



HAL
open science

Closed-form Filtering for Non-linear Systems

Théophile Cantelobre, Carlo Ciliberto, Benjamin Guedj, Alessandro Rudi

► **To cite this version:**

Théophile Cantelobre, Carlo Ciliberto, Benjamin Guedj, Alessandro Rudi. Closed-form Filtering for Non-linear Systems. 2024. hal-04458936

HAL Id: hal-04458936

<https://inria.hal.science/hal-04458936v1>

Preprint submitted on 15 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Closed-form Filtering for Non-linear Systems

Théophile Cantelobre¹

Carlo Ciliberto²

Benjamin Guedj^{2,3}

Alessandro Rudi¹

¹ Inria/PSL Research University, Paris, France

² AI Centre, Dept. of Computer Science, University College London, London, UK

³ Inria Lille - Nord Europe and Inria London, France

February 15, 2024

Abstract

Sequential Bayesian Filtering aims to estimate the current state distribution of a Hidden Markov Model, given the past observations. The problem is well-known to be intractable for most application domains, except in notable cases such as the tabular setting or for linear dynamical systems with gaussian noise. In this work, we propose a new class of filters based on Gaussian PSD Models, which offer several advantages in terms of density approximation and computational efficiency. We show that filtering can be efficiently performed in closed form when transitions and observations are Gaussian PSD Models. When the transition and observations are approximated by Gaussian PSD Models, we show that our proposed estimator enjoys strong theoretical guarantees, with estimation error that depends on the quality of the approximation and is adaptive to the regularity of the transition probabilities. In particular, we identify regimes in which our proposed filter attains a TV ϵ -error with memory and computational complexity of $O(\epsilon^{-1})$ and $O(\epsilon^{-3/2})$ respectively, including the offline learning step, in contrast to the $O(\epsilon^{-2})$ complexity of sampling methods such as particle filtering.

1 Introduction

Sequential Bayesian Filtering is the task of inferring the distribution of unobserved variable X_T from observations Y_1, \dots, Y_T where $(X_t, Y_t)_{n \geq 0}$ is a Hidden

Markov Model. The distribution of X_T given Y_1, \dots, Y_T is known as the filtering distribution (or optimal filter), denoted π_T^ν where $\nu \in \mathcal{P}(E)$ is the distribution of X_0 (or prior knowledge on X_0 more generally). In a Hidden Markov Model, the conditional distributions of X_n given X_{n-1} and Y_n given X_n are described by transition kernels Q and G respectively. In particular, $Q(u, x)$ identifies the conditional probability of transitioning to the state $X_n = x$ at time n given the fact that the system was in the state $X_{n-1} = u$ at time $n-1$, while $G(x, y)$ is the probability of observing $Y_n = y$ given the fact that the system is in the state $X_n = x$. If one has perfect knowledge of the initial distribution π_0 , the transition and observation kernels Q, G , then the filtering distribution can be computed recursively by applying Bayes rule:

$$\pi_n(dx) = \frac{\int Q(u, dx)G(x, y_n)\pi_{n-1}(du)}{\int \int Q(u, dx)G(x, y_n)\pi_{n-1}(du)} \quad (1)$$

The recursive expression of the filtering distribution can be seen as a combination of two steps: the prediction of X_{n+1} given belief on X_n given past observations, then the correction of this prediction based on the observation Y_{n+1} received. For complete introduction to Hidden Markov Models and inference, we refer the reader to [Cappé et al. \(2005\)](#).

Two quintessential problems in filtering are studying the stability and robustness of the optimal filter, and computing an approximation of it in practice.

The stability of the optimal filter is related the robustness of the sequence π_n^ν with respect to the initial ν , since this distribution is unknown in most application settings and estimates/priors need to be used in practice. The optimal filter is stable when its dependence on the initial distribution decreases as more observations are provided (also known as “forgetting”). More formally, a filter is stable when the distance between π_n^ν and π_n^μ goes to zero as $n \rightarrow \infty$ for any two initial distributions ν and μ . This problem has attracted considerable interest since the first contributions of [Ocone and Pardoux \(1996\)](#) and [Blankenship et al. \(1981\)](#). Some reference relevant to this work include [Gland and Oudjane \(2004\)](#); [McDonald and Yüksel \(2020\)](#); [LeGland \(1999\)](#); [De Castro et al. \(2017\)](#); [Mitrophanov et al. \(2005\)](#). A modern review of the literature and different approaches can be found in [Kim \(2022\)](#).

In general, the iteration in (1) is intractable. Two exceptions are when the state-space is finite or when the state-space is continuous but the transition kernels are Gaussian Linear Conditional Distributions. In the former case, the algorithm is known as the *forward algorithm*. In the latter, the algorithm is the well-known Kalman filter [Kálmán and Bucy \(1961\)](#). The Kalman filter is known to compute the recursion (1) exactly where the Hidden Markov

Model has linear dynamics and observations and independent Gaussian noise. If the state-space is finite, the Baum-Welch algorithm can compute filtering and smoothing distributions in closed-form using the forward-backward approach.

But many real-world systems do not have linear dynamics nor Gaussian uncertainty. On one hand, many algorithms have been devised to handle variations on the Kalman filter’s assumptions, including the Extended Kalman Filter or the Unscented Filter (Julier and Uhlmann, 2004) (see also Särkkä (2013)). These models approximate the state variable as a Gaussian, which excludes many systems where, for instance, multi-modality is present. On the other, Sequential Monte Carlo algorithms such as the Particle filter and variants were developed. These algorithms approximate the marginalization step above using sampling, and can handle multi-modality. This family of methods has strong theoretical guarantees (though under arguably stringent conditions), but do not give a closed-form expression of the approximate distribution and are known to be difficult to turn and prohibitively costly for online applications.

Gaussian PSD Models were introduced in Rudi and Ciliberto (2021) as models for probability distributions. Gaussian PSD Models are a special case of the family of models proposed in Marteau-Ferey et al. (2020). They generalize Gaussian Mixture Models by allowing for negative coefficients in the mixture. As originally highlighted in Rudi and Ciliberto (2021), Gaussian PSD Models enjoy appealing properties for applications involving Bayesian inference, with filtering as a special case: 1) They have optimal approximation guarantees with respect to a large family of probability densities 2) products and marginals of Gaussian PSD models can be efficiently computed in closed-form.

Approach & Contributions

In this paper, we study the problem of performing the iteration in (1) when knowledge of the transition Q and observation G probabilities are unknown and only an approximation in terms of two corresponding Gaussian PSD Models \hat{Q} and \hat{G} is available. We introduce a new algorithm to derive an estimator $\hat{\pi}_n(dx)$ and study its relation with the true $\pi_n(dx)$. The proposed estimator extends previous filtering strategies, such as the Kalman filter and offers strong theoretical guarantees on a large family of application settings.

Our main contributions are:

1. A novel algorithm to tackle Sequential Bayesian Filtering, which recovers previously proposed estimators and can be applied to any filtering

problem where the transition kernels admit a smooth density.

2. We show that the proposed estimator is both stable and robust with respect to a large family of application settings. These theoretical properties are adaptive to the regularity properties of the Hidden Markov Model.
3. The computational and space complexity of the proposed algorithm depends on the regularity of the transition kernels. For very regular kernels (e.g. infinitely differentiable) the algorithm has a computational complexity that is smaller than, for example, particle filtering.

Our paper is organized as follows: in [Section 2](#), we describe Gaussian PSD Models and their properties, in particular their stability with respect to probabilistic operations. In [Section 3](#), we devise an algorithm for learning Gaussian PSD Models from function evaluations and prove that optimal estimation rates are attained for smooth targets. In [Section 4](#), we introduce PSDFILTER, an approximate filtering algorithm which plugs \hat{Q} and \hat{G} in the iteration above. We prove this algorithm is robust to the choice of initial distribution and to the approximation error in Q and G . Sketches of the proofs of our main theorems [Theorem 4](#) and [Theorem 6](#) are presented in [Section 5](#). Finally, in [Section 6](#), we generalize Gaussian PSD Models to allow for a richer class of approximators, while retaining most of the desirable properties of Gaussian PSD Models.

Notation We denote $E = (-1, 1)^d$ and $F = (-1, 1)^{d'}$ the state and observation space. $\mathcal{P}(E)$ is the set of probability measures on E and $\mathcal{M}_+(E)$ the set of finite, positive measures on E . We assume that all measures admit a density with respect to the Lebesgue measure and use the abuse of notation $\mu(dx) = \mu(x)dx$. $\beta > 0$ is a smoothness parameter, σ is a mixing parameter for kernels and ϵ is the accuracy when doing function approximation. Denote \mathbb{R}_+^d the set of vectors in \mathbb{R}^d with all positive components and $\mathcal{S}^+(\mathbb{R}^M)$ the set of positive definite matrices of size M .

2 Gaussian PSD Models

Gaussian PSD Models, introduced in [Rudi and Ciliberto \(2021\)](#), is a family of models for non-negative functions and, in particular, probability densities specializing the PSD Models from [Marteau-Ferey et al. \(2020\)](#). They are non-negative everywhere, admit a linear parametrization and can be learned from samples and function evaluations. They are characterized by a linear

combination of kernels, with weights chosen such that the function is non-negative. In this section, we recall the definition of Gaussian PSD Models of [Rudi and Ciliberto \(2021\)](#), show that they extend most well-established probability models and present how to perform operations such as multiplication or marginalization.

Definition 1. *A Gaussian PSD Model of order M is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ which can be written:*

$$f(x) = \sum_{i=1}^M \sum_{j=1}^M A_{ij} k_{\eta}(x, x_i) k_{\eta}(x, x_j) \quad (2)$$

where $\eta \in \mathbb{R}_+^d$ is the precision vector, $X = (x_i)_{1 \leq i \leq M} \in \mathbb{R}^{M \times d}$ are the anchor points and $A \in \mathcal{S}^+(\mathbb{R}^d)$ is the weight matrix. Such a function is denoted $f(x; A, X, \eta)$ (or $f(x; \theta)$ for shorthand).

Adopting the point of view in [Marteau-Ferey et al. \(2020\)](#), Gaussian PSD Models can equivalently be defined as functions of the form $f(x) = \Phi_{\eta}(x)^{\top} A \Phi(x)$ where Φ_{η} is defined as $\Phi_{\eta}(x) = (k_{\eta}(x, x_1) \dots k_{\eta}(x, x_M))^{\top} \in \mathbb{R}^M$. When defined on the product of Euclidean spaces $\mathbb{R}^d \times \mathbb{R}^{d'}$ with anchor points $[X, Y]$ (the row-wise concatenation of X and Y) and precision vector $\eta = (\eta, \eta')$ (column-wise concatenation of η_1 and η_2), we denote the model $f(x, y; A, [X, Y], (\eta_1, \eta_2))$. This split notation is justified by the fact that $k_{\eta}((x, y), (u, v)) = k_{\eta_1}(x, u) k_{\eta_2}(y, v)$.

Example 1 (Gaussian Mixture Model). *Let $p(x) = \sum_{k=1}^M \alpha_k p(x|\mu_k, \eta)$ where $\mu_k \in \mathbb{R}^d$, $\eta_k \in \mathbb{R}_+^d$, and $\alpha \in \mathbb{R}_+^d$ with $\sum_{k=1}^M \alpha_k = 1$ and $p(x|\mu_k, \eta)$ is the Gaussian density with mean μ and precision vector η . p is known as a Gaussian Mixture Model. p can be written as a Gaussian PSD Model of order M $f(x; A, X, \eta/2)$ with $A = \text{diag}(a)$ and $X = (\mu_1 \dots \mu_M)^{\top}$.*

Example 2 (Squared linear Gaussian model). *Let $g(x) = w^{\top} \Phi_{\eta}(x)$ where $w \in \mathbb{R}^d$ and $\Phi_{\eta}(x) = (k_{\eta}(x, x_1), \dots, k_{\eta}(x, x_M))^{\top}$. Then, $f = g^2$ can be written as a Gaussian PSD Model of order M with $A = ww^{\top}$ and $X = (x_1, \dots, x_M)^{\top}$. Indeed,*

$$f(x) = (w^{\top} \Phi_{\eta}(x))^2 = w^{\top} \Phi_{\eta}(x) w^{\top} \Phi_{\eta}(x) = \Phi_{\eta}(x)^{\top} w w^{\top} \Phi_{\eta}(x). \quad (3)$$

As pointed out in the introduction of this section, because $k_{\eta}(x, u) k_{\eta}(x, v) \propto k_{2\eta}(x, \frac{u+v}{2})$, a Gaussian PSD Model can be seen as a linear combination of Gaussians. It is important to note that the coefficients of the components can be non-negative, which makes them much more expressive than Mixture models. Consider for instance $f(x) = (e^{-(x-2)^2} - e^{-(x-3)^2})^2$ which is clearly non-negative and can be written as a Gaussian PSD Model but not as a Mixture model.

2.1 Operations on Gaussian PSD Models

Gaussian PSD Models are compatible with operations on probabilistic models such as integration, partial evaluation, product and marginalization. The operations are summarized in [Proposition 1](#) and the algorithms, based on kernel evaluations and matrix-vector products are detailed in [Rudi and Ciliberto \(2021\)](#).

Proposition 1 (Closed form operations for Gaussian PSD Models). *Let $f(x, y; \theta_1)$ and $g(y, z; \theta_2)$ be two Gaussian PSD Models of order M_1 and M_2 respectively, as in [Definition 1](#). Then there exist some algorithms INTEGRAL, PARTIALEVAL, PRODUCT, MARGINAL such that*

- **Integral over \mathbb{R}^d or over a hypercube** $\int f(x, y; \theta_1) dx dy$ can be computed exactly and in closed form by the algorithm INTEGRAL(θ_1) with a computational cost of $O(M_1^2 d)$.
- **Partial evaluation** $f(x, y_0; \theta_1) = h(x; \theta')$ is a Gaussian PSD Model of order at most M_1 and θ' can be computed exactly and in closed form by the algorithm PARTIALEVAL(y_0, θ_1) with a computational cost of $O(M_1^2 d)$
- **Product** $f(x, y; \theta_1) g(y, z; \theta_2) = h(x, y, z; \theta')$ is a Gaussian PSD Model of order at most $M_1 \times M_2$ and θ' can be computed exactly and in closed form by the algorithm PRODUCT(θ_1, θ_2) with a computational cost of $O(M_1^2 M_2^2 d)$.
- **Marginalization** $\int f(x, y; \theta_1) dy = h(x; \theta')$ is a Gaussian PSD Model of order at most M_1 and θ' can be computed exactly and in closed form by the algorithm MARGINAL(y, θ_1) with a computational cost of $O(M_1^2 d)$.

The proof of the proposition above can be found in ([Rudi and Ciliberto, 2021](#), Appendix F). Note that a Markov transition $g(x) = \int Q(u, x) f(u) du$, when Q and f are Gaussian PSD models, can be decomposed in terms of product and marginalization and computed in closed form, with g again a Gaussian PSD Model. More importantly, g is of order M_1 (instead of the naïve $M_1 \times M_2$ according to [Proposition 1](#)). This is summarized in the following proposition

Proposition 2 (Constant order for Markov transition). *If $Q(u, x; \theta_Q)$ and $f(u; \theta_f)$ are two Gaussian PSD Models of order M_1 and M_2 respectively, then $g(x) = \int Q(u, x; \theta_Q) f(u; \theta_f) du$, computed via PRODUCT and MARGINAL is a Gaussian PSD Model of order M_1 .*

The proof of [Proposition 2](#) can be found in ([Rudi and Ciliberto, 2021](#), Appendix F.5).

3 Learning transition and observation kernels with Gaussian PSD Models

In this section, we show that General Gaussian Models can be used to efficiently approximate smooth, non-negative functions in L^∞ using function evaluations. Let $\Omega = (-1, 1)^d$ and $f : \Omega \rightarrow \mathbb{R}$ be the target function. We assume we can evaluate $f(x)$ at any point $x \in \Omega$. We assume that f is the sum of squared β -smooth functions. Formally, we introduce:

Assumption 1 (Smooth sum-of-squares assumption). *There exist $q \in \mathbb{N}$ and $\beta \geq 0$ and $f_i \in W_2^\beta(\Omega) \cap L^\infty(\Omega)$ such that $f(x) = \sum_{i=1}^q f_i(x)^2$.*

[Assumption 1](#) is verified for most continuous dynamical models of interest. For instance, any transition kernel $Q(x, y) \propto e^{-\|\Sigma^{-1/2}(y-h(x))\|^2}$ verifies the assumption. [Rudi and Ciliberto \(2021\)](#) provides a list of sufficient conditions, we recall in [Proposition 3](#).

Proposition 3 (Generality of [Assumption 1](#), Prop. 5 in [Rudi and Ciliberto \(2021\)](#)). *A function f satisfies [Assumption 1](#) on $\Omega = (-1, 1)^d$ as soon as:*

- *f is a probability density and $f \in W_2^\beta(\Omega) \cap L^\infty(\Omega)$, and strictly positive on $[-1, 1]^d$;*
- *f is an exponential model $f(x) = e^{-v(x)}$ with $v \in W_2^\beta(\Omega) \cap L^\infty(\Omega)$;*
- *f is a mixture of models from (b) ;*
- *f is $\beta + 2$ -times differentiable on $[-1, 1]^d$, with a finite set of zeroes all in $(-1, 1)^d$, and a positive definite Hessian in each zero.*

In [Section 3.1](#), we introduce the optimization problem we solve to learn f and present the learning algorithm. In [Section 3.2](#), we prove that the obtained estimator \hat{f} is

3.1 Learning algorithm

In [Section 2](#), we showed that the square of any Gaussian Linear Model is a Gaussian PSD Model and its weight matrix is of rank 1 and given by $A = aa^\top$ where a is the weight vector of the Gaussian Linear Model. We use

this insight to efficiently approximate a smooth sum-of-squares function f with a Gaussian PSD Model \hat{f} . This insight was first published by [Marteau-Ferey et al. \(2022\)](#) for probability densities. A full-rank estimator \hat{f} can be also learned, by solving a Semi-Definite Programming problem using e.g. Newton’s method.

In this work we propose to approximate f using $\hat{f} = \hat{g}^2$ where \hat{g} is a Gaussian Linear Model learned on g . Denoting the Linear Gaussian Model $\hat{g}(x; a, \eta, \tilde{X}) = \sum_{i=1}^M a_i k_\eta(x, \tilde{x}_i)$ where $a \in \mathbb{R}^M$, $\eta \in \mathbb{R}_+^d$ and $\tilde{X} \in \mathbb{R}^{M \times d}$, we introduce the optimization problem used to learn \hat{a} from data points $X \in \mathbb{R}^{n \times d}$:

$$\min_{a \in \mathbb{R}^M} \frac{1}{n} \sum_{k=1}^n |\sqrt{f(x_k)} - \hat{g}(x_k; a, \eta, \tilde{X})|^2 + \lambda a^\top K a \quad (4)$$

where K is described by $K_{ij} = k_\eta(\tilde{x}_i, \tilde{x}_j)$.

In [Appendix B](#), we cast (4) as a kernel ridge regression problem, which can be efficiently solved for large values of n and M in $O(n\sqrt{n})$ time using approximate kernel methods such as in [Rudi et al. \(2017\)](#).

Algorithm 1: LEARNRANKONE algorithm

Data: $f(x)$, M , N , η , λ
 $X \leftarrow \text{UNIFORMSAMPLE}(\Omega, N)$;
 $Y \leftarrow (\sqrt{f(x_i)} | x_i \in X)$;
 $\tilde{X} \leftarrow \text{UNIFORMSAMPLE}(\Omega, M)$;
 $\hat{a} \leftarrow \text{KERNELRIDGREGRESSION}(k_\eta, \tilde{X}, X, Y, \lambda)$;
 $\hat{f}(x) \leftarrow \text{GAUSSIANPSDMODEL}(a a^\top, X, \eta)$;
Result: $\hat{f}(x)$

3.2 Learning rates

The Gaussian PSD Model \hat{f} obtained from [Algorithm 1](#) using function evaluations at uniformly sampled training points X approximates f in $L^\infty(\Omega)$ with optimal learning rates for L^∞ norm ([Wendland, 2004](#)), if f is a β -smooth and bounded density, as formalized by [Assumption 1](#).

We build on the results in [Rudi and Ciliberto \(2021\)](#) and [Marteau-Ferey et al. \(2022\)](#). The former studies the L^2 convergence of \hat{f} to f when the training set is sampled from the target density, using the full-rank counterpart

to [Algorithm 1](#). The latter studies convergence in Hellinger distance using [Algorithm 1](#). Both works use the insights of [Rudi et al. \(2015\)](#).

[Algorithm 1](#) finds a solution to (4) in the reproducing kernel Hilbert space \mathcal{H}_η associated to k_η where η is chosen as a function of the desired precision ϵ . [Theorem 4](#) proves that \hat{f} converges to f in L^∞ , with optimal rates ([Wendland, 2004](#)). In particular, to learn f uniformly to precision ϵ , $N \approx \epsilon^{-\frac{2d}{2\beta-d}}$ function evaluations and a model with $M \approx N$ anchor points suffices.

Theorem 4. *Let $\beta > d/2$ and $\theta^{-1} < 1 + 2\beta/d$. Let $f : \Omega \rightarrow \mathbb{R}$ such that f verifies [Assumption 1](#). Set $M \geq C'(\log(\frac{1}{\epsilon}))^d \log(\frac{1}{\delta\epsilon})\epsilon^{-d/\beta}$ and $n \geq C'\epsilon^{-2d/\beta} \log \frac{1}{\delta}$. Consider the set of anchor points $\tilde{X} \in \mathbb{R}^{M \times d}$ and the set of training points $X \in \mathbb{R}^{n \times d}$ sampled independently and uniformly from $\Omega = (-1, 1)^d$. Let $\epsilon \leq \epsilon_0$. Let $\eta = \epsilon^{-2/\beta} 1_d$ and \mathcal{H}_η the RKHS associated to k_η . Let \hat{a} solution to the kernel ridge regression problem defined in (4). We denote $\hat{g}(\cdot) = \hat{g}(\cdot; \hat{a}, \tilde{X}, \eta)$ the estimator of g and $\hat{f} = \hat{g}^2$. With probability at least $1 - 3\delta$,*

$$\|\hat{f} - f\|_{L^\infty(\Omega)} \leq C \|\sqrt{f}\|_{W_2^\beta(\Omega)}^2 \epsilon^{1-\frac{d}{2\beta}} \quad (5)$$

where C, C' are constants depending only on β, d and independent of f and ϵ .

The proof of [Theorem 4](#) can be found in [Appendix B](#), and uses arguments from [Marteau-Ferey et al. \(2022\)](#). A sketch is given in [Section 5](#).

4 Gaussian PSD Models for filtering

Now that we know how to efficiently carry out closed-form filtering operations using Gaussian PSD Models and learn good approximations of non-negative functions using this family of models, we can compute an approximation of the filtering distribution.

4.1 Setting

Consider two discrete Markov Chains $(X_n) \in E^\mathbb{N}$ - the hidden state chain - and $(Y_n) \in F^\mathbb{N}$ - the observations. We assume that (X_t, Y_t) has a Hidden Markov Model structure described by (ν, Q, G) where $\nu \in \mathcal{P}(E)$ is an initial distribution, $Q : E \times E \rightarrow \mathbb{R}$ is a Markov kernel, and $G : E \times F \rightarrow \mathbb{R}$ is a transition kernel. Formally, this can be summarized as: (1) the law of X_n is fully-determined by the knowledge of X_{n-1} , i.e. $\mathbb{P}(X_n \in dx | X_{n-1} = x) = Q(x, dx)$; (2) the law of Y_n is fully-determined by the knowledge of X_n ; i.e.

$\mathbb{P}(Y_n \in dy | X_n = x) = G(x, dy)$; (3) the law of X_0 is given by ν . In particular, we have the usual Markovian structure: $X_n \perp\!\!\!\perp X_l | X_k$ and $Y_n \perp\!\!\!\perp Y_l | X_k$ for any $n > k > l$.

The goal of filtering is to compute the distribution of X_n conditionally on past observation Y_1, \dots, Y_n . We denote this distribution $\pi_n^\nu(z_{1:n}, dx)$ where ν is initial distribution and $z_{1:n}$ are the observations (not necessarily taken from the chain (Y_n)). This distribution is known as the filtering distribution or optimal filter. When clear from context, we drop the dependence of $z_{1:n}$ and ν .

Importantly, $\pi_n(dx)$ can be computed recursively using Q and G and beginning from ν :

$$\pi_n = \bar{R}_n(\pi_{n-1}), \quad \text{where } \bar{R}_n(\mu)(\cdot) := \frac{R_n\mu(\cdot)}{R_n\mu(E)} \quad \text{and} \quad R_n\mu(\cdot) := \int Q(u, \cdot)G(x, y_n)d\mu(u) \quad (6)$$

which recovers (1). Computing $\bar{R}_n(\pi_{n-1})$ is difficult in most circumstances since one must be able to compute products and marginals on probability distributions. Two notable exceptions include the Conditional Linear Gaussian Model (which corresponds to the Kalman filter) and when E is finite (which corresponds to the Baum-Welch algorithm).

4.2 PSD filter

To overcome these difficulties, we approximate Q and G from evaluations using Gaussian PSD Models then compute iteration (6) with these approximate kernels.

Given a sequence of observations $(z_k)_{k \geq 1}$, we define $\hat{R}_k(u, x) = \hat{Q}(u, x)\hat{G}(x, y_k)$ analogously to R_n . The non-linear transformation \tilde{R}_k is defined for any positive, finite measure μ by

$$\tilde{R}_k(\mu) = \frac{\hat{R}_k\mu}{\hat{R}_k\mu(E)}. \quad (7)$$

and summarized in [Algorithm 2](#). Note that at each step, $\hat{\pi}_k$ is a valid, normalized density. However, \hat{Q} and \hat{G} are not properly normalized, i.e. $\hat{G}(u, F) = 1$ is not guaranteed. As long as the initial distribution $\hat{\pi}_0$ is a General PSD Model and \hat{G} and \hat{Q} are valid Gaussian PSD Models, $\hat{\pi}_k$ is a Gaussian PSD Model for all $k \geq 0$ and moreover is a valid density. Importantly, [Corollary 5](#) shows that the order of $\hat{\pi}_k$ is constant for $k \geq 1$ and equal to the product of the orders of \hat{Q} and \hat{G} as shown in the following corollary which follows directly from [Proposition 2](#).

Corollary 5 ($\hat{\pi}_n$ has constant order for any n). *Let $\hat{\pi}_0(x)$ a Gaussian PSD Model on E of order M_0 , $\hat{Q}(u, x)$ a Gaussian PSD Model on $E \times E$ of order M_Q and $\hat{G}(x, y)$ a Gaussian PSD Model on $E \times E$ of order M_G . Let $(z_n) \in F^{\mathbb{N}}$. Let $(\hat{\pi}_n)_{n \in \mathbb{N}}$ the sequence of functions defined by the recursion (7). Then, for any $n \geq 0$, $\hat{\pi}_n$ is a Gaussian PSD Model of order at most $M_Q \times M_G$ and it is computed by [Algorithm 2](#).*

Algorithm 2: PSDFilter algorithm

Data: $z_1, \dots, z_T, \hat{\pi}_0, \hat{Q}, \hat{G}$
for $k = 1, \dots, T$ **do**
 $\beta \leftarrow \text{PRODUCT}(\hat{\pi}_{k-1}, \hat{Q});$
 $\hat{Q}\hat{\pi}_{k-1}(\cdot) \leftarrow \text{MARGINAL}(\beta(u, \cdot), [u]);$
 $\hat{G}_k(\cdot) \leftarrow \text{PARTIALEVAL}(\hat{G}(\cdot, y), y := y_k);$
 $\tilde{\pi}_k \leftarrow \text{PRODUCT}(\hat{Q}\hat{\pi}_{k-1}, \hat{G}_k);$
 $Z \leftarrow \text{INTEGRAL}(\tilde{\pi});$
 $\hat{\pi} \leftarrow \tilde{\pi}/Z;$
end
Result: $\hat{\pi}_1, \dots, \hat{\pi}_T$

4.3 Gaussian PSD Filter Stability and Robustness

In this section, we show that the algorithm introduced in [Section 4.2](#) estimates a distribution which is close to the true filtering distribution in total variation distance. [Theorem 6](#) combines stability properties of the optimal filter with robustness of the iteration (7) to the use of \hat{Q} and \hat{G} in place of Q and G . We introduce the following assumption.

Assumption 2 (R_n is mixing). *There exists $1 > \sigma > 0$ and a probability density $\xi \in \mathcal{P}(E)$ such that for any $n \in \mathbb{N}$, R_n is σ - ξ -mixing, i.e. for any $u, x \in E \times E$,*

$$\sigma \xi(x) \leq R_n(u, x) \leq \frac{1}{\sigma} \xi(x). \quad (8)$$

[Assumption 2](#) is a classical assumption for the study of filtering ([Cappé et al., 2005](#)). We are ready to state the main theorem of this work.

Theorem 6 (PSD filter robustness and stability). *Assume that Q and G verify [Assumptions 1](#) and [2](#). Let $\varepsilon > 0$. When \hat{G}, \hat{Q} are learned using [Algorithm 1](#) and $M, N \in \mathbb{N}$ are chosen such that*

$$M \geq C' (\log(1/\varepsilon))^D \log(1/\varepsilon\delta) \varepsilon^{-\frac{D}{\beta-D/2}}, \quad n \geq C' \varepsilon^{-\frac{D}{\beta-D/2}} \log(1/\delta)$$

then with probability at least $1 - 6\delta$, the following holds: for any $k \in \mathbb{N}$,

$$\|\pi_k - \hat{\pi}_k\|_{TV} \leq \frac{C}{\sigma^2} \left(\frac{1 - \sigma^2}{1 + \sigma^2} \right)^{k-1} \|\pi_0 - \hat{\pi}_0\|_{TV} + \frac{\varepsilon}{\sigma}, \quad (9)$$

where $\hat{\pi}_1, \dots, \hat{\pi}_k$ are computed using [Algorithm 2](#), moreover $C = \frac{2}{\log 3}$ and C' depends only on $\|\sqrt{Q}\|_{W_2^\beta}, \|\sqrt{G}\|_{W_2^\beta}, \beta, d$.

[Theorem 6](#) is proven in [Appendix C](#) and we include a sketch of the proof in [Section 5.1](#).

The theorem above shows that the distance between the probability $\hat{\pi}_k$ resulting from our algorithm and the optimal one π_k corresponding to k steps of sequential Bayesian filtering with the true Q, G is bounded by two terms: (1) the first term accounts for stability and goes to zero exponentially fast in the number of steps k and depends on how close we choose $\hat{\pi}_0$ with respect to π_0 ; (2) the second terms accounts for robustness, it does not increase in k and we can make it arbitrarily small by learning more precise \hat{Q}, \hat{G} , i.e. by increasing M, n . Assuming that $\hat{\pi}_0 = \pi_0$ the proposed algorithm [Algorithm 2](#) for any $k \in \mathbb{N}$ achieves a an error $\|\pi_k - \hat{\pi}_k\|_{TV} \leq \varepsilon/\sigma$ producing a Gaussian PSD model $\hat{\pi}_k$ that satisfies

$$\hat{\pi}_k \text{ of order } O(\varepsilon^{-\frac{2D}{\beta-D/2}}), \text{ and } O(\varepsilon^{-\frac{6D}{\beta-D/2}}),$$

for [Algorithm 2](#). The cost of learning \hat{Q}, \hat{G} via [Algorithm 1](#) is paid only once at the beginning and is of $O(\varepsilon^{-\frac{1.5D}{\beta-D/2}})$, if we use fast algorithms as the ones recalled in [Section 3.2](#). Note that the proposed approach is adaptive to the regularity of the kernels Q and G . In particular, if they are differentiable many times, i.e. $\beta \geq 4.5D$, then the order of $\hat{\pi}_k$ becomes only $O(\varepsilon^{-1/2})$, i.e.

$$\text{memory cost } O(\varepsilon^{-1}), \quad \text{computational cost } O(\varepsilon^{-1.5}).$$

This is remarkable since, for example, particle filter methods are bound to a computational complexity that cannot be smaller than ε^{-2} since they have to approximate an integral via sampling ([Gland and Oudjane, 2004](#)).

5 Sketch of the proof of Theorem 4 and Theorem 6

In this section, we give the main arguments for the proof of Theorem 4. The proof can be found in Appendix B. We denote f the target function and $g = \sqrt{g}$, and \hat{g} the solution to (4) and $\hat{f} = \hat{g}^2$.

As shown in Appendix B, by opening the proof of Proposition 11 in Marteau-Ferey et al. (2022), we can obtain bounds on $\|g - \hat{g}\|_{L^2(\Omega)}$ and (essentially) $\|\hat{g} - g\|_{W_2^m(\Omega)}$ as a function of $\|g\|_{W_2^\beta(\Omega) \cap L^\infty(\Omega)}$ and ϵ with optimal dependence on ϵ . We combine these results to bound $\|f - \hat{f}\|_{L^\infty(\Omega)}$ with optimal rates in ϵ , which is needed to control the error terms in Theorem 6. Since $\|f - \hat{f}\|_{L^\infty(\Omega)} \leq (2\|g\|_{L^\infty(\Omega)} + \|g - \hat{g}\|_{L^\infty(\Omega)}) \|g - \hat{g}\|_{L^\infty(\Omega)}$, we focus on controlling $\|g - \hat{g}\|_{L^\infty(\Omega)}$.

The first key argument is to separate $g - \hat{g}$ as the sum of an approximation error $g - g_{\tau,\epsilon}$ and an estimation error $g_{\tau,\epsilon} - \hat{g}$ where $g_{\tau,\epsilon} \in \mathcal{H}_\eta$ and $\|g - g_{\tau,\epsilon}\|_{L^\infty(\Omega)} \leq C\|g\|\epsilon^{1-\bar{\nu}}$. Using the triangle inequality, controlling $\|g - \hat{g}\|_{L^\infty(\Omega)}$ reduces to controlling $\|g_{\tau,\epsilon} - \hat{g}\|_{L^\infty(\Omega)}$. We apply the Gagliardo-Nirenberg inequality with well-chosen parameters to obtain a bound with the product of the Sobolev norm of the estimation error (which is controlled by the \mathcal{H}_η norm) and of the $L^2(\Omega)$ norm, both of which are bounded by the learning approach. Combining all the terms yields the result in Theorem 4.

5.1 Sketch of the proof of Theorem 6

We start by proving a lemma which controls how errors committed at each step accumulate and are eliminated.

Lemma 7. *Let $T > 1$.*

$$\|\pi_n^\nu - \hat{\pi}_n\|_{TV} \leq \underbrace{C_\sigma \tau_\sigma^n \|\hat{\pi}_0 - \nu\|_{TV}}_{A_n} + \underbrace{\delta_n + C_\sigma \sum_{k=1}^n \tau_\sigma^{n-k-1} \delta_k}_{B_n} \quad (10)$$

where $\delta_k = \|\tilde{R}_k(\hat{\pi}_{k-1}) - \bar{R}_n(\hat{\pi}_{k-1})\|_{TV}$.

The gist of Lemma 7 is that, using the triangle inequality, we can bound $\|\pi_n^\nu - \hat{\pi}_n\|$ by $A_n = \|\pi_n^\nu - \pi_n^{\hat{\pi}_0}\|$ and $B_n = \|\pi_n^{\hat{\pi}_0} - \hat{\pi}_n\|$. Lemma 7 is proven by combining Proposition 21 and Proposition 20 in Appendix C.

The first source of discrepancy is the initialization error. Indeed, $\hat{\pi}_k$ is initialized at $\hat{\pi}_0$ while π_k^ν is initialized at ν and we isolated this question above by considering the discrepancy between two optimal filters initialized

at ν and $\hat{\pi}_0$. The behavior of A_n is known as the stability or forgetting property, and is a property of the optimal filter, i.e. of the dynamical system considered, and not of the algorithms considered. In particular, since R_n is mixing for any $n \geq 0$, A_n decreased exponentially as the application of the filter is contractive (under the mixing assumption).

The second source of discrepancy is the accumulation of errors committed at each step by applying the approximate iteration (7) in place of (6). A telescopic argument in the proof of Lemma 7 shows that the accumulation is limited because the forgetting property of the optimal kernel tends to make past errors δ_j for $j < k$ disappear thanks to the τ^{n-k-1} term. Our argument generalizes the argument in Gland and Oudjane (2004) and relies on the projective properties of the Hilbert metric (Cohen and Fausti, 2023). We then combine Lemma 7 and Theorem 4 to obtain bound B_n by a constant ε as small as we want.

6 Extension: filtering with Generalized Gaussian PSD Models

Throughout this work, we have focused on Gaussian PSD Models. In fact, many of the properties studied above remain valid for a more general family of models we introduce and study in this section.

Seen as a mixture (with potentially negatively weighted components), Gaussian Mixture Models have components which are aligned with the axes of the space. In cases where $f(u, x) = Q(u, x)$ a transition kernel, we know that u and x are strongly "correlated" (indeed, $Q(u, x)$ is a conditional density) and having non-diagonal precision matrices can be useful. When $Q(u, x)$ is a Gaussian Linear Conditional Distribution, one Gaussian component with non-diagonal precision matrix is enough to approximate Q uniformly. We prove this in Theorem 10.

We introduce Generalized Gaussian PSD Models as the family of non-negative combinations of such components. In this section, we show that Generalized Gaussian PSD Models share many of the properties of Gaussian PSD Models. In addition, we show in Theorem 10 that they generalize Kalman filters to more general initial distributions such as multi-modal models.

In this section, for P a positive semi-definite matrix, we denote for any $x, y \in \mathbb{R}^d$, $k_P(x, y) = e^{-(x-y)^\top P(x-y)}$ and $C(P) = \int_{\mathbb{R}^d} k_P(x, 0) dx$.

Definition 2 (Generalized Gaussian PSD model of order M). *A Generalized Gaussian PSD model of order M is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ which can be*

written:

$$f(x) = \text{Tr}(AB(x)) \quad (11)$$

where A is a positive semi-definite matrix of size M and for any $x \in \mathbb{R}^d$, $B(x)$ is a positive semi-definite matrix with entries

$$B(x)_{ij} = e^{C_{ij}} k_{P_{ij}}(x, x_{ij}) \quad (12)$$

where $C_{ij} \in \mathbb{R}$, $x_{ij} \in \mathbb{R}^d$ and P_{ij} is a $d \times d$ positive semi-definite matrix, for any $1 \leq i, j \leq M$.

We use the notation $f(x; \theta)$ where $\theta = (X, A, P, C)$ and $X = (x_{ij})_{1 \leq i, j \leq M}$, $P = (P_{ij})_{1 \leq i, j \leq M}$, and $C = (C_{ij})_{1 \leq i, j \leq M}$.

Like Gaussian PSD Models, Generalized Gaussian PSD Model are a generalization of Gaussian Mixture Models allowing for negative weights. Indeed, by developing (11) with (12), f can be written : $f(x) = \sum_{i=1}^M \sum_{j=1}^M A_{ij} e^{C_{ij}} k_{P_{ij}}(x, x_{ij})$.

6.1 Examples of Generalized Gaussian PSD Models

Below we show that Generalized Gaussian PSD Models generalize Gaussian Mixture Models, Gaussian PSD Models from [Rudi and Ciliberto \(2021\)](#) and indeed, any squared linear combination of Gaussian functions.

Example 3 (Gaussian Mixture Models are Generalized Gaussian PSD Models). *Let $a \in \mathbb{R}_+^d$ such that $\sum_{i=1}^d a_i = 1$. If $f(x) = \sum_{i=1}^M a_i p(x|\mu_i, P_i)$ with $a_i \geq 0$ and $\sum_{i=1}^M a_i = 1$, then f verifies [Definition 2](#) with $A = \text{diag}(a)$, $P_{ii} = P_i/2$, $x_{ii} = x_i/2$ and $C_{ii} = -\frac{1}{2} \log(C(P_i))$.*

Example 4 (Gaussian PSD Models are Generalized Gaussian PSD Models). *Consider a Gaussian PSD Model $f(x) = \Phi_M(x)^\top A \Phi_M(x)$ with $\Phi_M(x) = (k_\eta(x, x_1), \dots, k_\eta(x, x_M))^\top \in \mathbb{R}^M$. Then, f is a Generalized Gaussian PSD Model. Indeed, $f(x) = \text{Tr}(A \Phi_M(x) \Phi_M(x)^\top)$. Each component of the positive semi-definite matrix $\Phi(x) \Phi(x)^\top$ is a function of the form $x \mapsto e^{C_{ij}} k_{2\eta}(x, \frac{x_i + x_j}{\sqrt{2}})$ with $C_{ij} = \|x_i\|^2 + \|x_j\|^2 - \frac{\|x_i + x_j\|^2}{2}$, and f verifies [Definition 2](#).*

Example 5 (Squared Gaussian Linear Models are Generalized Gaussian PSD Models). *Consider the function $f(x) = (w^\top \Phi_M(x))^2$ where $\Phi_M(x) = (k_{P_1}(x, x_1), \dots, k_{P_M}(x, x_M))^\top$ and $w \in \mathbb{R}^d$. Then, f is a Generalized Gaussian PSD Model. Indeed, $f(x) = \text{Tr}(\Phi_M(x)^\top w w^\top \Phi_M(x)) = \text{Tr}(w w^\top \Phi_M(x) \Phi_M(x)^\top)$. Since $w w^\top$ and $\Phi_M(x) \Phi_M(x)^\top$ are both positive semi-definite matrices, $f(x) \geq 0$ and $\Phi_M(x) \Phi_M(x)^\top = e_{ij}^C k_{P_{ij}}(x, x_{ij})$ with $P_{ij} = P_i + P_j$, $x_{ij} = P_{ij}^{-1/2} (P_i x_i + P_j x_j)$.*

6.2 Properties of Generalized Gaussian PSD Models

6.2.1 Closed-form stability with respect to probabilistic operations

Like Gaussian PSD Models, Generalized Gaussian PSD Models are closed under product, partial evaluation, and marginalization.

Proposition 8. *Let $f(x, y; \theta_1)$ and $g(y, z; \theta_2)$ be two Generalized Gaussian PSD Models of order M_1 and M_2 respectively, where all precision matrices are positive definite.*

- **Integral over** $\mathbb{R}^d \int f(x, y; \theta_1) dx dy$ is given by the algorithm `INTEGRAL`(θ_1).
- **Partial evaluation** $f(x, y_0; \theta_1) = h(x; \theta')$ is a Generalized Gaussian PSD Model of order M_1 and θ' is given by the algorithm `PARTIALEVAL`(y_0, θ_1).
- **Product** $f(x, y; \theta_1)g(y, z; \theta_2) = h(x, y, z; \theta')$ is a Generalized Gaussian PSD Model of order $M_1 \times M_2$ and θ' is given by the algorithm `PRODUCT`(θ_1, θ_2).
- **Marginalization** $\int f(x, y; \theta_1) dy = h(x; \theta')$ is a Generalized Gaussian PSD Model of order M_1 and θ' is given by the algorithm `MARGINAL`(y, θ_1).

Proof sketch for the product Using (11), $f(x, y)g(y, z)$ can be written $h(x, y, z) = \text{Tr}(A_1 \otimes A_2 \times B_1(x, y) \otimes B_2(y, z))$. The entries of $B_1(x, y) \otimes B_2(y, z)$ are products of Gaussian functions which can be simplified into the form (12). The complete proof and description of each operation can be found in [Appendix D](#).

6.2.2 Closed-form filtering iteration

An optimal filtering iteration (6) can be written using the four operations of [Proposition 8](#).

Proposition 9. *Let $\mu(x; \theta_\mu)$, $q(x, x'; \theta_q)$ and $g(x, y; \theta_g)$ be three Generalized Gaussian PSD models with order M , M_q and M_g respectively. Let $y \in \mathbb{R}^d$ such that $\int \int q(u, x)g(x, y)\mu(u)du dx > 0$. The density μ' defined by $\mu'(x) = \frac{\int q(u, x)g(x, y)\mu(u)du}{\int \int q(u, x)g(x, y)\mu(u)du dx}$ is a Generalized Gaussian PSD Model with order at most $M \times M_q \times M_g$ whose parameters are given by `FILTERSTEP`($y, \theta_\mu, \theta_q, \theta_g$).*

Applying [Proposition 9](#) recursively as in [Algorithm 2](#) to compute an approximate filter $\hat{\pi}_n$, the order of $\hat{\pi}_n$ increases exponentially with n . A constant number of anchor points can be used by compression $\hat{\pi}_n$ at each step,

for example by learning a Gaussian PSD Model with a given number of anchor points (indeed $\hat{\pi}_n$ is a smooth sum-of-squares), which is justified by [Theorem 4](#).

6.2.3 Generalized Gaussian PSD Models generalize Kalman filters

Conditional Gaussian Linear Distributions are widely used in filtering and dynamical modeling since they cover transition or observation state-space equations such as $X_{t+1} = FX_t + b + \Sigma U_t$ where U_t is Gaussian noise considered in the Kalman filter and extensions.

Theorem 10 (Approximating a Conditional Gaussian Linear Distribution). *Let p be a Conditional Gaussian Linear Distribution defined by $p(y|x) = \mathcal{N}(y|Fx + b, \Sigma)$ with $F \in \mathbb{R}^{d' \times d}$, $b \in \mathbb{R}^{d'}$ and $\Sigma \in \mathcal{S}_d^{++}(\mathbb{R})$. Then for any $\epsilon > 0$ and $R > 0$, there exists a Generalized Gaussian PSD Model of order 1 such that $|p(y|x) - \hat{p}(x, y)| \leq \epsilon, \forall x, y \in \mathbb{R}^d \times \mathbb{R}^{d'}$ such that $\|x\|_2^2 + \|y\|_2^2 \leq R^2$.*

This shows that a Generalized Gaussian PSD Model of order 1 can approximate a Conditional Gaussian Linear Model with arbitrary accuracy on a compact. Note that in this case, applying [Algorithm 2](#) with \hat{Q} and \hat{G} such approximations (each of order 1) and $\hat{\pi}_0$ of order M yields an approximation of π_n of constant order M . The proof of [Theorem 10](#) can be found in [Appendix D.5](#).

Generalized Gaussian PSD Models can be used to learn general transition kernels using non-convex optimization. We discuss this in [Appendix D.6](#).

Acknowledgments

T.C. gratefully acknowledges support from the French National Agency for Research, grant ANR-18-CE40-0016-01. C.C. acknowledges the support of the Royal Society (grant SPREM RGS\R1\201149) and Amazon.com Inc. (Amazon Research Award – ARA). B.G. acknowledges partial support by the U.S. Army Research Laboratory and the U.S. Army Research Office, and by the U.K. Ministry of Defence and the U.K. Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/R013616/1; B.G. also acknowledges partial support from the French National Agency for Research, grants ANR-18-CE40-0016-01 and ANR-18-CE23-0015-02. A.R. acknowledges partial support from the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), and support from the European Research Council (grant REAL 947908).

References

- G. L. Blankenship, W. E. Hopkins, and S. I. Marcus. Asymptotic behavior of nonlinear filters. In *1981 20th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes*, pages 89–89, 1981. doi: 10.1109/CDC.1981.269451.
- Olivier Cappé, Eric Moulines, and Tobias Ryden. *Inference in Hidden Markov Models (Springer Series in Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2005. ISBN 0387402640.
- Samuel N. Cohen and Eliana Fausti. Hyperbolic contractivity and the hilbert metric on probability measures, 2023.
- Yohann De Castro, Élisabeth Gassiat, and Sylvain Le Corff. Consistent estimation of the filtering and marginal smoothing distributions in nonparametric hidden markov models. *IEEE Transactions on Information Theory*, 63(8):4758–4777, 2017. doi: 10.1109/TIT.2017.2696959.
- François Le Gland and Nadia Oudjane. Stability and uniform approximation of nonlinear filters using the Hilbert metric and application to particle filters. *The Annals of Applied Probability*, 14(1):144 – 187, 2004. doi: 10.1214/aoap/1075828050.
- S.J. Julier and J.K. Uhlmann. Unscented filtering and nonlinear estimation.

- Proceedings of the IEEE*, 92(3):401–422, 2004. doi: 10.1109/JPROC.2003.823141.
- Rudolf E. Kálmán and Richard S. Bucy. New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83:95–108, 1961.
- Jin Won Kim. *Duality for nonlinear filtering*. PhD thesis, University of Illinois Urbana-Champaign, 2022.
- F. LeGland. Stability and approximation of nonlinear filters: an information theoretic approach. In *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No.99CH36304)*, volume 2, pages 1889–1894 vol.2, 1999. doi: 10.1109/CDC.1999.830910.
- Dong C. Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Math. Program.*, 45(1–3):503–528, aug 1989. ISSN 0025-5610.
- Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Non-parametric models for non-negative functions. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12816–12826. Curran Associates, Inc., 2020.
- Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Sampling from arbitrary functions via psd models. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 2823–2861. PMLR, 28–30 Mar 2022.
- Curtis McDonald and Serdar Yüksel. Exponential filter stability via Dobrushin’s coefficient. *Electronic Communications in Probability*, 25(none): 1 – 13, 2020. doi: 10.1214/20-ECP333.
- Alexander Yu. Mitrophanov, Alexandre Lomsadze, and Mark Borodovsky. Sensitivity of hidden markov models. *Journal of Applied Probability*, 42 (3):632–642, 2005. doi: 10.1239/jap/1127322017.
- Daniel Ocone and Etienne Pardoux. Asymptotic stability of the optimal filter with respect to its initial condition. *SIAM Journal on Control and Optimization*, 34(1):226–243, 1996. doi: 10.1137/S0363012993256617.

- Alessandro Rudi and Carlo Ciliberto. Psd representations for effective probability models. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 19411–19422. Curran Associates, Inc., 2021.
- Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. Falkon: An optimal large scale kernel method. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Simo Särkkä. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2013.
- Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.

A Markov kernels & Hidden Markov Models

A.1 Tools and notation

Let (E, \mathcal{E}) be a measurable space. We denote $\mathcal{M}(E)$ the set of finite signed measures on (E, \mathcal{E}) , $\mathcal{M}_+(E)$ the set of finite positive measures, $\mathcal{M}_0(E)$ the set of finite signed measures which sum to 0 and $\mathcal{P}(E)$ the set of probability distributions on (E, \mathcal{E}) . Let (F, \mathcal{F}) be a second measurable space.

A.2 Total variation

Definition 3 (Total variation norm). *Let ξ be a finite signed measure on (E, \mathcal{E}) . The total variation of ξ is $\|\xi\|_{TV} = \xi^+(E) + \xi^-(E)$ where ξ^+, ξ^- is the Jordan-Hahn decomposition of ξ . If ξ admits a density with respect to the Lebesgue measure, then $\|\xi\|_{TV} = \int |\xi(x)| dx$.*

Proposition 11. *Let μ and ν be two finite measures on (E, \mathcal{E}) . Then,*

$$\|\bar{\mu} - \bar{\nu}\| = \frac{\|\mu - \nu\|}{\mu(E)} + \frac{|\mu(E) - \nu(E)|}{\mu(E)} \quad (13)$$

In particular,

$$\|\bar{\mu} - \bar{\nu}\| \leq \frac{2\|\mu - \nu\|}{\mu(E)} \quad (14)$$

A.3 Transition kernels

We report the essential results relative to transition kernels taken [Cappé et al. \(2005\)](#).

Definition 4 (Transition kernel). *A function $Q : E \times \mathcal{F} \rightarrow \mathbb{R}^+$ is an unnormalized transition kernel if:*

- *for all $x \in E$, $Q(x, \cdot)$ is a positive measure on (F, \mathcal{F}) ;*
- *for all $A \in \mathcal{F}$, $x \mapsto Q(x, A)$ is measurable.*

Q is normalized if for any $x \in E$, $Q(x, F) = 1$. When $E = F$ and Q is normalized, Q is said to be a Markov transition kernel.

By abuse of notation, when Q admits a density with respect to the Lebesgue measure, we denote it Q as well, i.e. $Q(x, dy) = Q(x, y)dy$.

Note that $R_n(u, x) = Q(u, x)G(x, z_n)$ is indeed an unnormalized transition kernel.

Definition 5 (Effects of kernels). *Let K be an unnormalized kernel on $(E, \mathcal{E}) \times (F, \mathcal{F})$, $\mu \in \mathcal{M}_+(E)$ and f a bounded function $f : E \rightarrow \mathbb{R}$.*

Then, $K\mu \in \mathcal{M}_+(F)$ with for any $A \in \mathcal{F}$,

$$K\mu(A) = \int K(u, A)\mu(du), \quad (15)$$

and $Kf : E \rightarrow \mathbb{R}$ is a bounded function with for any $u \in E$,

$$Kf(u) = \int K(u, A)\mu(du). \quad (16)$$

A.4 Hidden Markov Models

Definition 6 (Hidden Markov Model). *Let (E, \mathcal{E}) and (F, \mathcal{F}) be two measurable spaces. Let Q and G denote a Markov transition kernel on (E, \mathcal{E}) and G denote a transition kernel from (E, \mathcal{E}) to (F, \mathcal{F}) . Let T be the Markov transition kernel defined on the product space $(E \times F, \mathcal{E} \otimes \mathcal{F})$ by*

$$\forall (x, y) \in E \times F, \forall C \in \mathcal{E} \otimes \mathcal{F}, T[(x, y), C] = \int \int 1_C((x', y'))Q(x, dx')G(x', dy') \quad (17)$$

*The Markov Chain $\{X_k, Y_k\}_{k \geq 0}$ with Markov transition kernel T and initial distribution $\nu \otimes G$, where ν is a probability distribution on (E, \mathcal{E}) is called a **Hidden Markov Model**.*

We denote \mathbb{P}_ν and \mathbb{E}_ν the probability measure and corresponding expectation associated with the process $\{(X_k, Y_k)\}$ over $((E \times F)^\mathbb{N}, (\mathcal{E} \otimes \mathcal{F})^{\otimes \mathbb{N}})$.

Throughout this work, we assume that for any $x \in E$, $Q(x, \cdot) \ll \lambda(\cdot)$ and $G(x, \cdot) \ll \lambda(\cdot)$ where λ is the Lebesgue measure over (E, \mathcal{E}) , and we denote $Q(x, dx') = Q(x, x')dx'$ and $G(x, dy) = G(x, y)dy$.

Definition 7 (Filtering distribution). *Let ν be a probability distribution over (E, \mathcal{E}) and $n \geq 0$. We denote π_n^ν the conditional distribution of X_n given $Y_{1:n}$, i.e.*

- π_n^ν is a transition kernel from F^n to E
- π_n^ν satisfies for any bounded function $f : E \rightarrow \mathbb{R}$,

$$\mathbb{E}_\nu[f(X_n)|Y_{1:n}] = \int f(x)\pi_n^\nu(Y_{1:n}, dx) \quad (18)$$

A.5 Mixing kernels

Definition 8 (Mixing kernel). *We say a kernel $K(x, dy)$ is mixing if there exists a positive constant $\sigma > 0$ and a non-negative measure ξ such that for any $x \in E$ and $A \in \mathcal{F}$,*

$$\sigma\xi(A) \leq K(x, A) \leq \frac{1}{\sigma}\xi(A). \quad (19)$$

If K is mixing for a measure ξ and a constant σ we write that K is σ - ξ -mixing.

Remark 1. *Note that we can add the constraint that ξ be normalized. Indeed, if K is σ - ξ -mixing with $\xi(E) \neq 1$, then K is $\bar{\sigma}$ - $\bar{\xi}$ -mixing with $\bar{\sigma} = \sigma \times \min\left(\xi(E), \frac{1}{\xi(E)}\right) \leq \sigma$.*

Proposition 12 (Sufficient condition for mixing when K admits a density). *If $K(x, dy) = \kappa(x, y)dy$ and there exists $\sigma > 0$ and a measure density ξ such that for any $x \in E$ and $y \in F$,*

$$\sigma\xi(y) \leq \kappa(x, y) \leq \frac{1}{\sigma}\xi(y). \quad (20)$$

then K is mixing with constant σ and $\xi(A) = \int 1_A(y)\xi(y)dy$.

Proposition 13. *If K is mixing with σ and ξ , then for any μ ,*

$$\sigma\xi(A) \leq K\mu(A) \leq \frac{1}{\sigma}\xi(A). \quad (21)$$

Proof.

$$\sigma\xi(A) \leq K\mu(A) = \int 1_A(x) \int K(u, dx)\mu(du) \leq \int 1_A(x) \int \frac{1}{\sigma}\xi(dx)\mu(du) = \frac{1}{\sigma}\xi(A) \quad (22)$$

□

A.6 Optimal kernel R_n

Definition 9. *Let $y_n \in F$. Let $R_n : E \times \mathcal{E} \rightarrow \mathbb{R}^+$ a kernel defined by: for any bounded function $f \in \mathcal{F}_b(E)$, and $u \in E$,*

$$R_n(u, f) = \int f(x)Q(u, x)G(x, y_n)dx \quad (23)$$

We call R_n the optimal kernel.

For alternative definitions see for example [Cappé et al., 2005](#), page 220.

A.7 Hilbert metric

Definition 10 (Comparable measures). *Let μ and ν be two measures on (E, \mathcal{E}) . μ and ν are said to be comparable if there exists $\infty > a, b > 0$ such that for any $A \in \mathcal{E}$,*

$$a\nu(A) \leq \mu(A) \leq b\nu(A) \quad (24)$$

Proposition 14. *Let μ and ν be two comparable measures on (E, \mathcal{E}) and let K be an unnormalized transition kernel on (E, \mathcal{E}) . Then, for any $n \geq 0$, $K^n\mu$ and $K^n\nu$ are comparable.*

Proof. By recursion,

$$bK\nu(A) \leq K\mu(A) = \int 1_A(x) \int K(u, dx)\mu(du) \leq aK\nu(A) \quad (25)$$

□

Definition 11 (Hilbert metric). *Let μ and ν be two comparable probability distributions on (E, \mathcal{E}) . The Hilbert metric between μ and ν is defined as*

$$h(\mu, \nu) = \log \left[\frac{\sup_{A \in \mathcal{E}, \nu(A) > 0} \frac{\mu(A)}{\nu(A)}}{\inf_{A \in \mathcal{E}, \nu(A) > 0} \frac{\mu(A)}{\nu(A)}} \right] \quad (26)$$

Proposition 15. *If μ and ν are two comparable probability distributions on (E, \mathcal{E}) and furthermore they both admit densities then,*

$$h(\mu, \nu) = \log \left[\frac{\text{ess sup}_x \frac{\mu(x)}{\nu(x)}}{\text{ess inf}_x \frac{\mu(x)}{\nu(x)}} \right] = \log \left[\left\| \frac{\mu(x)}{\nu(x)} \right\|_\infty \left\| \frac{\nu(x)}{\mu(x)} \right\|_\infty \right] \quad (27)$$

Definition 12 (Birkhoff contraction coefficient). *Let K be an unnormalized transition kernel. Define $\tau(K)$ such that*

$$\tau(K) = \sup_{0 < h(\mu, \nu) < \infty} \frac{h(K\mu, K\nu)}{h(\mu, \nu)} \quad (28)$$

where the supremum is taken over comparable, positive measures μ and ν .

Proposition 16 (Properties of the Birkhoff coefficient). *Let K be an unnormalized transition kernel.*

- τ is sub-multiplicative, i.e. $\tau(KL) \leq \tau(K)\tau(L)$
- $\tau \leq 1$

- if in addition K is σ - ξ -mixing, then $\tau(K) \leq \frac{1-\sigma^2}{1+\sigma^2}$

These properties are proven in [Cohen and Fausti \(2023\)](#), which studies the Hilbert metric in detail.

Proof. We have $\tau(K) = \tanh \left[\frac{1}{4} \Delta(K) \right]$ where $\Delta(K) = \sup_{\mu, \mu'} h(K\mu, K\mu')$ (Δ for diameter). Since K is σ - ξ -mixing, if μ and ν are two finite measures on (E, \mathcal{E}) ,

$$\sigma^2 K\mu(A) \leq \sigma\xi(A) \leq K\mu(A) \leq \frac{1}{\sigma}\xi(A) \leq \frac{1}{\sigma^2}K\mu'(A) \quad (29)$$

So, $h(K) \leq \log \left(\frac{1}{\sigma^4} \right) \leq \frac{1-\sigma^2}{1+\sigma^2}$. □

Proposition 17 (Total variation - Hilbert comparisons). *Without any hypotheses on μ, ν finite measures, if $\bar{\mu} = \mu/\mu(E)$ and $\bar{\nu} = \nu/\nu(E)$ are normalized counterparts to μ and ν , then*

$$\|\bar{\mu} - \bar{\nu}\| \leq \frac{2}{\log(3)} h(\mu, \nu) \quad (30)$$

If in addition, K is an σ -mixing kernel,

$$h(K\mu, K\nu) \leq \frac{1}{\sigma^2} \|\mu - \nu\| \quad (31)$$

These inequalities are proven in [Cohen and Fausti \(2023\)](#).

B Proof of [Theorem 4](#)

Notation and background results In this section, $\|\cdot\|$ without any subscript denotes $\|\cdot\|_{W_2^\beta(\Omega)}$. Let f a target function defined on $\Omega = (-1, 1)^d$ that verifies [Assumption 1](#), and let $g = \sqrt{f}$. Let $\nu = \min(1, d/2\beta)$. Let $\epsilon > 0$ and $\delta > 0$. In this proof all constants C and exponents $\alpha, \gamma, \rho, \dots$ are independent of f, g, \hat{g}, \hat{f} unless otherwise stated. Only $\beta, \nu, \tilde{n}u$ and θ have importance. We recall also the Gagliardo-Nirenberg inequality, that will be used later.

Lemma 18 (Gagliardo-Nirenberg inequality [Wendland \(2004\)](#)). *Let $u \in L^\infty(\mathbb{R}^d) \cap W^{m,2}(\mathbb{R}^d)$ then*

$$\|u\|_{L^\infty(\Omega)} \leq C \|u\|_{W^{m,2}(\Omega)}^\theta \|u\|_{L^2(\Omega)}^{1-\theta} \quad (32)$$

where $\theta = \frac{d}{m}$ and C is independent of u .

Setting the function space \mathcal{H}_η and $g_{\tau,\epsilon}$ Set $\tau = \epsilon^{-2/\beta}$ and $\lambda = \epsilon^{\frac{2\beta+d}{\beta}}$ and \mathcal{H}_η the reproducing kernel Hilbert space associated to k_η where $\eta = \tau 1_d$.

Existence and properties of $g_{\tau,\epsilon}$ As a consequence of the Stein extension theorem (see Corollary A.3 of [Rudi and Ciliberto \(2021\)](#)), there exists a function $\tilde{g} \in W_2^\beta(\Omega)$ such that $\tilde{g}|_\Omega = g$ and $\|\tilde{g}\|_{W_2^\beta(\Omega)} \leq \|g\|_{W_2^\beta(\Omega)}$ and $\|\tilde{g}\|_{L^\infty(\Omega)} \leq C\|g\|_{L^\infty(\Omega)}$. According to [Rudi and Ciliberto \(2021\)](#) and [Marteau-Ferey et al. \(2022\)](#) (Proposition 7) there exists $g_{\tau,\epsilon} \in \mathcal{H}_\eta$ and some constants C_1, C_2 depending only on β, d and independent of g such that:

$$\|g - g_\tau\|_{L^\infty(\Omega)} \leq C_1 \epsilon^{1-\nu} \|g\|_{W_2^\beta(\Omega)}, \quad (33)$$

$$\|g_\tau\|_{\mathcal{H}_\eta} \leq C_2 \|g\|_{W_2^\beta(\Omega)} \epsilon^{-d/2\beta}. \quad (34)$$

Learning \hat{g} and \hat{f} Let $M, n \in \mathbb{N}$, draw $X \in \mathbb{R}^{n \times d}$ the training set and $\tilde{X} \in \mathbb{R}^{M \times d}$ the set of anchor points. Define $y \in \mathbb{R}^n$ such that $y_i = g(x_i)$ for any $1 \leq i \leq n$ and x_i is the i -th row of X . We formalize (4) explicitly as a kernel ridge regression problem below:

$$\min_{a \in \mathbb{R}^d} \frac{1}{n} |a^T \Phi_\eta(X_i) - y_i|^2 + \lambda a^T K a \quad (35)$$

where $\Phi_\eta(x) = (k_\eta(x, \tilde{x}_1) \dots k_\eta(x, \tilde{x}_M))^T \in \mathbb{R}^M$ and $K \in \mathbb{R}^{M \times M}$ is given by $K_{ij} = k_\eta(\tilde{x}_i, \tilde{x}_j)$. Problem (35) is strongly convex and has a unique solution \hat{a} . We denote \hat{g} the Gaussian Linear Model defined by \hat{a}, η , and \tilde{X} , i.e. such that for any $x \in \mathbb{R}^d$, $\hat{g}(x) = \hat{a}^T \Phi_\eta(x)$. Define \hat{f} the Gaussian PSD Model $\hat{f} = \hat{g}^2$ where $\hat{f}(x) = \Phi_\eta(x)^T \hat{a} \hat{a}^T \Phi_\eta(x)$. The analysis of [Marteau-Ferey et al. \(2022\)](#) shows that when $M \geq C' \log^d(\frac{1}{\epsilon}) \log(\frac{1}{\delta\epsilon}) \epsilon^{-d/\beta}$, and $n \geq C' \epsilon^{-d/\beta} \log \frac{1}{\delta}$, then there exist two constants C_3 and C_4 independent of g and \hat{g} such that \hat{g} verifies the following inequalities each with probability at least $1 - \delta$,

$$\|\hat{g}\|_{\mathcal{H}} \leq C_3 \|g\|_{W_2^\beta(\Omega)} \epsilon^{-d/2\beta} \quad (36)$$

$$\|g - \hat{g}\|_{L^2(\Omega)} \leq C_4 \|g\|_{W_2^\beta(\Omega)} \epsilon. \quad (37)$$

Deriving the bound for $g - \hat{g}$ in $L^\infty(\Omega)$ In particular, using the triangle inequality and combining (33) and (37), there exists a constant C_5 such that with probability at least $1 - \delta$,

$$\|g_{\tau,\epsilon} - \hat{g}\|_{L^2(\Omega)} \leq C_5 \|g\|_{W_2^\beta(\Omega)} \epsilon. \quad (38)$$

We now have all the ingredients to bound $\|g - \hat{g}\|_{L^\infty(\Omega)}$ in high probability. First, notice that :

$$\|g - \hat{g}\|_{L^\infty(\Omega)} \leq \|g - g_{\tau,\epsilon}\|_{L^\infty(\Omega)} + \|g_{\tau,\epsilon} - \hat{g}\|_{L^\infty(\Omega)}. \quad (39)$$

We apply the Gagliardo-Nirenberg inequality (Lemma 18) to $\|g_{\tau,\epsilon} - \hat{g}\|_{L^\infty(\Omega)}$:

$$\|\hat{g} - g_{\tau,\epsilon}\|_{L^\infty(\Omega)} \leq C \|\hat{g} - g_{\tau,\epsilon}\|_{W_2^m(\Omega)}^\theta \|\hat{g} - g_{\tau,\epsilon}\|_{L^2(\Omega)}^{1-\theta}. \quad (40)$$

with $\theta = \frac{d}{2m}$ and $m \geq d/2$ (we fix m when we optimize the exponents below) and C is a constant independent of \hat{g} and $g_{\tau,\epsilon}$.

The Sobolev norm $\|h\|_{W_2^m(\mathbb{R}^d)}$ is upper bounded by the rkhs norm $\|h\|_{\mathcal{H}_\eta}$ for $h \in \mathcal{H}_\eta \subset W_2^m(\mathbb{R}^d)$. Thus, applying the triangle inequality and bounds (34), (36) and (38) there exists a constant C_6 such that with probability at least $1 - 2\delta$,

$$\|\hat{g} - g_{\tau,\epsilon}\|_{L^\infty(\mathbb{R}^d)} \leq C_6 \|g\| \epsilon^{1-\theta-\theta d/2\beta}. \quad (41)$$

Combining (33) and (41), there exist two constants $C, C' > 0$ such that with probability at least $1 - 2\delta$,

$$\|g - \hat{g}\|_{L^\infty(\Omega)} \leq C \|g\|_{W_2^\beta(\Omega)} \epsilon^{1-\nu} + C' \|g\|_{W_2^\beta(\Omega)} \epsilon^{1-\theta-\theta d/2\beta} \quad (42)$$

Choosing $m = \beta + \frac{d}{2}$, there exists a constant $C_7 > 0$ such that for ϵ small enough with probability at least $1 - 2\delta$,

$$\|g - \hat{g}\|_{L^\infty(\Omega)} \leq C \|g\|_{W_2^\beta(\Omega)} \epsilon^{1-d/2\beta}. \quad (43)$$

Bounding $\|g + \hat{g}\|_{L^\infty(\Omega)}$ Using the triangle inequality and (36), there exists a constant $C_8 > 0$ and $\rho > 0$ such that with probability at least $1 - \delta$,

$$\|g + \hat{g}\|_{L^\infty(\Omega)} \leq 2\|g\|_{L^\infty(\Omega)} + \|g - \hat{g}\|_{L^\infty(\Omega)} \leq 2\|g\|_{L^\infty(\Omega)} + C \|g\|_{W_2^\beta(\Omega)} \epsilon^{1-d/2\beta}. \quad (44)$$

Bounding $\|f - \hat{f}\|_\infty$ Notice that since $a^2 - b^2 = (a - b)(a + b)$,

$$\|\hat{f} - f\|_{L^\infty(\Omega)} \leq \|g - \hat{g}\|_{L^\infty(\Omega)} \|g + \hat{g}\|_{L^\infty(\Omega)}. \quad (45)$$

Combining the above bounds:

$$\|\hat{f} - f\|_{L^\infty(\Omega)} \leq C \|g\|_{W_2^\beta(\Omega)}^2 \epsilon^{2-d/\beta} + C' \|g\|_{W_2^\beta(\Omega)}^2 \epsilon^{1-d/2\beta} \quad (46)$$

Since $\beta > d/2$, under the conditions on M, n , there exists a constant $C'' = C + C'$ depending only on Ω, d, β and independent of f, g, \hat{g}, \hat{f} such that with probability at least $1 - 3\delta$,

$$\|\hat{f} - f\|_{L^\infty(\Omega)} \leq C'' \|g\|_{W_2^\beta(\Omega)}^2 \epsilon^{1-d/2\beta} \quad (47)$$

C Proof of Theorem 6

In this section $E = (-1, 1)^d$ and $\Omega = E \times F$. Without loss of generality, we assume that $F = E$. Of course, any compact can be considered.

C.1 Propagation of one-step errors

We generalize the proof technique in [Gland and Oudjane \(2004\)](#) to take into consideration general sequences of densities.

Proposition 19. *Let $\pi_0 \in \mathcal{P}(E)$. Let $(z_k)_{k \geq 1}$ and $(\pi_k)_{k \geq 1}$ the optimal filter sequence computed on the $(z_k)_{k \geq 1}$ and initialized at π_0 . Let $(\mu_k) \in \mathcal{P}(E)^{\mathbb{N}}$ a sequence of distributions such that $\mu_0 = \pi_0$. Then, for any $n \geq 0$:*

$$\mu_n - \pi_n = \sum_{k=1}^n \bar{R}_{n:k+1}(\mu_k) - \bar{R}_{n:k}(\mu_{k-1}) \quad (48)$$

with the notation that $\bar{R}_{n:k} = \bar{R}_n \circ \dots \circ \bar{R}_k$ and $R_{n+1:n} = id$ if $k > l$.

Proof. Telescopic sum:

$$\begin{aligned} \mu_n - \pi_n &= \sum_{k=1}^n \bar{R}_{n:k+1}(\mu_k) - \bar{R}_{n:k}(\mu_{k-1}) \quad (49) \\ &= \mu_n - \bar{R}_n(\mu_{n-1}) + \bar{R}_n(\mu_{n-1} - \bar{R}_{n:n-1}(\mu_{n-2})) \dots + \bar{R}_{n:2}(\mu_1) - \underbrace{\bar{R}_{n:1}(\mu_0)}_{=\bar{R}_{n:1}(\pi_0)=\pi_n} \quad (50) \end{aligned}$$

□

Proposition 20 (Optimal filter stability). *Let (π_n^ν) and (π_n^μ) two sequences of optimal filters initialized at ν and μ respectively, and computed on the same data sequence z_1, \dots, z_n . We assume that for all $n \geq 1$, R_n verifies [Assumption 2](#). Then,*

$$\|\pi_n^\nu - \pi_n^\mu\| \leq \frac{2}{\sigma^2 \log 3} \left(\frac{1 - \sigma^2}{1 + \sigma^2} \right)^{n-1} \|\mu - \nu\|_{TV} \quad (51)$$

Proof. By [\(30\)](#),

$$\|\pi_n^\mu - \pi_n^\nu\|_{TV} \leq h(R_{n:1}(\mu_n), R_{n:1}(\nu_n)) \quad (52)$$

Since R_n, \dots, R_1 are mixing, $R_1\mu$ and $R_1\nu$ are comparable and we can apply [Proposition 16](#) with the Hilbert contraction coefficient $\tau(R_{n:2}) \leq \tau_\sigma^{n-1} \leq \left(\frac{1-\sigma^2}{1+\sigma^2}\right)^{n-1}$.

$$\|\pi_n^\mu - \pi_n^\nu\|_{TV} \leq \frac{2}{\log 3} \tau_\sigma^{n-1} h(R_1(\mu), R_1(\nu)) \quad (53)$$

And finally, applying [\(31\)](#),

$$\|\pi_n^\mu - \pi_n^\nu\|_{TV} \leq \frac{2}{\sigma^2 \log(3)} \tau_\sigma^{n-1} \|\mu - \nu\|_{TV} \quad (54)$$

□

Proposition 21. *Let $E = (-1, 1)^d$. Let $\sigma > 0$ and $\xi \in \mathcal{P}(E)$ such that the optimal kernel $R_n : E \times E \rightarrow \mathbb{R}^+$ is σ - ξ mixing for all n . Let μ_n a sequence of positive, finite measures on E such that $\mu_0 = \pi_0$. Then,*

$$\|\mu_n - \pi_n\|_{TV} \leq \delta_n + \frac{2}{\log 3} \frac{1}{\sigma^2} \sum_{k=1}^{n-1} \tau^{n-k-1} \delta_k \quad (55)$$

where $\delta_n = \|\mu_n - \bar{R}_n(\mu_{n-1})\|_{TV}$.

Proof. Applying the triangle inequality with the total variation norm to [\(48\)](#),

$$\|\mu_n - \pi_n\|_{TV} \leq \sum_{k=1}^n \|\bar{R}_{n:k+1}(\mu_k) - \bar{R}_{n:k}(\mu_{k-1})\|_{TV} \quad (56)$$

$$\leq \|\mu_n - \bar{R}_n(\mu_{n-1})\|_{TV} + \sum_{k=1}^{n-1} \|\bar{R}_{n:k+1}(\mu_k) - \bar{R}_{n:k}(\mu_{k-1})\|_{TV} \quad (57)$$

We apply [\(30\)](#) from [Proposition 17](#),

$$\|\mu_n - \pi_n\|_{TV} \leq \|\mu_n - \bar{R}_n(\mu_{n-1})\|_{TV} + \frac{2}{\log 3} \sum_{k=1}^{n-1} h(R_{n:k+1}(\mu_k), R_{n:k}(\mu_{k-1})) \quad (58)$$

Since R_{k+1} and R_k are mixing, $R_{k+1}\mu_k$ and $R_{k+1}R_k\mu_{k-1}$ are comparable and we can apply [Proposition 16](#) with the Hilbert contraction coefficient $\tau(R_{n:k+2}) \leq \tau_\sigma^{n-k-1} = \left(\frac{1-\sigma^2}{1+\sigma^2}\right)^{n-k-1}$.

$$\|\mu_n - \pi_n\|_{TV} \leq \|\mu_n - \bar{R}_n(\mu_{n-1})\| + \frac{2}{\log 3} \sum_{k=1}^{n-1} \tau_\sigma^{n-k-1} h(R_{k+1}\mu_k, R_{k+1}R_k\mu_{k-1}) \quad (59)$$

Since R_{k+1} is mixing, we apply (31) from Proposition 17:

$$\|\mu_n - \pi_n\|_{TV} \leq \|\mu_n - \bar{R}_n(\mu_{n-1})\| + \frac{2}{\sigma^2 \log 3} \sum_{k=1}^{n-1} \tau_\sigma^{n-k-1} \frac{1}{\sigma^2} \|\mu_k - \bar{R}_k(\mu_{k-1})\|_{TV} \quad (60)$$

where we use that the Hilbert contraction coefficient is sub-multiplicative and can be bounded away from 1 as a function of σ (we denote it τ_σ this upper-bound given in Proposition 16). By denoting $\delta_n = \|\bar{\mu}_n - \bar{R}_n(\bar{\mu}_{n-1})\|_{TV}$ we obtain the result. \square

We now prove Lemma 7.

Proof. We apply the triangle inequality to $\|\pi_n^\nu - \hat{\pi}_0\|_{TV}$:

$$\|\pi_n^\nu - \hat{\pi}_0\|_{TV} \leq \|\pi_n^\nu - \pi_n^{\hat{\pi}_0}\|_{TV} + \|\pi_n^{\hat{\pi}_0} - \hat{\pi}_n\|_{TV} \quad (61)$$

and the result follows from Proposition 20 and Proposition 21. \square

C.2 Bound with smooth, bounded approximations

Let $\hat{Q} : E \times E \rightarrow \mathbb{R}^+$ and $\hat{G} : E \times F \rightarrow \mathbb{R}^+$ two bounded approximations of Q and G . Given a sequence $(z_k) \in F^{\mathbb{N}}$, define the approximate non-negative kernel $\hat{R}_n(u, x) = \hat{Q}(u, x)\hat{G}(x, z_n)$ defined on $E \times E$. We introduce $(\hat{\pi}_n)$ the sequence of probability distributions computed using the recursion $\hat{\pi}_n = \hat{R}_n \hat{\pi}_{n-1} / R_n \hat{\pi}_{n-1}(E) = \tilde{R}_n(\hat{\pi}_{n-1})$.

Proposition 22. *Let $\delta_n = \|\tilde{R}_n(\hat{\pi}_{n-1}) - \bar{R}_n(\hat{\pi}_{n-1})\|_{TV}$. Then,*

$$\delta_n \leq \frac{2}{\sigma \xi(E)} \left(\|G - \hat{G}\|_{L^\infty(E \times E)} + C_d \|\hat{G}\|_{L^\infty(E \times E)} \|Q - \hat{Q}\|_{L^\infty(E \times E)} \right) \quad (62)$$

where C_d is a constant independent of Q, G, \hat{G}, \hat{Q} .

Proof. δ_n is the total variation distance between two probability distributions. Recall that

$$\|\bar{\mu} - \bar{\nu}\|_{TV} \leq \frac{2}{\mu(E)} \|\mu - \nu\|_{TV} \quad (63)$$

for any μ, ν positive, finite measures on E and $\bar{\mu}, \bar{\nu}$ their normalized counterparts. Thus,

$$\delta_n \leq \frac{2}{R_n \hat{\pi}(E)} \|\hat{R}_n(\hat{\pi}_{n-1}) - R_n(\hat{\pi}_{n-1})\|_{TV}. \quad (64)$$

Recall that for any $\mu \in \mathcal{P}(E)$, $\sigma\xi(E) \leq R_n\mu(E) \leq \frac{1}{\sigma}\xi(E)$, which allows us to control the denominator:

$$\delta_n \leq \frac{2}{\sigma\xi(E)} \|\hat{R}_n(\hat{\pi}_{n-1}) - R_n(\hat{\pi}_{n-1})\|_{TV}. \quad (65)$$

Because we will be able to bound the quality of approximation between Q and \hat{Q} (and between G and \hat{G}), we split the above expression:

$$\|\hat{R}_n(\hat{\pi}_{n-1}) - R_n(\hat{\pi}_{n-1})\|_{TV} \leq \int \int \hat{\pi}_{n-1}(u) |Q(u, x)G(x, z_n) - \hat{Q}(u, x)\hat{G}(x, z_n)| dx du \quad (66)$$

$$\leq \underbrace{\int \int Q(u, x) \hat{\pi}_{n-1}(u) |G(x, z_n) - \hat{G}(x, z_n)| dx du}_{A_n} \quad (67)$$

$$+ \underbrace{\int \int \hat{G}(u, x) \hat{\pi}_{n-1}(u) |Q(u, x) - \hat{Q}(u, x)| dx du}_{B_n}. \quad (68)$$

First, let us bound A_n .

$$A_n = \int \int Q(u, x) \hat{\pi}_{n-1}(u) |G(x, y_n) - \hat{G}(x, y_n)| dx du \quad (69)$$

$$\leq \|G - \hat{G}\|_{L^\infty(E \times E)} \int \hat{\pi}_{n-1}(u) \int Q(u, x) dx du \quad (70)$$

$$= \|G - \hat{G}\|_{L^\infty(E \times E)}. \quad (71)$$

where we used that Q is a transition kernel (i.e. that $Q(u, E) = 1$ for all $u \in E$) and that $\hat{\pi}_{n-1}$ is a distribution.

Second, let us bound B_n :

$$B_n = \int \int \hat{G}(x, z_n) \hat{\pi}_{n-1}(u) |Q(u, x) - \hat{Q}(u, x)| dx du \quad (72)$$

$$\leq \|Q - \hat{Q}\|_{L^\infty(E \times E)} \int \hat{\pi}_{n-1}(u) du \int \hat{G}(x, z_n) dx \quad (73)$$

$$\leq \text{Vol}(\Omega) \|\hat{G}\|_{L^\infty(E \times E)} \|Q - \hat{Q}\|_{L^\infty(E \times E)} \quad (74)$$

where we again used that $\hat{\pi}_{n-1}$ is a probability distribution. \square

C.3 Putting everything together

We assume that $E = F$ without loss of generality (simply replace $D = 2d$ for Q and $D = d + d'$ for G). We choose $\beta > D/2$. Thus, here $\Omega = E \times E \subset \mathbb{R}^D$ where $D = 2d$.

Let $\gamma > 0$ and apply [Theorem 4](#) with its parameter $\epsilon = \gamma$, Then we have that there exist \hat{G} and \hat{Q} and two constants C_1, C_2 such that

$$\|\hat{G} - G\|_{L^\infty(\Omega)} \leq C_1 \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}}, \quad \|\hat{Q} - Q\|_{L^\infty(\Omega)} \leq C_2 \|\sqrt{Q}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}}$$

where C_1 and C_2 are independent of Q, \hat{Q}, \hat{G}, G .

By combining these inequalities with [Proposition 22](#), since $\xi(E) = 1$:

$$\delta_n \leq \frac{2}{\sigma} \left(C \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}} + C_d \|\hat{G}\|_{L^\infty(\Omega)}^2 \|\sqrt{Q}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}} \right). \quad (75)$$

Since $\|\hat{G}\|_{L^\infty(\Omega)} \leq \|G - \hat{G}\|_{L^\infty(\Omega)} + 2\|\sqrt{G}\|_{L^\infty(\Omega)}^2 \leq \|G - \hat{G}\|_{L^\infty(\Omega)} + 2\|\sqrt{G}\|_{W_2^\beta(\Omega)}^2$,

$$\delta_n \leq \frac{2}{\sigma} \left(C \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}} + \left(C + C\epsilon^{1-\frac{D}{2\beta}} \right) \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 \|\sqrt{Q}\|_{W_2^\beta(\Omega)}^2 \gamma^{1-\frac{D}{2\beta}} \right)$$

Then, by upper-bounding the negligible terms, there exist two constants $C_1, C_2 > 0$ independent of Q, \hat{Q}, G, \hat{G} , such that :

$$\delta_n \leq C_1 \underbrace{\left(C_2 \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 + C_2 \|\sqrt{G}\|_{W_2^\beta(\Omega)}^2 \|\sqrt{Q}\|_{W_2^\beta(\Omega)}^2 \right)}_{C'(Q,G)} \frac{\gamma^{1-\frac{D}{2\beta}}}{\sigma} \quad (76)$$

Note that C' only depends on parameters on $\|\sqrt{G}\|_{W_2^\beta(\Omega)}, \|\sqrt{Q}\|_{W_2^\beta(\Omega)}, d, \beta, \Omega$ and does not depend on $\hat{Q}, \hat{G}, \sigma, \gamma$.

We now apply [Theorem 4](#). Let $\epsilon > 0$ and $\delta > 0$. As a consequence of the development above, if γ is chosen such that $\gamma = \left(\frac{\epsilon}{C'}\right)^{\frac{2\beta}{2\beta-D}} / \sigma$ then (1) M, n correspond to the ones stated in the statement of the theorem (2) with probability at least $1 - 6\delta$,

$$\delta_n \leq \frac{\epsilon}{\sigma} \quad (77)$$

and then (3),

D Computations on Generalized Gaussian PSD Models

The stability properties of Generalized Gaussian PSD Models under probabilistic operations rely at a high-level on the fact that $Tr(AB)Tr(CD) =$

$Tr(A \otimes CB \otimes D)$ and that if A and B are positive semi-definite matrices then so is $A \otimes B$.

D.1 INTEGRAL

Proposition 23 (Integration of a Generalized Gaussian PSD Model). *Let $f(x) = Tr(AB(x))$ with parameters $\{A, C, (P_{ij}), (\mu_{ij})\}$. Then, $Z = \int f(x)dx$ where*

$$Z = Tr(A \circ \exp^\circ(C) \circ C(P)) \quad (78)$$

where \exp° is the element-wise exponential map and $C(P) \in \mathcal{S}(\mathbb{R}^M)$ is described by $C(P)_{ij} = C(P_{ij})$.

We denote \circ the Hadamard product.

Proof. The proof is clear by linearity of the trace. \square

Remark 2 (Computational complexity). *Because of the need to compute the determinant of P_{ij} the computational complexity of the partial evaluation operation is $O(M^2d!)$.*

D.2 PARTIALEVAL

Proposition 24 (Partial evaluation of a Generalized Gaussian PSD Model). *Let $f(x, y) = Tr(AB(x, y))$ with parameters $\{A, C, (P_{ij}), (\mu_{ij})\}$ and $y \in \mathbb{R}^d$. Then, $g(x) := f(x, y) = Tr(A'B'(x))$ with parameters $\{A', C', (P'_{ij}), (\mu'_{ij})\}$ where*

$$A' = A \quad (79)$$

$$P'_{ij} = P_{ijxx} \quad (80)$$

$$\mu'_{ij} = \mu_{ijx} + P_{ijxx}^{-1} P_{ijxy} (\mu_{ijy} - y) \quad (81)$$

$$C'_{ij} = C_{ij} + \nu P_{xx}^{-1} \nu - y P_{yy} y + 2\mu_x P_{xy} y + 2\mu_y P_{yy} y - \mu P \mu \quad (82)$$

where

$$P = \begin{pmatrix} P_{xx} & P_{xy} \\ P_{xy}^T & P_{yy} \end{pmatrix}. \quad (83)$$

Proof. We can compute $B'(x)$ by expanding $C - \log(B(x, y)_{ij})$ for any i, j .

Dropping the i, j dependence:

$$\left\| P^{1/2} \begin{pmatrix} x \\ y \end{pmatrix} - \mu \right\|^2 = x^T P_{xx} x + 2y P_{xy}^T x + y P_{yy} y - 2\mu P \begin{bmatrix} x \\ y \end{bmatrix} + \mu P \mu \quad (84)$$

$$= x^T P_{xx} x + 2y P_{xy}^T x + y P_{yy} y - 2\mu_x P_{xx} x - 2\mu_x P_{xy} y - 2\mu_y P_{xy} x - 2\mu_y P_{yy} y + \mu P \mu \quad (85)$$

$$= x^T P_{xx} x - 2 \underbrace{(\mu_x P_{xx} + \mu_y P_{xy} - y P_{xy}^T)}_{\nu} x + y P_{yy} y - 2\mu_x P_{xy} y - 2\mu_y P_{yy} y + \mu P \mu \quad (86)$$

$$= \left\| P_{xx}^{1/2} (x - P_{xx}^{-1} \nu) \right\|^2 - \nu P_{xx}^{-1} \nu + y P_{yy} y - 2\mu_x P_{xy} y - 2\mu_y P_{yy} y + \mu P \mu. \quad (87)$$

□

Remark 3 (Computational complexity). *Because of the need to compute the inverse of P_{xx} , the computational complexity of the partial evaluation operation is $O(M^2 d_x^3)$.*

D.3 MARGINALIZATION

Proposition 25 (Marginalization of a Generalized Gaussian PSD Model).

Let $f(x, y) = \text{Tr}(AB(x, y))$ with parameters $\{A, C, (P_{ij}), (\mu_{ij})\}$. Then, $h(x) := \int f(x, y) dy = \text{Tr}(A'B'(x))$ with parameters $\{A', C', (P'_{ij}), (\mu'_{ij})\}$ where

$$A' = A \quad (88)$$

$$P'_{ij} = \left([P_{ij}^{-1}]_{xx} \right)^{-1} \quad (89)$$

$$\mu'_{ij} = [\mu_{ij}]_x \quad (90)$$

$$C'_{ij} = C_{ij} + \log(C(P_{ij})) - \log(C(P'_{ij})) \quad (91)$$

Proof. We compute the integral component-wise, denoting $\Sigma = (2P)^{-1}$:

$$\int B(x, y)_{ij} dy = e^{C_{ij}} \int \exp \left(-\frac{\left\| \sqrt{2} P_{ij}^{1/2} \begin{bmatrix} x \\ y \end{bmatrix} - \mu_{ij} \right\|^2}{2} \right) dy \quad (92)$$

$$= e^{C_{ij}} \int \exp \left(-\frac{\left\| \Sigma^{-1/2} \begin{bmatrix} x \\ y \end{bmatrix} - \mu_{ij} \right\|^2}{2} \right) dy \quad (93)$$

$$= e^{C_{ij}} \sqrt{(2\pi)^{d_x+d_y} |\Sigma|} \int \frac{1}{\sqrt{(2\pi)^{d_x+d_y} |\Sigma|}} \exp \left(-\frac{\left\| \Sigma^{-1/2} \begin{bmatrix} x \\ y \end{bmatrix} - \mu_{ij} \right\|^2}{2} \right) dy \quad (94)$$

$$= e^{C_{ij}} \sqrt{(2\pi)^{d_x+d_y} |\Sigma|} \frac{1}{\sqrt{(2\pi)^{d_x} |\Sigma]_{xx}|}} \exp \left(-\frac{\left\| [\Sigma]_{xx}^{-1/2} (x - [\mu_{ij}]_x) \right\|^2}{2} \right) \quad (95)$$

$$= e^{C_{ij}} \frac{C_{P_{ij}}}{C_{P'_{ij}}} \exp \left(-\left\| P'_{ij}{}^{1/2} (x - [\mu_{ij}]_x) \right\|^2 \right) \quad (96)$$

$$= e^{C'_{ij}} \exp \left(-\left\| P'_{ij}{}^{1/2} (x - \mu'_{ij}) \right\|^2 \right) \quad (97)$$

where $C'_{ij} = C_{ij} + \log(C_{P_{ij}}) - \log(C_{P'_{ij}})$, $P'_{ij} = \left([P_{ij}^{-1}]_{xx} \right)^{-1}$ and $\mu'_{ij} = [\mu_{ij}]_x$. \square

Remark 4 (Computational complexity). *Because of the need to compute the determinant of P as well as invert it, the computational complexity of the partial evaluation operation is $O(M^2 \max(d!, d^3))$.*

D.4 PRODUCT

Proposition 26 (Product of two Generalized Gaussian PSD Models). *Let $f(x, y) = \text{Tr}(AB(x, y))$ a generalized PSD model of order M with parameters $\{A, C, (P_{ij}), (\mu_{ij})\}$. Let $g(x, y) = \text{Tr}(\tilde{A}\tilde{B}(x, y))$ a generalized PSD model of order m with parameters $\{\tilde{A}, \tilde{C}, (\tilde{P}_{kl}), (\tilde{\mu}_{kl})\}$.*

Then, $h(x) := f(x)g(x) = \text{Tr}(A'B'(x))$ is a generalized PSD model of order Mm with parameters $\{A', C', (P'_{ij}), (\mu'_{ij})\}$ where

$$A' = A \otimes \tilde{A} \quad (98)$$

$$P'_{ijkl} = \begin{bmatrix} P_{ijxx} & P_{ijxy}^T & 0 \\ P_{ijxy} & P_{ijyy} + \tilde{P}_{klyy} & \tilde{P}_{klyz} \\ 0 & \tilde{P}_{klyz} & \tilde{P}_{klzz} \end{bmatrix} \quad (99)$$

$$\mu'_{ijkl} = P'_{ijkl}{}^{-1} \hat{\mu}_{ijkl} \quad (100)$$

$$C'_{ijkl} = C_{ij}^f + C_{kl}^g + \hat{\mu}_{ijkl} P_{ijkl} \hat{\mu}_{ijkl} - \mu_{ij} P_{ij} \mu_{ij} - \tilde{\mu}_{kl} \tilde{P}_{kl} \tilde{\mu}_{kl} \quad (101)$$

where

$$\hat{\mu}_{ijkl} = \begin{pmatrix} P_{ij} \mu_{ij} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \tilde{P}_{kl} \tilde{\mu}_{kl} \end{pmatrix} = \begin{pmatrix} [P_{ij} \mu_{ij}]_x \\ [P_{ij} \mu_{ij}]_y + [\tilde{P}_{kl} \tilde{\mu}_{kl}]_y \\ [\tilde{P}_{kl} \tilde{\mu}_{kl}]_z \end{pmatrix} \quad (102)$$

Proof. Notice that:

$$\log(B_f(x, y) \otimes B_g(y, z)_{ijkl}) = C_{ij}^f + C_{kl}^g - \left\| P_{ij}^{1/2} \begin{bmatrix} x \\ y \end{bmatrix} - \mu_{ij} \right\|^2 - \left\| \tilde{P}_{kl}^{1/2} \begin{bmatrix} x \\ y \\ z \end{bmatrix} - \tilde{\mu}_{kl} \right\|^2 \quad (103)$$

Let us compute the following term by computing the square:

$$\left\| P_{ij}^{1/2} \begin{bmatrix} x \\ y \end{bmatrix} - \mu_{ij} \right\|^2 + \left\| \tilde{P}_{kl}^{1/2} \begin{bmatrix} y \\ z \end{bmatrix} - \tilde{\mu}_{kl} \right\|^2 \quad (104)$$

$$= \begin{pmatrix} x \\ y \\ z \end{pmatrix}^T P_{ijkl} \begin{pmatrix} x \\ y \\ z \end{pmatrix} - 2 \hat{\mu}_{ijkl}^T \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \mu_{ij} P_{ij} \mu_{ij} + \tilde{\mu}_{kl} \tilde{P}_{kl} \tilde{\mu}_{kl} \quad (105)$$

$$= \left\| P_{ijkl}^{1/2} \begin{pmatrix} x \\ y \\ z \end{pmatrix} - P_{ijkl}^{-1} \hat{\mu}_{ijkl} \right\|^2 - \hat{\mu}_{ijkl} P_{ijkl}^{-1} \hat{\mu}_{ijkl} + \mu_{ij} P_{ij} \mu_{ij} + \tilde{\mu}_{kl} \tilde{P}_{kl} \tilde{\mu}_{kl} \quad (106)$$

$$(107)$$

where

$$P_{ijkl} = \begin{bmatrix} P_{ijxx} & P_{ijxy}^T & 0 \\ P_{ijxy} & P_{ijyy} + \tilde{P}_{klyy} & \tilde{P}_{klyz} \\ 0 & \tilde{P}_{klyz} & \tilde{P}_{klzz} \end{bmatrix}, \quad (108)$$

$$\hat{\mu}_{ijkl} = \begin{pmatrix} P_{ij}\mu_{ij} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \tilde{P}_{kl}\tilde{\mu}_{kl} \end{pmatrix} = \begin{pmatrix} [P_{ij}\mu_{ij}]_x \\ [P_{ij}\mu_{ij}]_y + [\tilde{P}_{kl}\tilde{\mu}_{kl}]_y \\ [\tilde{P}_{kl}\tilde{\mu}_{kl}]_z \end{pmatrix} \quad (109)$$

So

$$(B_f(x, y) \otimes B_g(y, z))_{ijkl} = \exp\left(C_{ijkl} - \left\| P_{ijkl}^{1/2}((x, y, z) - \mu_{ijkl}) \right\|^2\right) \quad (110)$$

where

$$\mu_{ijkl} = P_{ijkl}^{-1}\hat{\mu}_{ijkl} \quad (111)$$

$$C_{ijkl} = C_{ij}^f + C_{kl}^g + \hat{\mu}_{ijkl}P_{ijkl}^{-1}\hat{\mu}_{ijkl} - \mu_{ij}P_{ij}\mu_{ij} - \tilde{\mu}_{kl}\tilde{P}_{kl}\tilde{\mu}_{kl} \quad (112)$$

□

Remark 5 (Computational complexity). *Because of the need to compute the inverse of P' the computational complexity of the product operation between models of order M and m is $O(M^2m^2d^3)$.*

D.5 Proof of Theorem 10

Proof. Let $P = L^T\Sigma^{-1}L$ where $L = (F \ -I)$ and $P_\lambda = P + \lambda I$.

We have

$$-\log p(y|x) = -C_\Sigma + \|\Sigma^{-1/2}(Fx + b - y)\|^2 = -C_\Sigma + \|\Sigma^{-1/2}(Lu + b)\|^2 \quad (113)$$

$$= -C_\Sigma + uL^T\Sigma^{-1}Lu - 2b^T\Sigma^{-1}Lu + b^T\Sigma^{-1}b \quad (114)$$

If we define $\mu = P_\lambda^{-1}\beta$ where $\beta = L^T\Sigma^{-1}b$ then,

$$-\log p(y|x) = -C + \|P_\lambda^{1/2}(u - \mu)\|^2 \quad (115)$$

$$= -C + u^T P_\lambda u - 2\mu^T P_\lambda u + \mu^T P_\lambda \mu \quad (116)$$

$$= -C + u^T P u + \lambda\|u\|^2 - 2\beta^T u + \beta^T P_\lambda^{-1} P_\lambda P_\lambda^{-1} \beta \quad (117)$$

$$= -C + u^T P u + \lambda\|u\|^2 - 2\beta^T u + \beta^T P_\lambda^{-1} \beta \quad (118)$$

$$(119)$$

And so,

$$-\log(p(y|x)/\hat{p}(y|x)) = -C_\Sigma + uPu - 2\beta^T u + b^T\Sigma^{-1}b + C - u^T Pu - \lambda\|u\|^2 + 2\beta^T u - \beta^T P_\lambda^{-1}\beta \quad (120)$$

$$= C - C_\Sigma + b^T\Sigma^{-1}b - \lambda\|u\|^2 - \beta^T P_\lambda^{-1}\beta \quad (121)$$

Using Woodbury,

$$b^T \Sigma^{-1} b - \beta^T P_\lambda^{-1} \beta = b^T \left(\Sigma^{-1} - \Sigma^{-1} L [L^T \Sigma^{-1} L + \lambda I]^{-1} L^T - \Sigma^{-1} \right) b \quad (122)$$

$$= \lambda b^T (\lambda \Sigma + LL^T)^{-1} b \quad (123)$$

and

$$-\log(p(y|x)/\hat{p}(y/x)) = C - C_\Sigma + \lambda b^T (\lambda \Sigma + LL^T)^{-1} b - \lambda \|u\|^2 \quad (124)$$

With $C = C_\Sigma - \lambda b^T (\lambda \Sigma + LL^T)^{-1} b$, $\frac{p(y|x)}{\hat{p}(x,y)} = e^{\lambda \|u\|^2}$ and the result follows. \square

D.6 Learning Generalized Gaussian PSD Models

From an approximation perspective, a Generalized Gaussian PSD Model is a Gaussian PSD Model in which one can optimize the anchor points \tilde{x} and precision matrices P of each kernel function. In the case of approximating transition kernels, this can yield significant improvements in model order. Indeed, a transition kernel $Q(u, x)$ is a conditional probability distribution which depends in which the probability of the value x depends on the value u . This dependence is encoded in the combination of kernel evaluations but not in the kernel evaluations themselves.

To approximate a function f with a Generalized Gaussian PSD Model, we implicitly approximate the square-root of f using a Gaussian Linear Model:

$$\min_{\hat{g} \in \mathcal{G}_M} \frac{1}{n} \sum_{i=1}^n |f(x_i) - \hat{g}(x_i)|^2, \quad (125)$$

where x_i are sampled or chosen on a grid, and $\mathcal{G}_M = \{\sum_{j=1}^M \alpha_j k_{P_j}(x, \mu_j) \mid \mu_j \in \mathbb{R}^d, P = R_j^\top R_j, R_j \in \mathbb{R}^{d \times d}, \alpha_j \in \mathbb{R} \text{ for } 1 \leq j \leq M\}$. (125) is a smooth, non-convex problem which can be solved approximately using off-the-shelf solver like L-BFGS (Liu and Nocedal, 1989).

In practice, we initialize the model by placing μ_i in regions where $f(\mu_i)$ is large. In the case where $f(u, x)$ is a transition kernel $Q(u, x)$, one strategy is to choose $[\mu_i]_u$ on a grid (or sampled uniformly) and then choose $[\mu_i]_v$ such that $f(\mu_i) = \sup_x f([\mu_i]_u, v)$. This is particularly interesting when Q is a non-linear Gaussian model $Q(u, x) \propto e^{-\|\Sigma^{-1/2}(x-h(u))\|^2}$ for some non-linear transition model h .