that this approximation can be computed efficiently.

Consider $P_{p,q}$. Recall that by Equation (E.4), $\mathcal{F}^{(K)}\mathcal{F}^{(K)H}$ is approximately block diagonal with blocks $\boldsymbol{B}_{m_1,\ldots,m_K}$; hence $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{K}^{-(K)}\boldsymbol{\mathcal{F}}^{(K)H}$ is approximately block diagonal with blocks $\boldsymbol{B}_{m_1,\ldots,m_K}^{-1}$. Thus

$$P_{p,q} = \underbrace{(\boldsymbol{U}_{:,p})^{T}}_{\boldsymbol{u}^{T}} \boldsymbol{K}^{-(K)} \underbrace{\boldsymbol{V}_{:,q}}_{\boldsymbol{v}}$$
$$= (\boldsymbol{\mathcal{F}}^{(K)} \boldsymbol{u})^{H} (\boldsymbol{\mathcal{F}}^{(K)} \boldsymbol{K}^{-(K)} \boldsymbol{\mathcal{F}}^{(K)H}) (\boldsymbol{\mathcal{F}}^{(K)} \boldsymbol{v})$$
$$\approx \sum_{d=1}^{Y} \begin{bmatrix} u_{(d-1)N+1} & \cdots & u_{dN} \end{bmatrix} \boldsymbol{B}_{m^{(1)}(d),\dots,m^{(K)}(d)}^{-1} \begin{bmatrix} v_{(d-1)N+1} & \cdots & v_{dN} \end{bmatrix}^{T}. \quad (E.7)$$

We now show that Equation (E.7) can be computed efficiently. First, we show that all $\boldsymbol{B}_{m_1,\dots,m_K}^{-1}$ can be computed efficiently. Appendix E.4 showed that all $\boldsymbol{B}_{m_1,\dots,m_K}$ can be computed in $\mathcal{O}(N^2Y\log Y)$ time. Therefore all $\boldsymbol{B}_{m_1,\dots,m_K}^{-1}$ can be computed in

 $\mathcal{O}(N^2 Y \log Y + N^3 Y)$ time.

Second, we show that $\mathcal{F}^{(K)}\boldsymbol{u}$ and thereby $\mathcal{F}^{(K)}\boldsymbol{v}$ from Equation (E.7) can be computed efficiently. It holds that

$$\mathcal{F}^{(K)}\boldsymbol{u} = (\mathcal{F}_{Y_K} \otimes \cdots \otimes \mathcal{F}_{Y_1} \otimes \boldsymbol{I})\boldsymbol{u} \\ = \underbrace{(\mathcal{F}_{Y_K} \otimes \boldsymbol{I}_{Y_{K-1}} \otimes \cdots \otimes \boldsymbol{I}_{Y_1} \otimes \boldsymbol{I}_N)}_{T_{Y_K}} \cdots \underbrace{(\boldsymbol{I}_{Y_K} \otimes \cdots \otimes \boldsymbol{I}_{Y_2} \otimes \mathcal{F}_{Y_1} \otimes \boldsymbol{I}_N)}_{T_{Y_1}} \boldsymbol{u}.$$

Observe that each T_{Y_i} is of the form $I_{Z_i^{(1)}} \otimes \mathcal{F}_{Y_i} \otimes I_{Z_i^{(2)}}$ where $Z_i^{(1)} = Y_K \cdots Y_{i+1}$ and $Z_i^{(2)} = Y_{i-1} \cdots Y_1 N$. Inspection of

$$(\boldsymbol{\mathcal{F}}_{Y_i} \otimes \boldsymbol{I}_{Z_i^{(2)}}) \boldsymbol{x} = egin{bmatrix} \mathcal{F}_{Y_i,1,1} \boldsymbol{I}_{Z_i^{(2)}} & \cdots & \mathcal{F}_{Y_i,1,Y_i} \boldsymbol{I}_{Z_i^{(2)}} \ dots & \ddots & dots \ \mathcal{F}_{Y_i,Y_i,1} \boldsymbol{I}_{Z_i^{(2)}} & \cdots & \mathcal{F}_{Y_i,Y_i,Y_i} \boldsymbol{I}_{Z_i^{(2)}} \end{bmatrix} \boldsymbol{x} = \boldsymbol{y}$$

shows that

$$\boldsymbol{\mathcal{F}}_{Y_i}\underbrace{\begin{bmatrix} x_z & x_{z+Z_i^{(2)}} & \cdots & x_{z+(Y_i-1)Z_i^{(2)}} \end{bmatrix}^T}_{\text{every } Z_i^{(2)}, \text{th element of } \boldsymbol{x}, \text{ starting at } z} = \underbrace{\begin{bmatrix} y_z & y_{z+Z_i^{(2)}} & \cdots & y_{z+(Y_i-1)Z_i^{(2)}} \end{bmatrix}^T}_{\text{every } Z_i^{(2)}, \text{th element of } \boldsymbol{y}, \text{ starting at } z}$$
(E.8)

for $z = 1, \ldots, Z_i^{(2)}$, which completely specifies \boldsymbol{y} . The fast Fourier transform algorithm can be used to compute Equation (E.8) in $\mathcal{O}(Y_i \log Y_i)$ time, which means that \boldsymbol{y} can be computed in $\mathcal{O}(Z_i^{(2)}Y_i \log Y_i)$ time. Similarly, by directly applying the definition of the Kronecker product we see that $(\boldsymbol{I}_{Z_i^{(1)}} \otimes \boldsymbol{\mathcal{F}}_{Y_i} \otimes \boldsymbol{I}_{Z_i^{(2)}})\boldsymbol{x} = \boldsymbol{y}$ yields the \boldsymbol{y} such that

$$(\boldsymbol{\mathcal{F}}_{Y_{i}} \otimes \boldsymbol{I}_{Z_{i}^{(2)}}) \underbrace{ \begin{bmatrix} x_{(z-1)Z_{i}^{(2)}Y_{i}+1} & x_{(z-1)Z_{i}^{(2)}Y_{i}+2} & \cdots & x_{zZ_{i}^{(2)}Y_{i}} \end{bmatrix}^{T}}_{z'\text{th consecutive group of } Z_{i}^{(2)}Y_{i} \text{ consecutive elements of } x} \\ = \underbrace{ \begin{bmatrix} y_{(z-1)Z_{i}^{(2)}Y_{i}+1} & y_{(z-1)Z_{i}^{(2)}Y_{i}+2} & \cdots & y_{zZ_{i}^{(2)}Y_{i}} \end{bmatrix}^{T}}_{z'\text{th consecutive group of } Z_{i}^{(2)}X_{i} \text{ consecutive elements of } x}$$

z'th consecutive group of $Z_i^{(2)} Y_i$ consecutive elements of y

for $z = 1, \ldots, Z_{Y_i}^{(1)}$, which again completely specifies \boldsymbol{y} . Hence multiplication by any \boldsymbol{T}_{Y_i} can be done in $\mathcal{O}(Z_{Y_i}^{(1)}Z_{Y_i}^{(2)}Y_i\log Y_i) = \mathcal{O}(NY\log Y_i)$ time. Thus $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{u}$ and $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{v}$ can be computed in

$$\mathcal{O}[NY(\log Y_1 + \ldots + \log Y_K)].$$

time. This complexity is bounded by $\mathcal{O}(NY \log Y)$.

Finally, Equation (E.7) shows that after all $\boldsymbol{B}_{m_1,\dots,m_K}^{-1}$, $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{u}$ and $\boldsymbol{\mathcal{F}}^{(K)}\boldsymbol{v}$ have been computed, $P_{p,q}$ can be computed in $\mathcal{O}(N^2Y)$ time. Therefore, since all $\boldsymbol{B}_{m_1,\dots,m_K}^{-1}$ have to be computed only once, the resulting complexity of computing \boldsymbol{P} is $\mathcal{O}[N^2Y\log Y + N^3Y + PQ(NY\log Y + N^2Y)]$. The complete procedure to compute \boldsymbol{P} is outlined in Algorithm 2.

E.6 Conclusion

We have derived procedures to efficiently approximate the determinant of a stationary multi-output kernel matrix and a product involving an inverse stationary multi-output kernel matrix. These procedures have respectively time complexities $\mathcal{O}(N^2Y \log Y + N^3Y)$ and $\mathcal{O}[N^2Y \log Y + N^3Y + PQ(NY \log Y + N^2Y)]$. Importantly, these complexities scale favourably in Y.

Algorithm 2 Efficient approximation of a product involving an inverse stationary multioutput kernel matrix. Runs in $\mathcal{O}[N^2Y \log Y + N^3Y + PQ(NY \log Y + N^2Y)]$ time.

F | Exponentiated Quadratic Forms

F.1 Introduction

This chapter develops notation that makes working with exponentiated quadratic forms more convenient.

F.2 General Form

Let $\boldsymbol{x} \in \mathbb{R}^N$. Denote an exponentiated quadratic form

$$p(\boldsymbol{x}) = C \exp\left(-\frac{1}{2}\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} + \boldsymbol{x}^T \boldsymbol{b} + c\right)$$

by $(C, \mathbf{A}, \mathbf{b}, c)$. We call \mathbf{x} the composite vector of variables, or composite vector in short.

Lemma 4 (Product Identity (General Form)). Let p_1 and p_2 be two exponentiated quadratic forms. Then

$$p_1p_2 = (C_1, A_1, b_1, c_1)(C_2, A_2, b_2, c_2) = (C_1C_2, A_1 + A_2, b_1 + b_2, c_1 + c_2)$$

Proof. Direct computation yields that

$$p_1 p_2 = C_1 C_2 \exp\left(-\frac{1}{2} \boldsymbol{x}^T \boldsymbol{A}_1 \boldsymbol{x} + \boldsymbol{x}^T \boldsymbol{b}_1 - \frac{1}{2} \boldsymbol{x}^T \boldsymbol{A}_2 \boldsymbol{x} + \boldsymbol{x}^T \boldsymbol{b}_2 + c\right)$$

= $C_1 C_2 \exp\left[-\frac{1}{2} \boldsymbol{x}^T (\boldsymbol{A}_1 + \boldsymbol{A}_2) \boldsymbol{x} + \boldsymbol{x}^T (\boldsymbol{b}_1 + \boldsymbol{b}_2) + (c_1 + c_2)\right]$
= $(C_1 C_2, \boldsymbol{A}_1 + \boldsymbol{A}_2, \boldsymbol{b}_1 + \boldsymbol{b}_2, c_1 + c_2).$

Lemma 5 (Integration Identity (General Form)). Let the composite vector be $\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix}^T$ and consider

$$p = (C, \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix}, \begin{bmatrix} \boldsymbol{b}_1 \\ \boldsymbol{b}_2 \end{bmatrix}^T, c).$$

Denote integration over the subset of variables x_1 by $I_{x_1}(p)$:

$$I_{\boldsymbol{x}_1}(p)(\boldsymbol{x}_2) = \int (C, \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix}, \begin{bmatrix} \boldsymbol{b}_1 \\ \boldsymbol{b}_2 \end{bmatrix}, c) \left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix} \right) \mathrm{d}\boldsymbol{x}_1.$$

Then

$$I_{\boldsymbol{x}_1}(p) = \left(C \frac{(2\pi)^{N/2}}{|\boldsymbol{A}_{1,1}|^{1/2}}, \boldsymbol{A}_{2,2} - \boldsymbol{A}_{2,1} \boldsymbol{A}_{1,1}^{-1} \boldsymbol{A}_{1,2}, \boldsymbol{b}_2 - \boldsymbol{A}_{2,1} \boldsymbol{A}_{1,1}^{-1} \boldsymbol{b}_1, c + \frac{1}{2} \boldsymbol{b}_1^T \boldsymbol{A}_{1,1}^{-1} \boldsymbol{b}_1 \right).$$

Proof. An exponentiated quadratic form can be integrated as follows:

$$\begin{split} &\int_{\mathbb{R}^N} \exp\left(-\frac{1}{2}\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} + \boldsymbol{b}^T \boldsymbol{x} + c\right) \mathrm{d} \boldsymbol{x} \\ &= \exp\left(\frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} + c\right) \int_{\mathbb{R}^N} \exp\left(-\frac{1}{2}\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} + \boldsymbol{b}^T \boldsymbol{x} - \frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) \mathrm{d} \boldsymbol{x} \\ &= \frac{(2\pi)^{N/2}}{|\boldsymbol{A}|^{1/2}} \exp\left(\frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} + c\right) \int_{\mathbb{R}^N} \frac{(2\pi)^{-N/2}}{|\boldsymbol{A}|^{-1/2}} \exp\left(-\frac{1}{2}\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} + \boldsymbol{b}^T \boldsymbol{x} - \frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) \mathrm{d} \boldsymbol{x} \\ &= \frac{(2\pi)^{N/2}}{|\boldsymbol{A}|^{1/2}} \exp\left(\frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} + c\right) \int_{\mathbb{R}^N} \mathcal{N}(\boldsymbol{x}; \boldsymbol{A}^{-1} \boldsymbol{b}, \boldsymbol{A}) \mathrm{d} \boldsymbol{x} \\ &= \frac{(2\pi)^{N/2}}{|\boldsymbol{A}|^{1/2}} \exp\left(\frac{1}{2}\boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} + c\right). \end{split}$$

Now,

$$-\frac{1}{2}\begin{bmatrix}\boldsymbol{x}_{1}\\\boldsymbol{x}_{2}\end{bmatrix}^{T}\begin{bmatrix}\boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2}\\\boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2}\end{bmatrix}\begin{bmatrix}\boldsymbol{x}_{1}\\\boldsymbol{x}_{2}\end{bmatrix} + \begin{bmatrix}\boldsymbol{x}_{1}\\\boldsymbol{x}_{2}\end{bmatrix}^{T}\begin{bmatrix}\boldsymbol{b}_{1}\\\boldsymbol{b}_{2}\end{bmatrix} + c$$

$$= -\frac{1}{2}\boldsymbol{x}_{1}^{T}\boldsymbol{A}_{1,1}\boldsymbol{x}_{1} + \boldsymbol{x}_{1}^{T}\begin{bmatrix}\boldsymbol{b}_{1} - \underbrace{\frac{1}{2}(\boldsymbol{A}_{1,2} + \boldsymbol{A}_{2,1}^{T})}_{\boldsymbol{A}_{1,2}}\boldsymbol{x}_{2}\end{bmatrix} + (\boldsymbol{x}_{2}^{T}\boldsymbol{b}_{2} - \frac{1}{2}\boldsymbol{x}_{2}^{T}\boldsymbol{A}_{2,2}\boldsymbol{x}_{2} + c).$$

F.2 General Form

Thus

$$\begin{split} I_{\boldsymbol{x}_{1}}(p)(\boldsymbol{x}_{2}) &= C \frac{(2\pi)^{N/2}}{|\boldsymbol{A}_{1,1}|^{1/2}} \exp\left[\frac{1}{2}(\boldsymbol{b}_{1} - \boldsymbol{A}_{1,2}\boldsymbol{x}_{2})^{T}\boldsymbol{A}_{1,1}^{-1}(\boldsymbol{b}_{1} - \boldsymbol{A}_{1,2}\boldsymbol{x}_{2}) \right. \\ &+ (\boldsymbol{x}_{2}^{T}\boldsymbol{b}_{2} - \frac{1}{2}\boldsymbol{x}_{2}^{T}\boldsymbol{A}_{2,2}\boldsymbol{x}_{2} + c)\right] \\ &= C \frac{(2\pi)^{N/2}}{|\boldsymbol{A}_{1,1}|^{1/2}} \exp\left[-\frac{1}{2}\boldsymbol{x}_{2}^{T}(\boldsymbol{A}_{2,2} - \boldsymbol{A}_{2,1}\boldsymbol{A}_{1,1}^{-1}\boldsymbol{A}_{1,2})\boldsymbol{x}_{2} \right. \\ &+ \boldsymbol{x}_{2}^{T}(\boldsymbol{b}_{2} - \boldsymbol{A}_{2,1}\boldsymbol{A}_{1,1}^{-1}\boldsymbol{b}_{1}) + (c + \frac{1}{2}\boldsymbol{b}_{1}^{T}\boldsymbol{A}_{1,1}^{-1}\boldsymbol{b}_{1})\right] \end{split}$$

and so the result follows.

Lemma 6 (Linear Expansion Property (General Form)). Let the composite vector be $\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix}^T$ and consider

$$p_0 = (C, \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} \end{bmatrix}, \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}, c).$$

Then the exponentiated quadratic forms

$$p_1\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_3^T \end{bmatrix}^T\right) = p_0\left(\begin{bmatrix} (\boldsymbol{x}_1 + \boldsymbol{B}\boldsymbol{x}_3)^T & \boldsymbol{x}_2^T \end{bmatrix}^T\right),$$
$$p_2\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_3^T \end{bmatrix}^T\right) = p_0\left(\begin{bmatrix} \boldsymbol{x}_1^T & (\boldsymbol{x}_2 + \boldsymbol{B}\boldsymbol{x}_3)^T \end{bmatrix}^T\right)$$

are given by

$$p_i = (C, \begin{bmatrix} \mathbf{A} & \mathbf{A}_{1,i}\mathbf{B} \\ \mathbf{A} & \mathbf{A}_{2,i}\mathbf{B} \\ \mathbf{B}^T\mathbf{A}_{i,1} & \mathbf{B}^T\mathbf{A}_{i,2} & \mathbf{B}^T\mathbf{A}_{i,i}\mathbf{B} \end{bmatrix}, \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{B}^T\mathbf{b}_i \end{bmatrix}, c)$$

for respectively i = 1, 2.

Proof. The exponent of p_1 can be simplified as follows:

$$\begin{split} &-\frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} + \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} + \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{x}_{2} \end{bmatrix} + \begin{bmatrix} \boldsymbol{x}_{1} + \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \end{bmatrix} + c \\ &= -\frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix} \\ &+ \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \end{bmatrix} + \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \end{bmatrix} + \begin{bmatrix} \boldsymbol{B}\boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \end{bmatrix} + c \\ &= -\frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \boldsymbol{0} & \boldsymbol{B}^{T} \boldsymbol{A}_{1,2} \\ \boldsymbol{0} & \boldsymbol{B}^{T} \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} \\ &- \frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{B}^{T} \boldsymbol{A}_{1,2} \\ \boldsymbol{0} & \boldsymbol{B}^{T} \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} \\ &- \frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} \boldsymbol{B} & \boldsymbol{0} \\ \boldsymbol{A}_{2,1} \boldsymbol{B} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix} - \frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{B}^{T} \boldsymbol{A}_{1,1} \\ \boldsymbol{B} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{3} \\ \boldsymbol{0} \end{bmatrix} \\ &+ \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} & \boldsymbol{A}_{2,1} \boldsymbol{B} \\ \boldsymbol{0} \end{bmatrix} + c \\ &= -\frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{3} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} & \boldsymbol{A}_{1,1} \boldsymbol{B} \\ \boldsymbol{B}^{T} \boldsymbol{A}_{1,1} & \boldsymbol{B}^{T} \boldsymbol{A}_{1,2} & \boldsymbol{B}^{T} \boldsymbol{A}_{1,1} \boldsymbol{B} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{3} \end{bmatrix} + \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{3} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \end{bmatrix} + c. \\ &= -\frac{1}{2} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{3} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} & \boldsymbol{A}_{1,1} \boldsymbol{B} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix} \begin{bmatrix} \boldsymbol{A}_{1} & \boldsymbol{A}_{1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2} \end{bmatrix} + c. \\ &= -\frac{1}{2} \begin{bmatrix} \boldsymbol$$

Thus

$$p_1 = (C, \begin{bmatrix} \mathbf{A} & \mathbf{A}_{1,1}\mathbf{B} \\ \mathbf{A} & \mathbf{A}_{2,1}\mathbf{B} \\ \mathbf{B}^T\mathbf{A}_{1,1} & \mathbf{B}^T\mathbf{A}_{1,2} & \mathbf{B}^T\mathbf{A}_{1,1}\mathbf{B} \end{bmatrix}, \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{B}^T\mathbf{b}_1 \end{bmatrix}, c).$$

The case p_2 follows similarly.

Lemma 7 (Compression Property (General Form)). Let the composite vector be $\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_2^T \end{bmatrix}^T$ and consider

$$p_{1} = (C, \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} & \mathbf{A}_{1,3} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} & \mathbf{A}_{2,3} \\ \mathbf{A}_{3,1} & \mathbf{A}_{3,2} & \mathbf{A}_{3,3} \end{bmatrix}, \begin{bmatrix} \mathbf{b}_{1} \\ \mathbf{b}_{2} \\ \mathbf{b}_{3} \end{bmatrix}, c).$$

F.3 Kronecker-Structured Form

Then the exponentiated quadratic form

$$p_2\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T\end{bmatrix}^T\right) = p_1\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_2^T\end{bmatrix}^T\right)$$

is given by

$$p_{2} = (C, \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} + \mathbf{A}_{1,3} \\ \mathbf{A}_{2,1} + \mathbf{A}_{3,1} & \mathbf{A}_{2,2} + \mathbf{A}_{2,3} + \mathbf{A}_{3,2} + \mathbf{A}_{3,3} \end{bmatrix}, \begin{bmatrix} \mathbf{b}_{1} \\ \mathbf{b}_{2} + \mathbf{b}_{3} \end{bmatrix}, c).$$

Proof. The exponent of p_1 can be simplified as follows:

$$\begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} & \boldsymbol{A}_{1,3} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} & \boldsymbol{A}_{2,3} \\ \boldsymbol{A}_{3,1} & \boldsymbol{A}_{3,2} & \boldsymbol{A}_{3,3} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} \\ \boldsymbol{b}_{3} \end{bmatrix} + c$$

$$= \boldsymbol{x}_{1}^{T} \boldsymbol{A}_{1,1} \boldsymbol{x} + \boldsymbol{x}_{1}^{T} (\boldsymbol{A}_{1,2} + \boldsymbol{A}_{1,3}) \boldsymbol{x}_{2} + \boldsymbol{x}_{2}^{T} (\boldsymbol{A}_{2,1} + \boldsymbol{A}_{2,3}) \boldsymbol{x}_{1}$$

$$+ \boldsymbol{x}_{2}^{T} (\boldsymbol{A}_{2,2} + \boldsymbol{A}_{2,3} + \boldsymbol{A}_{3,2} + \boldsymbol{A}_{3,3}) \boldsymbol{x}_{2} + \boldsymbol{x}_{1}^{T} \boldsymbol{b}_{1} + \boldsymbol{x}_{2}^{T} (\boldsymbol{b}_{2} + \boldsymbol{b}_{3}) + c$$

$$= \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} + \boldsymbol{A}_{1,3} \\ \boldsymbol{A}_{2,1} + \boldsymbol{A}_{3,1} & \boldsymbol{A}_{2,2} + \boldsymbol{A}_{2,3} + \boldsymbol{A}_{3,2} + \boldsymbol{A}_{3,3} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix} + \begin{bmatrix} \boldsymbol{x}_{1} \\ \boldsymbol{x}_{2} \end{bmatrix}^{T} \begin{bmatrix} \boldsymbol{b}_{1} \\ \boldsymbol{b}_{2} + \boldsymbol{b}_{3} \end{bmatrix} + c.$$

Hence the result follows.

F.3 Kronecker-Structured Form

Let the composite vector be $\begin{bmatrix} \boldsymbol{x}_1^T & \cdots & \boldsymbol{x}_M^T \end{bmatrix}^T \in \mathbb{R}^{MN}$. Denote a exponentiated quadratic form of the form $(C, \boldsymbol{A} \otimes \boldsymbol{I}_N, \boldsymbol{0}, 0)$ by (C, \boldsymbol{A}) . This is a *Kronecker-structured* exponentiated quadratic form.

Lemma 8 (Product Identity (Kronecker-Structured Form)). Let p_1 and p_2 be two Kronecker-structured exponentiated quadratic forms. Then

$$p_1p_2 = (C_1, A_1)(C_2, A_2) = (C_1C_2, A_1 + A_2).$$

Proof. Follows from Lemma 4 and bilinearity of the Kronecker product. \Box

Lemma 9 (Integration Identity (Kronecker-Structured Form)). Let the composite vector

be $\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix}^T$ where $\boldsymbol{x}_1 \in \mathbb{R}^{KN}$ and $\boldsymbol{x}_2 \in \mathbb{R}^{MN}$. Consider $p = (C, \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix}).$

Denote integration over the subset of variables x_1 by $I_{x_1}(p)$:

$$I_{\boldsymbol{x}_1}(p)(\boldsymbol{x}_2) = \int (C, \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix}, \begin{bmatrix} \boldsymbol{b}_1 \\ \boldsymbol{b}_2 \end{bmatrix}, c) \left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix}^T \right) \mathrm{d}\boldsymbol{x}_1.$$

Then

$$I_{\boldsymbol{x}_1}(p) = \left(C \frac{(2\pi)^{KN/2}}{|\boldsymbol{A}_{1,1}|^{N/2}}, \boldsymbol{A}_{2,2} - \boldsymbol{A}_{2,1} \boldsymbol{A}_{1,1}^{-1} \boldsymbol{A}_{1,2} \right).$$

Proof. Direct computation yields that

$$I_{x_{1}}[(C, \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix})] = I_{x_{1}}[(C, \begin{bmatrix} A_{1,1} \otimes I & A_{1,2} \otimes I \\ A_{2,1} \otimes I & A_{2,2} \otimes I \end{bmatrix}, \mathbf{0}, \mathbf{0})]$$

$$= \left(C \frac{(2\pi)^{KN/2}}{|A_{1,1} \otimes I|^{1/2}}, A_{2,2} \otimes I - (A_{2,1} \otimes I)(A_{1,1} \otimes I)^{-1}(A_{1,2} \otimes I), \mathbf{0}, \mathbf{0})\right)$$

$$= \left(C \frac{(2\pi)^{KN/2}}{|A_{1,1}|^{N/2}}, (A_{2,2} - A_{2,1}A_{1,1}^{-1}A_{1,2}) \otimes I, \mathbf{0}, \mathbf{0})\right)$$

$$= \left(C \frac{(2\pi)^{KN/2}}{|A_{1,1}|^{N/2}}, A_{2,2} - A_{2,1}A_{1,1}^{-1}A_{1,2}\right).$$

Lemma 10 (Linear Expansion Property (Kronecker-Structured Form)). Let the composite vector be $\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T \end{bmatrix}^T$ and consider

$$p_0 = (C, \begin{bmatrix} \boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2} \\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2} \end{bmatrix}).$$

F.3 Kronecker-Structured Form

Then the exponentiated quadratic forms

$$p_1\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_3^T \end{bmatrix}^T\right) = p_0\left(\begin{bmatrix} (\boldsymbol{x}_1 + \boldsymbol{B}\boldsymbol{x}_3)^T & \boldsymbol{x}_2^T \end{bmatrix}^T\right),$$
$$p_2\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_3^T \end{bmatrix}^T\right) = p_0\left(\begin{bmatrix} \boldsymbol{x}_1^T & (\boldsymbol{x}_2 + \boldsymbol{B}\boldsymbol{x}_3)^T \end{bmatrix}^T\right)$$

are given by

$$p_i = (C, \begin{bmatrix} \mathbf{A} & \mathbf{A}_{1,i}\mathbf{B} \\ \mathbf{A} & \mathbf{A}_{2,i}\mathbf{B} \\ \mathbf{B}^T \mathbf{A}_{i,1} & \mathbf{B}^T \mathbf{A}_{i,2} & \mathbf{B}^T \mathbf{A}_{i,i}\mathbf{B} \end{bmatrix})$$

for respectively i = 1, 2.

Proof. Follows from Lemma 6 and bilinearity of the Kronecker product.

Lemma 11 (Compression Property (Kronecker-Structured Form)). Let the composite vector be $\begin{bmatrix} \mathbf{x}_1^T & \mathbf{x}_2^T & \mathbf{x}_2^T \end{bmatrix}^T$ and consider

$$p_1 = (C, \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} & \mathbf{A}_{1,3} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} & \mathbf{A}_{2,3} \\ \mathbf{A}_{3,1} & \mathbf{A}_{3,2} & \mathbf{A}_{3,3} \end{bmatrix}).$$

Then the exponentiated quadratic form

$$p_2\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T\end{bmatrix}^T\right) = p_1\left(\begin{bmatrix} \boldsymbol{x}_1^T & \boldsymbol{x}_2^T & \boldsymbol{x}_2^T\end{bmatrix}^T\right)$$

is given by

$$p_2 = (C, \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} + \mathbf{A}_{1,3} \\ \mathbf{A}_{2,1} + \mathbf{A}_{3,1} & \mathbf{A}_{2,2} + \mathbf{A}_{2,3} + \mathbf{A}_{3,2} + \mathbf{A}_{3,3} \end{bmatrix}).$$

Proof. Follows from Lemma 7 and bilinearity of the Kronecker product.

Finally, we prove one additional identity.

Lemma 12 (Direct Integration Identity (Kronecker-Structured Form)). Let the com-

Exponentiated Quadratic Forms

posite vector be
$$\begin{bmatrix} \mathbf{x}_{1}^{T} & \mathbf{x}_{2}^{T} \end{bmatrix}^{T}$$
 where $\mathbf{x}_{1} \in \mathbb{R}^{KN}$ and $\mathbf{x}_{2} = \begin{bmatrix} \mathbf{y}_{1}^{T} & \cdots & \mathbf{y}_{M}^{T} \end{bmatrix}^{T} \in \mathbb{R}^{MN}$. Then
 $I_{\mathbf{x}_{1}}[(C, 2 \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} \end{bmatrix})](\mathbf{x}_{2}) = C \frac{\pi^{KN/2}}{|\mathbf{A}_{1,1}|^{N/2}} \exp \left\{ -\sum_{m_{1}=1,m_{2}=1}^{M,M} A_{2,2,m_{1},m_{2}} \mathbf{y}_{m_{1}}^{T} \mathbf{y}_{m_{2}} + \sum_{k_{1}=1,k_{2}=1}^{K,K} A_{1,1,k_{1},k_{2}}^{-1} \left(\sum_{m=1}^{M} A_{2,1,k_{1},m} \mathbf{y}_{m} \right)^{T} \left(\sum_{m=1}^{M} A_{2,1,k_{2},m} \mathbf{y}_{m} \right) \right\}.$

Proof. Direct computation yields that

$$\begin{split} I_{\boldsymbol{x}_{1}}[(C, 2\begin{bmatrix}\boldsymbol{A}_{1,1} & \boldsymbol{A}_{1,2}\\ \boldsymbol{A}_{2,1} & \boldsymbol{A}_{2,2}\end{bmatrix})] &= \left(C\frac{\pi^{KN/2}}{|\boldsymbol{A}_{1,1}|^{N/2}}, 2(\boldsymbol{A}_{2,2} - \boldsymbol{A}_{2,1}\boldsymbol{A}_{1,1}^{-1}\boldsymbol{A}_{1,2})\right) \\ &= C\frac{\pi^{KN/2}}{|\boldsymbol{A}_{1,1}|^{N/2}} \exp\{-\boldsymbol{x}_{2}^{T}[(\boldsymbol{A}_{2,2} - \boldsymbol{A}_{2,1}\boldsymbol{A}_{1,1}^{-1}\boldsymbol{A}_{1,2}) \otimes \boldsymbol{I}_{N}]\boldsymbol{x}_{2}\} \\ &= C\frac{\pi^{KN/2}}{|\boldsymbol{A}_{1,1}|^{N/2}} \exp\{-\boldsymbol{x}_{2}^{T}(\boldsymbol{A}_{2,2} \otimes \boldsymbol{I}_{N})\boldsymbol{x}_{2} \\ &+ \boldsymbol{x}_{2}^{T}(\boldsymbol{A}_{2,1} \otimes \boldsymbol{I}_{N})(\boldsymbol{A}_{1,1}^{-1} \otimes \boldsymbol{I}_{N}) \\ &[(\boldsymbol{A}_{1,2} \otimes \boldsymbol{I}_{N})\boldsymbol{x}_{2}]\}. \end{split}$$

Hence the result follows.

F.4 Conclusion

We have developed notation that makes working with exponentiated quadratic forms more convenient. Importantly, the notation reduces manipulation of exponentiated quadratic forms to relatively simple matrix operations.

G | Roots of Kernels

First, using the notation and identities from Appendix F with composite vector $\begin{bmatrix} t^T & \tau^T \end{bmatrix}^T$, any kernel of exponentiated quadratic form has a root:

$$R\left[\left(\sigma\frac{|2\boldsymbol{A}|^{1/4}}{\pi^{K/4}}, 2\boldsymbol{A}, \boldsymbol{0}, 0\right)\right] * \left(\sigma\frac{|2\boldsymbol{A}|^{1/4}}{\pi^{K/4}}, 2\boldsymbol{A}, \boldsymbol{0}, 0\right)$$

$$= I_{\tau}\left[\left(\sigma\frac{|2\boldsymbol{A}|^{1/4}}{\pi^{K/4}}, 2\begin{bmatrix}\boldsymbol{A} & -\boldsymbol{A}\\ -\boldsymbol{A} & \boldsymbol{A}\end{bmatrix}, \boldsymbol{0}, 0\right)\left(\sigma\frac{|2\boldsymbol{A}|^{1/4}}{\pi^{K/4}}, 2\begin{bmatrix}\boldsymbol{0} & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{A}\end{bmatrix}, \boldsymbol{0}, 0\right)\right]$$

$$= I_{\tau}\left[\left(\sigma^{2}\frac{|2\boldsymbol{A}|^{1/2}}{\pi^{K/2}}, 2\begin{bmatrix}\boldsymbol{A} & -\boldsymbol{A}\\ -\boldsymbol{A} & 2\boldsymbol{A}\end{bmatrix}, \boldsymbol{0}, 0\right)\right]$$

$$= \left(\sigma^{2}\frac{|2\boldsymbol{A}|^{1/2}}{\pi^{K/2}}\frac{\pi^{K/2}}{|2\boldsymbol{A}|^{1/2}}, 2[\boldsymbol{A} - \boldsymbol{A}(2\boldsymbol{A})^{-1}\boldsymbol{A}], \boldsymbol{0}, 0\right)$$

$$= (\sigma^{2}, \boldsymbol{A}, \boldsymbol{0}, 0).$$

Second, consider the diagonal multi-output exponentiated-quadratic kernel. Its diagonal entries are of exponentiated quadratic form and thus have roots. Now, for any diagonal H it holds that

$$R(\mathbf{H}) * \mathbf{H}^{T} = \text{diag}[R(H_{1,1}) * H_{1,1}, \dots, R(H_{N,N})].$$

Thus the diagonal entries having roots implies that the diagonal multi-output exponentiatedquadratic kernel has a root. Specifically, if

$$\mathcal{K}(\boldsymbol{t}_1, \boldsymbol{t}_2) = \sigma^2 \exp(-\gamma \|\boldsymbol{t}_1 - \boldsymbol{t}_2\|^2) \boldsymbol{I},$$

then $\mathcal{K}(t_1, t_2) = [R(\mathbf{R}) * \mathbf{R}](t_2 - t_1)$ where

$$\boldsymbol{R}(\boldsymbol{t}) = \sigma \frac{(4\gamma)^{K/4}}{\pi^{K/4}} \exp(-2\gamma \|\boldsymbol{t}\|^2) \boldsymbol{I}.$$

Finally, the white noise kernel is its own root:

$$[R(\delta \boldsymbol{I}) * \delta \boldsymbol{I}](\boldsymbol{t}) = \int_{\mathbb{R}^K} \delta(\boldsymbol{t}' - \boldsymbol{t}) \boldsymbol{I} \delta(\boldsymbol{t}') \boldsymbol{I} \, \mathrm{d} \boldsymbol{t}' = \delta(\boldsymbol{t}) \boldsymbol{I}.$$

H | Approximate Kernel Model

We derive an analytical expression for $\tilde{\mathcal{K}}_{f|H(T)}$. Recall that H's components are modelled independently. Then

$$\tilde{\mathcal{K}}_{f_i,f_j \mid \boldsymbol{H}(\boldsymbol{T})} = \sum_{k=1}^M \int_{\mathbb{R}^K} \mathbb{E}[H_{i,k}(\boldsymbol{t}_1 - \boldsymbol{\tau}) H_{j,k}(\boldsymbol{t}_2 - \boldsymbol{\tau}) \mid \boldsymbol{H}(\boldsymbol{T})] \, \mathrm{d}\boldsymbol{\tau}$$

where, denoting $h_i = H_{i,k}$, $h_j = H_{j,k}$, $\boldsymbol{u}_{h_i} = h_i(\boldsymbol{T})$ and $\boldsymbol{u}_{h_j} = h_j(\boldsymbol{T})$,

$$\int_{\mathbb{R}^{K}} \mathbb{E}[H_{i,k}(\boldsymbol{t}_{1}-\boldsymbol{\tau})H_{j,k}(\boldsymbol{t}_{2}-\boldsymbol{\tau}) \mid \boldsymbol{H}(\boldsymbol{T})] d\boldsymbol{\tau} \\
= \int_{\mathbb{R}^{K}} \{\mathcal{K}_{h_{i},h_{j}\mid\boldsymbol{u}_{h_{i}},\boldsymbol{u}_{h_{j}}}(\boldsymbol{t}_{1}-\boldsymbol{\tau},\boldsymbol{t}_{2}-\boldsymbol{\tau}) \\
- \mathbb{E}[h_{i}(\boldsymbol{t}_{1}-\boldsymbol{\tau}) \mid \boldsymbol{u}_{h_{i}}]\mathbb{E}[h_{j}(\boldsymbol{t}_{1}-\boldsymbol{\tau}) \mid \boldsymbol{u}_{h_{j}}]\} d\boldsymbol{\tau} \\
= \int_{\mathbb{R}^{2K}} [\mathbb{1}(i-j)\mathcal{K}_{h_{i}}(\boldsymbol{t}_{1}-\boldsymbol{\tau},\boldsymbol{t}_{2}-\boldsymbol{\tau}) \\
+ \mathcal{K}_{h_{i}}(\boldsymbol{t}_{1}-\boldsymbol{\tau},\boldsymbol{T})\boldsymbol{M}^{(h_{i},h_{j})}\mathcal{K}_{h_{j}}(\boldsymbol{T},\boldsymbol{t}_{2}-\boldsymbol{\tau}) d\boldsymbol{\tau} \\
= \mathbb{1}(i-j)\underbrace{\int_{\mathbb{R}^{K}} \mathcal{K}_{h_{i}}(\boldsymbol{t}_{1}-\boldsymbol{\tau},\boldsymbol{t}_{2}-\boldsymbol{\tau}) d\boldsymbol{\tau}}_{I^{(1,h_{i})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2})} \tag{H.1}$$

and

$$M^{(h_i,h_j)} = K_{u_{h_i}}^{-1} u_{h_i} u_{h_j}^T K_{u_{h_j}}^{-1} - \mathbb{1}(i-j) K_{u_{h_i}}^{-1}$$

Hence,

$$\begin{split} \tilde{\mathcal{K}}_{f_i,f_j \mid \boldsymbol{H}(\boldsymbol{T})} &= \sum_{k=1}^{M} \left[\mathbbm{1}(i-j) I^{(1,H_{i,k})} + \operatorname{tr}(\boldsymbol{M}^{(H_{i,k},H_{j,k})T} \boldsymbol{I}^{(2,H_{i,k},H_{j,k})}) \right] \\ &= \sum_{k=1}^{M} \left\{ \mathbbm{1}(i-j) [I^{(1,H_{i,k})} - \operatorname{tr}(\boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1} \boldsymbol{I}^{(2,H_{i,k},H_{i,k})})] \right. \\ &+ H_{i,k}^T(\boldsymbol{T}) \boldsymbol{K}_{H_{i,k}(\boldsymbol{T})}^{-1} \boldsymbol{I}^{(2,H_{i,k},H_{j,k})} \boldsymbol{K}_{H_{j,k}(\boldsymbol{T})}^{-1} H_{j,k}(\boldsymbol{T}) \right\}. \end{split}$$

I | Variational Free Energy of the Nonparametric Kernel Model

I.1 Introduction

This chapter derives an analytical expression for the variational free energy of Model 5. We also determine the asymptotic time complexity of computing the free energy.

Recall that the free energy is given by

$$\mathcal{F}(q) = \underbrace{-\frac{1}{2} \log[(2\pi)^{N} |\mathbf{\Lambda}|^{2}] - \frac{1}{2} \sum_{i=1}^{Y} \mathbb{E}_{q} \{ \|\mathbf{\Lambda}^{-1} [\mathbf{Y}_{i,:} - \mathbf{y}(\mathbf{T}_{i,:})] \|^{2} \}}_{\mathbb{E}_{q} [\log p(\mathbf{Y} | \mathbf{H}, \mathbf{x})]} - \sum_{i=1,j=1}^{N,M} D_{KL} [q(\mathbf{u}_{H_{i,j}}) \| p(\mathbf{u}_{H_{i,j}})] - \sum_{j=1}^{M} D_{KL} [q(\mathbf{u}_{\tilde{x}_{j}}) \| p(\mathbf{u}_{\tilde{x}_{j}})].$$

In more detail, the Kullback-Leibler divergences are given by (Appendix B.2)

$$D_{KL}[q(\boldsymbol{u}_{H_{i,j}}) \| p(\boldsymbol{u}_{H_{i,j}})] = \frac{1}{2} \left[\operatorname{tr}(\boldsymbol{K}_{\boldsymbol{u}_{H_{i,j}}}^{-1} \boldsymbol{\Sigma}_{H_{i,j}}) + \boldsymbol{\mu}_{H_{i,j}}^{T} \boldsymbol{K}_{\boldsymbol{u}_{H_{i,j}}}^{-1} \boldsymbol{\mu}_{H_{i,j}} + \log \frac{|\boldsymbol{K}_{\boldsymbol{u}_{H_{i,j}}}|}{|\boldsymbol{\Sigma}_{H_{i,j}}|} - N \right],$$
(I.1)

$$D_{KL}[q(\boldsymbol{u}_{\tilde{x}_j}) \| p(\boldsymbol{u}_{\tilde{x}_j})] = \frac{1}{2} \left[\operatorname{tr}(\boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_j}}^{-1} \boldsymbol{\Sigma}_{\tilde{x}_j}) + \boldsymbol{\mu}_{\tilde{x}_j}^T \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_j}}^{-1} \boldsymbol{\mu}_{\tilde{x}_j} + \log \frac{|\boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_j}}|}{|\boldsymbol{\Sigma}_{\tilde{x}_j}|} - N \right]$$
(I.2)

and the likelihood is given by

$$\mathbb{E}_{q}[\log p(\boldsymbol{Y} \mid \boldsymbol{H}, \boldsymbol{x})] = -\frac{1}{2}\log[(2\pi)^{N}|\boldsymbol{\Lambda}|^{2}] - \frac{1}{2}\sum_{i=1}^{Y} \left\{ \boldsymbol{Y}_{i,:}^{T}\boldsymbol{\Lambda}^{-2}\boldsymbol{Y}_{i,:} - 2\boldsymbol{Y}_{i,:}^{T}\boldsymbol{\Lambda}^{-2}\mathbb{E}_{q}[\boldsymbol{y}(\boldsymbol{T}_{i,:})] + \operatorname{tr}\left\{\boldsymbol{\Lambda}^{-2}\mathbb{E}_{q}[\boldsymbol{y}(\boldsymbol{T}_{i,:})\boldsymbol{y}^{T}(\boldsymbol{T}_{i,:})]\right\} \right\}.$$

It thus remains to determine the predictive mean $\mathbb{E}_q[\boldsymbol{y}(\boldsymbol{t})]$ and predictive autocovariance $\mathbb{E}_q[\boldsymbol{y}(\boldsymbol{t}_1)\boldsymbol{y}^T(\boldsymbol{t}_2)]$. Note that the latter is only evaluated on its diagonal.

Recall that $H_{i,j}$ and \tilde{x}_j denote the number of inducing points for the processes $T_{H_{i,j}}$ and $T_{\tilde{x}_j}$ respectively. Also recall that $T_H = \max_{i,j} T_{H_{i,j}}$ and $T_{\tilde{x}} = \max_j T_{\tilde{x}_j}$. Finally, recall that Y denotes the number of observed data points. Then Equations (I.1) and (I.2) have respectively time complexities $\mathcal{O}(T_H^3)$ and $\mathcal{O}(T_{\tilde{x}}^3)$ (Appendix D). Thus $\mathcal{F}(q)$ has time complexity $\mathcal{O}[NMT_H^3 + MT_{\tilde{x}}^3 + YN(C_\mu + C_\Sigma)]$ where C_μ and C_Σ represent the number of operations required to compute respectively $\mathbb{E}_q[y_i(t)]$ and $\mathbb{E}_q[y_i(t)y_i(t)]$.

In the remainder of this chapter we utilise the notation and identities from Appendix F without further notion.

I.2 Predictive Mean

We have that

$$\mathbb{E}_q[y_i(\boldsymbol{t})] = \mathbb{E}_q\{[(\boldsymbol{H} \ast \boldsymbol{x})(\boldsymbol{t})]_i\} = \sum_{j=1}^M \underbrace{\int_{\mathbb{R}^K} \mathbb{E}_q[H_{i,j}(\boldsymbol{t} - \boldsymbol{\tau})x_j(\boldsymbol{\tau})] \,\mathrm{d}\boldsymbol{\tau}}_{T_{i,j}^{(1)}}$$

where, denoting $h = H_{i,j}$ and $x = x_j$,

$$T_{i,j}^{(1)} = \int h(\boldsymbol{t} - \boldsymbol{\tau}) p(h \mid \boldsymbol{u}_h) \, \mathrm{d}h q(\boldsymbol{u}_h) \, \mathrm{d}\boldsymbol{u}_h$$
$$x(\boldsymbol{\tau}) p(\tilde{x} \mid \boldsymbol{u}_{\tilde{x}}) \, \mathrm{d}x q(\boldsymbol{u}_{\tilde{x}}) \, \mathrm{d}\boldsymbol{u}_{\tilde{x}} \, \mathrm{d}\boldsymbol{\tau}$$
$$= \int_{\mathbb{R}^2} \mathcal{K}_h(\boldsymbol{t} - \boldsymbol{\tau}, \boldsymbol{T}_h) \boldsymbol{K}_{\boldsymbol{u}_h}^{-1} \boldsymbol{\mu}_h \mathcal{K}_{x,\tilde{x}}(\boldsymbol{\tau}, \boldsymbol{T}_{\tilde{x}}) \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}}}^{-1} \boldsymbol{\mu} \, \mathrm{d}\boldsymbol{\tau}$$
$$= \boldsymbol{\mu}_h^T \boldsymbol{K}_{\boldsymbol{u}_h}^{-1} \underbrace{\int_{\mathbb{R}^K} \mathcal{K}_h(\boldsymbol{T}_h, \boldsymbol{t} - \boldsymbol{\tau}) \mathcal{K}_{x,\tilde{x}}(\boldsymbol{\tau}, \boldsymbol{T}_{\tilde{x}}) \, \mathrm{d}\boldsymbol{\tau}}_{\boldsymbol{I}^{(L,h,x)}(\boldsymbol{t})} \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}}}^{-1} \boldsymbol{\mu}_{\tilde{x}} \boldsymbol{\mu}_{\tilde{x}}$$
I.3 Predictive Autocovariance

In summary,

$$\mathbb{E}_{q}[y_{i}(m{t})] = \sum_{j=1}^{M} m{\mu}_{H_{i,j}}^{T} m{K}_{m{u}_{H_{i,j}}}^{-1} m{I}^{(L,H_{i,j},x_{j})}(m{t}) m{K}_{m{u}_{ ilde{x}_{j}}}^{-1} m{\mu}_{ ilde{x}_{j}}.$$

The integral $I^{(L,\cdot,\cdot)}$ is computed in Appendix I.4.

To begin with, there is a one-time cost of constructing and inverting all $K_{u_{H_{i,j}}}$ and $K_{u_{\tilde{x}_j}}$, which has time complexity $\mathcal{O}[M(T_H^3 + KT_H^2 + KT_{\tilde{x}}^2 + T_{\tilde{x}}^3)]$. Afterwards, all $I^{(L,H_{i,j},\tilde{x}_j)}$ are constructed in $\mathcal{O}(MKT_HT_{\tilde{x}})$ time and the predictive mean is computed in $\mathcal{O}[M(T_H^2 + T_HT_{\tilde{x}} + T_{\tilde{x}}^2)]$. Hence $\mathcal{O}(C_{\mu}) = \mathcal{O}[M(T_H^2 + KT_HT_{\tilde{x}} + T_{\tilde{x}}^2)]$.

I.3 Predictive Autocovariance

We have that

$$\begin{split} \mathbb{E}_{q}[y_{i}(\boldsymbol{t}_{1})y_{j}(\boldsymbol{t}_{2})] &= \sum_{k=1,l=1}^{M,M} \int_{\mathbb{R}^{2K}} \mathbb{E}_{q}[H_{i,k}(\boldsymbol{t}_{1}-\boldsymbol{\tau}_{1})x_{k}(\boldsymbol{\tau}_{1})H_{j,l}(\boldsymbol{t}_{2}-\boldsymbol{\tau}_{2})x_{l}(\boldsymbol{\tau}_{2})] \,\mathrm{d}\boldsymbol{\tau}_{1} \,\mathrm{d}\boldsymbol{\tau}_{2} \\ &= \sum_{k=1}^{M} \int_{\mathbb{R}^{K}} \mathbb{E}_{q}[H_{i,k}(\boldsymbol{t}_{1}-\boldsymbol{\tau}_{1})x_{k}(\boldsymbol{\tau})] \,\mathrm{d}\boldsymbol{\tau} \sum_{l=1,l\neq k}^{M} \int_{\mathbb{R}^{K}} \mathbb{E}_{q}[H_{j,l}(\boldsymbol{t}_{2}-\boldsymbol{\tau})x_{l}(\boldsymbol{\tau})] \,\mathrm{d}\boldsymbol{\tau} \\ &+ \sum_{k=1}^{M} \underbrace{\int_{\mathbb{R}^{2K}} \mathbb{E}_{q}[H_{i,k}(\boldsymbol{t}_{1}-\boldsymbol{\tau}_{1})x_{k}(\boldsymbol{\tau}_{1})H_{j,k}(\boldsymbol{t}_{2}-\boldsymbol{\tau}_{2})x_{k}(\boldsymbol{\tau}_{2})] \,\mathrm{d}\boldsymbol{\tau}_{1} \,\mathrm{d}\boldsymbol{\tau}_{2}}{\boldsymbol{\tau}_{i,j,k}^{(2)}} \\ &= \sum_{k=1}^{M} T_{i,k}^{(1)} \sum_{l=1:l\neq k}^{M} T_{j,l}^{(1)} + \sum_{k=1}^{M} T_{i,j,k}^{(2)}. \end{split}$$

Then, denoting $h_i = H_{i,k}$, $h_j = H_{j,k}$ and $x = x_k$,

$$T_{i,j,k}^{(2)} = \int h_i(\boldsymbol{t}_1 - \boldsymbol{\tau}_1) h_j(\boldsymbol{t}_2 - \boldsymbol{\tau}_2) p(h_i, h_j | \boldsymbol{u}_{h_i}, \boldsymbol{u}_{h_j}) \, \mathrm{d}(h_i, h_j) q(\boldsymbol{u}_{h_i}, \boldsymbol{u}_{h_j}) \, \mathrm{d}(\boldsymbol{u}_{h_i}, \boldsymbol{u}_{h_j}) \\ x(\boldsymbol{\tau}_1) x(\boldsymbol{\tau}_2) p(\tilde{x} | \boldsymbol{u}_{\tilde{x}}) \, \mathrm{d}x q(\boldsymbol{u}_{\tilde{x}}) \, \mathrm{d}\boldsymbol{u}_{\tilde{x}} \, \mathrm{d}\boldsymbol{\tau}_1 \, \mathrm{d}\boldsymbol{\tau}_2 \\ = \int_{\mathbb{R}^{2K}} [\mathbb{1}(i - j) \mathcal{K}_{h_i}(\boldsymbol{t}_1 - \boldsymbol{\tau}_1, \boldsymbol{t}_2 - \boldsymbol{\tau}_2) + \mathcal{K}_{h_i}(\boldsymbol{t}_1 - \boldsymbol{\tau}_1, \boldsymbol{T}_{h_i}) \boldsymbol{M}^{(h_i, h_j)} \mathcal{K}_{h_j}(\boldsymbol{T}_{h_j}, \boldsymbol{t}_2 - \boldsymbol{\tau}_2)] \\ [\mathcal{K}_x(\boldsymbol{\tau}_1, \boldsymbol{\tau}_2) + \mathcal{K}_{x,\tilde{x}}(\boldsymbol{\tau}_1, \boldsymbol{T}_{\tilde{x}}) \boldsymbol{M}^{(\tilde{x})} \mathcal{K}_{\tilde{x},x}(\boldsymbol{T}_{\tilde{x}}, \boldsymbol{\tau}_2)] \, \mathrm{d}\boldsymbol{\tau}_1 \, \mathrm{d}\boldsymbol{\tau}_2$$

where

$$\boldsymbol{M}^{(h_{i},h_{j})} = \boldsymbol{K}_{\boldsymbol{u}_{h_{i}}}^{-1} [\mathbb{1}(i-j)\boldsymbol{\Sigma}_{h_{i}} + \boldsymbol{\mu}_{h_{i}}\boldsymbol{\mu}_{h_{j}}^{T}] \boldsymbol{K}_{\boldsymbol{u}_{h_{j}}}^{-1} - \mathbb{1}(i-j)\boldsymbol{K}_{\boldsymbol{u}_{h_{i}}}^{-1},$$
$$\boldsymbol{M}^{(\tilde{x})} = \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}}}^{-1} [\boldsymbol{\Sigma}_{\tilde{x}} + \boldsymbol{\mu}_{\tilde{x}}\boldsymbol{\mu}_{\tilde{x}}^{T}] \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}}}^{-1} - \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}}}^{-1}.$$

Thus $T_{i,j,k}^{(2)} = Q_1 + Q_2 + Q_3 + Q_4$ where

In summary,

$$\mathbb{E}_{q}[y_{i}(\boldsymbol{t}_{1})y_{j}(\boldsymbol{t}_{2})] = \sum_{k=1}^{M} \boldsymbol{\mu}_{H_{i,k}}^{T} \boldsymbol{K}_{\boldsymbol{u}_{H_{i,k}}}^{-1} \boldsymbol{I}^{(L,H_{i,k},x_{k})}(\boldsymbol{t}_{1}) \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_{k}}}^{-1} \boldsymbol{\mu}_{\tilde{x}_{k}} \\ \underbrace{\sum_{l=1,l\neq k}^{M} \boldsymbol{\mu}_{H_{j,l}}^{T} \boldsymbol{K}_{\boldsymbol{u}_{H_{j,l}}}^{-1} \boldsymbol{I}^{(L,H_{j,l},x_{l})}(\boldsymbol{t}_{2}) \boldsymbol{K}_{\boldsymbol{u}_{\tilde{x}_{l}}}^{-1} \boldsymbol{\mu}_{\tilde{x}_{l}}}{T^{(1)}} \\ + \mathbbm{1}(i-j) \sum_{k=1}^{M} I^{(Q_{1},H_{i,k},x_{k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2}) \\ + \mathbbm{1}(i-j) \sum_{k=1}^{M} \operatorname{tr}[\boldsymbol{M}^{(\tilde{x}_{k})T} \boldsymbol{I}^{(Q_{2},H_{i,k},x_{k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2})] \\ + \sum_{k=1}^{M} \operatorname{tr}[\boldsymbol{M}^{(H_{i,k})T} \boldsymbol{I}^{(Q_{3},H_{i,k},H_{j,k},x_{k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2})] \\ + \sum_{k=1}^{M} \operatorname{tr}[\boldsymbol{M}^{(H_{i,k},H_{j,k})T} \boldsymbol{I}^{(Q_{4},H_{i,k},x_{k})}(\boldsymbol{t}_{1},\boldsymbol{t}_{2})] \\ + \sum_{k=1}^{M} \operatorname{tr}[\boldsymbol{M}^{(H_{i,k},H_{j,k})T} \boldsymbol{I}^{(Q_{4},H_{i,k},x_{k})}(\boldsymbol{t}_{1}) \boldsymbol{M}^{(\tilde{x}_{k})} \boldsymbol{I}^{(Q_{4},H_{j,k},x_{k})T}(\boldsymbol{t}_{2})]. \\ \end{bmatrix}_{T^{(2)}}$$

The integrals $I_{\cdot}^{(\cdot)}$ are computed in Appendix I.4.

To begin with, there is a one-time cost of constructing and inverting all $\mathbf{K}_{u_{H_{i,j}}}$ and $\mathbf{K}_{u_{\tilde{x}_j}}$, which has time complexity $\mathcal{O}[M(T_H^3 + KT_H^2 + KT_{\tilde{x}}^2 + T_{\tilde{x}}^3)]$, and a one-time cost of computing all $\mathbf{M}^{(H_{i,j},H_{i,j})}$ and $\mathbf{M}^{(\tilde{x}_j)}$, which has time complexity $\mathcal{O}[M(T_H^3 + T_{\tilde{x}}^3)]$. Thus, the resulting one-time cost has time complexity $\mathcal{O}[M(T_H^3 + KT_H^2 + KT_{\tilde{x}}^2 + T_{\tilde{x}}^3)]$. Afterwards, all $\mathbf{I}^{(Q_2,H_{i,j},x_j)}$, $\mathbf{I}^{(Q_3,H_{i,j},H_{i,j},x_j)}$ and $\mathbf{I}^{(Q_4,H_{i,j},x_j)}$ are constructed in $\mathcal{O}[M(KT_H^2 + KT_HT_{\tilde{x}} + KT_{\tilde{x}}^2) + M^2]$ and $\mathcal{O}(M[T_{\tilde{x}}^2T_{\tilde{x}} + T_HT_{\tilde{x}}^2])$ time. Hence

$$\mathcal{O}(C_{\Sigma}) = \mathcal{O}[M(T_H^2 T_{\tilde{x}} + K T_H^2 + K T_H T_{\tilde{x}} + K T_{\tilde{x}}^2 + T_H T_{\tilde{x}}^2) + M^2]$$

I.4 Integrals $I_{\cdot}^{(\cdot)}$

In this section we compute the integrals $I_{\cdot}^{(\cdot)}$ from Appendix I.3.

I.4.1 Kernels

First, let the composite vector be $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T & \boldsymbol{\tau} \end{bmatrix}^T$. Then

$$\begin{split} \mathcal{K}_{\tilde{x}_{j}}(\boldsymbol{t}_{1},\boldsymbol{t}_{2}) &= \int_{\mathbb{R}^{2K}} \exp(-\omega \|\boldsymbol{t}_{1}-\boldsymbol{\tau}_{1}\|^{2}) \underbrace{\mathbb{E}[x_{j}(\boldsymbol{\tau}_{1})x_{j}(\boldsymbol{\tau}_{2})]}_{\delta(\boldsymbol{\tau}_{1}-\boldsymbol{\tau}_{2})} \exp(-\omega \|\boldsymbol{t}_{2}-\boldsymbol{\tau}_{2}\|^{2}) \, \mathrm{d}\boldsymbol{\tau}_{1} \, \mathrm{d}\boldsymbol{\tau}_{2} \\ &= \int_{\mathbb{R}^{K}} \exp(-\omega \|\boldsymbol{t}_{1}-\boldsymbol{\tau}\|^{2} - \omega \|\boldsymbol{t}_{2}-\boldsymbol{\tau}\|^{2}) \, \mathrm{d}\boldsymbol{\tau} \\ &= I_{\boldsymbol{\tau}}[(1,2\begin{bmatrix} \omega & 0 & | -\omega \\ 0 & \omega & -\omega \\ | -\omega & -\omega & | 2\omega \end{bmatrix})] \\ &= \left(\frac{\pi^{K/2}}{(2\omega)^{K/2}}, 2\begin{bmatrix} \omega & 0 \\ 0 & \omega \end{bmatrix} - \frac{1}{\omega} \begin{bmatrix} \omega \\ \omega \end{bmatrix} \begin{bmatrix} \omega \\ \omega \end{bmatrix}^{T} \right) \\ &= \left(\frac{\pi^{K/2}}{(2\omega)^{K/2}}, \begin{bmatrix} \omega & -\omega \\ -\omega & \omega \end{bmatrix} \right) \\ &= \frac{\pi^{K/2}}{(2\omega)^{K/2}} \exp\left(-\frac{1}{2}\omega \|\boldsymbol{t}_{1}-\boldsymbol{t}_{2}\|^{2}\right) \\ &= \mathcal{K}_{\tilde{x}_{j}}(\boldsymbol{t}_{1}-\boldsymbol{t}_{2}). \end{split}$$

Second, let the composite vector be $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T \end{bmatrix}^T$. Then

$$\begin{aligned} \mathcal{K}_{x_j,\tilde{x}_j}(\boldsymbol{t}_1,\boldsymbol{t}_2) &= \int_{\mathbb{R}^K} \underbrace{\mathbb{E}[x_j(\boldsymbol{t}_1)x_j(\boldsymbol{t}_2-\boldsymbol{\tau})]}_{\delta(\boldsymbol{t}_1-\boldsymbol{t}_2+\boldsymbol{\tau})} \exp(-\omega \|\boldsymbol{\tau}\|^2) \,\mathrm{d}\boldsymbol{\tau} \\ &= \exp(-\omega \|\boldsymbol{t}_1-\boldsymbol{t}_2\|^2) \\ &= (1,2 \begin{bmatrix} \omega & -\omega \\ -\omega & \omega \end{bmatrix}) \\ &= \mathcal{K}_{x_j,\tilde{x}_j}(\boldsymbol{t}_1-\boldsymbol{t}_2) \end{aligned}$$

and $\mathcal{K}_{\tilde{x}_j,x_j}(\boldsymbol{t}_1,\boldsymbol{t}_2) = \mathcal{K}_{x_j,\tilde{x}_j}(\boldsymbol{t}_1,\boldsymbol{t}_2).$

Finally, let the composite vector be $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T \end{bmatrix}^T$. Then

$$\begin{aligned} \mathcal{K}_{H_{i,j}}(\boldsymbol{t}_1, \boldsymbol{t}_2) &= \sigma_h^2 \exp(-\alpha \|\boldsymbol{t}_1\|^2 - \alpha \|\boldsymbol{t}_2\|^2 - \gamma \|\boldsymbol{t}_2 - \boldsymbol{t}_1\|^2) \\ &= (\sigma_h^2, 2 \begin{bmatrix} \alpha + \gamma & -\gamma \\ -\gamma & \alpha + \gamma \end{bmatrix}). \end{aligned}$$

I.4.2 Integral $I_{\cdot,\cdot}^{(L,\cdot,\cdot)}$

Let the composite vector be $\begin{bmatrix} \boldsymbol{t}^T & \boldsymbol{T}_{h,i,:}^T & \boldsymbol{T}_{\tilde{x},j,:}^T & \boldsymbol{\tau}^T \end{bmatrix}^T$ and let $L = \alpha + \gamma + \omega$. Then

I.4.3 Integral $I^{(Q_1,\cdot,\cdot)}$

Let the composite vector be $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T & \boldsymbol{\tau}^T \end{bmatrix}^T$. Then

$$\begin{split} I^{(Q_1,\cdot,\cdot)}(\boldsymbol{t}_1,\boldsymbol{t}_2) &= \int_{\mathbb{R}^K} \mathcal{K}_h(\boldsymbol{t}_1 - \boldsymbol{\tau}, \boldsymbol{t}_2 - \boldsymbol{\tau}) \, \mathrm{d}\boldsymbol{\tau} \\ &= (\sigma_h^2, 2 \begin{bmatrix} \gamma & -\gamma \\ -\gamma & \gamma \end{bmatrix}) I_{\boldsymbol{\tau}} [(1, 2 \begin{bmatrix} \alpha & 0 & | -\alpha \\ 0 & \alpha & | -\alpha \\ | -\alpha & -\alpha & | 2\alpha \end{bmatrix})] \\ &= (\sigma_h^2, 2 \begin{bmatrix} \gamma & -\gamma \\ -\gamma & \gamma \end{bmatrix}) \left(\frac{\pi^{K/2}}{(2\alpha)^{K/2}}, 2 \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} - \frac{1}{\alpha} \begin{bmatrix} -\alpha \\ -\alpha \end{bmatrix} \begin{bmatrix} -\alpha \\ -\alpha \end{bmatrix}^T \right) \\ &= (\sigma_h^2, 2 \begin{bmatrix} \gamma & -\gamma \\ -\gamma & \gamma \end{bmatrix}) \left(\frac{\pi^{K/2}}{(2\alpha)^{K/2}}, \begin{bmatrix} \alpha & -\alpha \\ -\alpha & \alpha \end{bmatrix} \right) \\ &= \left(\sigma_h^2 \frac{\pi^{K/2}}{(2\alpha)^{K/2}}, \begin{bmatrix} \alpha + 2\gamma & -\alpha - 2\gamma \\ -\alpha - 2\gamma & \alpha + 2\gamma \end{bmatrix} \right) \\ &= \sigma_h^2 \frac{\pi^{K/2}}{(2\alpha)^{K/2}} \exp\left[-\frac{1}{2} (\alpha + 2\gamma) \| \boldsymbol{t}_1 - \boldsymbol{t}_2 \|^2 \right]. \end{split}$$

I.4.4 Integral $I^{(Q_2,\cdot,\cdot)}_{\cdot,\cdot}$

Let the composite vector be = $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T & \boldsymbol{T}_{\tilde{x},i,:}^T & \boldsymbol{T}_1^T & \boldsymbol{\tau}_2^T \end{bmatrix}^T$. Then

I.4.5 Integral $I^{(Q_3,\cdot,\cdot,\cdot)}_{\cdot,\cdot}$

Let the composite vector be = $\begin{bmatrix} \boldsymbol{t}_1^T & \boldsymbol{t}_2^T & \boldsymbol{T}_{h_k,i,:}^T & \boldsymbol{T}_{h_l,j,:}^T & \boldsymbol{\tau}^T \end{bmatrix}^T$. Then

$$\begin{split} I_{i,j}^{(Q_3,h_k,h_l,x)}(\boldsymbol{t}_1,\boldsymbol{t}_2) \\ &= \int_{\mathbb{R}^K} \mathcal{K}_h(\boldsymbol{t}_1 - \boldsymbol{\tau},\boldsymbol{T}_{h_k,i,:}) \mathcal{K}_h(\boldsymbol{T}_{h_l,j,:},\boldsymbol{t}_2 - \boldsymbol{\tau}) \, \mathrm{d}\boldsymbol{\tau} \\ &= I_{\boldsymbol{\tau}} [(\sigma_h^2, 2 \begin{bmatrix} \alpha + \gamma & 0 & -\gamma & 0 & -\alpha - \gamma \\ 0 & 0 & 0 & 0 & 0 \\ -\gamma & 0 & \alpha + \gamma & 0 & \gamma \\ 0 & 0 & 0 & 0 & 0 \\ -\alpha - \gamma & 0 & \gamma & 0 & \alpha + \gamma \end{bmatrix}) (\sigma_h^2, 2 \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & \alpha + \gamma & 0 & -\gamma & -\alpha - \gamma \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -\gamma & 0 & \alpha + \gamma & \gamma \\ 0 & -\alpha - \gamma & 0 & \gamma & \alpha + \gamma \end{bmatrix})] \\ &= I_{\boldsymbol{\tau}} [(\sigma_h^4, 2 \begin{bmatrix} \alpha + \gamma & 0 & -\gamma & 0 & | & -\alpha - \gamma \\ 0 & \alpha + \gamma & 0 & -\gamma & | & -\alpha - \gamma \\ 0 & \alpha + \gamma & 0 & -\gamma & | & -\alpha - \gamma \\ -\gamma & 0 & \alpha + \gamma & 0 & | & \gamma \\ 0 & -\gamma & 0 & \alpha + \gamma & \gamma \\ -\alpha - \gamma & -\alpha - \gamma & \gamma & \gamma & | & 2\alpha + 2\gamma \end{bmatrix})] \\ &= \left(\sigma_h^4 \frac{\pi^{K/2}}{(2\alpha + 2\gamma)^{K/2}}, 2 \begin{bmatrix} \alpha + \gamma & 0 & -\gamma & 0 \\ 0 & \alpha + \gamma & 0 & -\gamma \\ -\gamma & 0 & \alpha + \gamma & 0 \\ 0 & -\gamma & 0 & \alpha + \gamma \end{bmatrix} - \frac{1}{\alpha + \gamma} \begin{bmatrix} \alpha + \gamma \\ \alpha + \gamma \\ -\gamma \\ -\gamma \\ -\gamma \end{bmatrix} \right] \right). \end{split}$$

Now

and so

$$\begin{split} I_{i,j}^{(Q_3,h_k,h_l,x)}(\boldsymbol{t}_1,\boldsymbol{t}_2) &= \sigma_h^4 \frac{\pi^{K/2}}{(2\alpha+2\gamma)^{K/2}} \exp\left[-\frac{1}{2}\alpha \|\boldsymbol{t}_1 - \boldsymbol{t}_2\|^2 - \frac{1}{2}\alpha \|\boldsymbol{T}_{h_k,i,:} - \boldsymbol{T}_{h_l,j,:}\|^2 \\ &- \frac{1}{2}\gamma \|(\boldsymbol{t}_1 - \boldsymbol{t}_2) - (\boldsymbol{T}_{h_k,i,:} - \boldsymbol{T}_{h_l,j,:})\|^2 \\ &+ \frac{\gamma^2 - (\alpha+\gamma)^2}{\alpha+\gamma} \|\boldsymbol{T}_{h_k,i,:} + \boldsymbol{T}_{h_l,j,:}\|^2\right]. \end{split}$$

Furthermore, denote $\boldsymbol{t}_1 - \boldsymbol{t}_2 = \boldsymbol{t}$ and $\boldsymbol{T}_{h_k,i,;} - \boldsymbol{T}_{h_l,j,:} = \boldsymbol{T}$. Then

$$\alpha \|\boldsymbol{t}\|^{2} + \alpha \|\boldsymbol{T}\|^{2} + \gamma \|\boldsymbol{t} + \boldsymbol{T}\|^{2} = (\alpha + \gamma) \|\boldsymbol{t}\|^{2} - 2\gamma \boldsymbol{t}^{T} \boldsymbol{T} + (\alpha + \gamma) \|\boldsymbol{T}\|^{2}$$
$$= (\alpha + \gamma) \left\|\boldsymbol{t} - \frac{\gamma}{\alpha + \gamma} \boldsymbol{T}\right\|^{2} + \frac{(\alpha + \gamma)^{2} - \gamma^{2}}{\alpha + \gamma} \|\boldsymbol{T}\|^{2}$$

so that

$$\begin{split} I_{i,j}^{(Q_3,h_k,h_l,x)}(\boldsymbol{t}_1,\boldsymbol{t}_2) &= \sigma_h^4 \frac{\pi^{K/2}}{(2\alpha+2\gamma)^{K/2}} \exp\left[-\frac{1}{2}(\alpha+\gamma) \Big\| (\boldsymbol{t}_1 - \boldsymbol{t}_2) - \frac{\gamma}{\alpha+\gamma} (\boldsymbol{T}_{h_k,i,:} - \boldsymbol{T}_{h_l,j,:}) \Big\|^2 \\ &- \frac{1}{2} \frac{\gamma^2 - (\alpha+\gamma)^2}{\alpha+\gamma} (\|\boldsymbol{T}_{h_k,i,:} - \boldsymbol{T}_{h_l,j,:}\|^2 \\ &- 2\|\boldsymbol{T}_{h_k,i,:} + \boldsymbol{T}_{h_l,j,:}\|^2) \right]. \end{split}$$

I.4.6 Integral $I^{(Q_4,\cdot,\cdot)}_{\cdot,\cdot}$

Let the composite vector be $\begin{bmatrix} \boldsymbol{t}^T & \boldsymbol{T}_{h,i,:}^T & \boldsymbol{T}_{\tilde{x},j,:}^T & \boldsymbol{\tau}^T \end{bmatrix}^T$. Then

I.5 Conclusion

We have derived an analytical expression for the variational free energy of Model 5. The asymptotic time complexity of computing the free energy including one-time costs is I.5 Conclusion

given by

$$\mathcal{O}[NMT_{H}^{3} + MT_{\tilde{x}}^{3} + YN(C_{\mu} + C_{\Sigma}) + \text{one-time costs}]$$

$$= \mathcal{O}[NMT_{H}^{3} + MT_{\tilde{x}}^{3} + YNM(T_{H}^{2}T_{\tilde{x}} + KT_{H}^{2} + KT_{H}T_{\tilde{x}} + KT_{\tilde{x}}^{2} + T_{H}T_{\tilde{x}}^{2}) + YNM^{2}$$

$$+ M(T_{H}^{3} + KT_{H}^{2} + KT_{\tilde{x}}^{2} + T_{\tilde{x}}^{3})]$$

$$= \mathcal{O}[NMT_{H}^{3} + MT_{\tilde{x}}^{3} + YNM(T_{H}^{2}T_{\tilde{x}} + KT_{H}^{2} + KT_{H}T_{\tilde{x}} + KT_{\tilde{x}}^{2} + T_{H}T_{\tilde{x}}^{2}) + YNM^{2}].$$

References

- Alvarez, M. A. and Lawrence, N. D. (2011). Computationally Efficient Convolved Multiple Output Gaussian Processes. Journal of Machine Learning Research, (12):1459–1500.
 (Pages xxii, 31, 33, and 34)
- Alvarez, M. A., Luengo, D., and Lawrence, N. D. (2009). Latent Force Models. Artificial Intelligence and Statistics, 5:9–16. (Pages xxiii, 31, 33, and 34)
- Bonilla, E. V., Chai, K. M., and Williams, C. K. I. (2008). Multi-Task Gaussian Process Prediction. Advances in Neural Information Processing Systems, 20:153–160. (Pages xxiii, 33, and 34)
- Duvenaud, D. (2014). Automatic Model Construction with Gaussian Processes. PhD thesis, Computational and Biological Learning Laboratory, University of Cambridge. (Page 1)
- Goovaerts, P. (1997). Geostatistics for Natural Resources Evaluation. Oxford University Press. (Pages xxiii, 33, and 34)
- Gray, R. M. (2006). Toeplitz and Circulant Matrices: A Review. Foundations and Trends in Communications and Information Theory, 2(3):155–239. (Pages 62, 63, and 66)
- Higham, N. J. (1988). Computing a Nearest Symmetric Positive Semidefinite Matrix. Linear Algebra and its Applications, 103:103–118. (Page 57)
- MacKay, D. J. C. (2002). Information Theory, Inference & Learning Algorithms. Cambridge University Press. (Pages 10 and 16)
- Minka, T. (2000). Deriving Quadrature Rules from Gaussian Processes. Technical report, Statistics Department, Carnegie Mellon University. (Page 21)

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press. (Page 51)

Nguyen, T. V. and Bonilla, E. V. (2014). Collaborative Multi-Output Gaussian Processes. *Conference on Uncertainty in Artificial Intelligence*, 30.

(Pages xxii, 33, and 34)

- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning.* MIT Press. (Pages 1 and 10)
- Teh, Y. W. and Seeger, M. (2005). Semiparametric Latent Factor Models. International Workshop on Artificial Intelligence and Statistics, 10. (Pages xxiii, 33, and 34)
- Titsias, M. K. (2009). Variational Learning of Inducing Variables in Sparse Gaussian Processes. Artificial Intelligence and Statistics, 12:567–574. (Pages 23, 26, and 27)
- Tobar, F., Bui, T. D., and Turner, R. E. (2015a). Design of Covariance Functions
 Using Inter-Domain Inducing Variables. *Time Series Workshop on Advances in Neural Information Processing Systems*. (Page 27)
- Tobar, F., Bui, T. D., and Turner, R. E. (2015b). Learning Stationary Time Series using Gaussian Processes with Nonparametric Kernels. Advances in Neural Information Processing Systems, 29:3501–3509. (Pages vii, xvii, xxii, 3, 15, 21, 25, 26, and 37)
- Tobar, F. and Turner, R. E. (2016). Modelling Time Series via Automatic Learning of Basis Functions. Ninth IEEE Sensor Array and Multichannel Signal Processing Workshop. (Pages 29 and 30)
- Trefethen, L. N. and Bau, D. (1997). *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics. (Pages 55 and 56)
- Ulrich, K., Carlson, D. E., Dzirasa, K., and Carin, L. (2015). GP Kernels for Cross-Spectrum Analysis. Advances in Neural Information Processing Systems, 28:1999–2007. (Pages vii, xxii, 3, 25, 26, 33, 34, and 61)
- Wilson, A. G. and Adams, R. P. (2013). Gaussian Process Kernels for Pattern Discovery and Extrapolation. International Conference on Machine Learning, 3:1067–1075. (Pages xxiii, 1, 25, and 26)

Wilson, A. G., Knowles, D. A., and Ghahramani, Z. (2012). Gaussian process Regression Networks. International Conference on Machine Learning, 29.

(Pages xxii, 31, and 35)

Index

Approximate Deep Kernel Model, 43 approximate inference, 27 Approximate Kernel Model, 22 automatic relevance determination, 16

Basis Function Model, 29 Bayesian numerical approximation, 21

Cholesky decomposition, 56 circulant approximation, 62 Convolutional Mixing Model, 32 cross-spectral mixture kernel, 25

Deep Gaussian Process Convolution Model, 39 Deep Kernel Model, 40

evidence, 10

free energy, 28

Gaussian process, 9 Gaussian Process Convolution Model, 25, 38 explicit decay, 38 Gaussian process regression, 10 Generalised Gaussian Process Convolution Model, 11 Instantaneous Mixing Model, 32 inter-domain transformation, 27 kernel, 1 decaying exponentiated-quadratic kernel, 15 diagonal multi-output decaying automatic relevance determination kernel, 16 diagonal multi-output decaying exponentiated-quadratic kernel, 15diagonal multi-output exponentiated-quadratic kernel, 15exponentiated-quadratic kernel, 15root, 19 kernel design problem, 1 Kullback-Leibler divergence, 51 linear state-space model, 6 marginal likelihood, 10 mixing model hierarchy, 33

Index

multi-task learning, 31	nonparametric model, 1
Multidimensional Time-Invariant Model, 9	Occam's razor, 10
	spectral mixture kernel, 25
nearest symmetric positive-semidefinite matrix, 57 neural network, 41	Time-Invariant Model, 6 Time-Variant Model, 6
Nonparametric Kernel Model, 12	variational free energy, 28