

AALTO UNIVERSITY
SCHOOL OF SCIENCE

Arno Solin

**TRACKING AND ELIMINATION OF
PERIODIC NOISE IN FMRI USING
BAYESIAN INFERENCE**

Semester project undertaken as the course *Mat-2.4108 –
Independent research project in applied mathematics* in the Degree
Programme in Engineering Physics and Mathematics.

Espoo 2012

Supervisor:

Prof. Harri Ehtamo

Instructor:

D.Sc. (Tech.) Simo Särkkä

The document can be stored and made available to the public on the
open internet pages of Aalto University. All other rights are reserved.

AALTO UNIVERSITY SCHOOL OF SCIENCE P.O. Box 1100, FI-00076 AALTO http://www.aalto.fi		ABSTRACT OF THE PROJECT	
Author: Arno Solin			
Title: Tracking and Elimination of Periodic Noise in fMRI Using Bayesian Inference			
Degree programme: Engineering Physics and Mathematics			
Major subject: Systems and Operations Research		Minor subject: Computational Science and Engineering	
Chair (code): Mat-2			
Supervisor: Prof. Harri Ehtamo			
Instructor: D.Sc. (Tech.) Simo Särkkä			
<p>In this work the formulation of DRIFTER is studied. It is a model-based Bayesian method for estimation and removal of physiological noise, such as cardiac- and respiration-induced effects, in functional magnetic resonance imaging (fMRI). The method is due to Särkkä, Solin and colleagues, and this study aims to broaden some aspects discussed in the original article.</p> <p>The background of the DRIFTER method is presented by providing some insight in stochastic resonator models and modeling of quasi-periodic signals. The method is based on first estimating frequency trajectories of physiological noise components by using the interactive multiple models (IMM) algorithm and reference signals. A retrospective image-based state space formulation is used to estimate the noise-induced components in the fMRI signal with Kalman filtering and Rauch–Tung–Striebel smoothing. Separate estimates are gained for cardiac- and respiration-induced noise components, the cleaned blood oxygenation level dependent (BOLD) brain signal and a white measurement noise estimate.</p> <p>In this study, two aspects of using the DRIFTER method are studied in more detail: the effect of slow sampling rates and signal aliasing, and an example of estimation of frequencies without physiological reference signals. A brief analysis of these questions is provided and the results are discussed.</p>			
Date: May 21, 2012	Language: English	Number of pages: 25 + 4	
Keywords: Quasi-periodic signal, Physiological noise, fMRI, Kalman filtering			

Preface

This work was carried out in the Bayesian Statistical Methods group in the Department of Biomedical Engineering and Computational Science at Aalto University, Finland.

This study acts as an extension to the recently published DRIFTER article Särkkä *et al.* (2012) in which the author acts as second writer and has been responsible for implementing the methods in Matlab. I wish to thank all the contributors of the DRIFTER project. First and foremost my instructor Dr. Simo Särkkä, and Drs. Aki Vehtari, Aapo Nummenmaa, Toni Auranen, Simo Vanni, and profs. Fa-Hsuan Lin and Jouko Lampinen for expertise, support and providing data for analysis. Additionally prof. Harri Ehtamo deserves a warm thank for acting as the supervisor of this semester project.

Otaniemi, 2012

Arno Solin

Contents

Abstract	ii
Preface	iii
Contents	iv
Symbols and Abbreviations	v
1 Introduction	1
2 Materials and Methods	4
2.1 Optimal Estimation	4
2.2 Modeling Periodic Signals	6
2.3 Tracking of Frequency in Reference Signals	10
2.4 Modeling the Components of a BOLD Signal	13
2.5 Kalman Filter and RTS Smoother Implementation	15
3 Results	17
3.1 Sampling Rates and Signal Aliasing	17
3.2 Estimation of Frequencies Without References	19
4 Discussion and Conclusions	21
References	23
Appendices	1
Appendices	1
Appendix A: Simulation Code	1

Symbols and Abbreviations

Matrices are capitalized and vectors are in bold type. We do not generally distinguish between probabilities and probability densities.

Operators and miscellaneous notation

$1 : k$	$1, 2, \dots, k$
$p(\mathbf{x} \mathbf{y})$	Conditional probability density of \mathbf{x} given \mathbf{y}
$\mathbf{x}_{k k-1}$	Conditional value of \mathbf{x}_k given values up to step $k - 1$
\mathbb{R}	The real numbers
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$
\mathbf{I}	Identity matrix
\mathbf{A}^\top	Matrix transpose
$\text{diag}(\mathbf{a})$	A diagonal matrix with elements of \mathbf{a} on its diagonal

General notation

\mathbf{x}	System state
\mathbf{y}	Observation
k	Time step
T	Final time step
\mathbf{q}_k	Zero-mean (Gaussian) Process noise
\mathbf{r}_k	Zero-mean (Gaussian) Measurement noise
\mathbf{Q}_k	Process noise covariance
\mathbf{R}_k	Measurement noise covariance

Abbreviations

RTS	Rauch–Tung–Striebel (smoother)
MRI	Magnetic Resonance Imaging
fMRI	Functional MRI
EPI	Echo Planar Imaging
InI	Dynamical Inverse Imaging
BOLD	Blood oxygenation level dependent
TR	Repetition time, interval between subsequent scans (T_R)
RMSE	Root mean square error
SNR	Signal-to-noise ratio

1 Introduction

The field of functional human brain mapping is dominated by three concurrent imaging methods: *Positron Emission Tomography* (PET) uses short-half-life radiotracers and provides high spatial resolution by measuring changes in blood flow. *Electroencephalography* (EEG) and *magnetoencephalography* (MEG) are based on recording of electrical/magnetic activity along the scalp, which provides high temporal resolution, but are subject to artefacts and noise. The third method is *functional Magnetic Resonance Imaging* (fMRI) which has high spatial resolution but relatively poor temporal resolution. (see, e.g., Huettel *et al.*, 2004; Buxton, 2009)

In fMRI (Ogawa *et al.*, 1990; Kwong *et al.*, 1992), the field of functional neuroimaging has recently primarily focused on a phenomenon called the blood oxygenation level dependent (BOLD) effect. Neurons in the brain require energy in form of glucose and oxygen to function. Hemoglobin in blood carries O_2 and when it loses it, the magnetic properties change in a subtle way — the signal being stronger the more the blood becomes oxygenated (Buxton, 2009). When an area in the brain becomes activated, the blood flow however actually increases more than the neurons consume energy, because the flow is haemodynamically regulated to give active neural assemblies more energy. This produces the BOLD effect, which accounts for a local increase in the T_2^* contrast based functional MR signal during neural activity (Buxton, 2009).

To fight the disadvantageous poor temporal resolution and high noise levels in fMRI, more effective MRI hardware has been developed. As the spatio-temporal resolution and signal-to-noise-ratio (SNR) of fMRI increases, accurate treatment of various noise sources in measurements becomes more and more important (as has been shown, e.g., by Hutton *et al.*, 2011). The interpretation and identification of these noises in the signal is important, because not all noises can be modeled as white noise (Lund *et al.*, 2006). These non-white noise sources cannot be eliminated by improving the data acquisition hardware, because they are not actual ‘noise’ as such but a part of the measured phenomenon.

Physiological noise, most importantly neurovascular fluctuation together with cardiac- and respiration-induced quasi-periodic oscillations can easily account for up to one third of the signal variation already at a field of 3 T

(Krüger and Glover, 2001). At higher field strengths these phenomena are even more dominant (Krüger and Glover, 2001; Triantafyllou *et al.*, 2005; Hutton *et al.*, 2011).

This problem has been recognized and there exists several approaches to elimination of cardiac and respiration related physiological noise from fMRI measurements. Biswal *et al.* (1996) use notch filters (band-stop filters) to eliminate the frequency bands corresponding to cardiac and respiratory activity. This approach requires the fMRI data to be sampled at a high frequency and cannot cope well with aliasing. This approach also assumes stationarity of the signal, and it has troubles adapting to fluctuations in the heart beat rate and changes in respiration cycles.

A popular approach is the *Image-Based Method for Retrospective Correction of Physiological Motion Effects in fMRI*, in short RETROICOR, due to Glover *et al.* (2000). It is based on fitting low-order Fourier series to the image data, where the periodical signals match the cardiac and respiratory signal phases. These phases are estimated with the help of external reference signals by peak-detection and histogram-based methods. The use of external references help to avoid the aliasing related problems in the notch filter.

Other image-based physiological noise reduction approaches include adaptive filtering (Deckers *et al.*, 2006), Principal Component Analysis (PCA) and Independent Component Analysis (ICA) (Thomas *et al.*, 2002), and IMPACT (Chuang and Chen, 2001). It is also possible to do retrospective noise reduction in k -space (the scanner's image-acquisition Fourier domain) (Hu *et al.*, 1995; Le and Hu, 1996; Frank *et al.*, 2001) or by utilizing the phase information (Cheng and Li, 2010).

Recently Särkkä *et al.* (2012) proposed a new Bayesian method for physiological noise modeling and removal. The method is named DRIFTER and it allows accurate dynamical tracking of the variations in the cardiac and respiratory frequencies by using Interacting Multiple Models (IMM), Kalman filter (KF) and Rauch–Tung–Striebel (RTS) smoother algorithms (Bar-Shalom *et al.*, 2001; Grewal and Andrews, 2001).

In this study we will go through the formulation of DRIFTER in detail with emphasis on how it can be implemented in practice. The method is flexible and the frequency trajectories can be either estimated from reference signals (e.g. by using pulse meters and respiratory belts, as in RETROICOR),

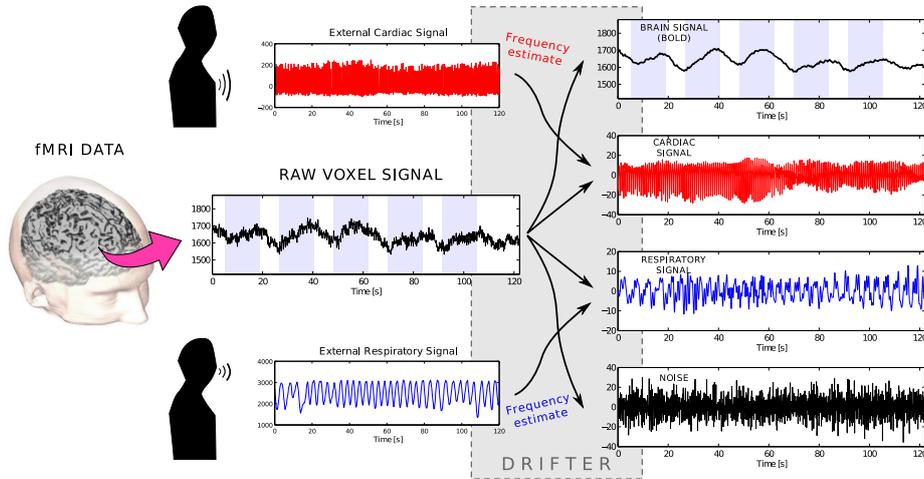


Figure 1: Illustration of the idea in DRIFTER where the frequencies of the external cardiac and respiratory signals are estimated and then the respiration- and cardiac-induced noises are separated from the raw voxel (volumetric pixel) signal.

or if the time resolution allows, directly from the fMRI signal (e.g., from the spatially averaged fMRI signal). The estimated frequency trajectory is used for accurate model based separation of the fMRI signal into activation, physiological noise and white noise components using Kalman filter and RTS smoother algorithms. This separation is done for each voxel (volumetric pixel) in the image. The basic idea of the method is to build a stochastic model for each component of the signal: the BOLD signal is a relatively slowly varying signal, cardiac and respiration are stochastic resonators with multiple harmonics, and the rest of the signal is assumed to be white noise. The basic idea of the method is presented in Figure 1, where the components are visualized on the far right.

The mathematical models and algorithms are presented in Section 2, starting from concepts in optimal estimation and extending the perusal to stochastic resonator models and how to implement the inference. In Section 3 we address the questions related to sampling rates and signal aliasing and provide an example of estimating the signal frequencies directly from the fMRI data without reference signals. These two questions have a practical impact on the usability of the method, because whole-brain fMRI data typically has sampling rates of approximately two seconds and physiological signals are often not included in routine clinical fMRI studies.

2 Materials and Methods

In this section we present the DRIFTER method as a composition of means to model the signal components in fMRI data, methods to estimate the frequency time series and track the values of these components, and an effective implementation scheme of all this into a Kalman filter compatible manner. The theory follows the same structure as Särkkä *et al.* (2012).

Kalman filters, Rauch–Tung–Striebel Smoothers and Interactive multiple model (IMM) algorithms are gone through in brief in Section 2.1. For modeling the quasi-periodic noise components in the fMRI data, an IMM scheme for the frequency estimation is presented in Sections 2.2–2.3. In Section 2.4 we split the fMRI signal into components that fit into the Kalman filtering state space form, and where the periodic noises can be tracked using the frequency estimates. Finally this is combined into an effective Kalman filtering formulation.

2.1 Optimal Estimation

The term *optimal estimation* refers to the methods used to estimate the underlying state of a time-varying system of which there exist only indirectly observed noisy measurements. In many cases Kalman filter and Rauch–Tung–Striebel smoother (see, e.g., Grewal and Andrews, 2001; Särkkä, 2006; Solin, 2010) algorithms are the ones referred to with optimal estimation. These two algorithms can be used for computing the exact Bayesian posterior filtering distributions of the state in discrete-time linear Gaussian state space models of the form

$$\begin{aligned}\mathbf{x}(t_{k+1}) &= \mathbf{A}_k \mathbf{x}(t_k) + \mathbf{q}_k \\ \mathbf{y}(t_k) &= \mathbf{H}_k \mathbf{x}(t_k) + \mathbf{r}_k,\end{aligned}\tag{1}$$

where $\mathbf{x}(t_k) \in \mathbb{R}^n$ is the state at time t_k , where $k = 0, 1, 2, \dots$, $\mathbf{y}(t_k) \in \mathbb{R}^d$ is the measurement at time t_k , $\mathbf{q}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$ is the Gaussian process noise, and $\mathbf{r}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_k)$ is the Gaussian measurement noise. Matrix \mathbf{A}_k is the state transition matrix and \mathbf{H}_k is the measurement model matrix.

Continuous-time models of equivalent linear kind can be handled by first discretizing the dynamics (see, e.g., Grewal and Andrews, 2001) of the

model.

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{F} \mathbf{x}(t) + \mathbf{L} \mathbf{e}(t), \quad (2)$$

where $\mathbf{e}(t)$ is a white noise process with a given spectral density matrix \mathbf{Q}_c . If we assume that the sampling period is Δt , and we define $t_k = k \Delta t$, then the (weak) solution (Øksendal, 2003) to this continuous-time stochastic differential equation can be expressed as

$$\mathbf{x}(t_{k+1}) = \exp(\Delta t \mathbf{F}) \mathbf{x}(t_k) + \int_{t_k}^{t_{k+1}} \exp((t_{k+1} - s) \mathbf{F}) \mathbf{L} \mathbf{e}(s) ds. \quad (3)$$

The second integral above is just a Gaussian random variable with covariance

$$\mathbf{Q}_k = \int_0^{\Delta t} \exp((\Delta t - \tau) \mathbf{F}) \mathbf{L} \mathbf{Q}_c \mathbf{L}^\top \exp((\Delta t - \tau) \mathbf{F})^\top d\tau. \quad (4)$$

Thus, if we define $\mathbf{A}_k = \exp(\Delta t \mathbf{F})$, the model becomes equivalent to the discrete model in Equations (1).

The *Kalman filter* (Kalman, 1960) is a closed-form solution to the linear discrete-time filtering problem in Equation (1). As the Kalman filter is conditional to all measurements up to time step k , the recursive filtering algorithm can be seen as a two-step process that first includes calculating the marginal distribution of the next step using the known system dynamics. After this the information is updated using new observations.

Similarly as the discrete-time linear Kalman filter gives a closed-form filtering solution, the discrete-time *Rauch–Tung–Striebel (RTS) smoother* (Rauch *et al.*, 1965) gives a closed-form solution to the linear smoothing problem, which is conditional to all the measurements $\mathbf{y}(t_k)$, where $k \in \{1, \dots, T\}$ is a fixed interval. The Kalman filtering and Rauch–Tung–Striebel smoothing equations are presented and discussed in detail in Section 2.5.

The state space model can be extended by including an additional latent variable into the formulation. The *Interacting Multiple Models* (IMM) algorithm (Bar-Shalom *et al.*, 2001) is a method that can be used for computing posterior distributions of models, where the model matrices depend on an additional latent variable θ_k such that

$$\begin{aligned} \mathbf{x}(t_{k+1}) &= \mathbf{A}_k(\theta_k) \mathbf{x}(t_k) + \mathbf{q}_k \\ \mathbf{y}(t_k) &= \mathbf{H}_k(\theta_k) \mathbf{x}(t_k) + \mathbf{r}_k. \end{aligned} \quad (5)$$

This variable takes values in a finite set $\theta_k \in \Omega = \{\theta^{(1)}, \dots, \theta^{(S)}\}$ and it is assumed that we can model its dynamics using a Markov chain with a transition matrix $\mathbf{\Pi}$, where the transition probabilities are given by

$$P(\theta_k^{(i)} | \theta_{k-1}^{(j)}) = \mathbf{\Pi}_{i,j}. \quad (6)$$

The IMM algorithm provides efficient means for computing an S -component Gaussian mixture approximation to the joint posterior distribution of the latent variables and states.

Kalman filters are restricted to model systems with known dynamical and measurement models. The IMM method is just one possibility to deal with this restriction. In practice, in the IMM formulation multiple models — typically matching discrete prior assumptions of different model possibilities — are run in parallel, and the most probable of them is chosen. Typical applications include tracking of manoeuvring targets, or simplification of non-linear dynamics into a few separate cases.

2.2 Modeling Periodic Signals

Various kind of noise is present in BOLD fMRI data, and these can be divided into scanner hardware related thermal noise and scanner signal drifting, head movement artefacts, and physiological noise. However physiological noise has the pleasant property that cardiac- and respiration-induced noises follow the same periodic structure as the actual beating of the heart and breathing cycles. This structure makes it possible to track these physiological noise components in the fMRI data on a voxel level.

A band-limited zero-mean periodic signal with period frequency f can be approximated to an arbitrary precision with a truncated Fourier series

$$c(t) = \sum_{n=1}^N a_n \cos(2\pi n f t) + b_n \sin(2\pi n f t), \quad (7)$$

where a_n and b_n are the Fourier coefficients. The frequency f of the base sinusoidal is often referred to as the *fundamental* frequency of the signal. The rest of the components in the superposition are referred to as the harmonic components. The harmonics have frequencies that are multiples of the fundamental frequency, $n f$. To model any arbitrary signal, the number

of harmonics would need to be infinite. However, in the case of real-world noise, a finite N gives a qualified approximation.

We are interested in modeling of signals which are almost periodic, that is, quasi-periodic. The strongest source of aperiodicity in the signals is caused by time variance of the base frequency, which means that the frequency is actually a function of time $f(t)$.

The most simple approach to modeling signals with a changing frequency would be to simply plug in the frequency function $f(t)$ into Equation (7). However, this would not be a wise choice. This model is very sensitive to changes in frequency. For example, when t is large, even a tiny change in the frequency causes a large change in signal $c(t)$. Discontinuities in the frequency would also cause the signal $c(t)$ to be discontinuous.

The RETROICOR method (Glover *et al.*, 2000) solves this problem by estimating the phase $\phi_c(t)$ of the signal instead of the frequency. That is

$$\phi_c(t) = \int_0^t 2\pi f(t) dt.$$

This method deals effectively with the problems regarding time-varying frequencies. However, there is another problem: the Fourier-like coefficients a_n and b_n are assumed to be constant in time, which implies that the amplitude of the phenomenon is assumed to be constant. This is a quite unrealistic assumption in real data, as was demonstrated by Särkkä *et al.* (2012).

The approach that is used in DRIFTER is based on the observation that the Fourier series (7) can also be represented in an alternative form by modeling the signal as an oscillator

$$\frac{d^2 c_n(t)}{dt^2} = -(2\pi n f)^2 c_n(t), \quad (8)$$

that has the solution

$$c_n(t) = a_n \cos(2\pi n f t) + b_n \sin(2\pi n f t), \quad (9)$$

where the constants a_n and b_n are set by the initial conditions of the differential equation. The Fourier series (7) is thus equivalent to the representation

$$c(t) = \sum_{n=1}^N c_n, \quad (10)$$

where the initial conditions of the oscillators implicitly define the Fourier coefficients. We can now replace the constant frequency with a time series $f(t)$ of frequencies, which leads to the following differential equation model for the n th harmonic

$$\frac{d^2 c_n(t)}{dt^2} = - (2\pi n f(t))^2 c_n(t). \quad (11)$$

Unlike the extended Fourier series (10), this signal has the pleasant property that it is continuous even when frequency has discontinuities.

Another source of aperiodicity of the signal are small changes in the shape of the signal, which correspond to changes in amplitudes and phases in the harmonics. These changes can be modeled by adding a white noise component $e_n(t)$ with spectral density q_n to the differential equation of each harmonic component

$$\frac{d^2 c_n(t)}{dt^2} = - (2\pi n f(t))^2 c_n(t) + e_n(t). \quad (12)$$

The full quasi-periodic signal then has the representation

$$c(t) = \sum_{n=1}^N c_n. \quad (13)$$

Figure 2 shows an oscillator with one harmonic. The amplitude is subject to change at 10s, the frequency at 20s, and the model changes to a stochastic oscillator at 30s. This demonstrates the robustness of the approach, and the outcome of this model seems very intuitive.

The model can also be represented in canonical state space form by defining the state and vector of noises respectively as

$$\begin{aligned} \mathbf{x} &= \left[c_1 \quad \frac{dc_1}{dt} \quad c_2 \quad \frac{dc_2}{dt} \quad \cdots \quad c_N \quad \frac{dc_N}{dt} \right]^T \\ \mathbf{e} &= \left[e_1 \quad e_2 \quad \cdots \quad e_N \right]^T. \end{aligned} \quad (14)$$

Furthermore, we want to preserve the norm of the discretized signal even if the frequency of the oscillatory components change. Therefore we violate the requirement that we want the solution to be of form (12) and formulate

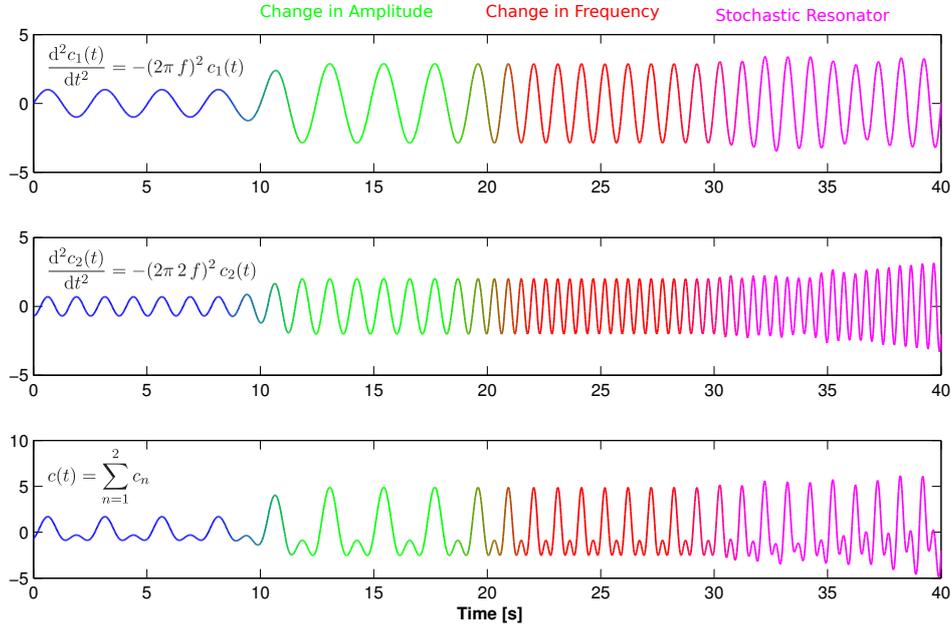


Figure 2: Decomposition of a quasi-periodic signal $c(t)$ into sinusoidal components $c_1(t)$ and $c_2(t)$. The effect of an amplitude change, a change in frequency and adding a stochastic part to the equation, are all demonstrated.

the state space dynamics as in Särkkä *et al.* (2012). If we now define

$$\mathbf{G}(f) = \begin{bmatrix} 0 & 2\pi f \\ -2\pi f & 0 \end{bmatrix} \quad (15a)$$

$$\mathbf{F}_o(f) = \begin{bmatrix} \mathbf{G}(f) & & & \\ & \mathbf{G}(2f) & & \\ & & \ddots & \\ & & & \mathbf{G}(Nf) \end{bmatrix} \quad (15b)$$

where $\mathbf{F}_o(f)$ can as well be denoted by ‘blockdiag ($\mathbf{G}(f), \mathbf{G}(2f), \dots, \mathbf{G}(Nf)$)’, the stochastic state space model for the quasi-periodic signal can be written as follows

$$\begin{aligned} \frac{d\mathbf{x}(t)}{dt} &= \mathbf{F}_o(f(t)) \mathbf{x}(t) + \mathbf{L} \mathbf{e}(t) \\ c(t) &= \mathbf{H} \mathbf{x}(t), \end{aligned} \quad (16)$$

where the matrix \mathbf{L} has elements $\mathbf{L}_{2n,n} = 1$ for $n = 1, \dots, N$, and all other are zero, and $\mathbf{H} = [1 \ 0 \ 1 \ 0 \ \dots \ 1 \ 0]$. When the frequency trajectory $f(t)$ is known, the model above is a time-varying linear state space model which is directly compatible for Kalman filters. With unknown $f(t)$ we can use adaptive Kalman filters for inferring the state and frequency trajectories as will be shown in the next section.

2.3 Tracking of Frequency in Reference Signals

In modeling the periodic noise components, it is essential to have good approximations of the time-depending frequencies of the signals. Most methods for removing physiological noise in fMRI use reference signal data to track the frequencies, or equivalently the phases of the signals. We assume that we have some reference sensor, which measures the cardiac cycle (e.g. an ECG sensor or pulse oximeter). Equivalently we assume that also the respiration cycles are monitored, typically with the help of a respiratory belt that is put around the subject's chest or abdomen.

These signals can be used in different ways. For example, the formulation of RETROICOR in Glover *et al.* (2000) presents separate methods for estimating the phase of a cardiac or respiration reference signal. The cardiac phase is estimated using a peak-detection method and comparing the R-R-times in the signal. This is not an effective way to deal with the more unpredictable respiratory signal, and thus a histogram-based method is used for this in RETROICOR.

In DRIFTER an IMM approach (adaptive Kalman filtering) is used for both the respiration and cardiac activity. The cardiac signal can now be modeled with the quasi-periodic signal model described in the previous section. Typically the references — especially the respiration reference signal — are not zero mean signals, but there are level changes in the signals over time. To account for this possible drifting of the reference signal, we include a time-varying bias $b(t)$, and model it using a Wiener velocity model (see, e.g. Särkkä, 2006)

$$\frac{d^2b(t)}{dt^2} = e_b(t), \quad (17)$$

where $e_b(t)$ is a white noise process with spectral density q_b . Note that this actually corresponds to the zero-frequency stochastic oscillator, but we could similarly use any other model for the bias term.

If we define the joint state consisting of the bias and a quasi-periodic signal with N_c harmonics as

$$\mathbf{x}_c = \left[b \quad \frac{db}{dt} \quad c_1 \quad \frac{dc_1}{dt} \quad \dots \quad c_{N_c} \quad \frac{dc_{N_c}}{dt} \right]^\top, \quad (18)$$

then the measured reference cardiac signal y_c , which is sampled at times t_k can be modeled as

$$\begin{aligned} \frac{d\mathbf{x}_c(t)}{dt} &= \mathbf{F}_c(f_c(t)) \mathbf{x}_c + \mathbf{L}_c \mathbf{e}_c \\ y_c(t_k) &= \mathbf{H}_c \mathbf{x}_c(t_k) + v_c, \end{aligned} \quad (19)$$

where where $v_c \sim \mathcal{N}(0, \sigma_c^2)$ is Gaussian measurement noise (residual noise) with zero mean and variance σ^2 , which accounts for the physical noise, uncertainties and the differences between the model and the reality. Above, the matrices are defined as

$$\mathbf{G}_b = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad (20)$$

$$\mathbf{F}_c(f) = \text{blockdiag}(\mathbf{G}_b, \mathbf{G}(f), \mathbf{G}(2f), \dots, \mathbf{G}(N_c f)),$$

and $\mathbf{G}(nf)$ is defined as in Equation (15a). The $N_c + 1$ dimensional white noise process \mathbf{e}_c consists of the bias noise process and the oscillator noise processes, and the matrices L_c and H_c are defined in analogous manner as in Equations (16).

If we assume that the frequency $f_c(t)$ is constant between the measurements, say, has value $f_c(t_k)$ on the interval $t \in [t_k, t_{k+1})$, then we can use the discretization procedure presented in Section 2.1 to convert the dynamic model into the form

$$\mathbf{x}_c(t_{k+1}) = \mathbf{A}_c(f_c(t_k)) \mathbf{x}_c(t_k) + \mathbf{q}_c, \quad (21)$$

where $\mathbf{q}_c \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_c(f_c(t_k)))$. Because the frequency trajectory $f_c(t)$ is unknown, we model it as a stochastic process as well. We assume that the frequency is constant between the measurements and that it can only take values from a given discrete set $f_c \in \{f_c^{(1)}, \dots, f_c^{(M_c)}\}$. If we model the time behavior of the discrete set of frequencies as a Markov chain, we obtain the

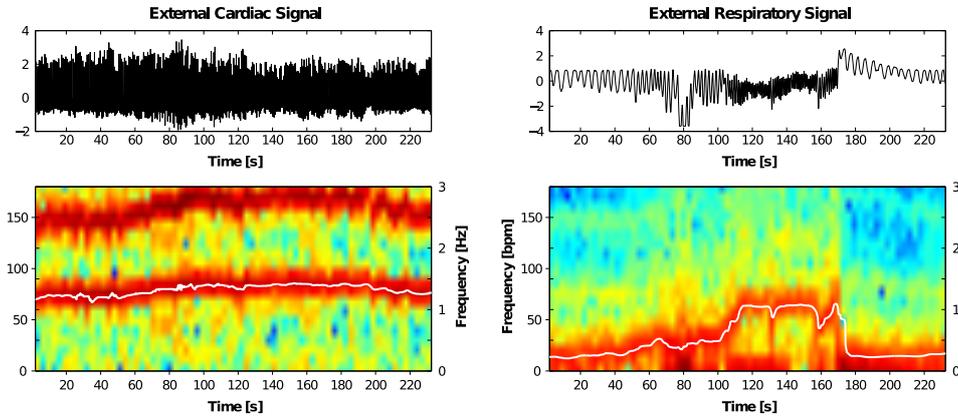


Figure 3: Examples of reference cardiac and respiration signals, their spectrograms and the estimated frequency trajectories. The fluctuations in the respiratory signal are considerable, because the test subject was instructed to breathe heavily during the scan.

following switching linear state space model:

$$\begin{aligned}
 P(f_c^{(i)} | f_c^{(j)}) &= \Pi_{i,j}^c \\
 \mathbf{x}_c(t_{k+1}) &= \mathbf{A}_c(f_c(t_k)) \mathbf{x}_c(t_k) + \mathbf{q}_c \\
 y_c(t_k) &= \mathbf{H}_c \mathbf{x}_c(t_k) + v_c,
 \end{aligned} \tag{22}$$

where Π^c is the transition matrix of the discrete frequency Markov chain and $\mathbf{q}_c \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_c(f_c(t_k)))$, $v_c \sim \mathcal{N}(0, \sigma_c^2)$. Comparing to the Section 2.1 it is easy to see that this model is of the form that can be treated with the IMM algorithm.

The respiratory reference signal and its frequency trajectory $f_r(t)$ can be modeled in a completely analogous manner as the cardiac signal. Obviously, the discrete set of frequencies $f_r^{(i)}$ needs to be different and the spectral densities \mathbf{q}_r, v_r of the noises need to be adjusted.

Although here we have used a fairly simple set of modes consisting of cardiac or respiratory signals in different frequencies, it would also be possible to model and include other types of modes to the switching model. We could, for example, detect extra or missing heart beats, arrhythmia, breath holding or various other conditions and pass them to the fMRI estimation stage along with the frequencies. (Särkkä *et al.*, 2012)

Figure 3 shows examples of cardiac and respiration reference signals and their spectrograms. The presence of multiple harmonics can be clearly seen

in both of the images. The frequencies of the cardiac and respiration change quite rapidly in time and thus it is crucial to explicitly account for the time variance in the frequencies.

2.4 Modeling the Components of a BOLD Signal

The fMRI data is a four-dimensional signal with three spatial dimensions and one temporal dimension, where we have separate time series for each voxel in three-dimensional space. In the most simple variant of DRIFTER we do not assume anything about the spatial structure of the signal, but all the voxel time series are treated independently. Different types of spatial priors could be applied at this stage to the model.

We assume that the measured signal consists of (i) a cardiac-induced signal with known base frequency $f_c(t)$ from the analysis described in the previous section, but unknown amplitudes and phases of the harmonics, (ii) a respiration-induced signal with known base frequency $f_r(t)$ as earlier, but unknown amplitudes and phases of the harmonics, and (iii) a brain signal (cleaned BOLD signal) that accounts for the fluctuation of the signal caused for example by haemodynamical responses. The rest of the voxel time series, that does not fit into these three categories, is assumed to be white noise.

The model is as follows:

1. The cardiac-induced signal is modeled as a zero mean periodic signal with a given frequency trajectory $f_c(t)$ as estimated from the reference signal. The model is of the form

$$\frac{d\mathbf{x}_c(t)}{dt} = \mathbf{F}_c(f_c(t)) \mathbf{x}_c(t) + \mathbf{L}_c \mathbf{e}_c(t), \quad (23)$$

where

$$\mathbf{F}_c(f) = \text{blockdiag}(\mathbf{G}(f), \mathbf{G}(2f), \dots, \mathbf{G}(N_c f)), \quad (24)$$

and $G(\cdot)$ is defined as in Equation (15a) and N_c is the number of modeled harmonics in the cardiac signal.

2. The respiration-induced signal is modeled as a zero mean periodic signal with the given frequency trajectory $f_r(t)$:

$$\frac{d\mathbf{x}_r(t)}{dt} = \mathbf{F}_r(f_r(t)) \mathbf{x}_r(t) + \mathbf{L}_r \mathbf{e}_r(t), \quad (25)$$

where

$$\mathbf{F}_r(f) = \text{blockdiag}(\mathbf{G}(f), \mathbf{G}(2f), \dots, \mathbf{G}(N_r f)). \quad (26)$$

3. The brain signal (BOLD signal) is assumed to be smooth and thus it is modeled using a simple Wiener velocity model

$$\frac{d\mathbf{x}_b(t)}{dt} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \mathbf{x}_b(t) + \mathbf{L}_b \mathbf{e}_b(t). \quad (27)$$

We define the full state of the signal in a single voxel as a combination of all the above states

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_c \\ \mathbf{x}_r \\ \mathbf{x}_b \end{bmatrix}. \quad (28)$$

This state is different in each spatial location \mathbf{r} , that is, has the form $\mathbf{x}(t, \mathbf{r})$. The full differential equation for the dynamics of the state and the corresponding measurements can be expressed as

$$\begin{aligned} \frac{\partial \mathbf{x}(t, \mathbf{r})}{\partial t} &= \mathbf{F} \mathbf{x}(t, \mathbf{r}) + \mathbf{L} \mathbf{e}(t, \mathbf{r}) \\ y(t_k, \mathbf{r}) &= \mathbf{H} \mathbf{x}(t_k, \mathbf{r}) + v_f(\mathbf{r}), \end{aligned} \quad (29)$$

where v_f is a zero mean Gaussian sequence with variance σ^2 , which is independent in each voxel location \mathbf{r} , models the measurement noise in fMRI images. The white noises $\mathbf{e}(t, \mathbf{r})$ are assumed to be independent in each voxel and have a joint spectral density \mathbf{Q}_c , which is independent of the position \mathbf{r} .

The continuous-time model is defined by the sparse block diagonal matrices $\mathbf{F} \in \mathbb{R}^{2N \times 2N}$, $\mathbf{L} \in \mathbb{R}^{2N \times N}$, $\mathbf{Q}_c \in \mathbb{R}^{N \times N}$, and $\mathbf{H} \in \mathbb{R}^{2N}$, where $N = N_c + N_r + 1$. These matrices take the forms $\mathbf{F}(f_c, f_r) = \text{blockdiag}(\mathbf{F}_c(f_c), \mathbf{F}_r(f_r), \mathbf{F}_b)$, $\mathbf{L} = \text{blockdiag}([0 \ 1]^\top, \dots, [0 \ 1]^\top, \dots, [0 \ 1]^\top)$, $\mathbf{Q}_c = \text{diag}(q_c/1, q_c/2, \dots, q_r/1, q_r/2, \dots, q_b)$ and $\mathbf{H} = [1 \ 0 \ \dots \ 1 \ 0 \ \dots \ 1 \ 0]$.

The discretization procedure presented in Section 2.1 now results in the model of the form

$$\begin{aligned} \mathbf{x}(t_{k+1}, \mathbf{r}) &= \mathbf{A}_k \mathbf{x}(t_k, \mathbf{r}) + \mathbf{q}_k(\mathbf{r}) \\ y(t_k, \mathbf{r}) &= \mathbf{H} \mathbf{x}(t_k, \mathbf{r}) + v_f(\mathbf{r}). \end{aligned} \quad (30)$$

where $\mathbf{q}_k(\mathbf{r}) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$. This discrete-time state space equation can now be used in the Kalman filtering formulation.

2.5 Kalman Filter and RTS Smoother Implementation

The four-dimensional fMRI signal $\mathbf{y}(t_k) \in \mathbb{R}^{n_x \times n_y \times n_z}$, where $k = 1, \dots, n_t$, can be handled independently in each voxel. As described in the previous subsection, we now assume each voxel measurement is a superposition of all the components in the model, $\mathbf{x}(t) = (\mathbf{x}_c(t), \mathbf{x}_r(t), \mathbf{x}_b(t))$.

The actual state cannot be observed and is thus unknown to us. We define the process to be a realisation of a prior distribution and determine the trajectory subject to observations provided by the fMRI data. To put this under a standard Kalman filter compatible formulation, the probability distribution is assumed to be approximately a multivariate Gaussian normal distribution. This means that the state $\mathbf{x}(t_k, \mathbf{r}) \sim \mathcal{N}(\mathbf{m}(t_k, \mathbf{r}), \mathbf{P}(t_k, \mathbf{r}))$ is governed by the mean $\mathbf{m}(t_k, \mathbf{r})$ and covariance $\mathbf{P}(t_k, \mathbf{r})$ of the distribution at each time step t_k in each voxel position \mathbf{r} .

The Kalman filtering equations (Grewal and Andrews, 2001) for this model can be written in the form of a *prediction step*

$$\begin{aligned} \mathbf{m}^-(t_{k+1}, \mathbf{r}) &= \mathbf{A}_k \mathbf{m}(t_k, \mathbf{r}) \\ \mathbf{P}^-(t_{k+1}, \mathbf{r}) &= \mathbf{A}_k \mathbf{P}(t_k, \mathbf{r}) \mathbf{A}_k^\top + \mathbf{Q}_k. \end{aligned} \quad (31)$$

and a subsequent *update step*

$$\begin{aligned} S_{k+1}(\mathbf{r}) &= \mathbf{H} \mathbf{P}^-(t_{k+1}, \mathbf{r}) \mathbf{H}^\top + \sigma^2 \\ \mathbf{K}_{k+1}(\mathbf{r}) &= \mathbf{P}^-(t_{k+1}, \mathbf{r}) \mathbf{H}^\top S_{k+1}^{-1}(\mathbf{r}) \\ \mathbf{m}(t_{k+1}, \mathbf{r}) &= \mathbf{m}^-(t_{k+1}, \mathbf{r}) + \mathbf{K}_{k+1}(\mathbf{r}) [y(t_k, \mathbf{r}) - \mathbf{H} \mathbf{m}^-(t_{k+1}, \mathbf{r})] \\ \mathbf{P}(t_{k+1}, \mathbf{r}) &= \mathbf{P}^-(t_{k+1}, \mathbf{r}) - \mathbf{K}_{k+1}(\mathbf{r}) S_{k+1}(\mathbf{r}) \mathbf{K}_{k+1}^\top(\mathbf{r}). \end{aligned} \quad (32)$$

However, the evaluation of the filtering (and smoothing) equations can be made efficient by assuming that the initial covariance of the signals is independent of the voxel position \mathbf{r} such that $\mathbf{P}(t_0, \mathbf{r}) = \mathbf{P}(t_0)$. This is a reasonable assumption which lets us formulate the Kalman filter equations in a parallel manner. The terms \mathbf{P} , S , and \mathbf{K} become independent of the position \mathbf{r} . Thus we only need to compute the following once per measurement time:

For a practical implementation, we can define the mean $\mathbf{m}(t_k) = [\mathbf{m}(t_k, \mathbf{r}_1) \mathbf{m}(t_k, \mathbf{r}_2) \dots \mathbf{m}(t_k, \mathbf{r}_{N_{xyz}})] \in \mathbb{R}^{2N \times N_{xyz}}$, where $N = N_c + N_r + 1$ and $N_{xyz} = n_x \cdot n_y \cdot n_z$, whereas the covariance dimensions are conveniently $\mathbf{P}(t_k) \in \mathbb{R}^{2N \times 2N}$. Similarly as for the state mean we combine the voxel measurements into a vector $\mathbf{y}(t_k) = [y(t_k, \mathbf{r}_1) y(t_k, \mathbf{r}_2) \dots y(t_k, \mathbf{r}_{N_{xyz}})] \in \mathbb{R}^{1 \times N_{xyz}}$.

Thus the *prediction step* becomes simply

$$\begin{aligned} \mathbf{m}^-(t_{k+1}) &= \mathbf{A}_k \mathbf{m}(t_k) \\ \mathbf{P}^-(t_{k+1}) &= \mathbf{A}_k \mathbf{P}(t_k) \mathbf{A}_k^\top + \mathbf{Q}_k \end{aligned} \quad (33)$$

and the subsequent *update step*

$$\begin{aligned} \mathbf{S}_{k+1} &= \mathbf{H} \mathbf{P}^-(t_{k+1}) \mathbf{H}^\top + \sigma^2 \mathbf{I} \\ \mathbf{K}_{k+1} &= \mathbf{P}^-(t_{k+1}) \mathbf{H}^\top \mathbf{S}_{k+1}^{-1} \\ \mathbf{m}(t_{k+1}) &= \mathbf{m}^-(t_{k+1}) + \mathbf{K}_{k+1} [\mathbf{y}(t_k) - \mathbf{H} \mathbf{m}^-(t_{k+1})] \\ \mathbf{P}(t_{k+1}) &= \mathbf{P}^-(t_{k+1}) - \mathbf{K}_{k+1} \mathbf{S}_{k+1} \mathbf{K}_{k+1}^\top. \end{aligned} \quad (34)$$

The Rauch–Tung–Striebel smoother covariance and gain are also independent of the position. This makes it possible to write the backward sweep of the smoother in a similar manner to the filtering equations

$$\begin{aligned} \mathbf{m}^-(t_{k+1}) &= \mathbf{A}_k \mathbf{m}(t_k) \\ \mathbf{P}^-(t_{k+1}) &= \mathbf{A}_k \mathbf{P}(t_k) \mathbf{A}_k^\top + \mathbf{Q}_k \\ \mathbf{C}_{k+1} &= \mathbf{P}(t_k) \mathbf{A}_k^\top [\mathbf{P}^-(t_{k+1})]^{-1} \\ \mathbf{m}^s(t_k) &= \mathbf{m}(t_k) + \mathbf{C}_{k+1} [\mathbf{m}^s(t_{k+1}) - \mathbf{m}^-(t_{k+1})] \\ \mathbf{P}^s(t_k) &= \mathbf{P}(t_k) + \mathbf{C}_{k+1} [\mathbf{P}^s(t_{k+1}) - \mathbf{P}^-(t_{k+1})] \mathbf{C}_{k+1}^\top. \end{aligned} \quad (35)$$

The values for $\mathbf{m}^-(t_{k+1}, \mathbf{r})$ and $\mathbf{P}^-(t_{k+1})$ were already calculated during the ‘forward sweep’ by the filtering equations (33–34).

3 Results

Särkkä *et al.* (2012) present many aspects of the DRIFTER method both using simulated and empirical fMRI data. The implementation in the article was published online as the ‘DRIFTER toolbox’ for Matlab¹, which follows the formulation that was presented in the article and coincides with the previous section. The results in the article show that in most cases DRIFTER performs well in identifying the oscillatory components and outperforms the competing RETROICOR method.

However, some questions were left open to discussion in Särkkä *et al.* (2012). These were, for example, the efficiency of the IMM implementation in identifying frequency trajectories without reference signals. This is of practical importance, because in clinical practice these are often not recorded on regular basis, and typically no such recordings exist for previous datasets. Another question regarding DRIFTER is the efficiency of the method in slowly sampled fMRI data with TRs ranging up to several seconds. These two questions are addressed in this section by providing some illustrative simulation study results.

3.1 Sampling Rates and Signal Aliasing

As the repetition times in MR studies often are significantly slower than 100 ms which was mainly used in the proof-of-concept examples in Särkkä *et al.* (2012), the effects of sampling rate and signal aliasing are of important nature. According to the sampling theorem by Nyquist and Shannon (see, e.g., Oppenheim *et al.*, 1999), a signal sampled at a frequency of f_s can be reconstructed only if the original signal is band-limited to contain only frequencies at and below the Nyquist frequency $f_s/2$. Higher frequency components become aliased into lower frequency components, and they cannot be observed with certainty.

However, if the signal consists of a finite number of sinusoids, then the recovery of the original signal is sometimes possible even when the signal has frequencies above the Nyquist limit (see Candès and Wakin, 2008). Fortunately the periodic physiological noise components in this study fall

¹The DRIFTER toolbox by Arno Solin and Simo Särkkä is available for download on <http://becs.aalto.fi/en/research/bayes/drifter/> (version 2012-04-25 used).

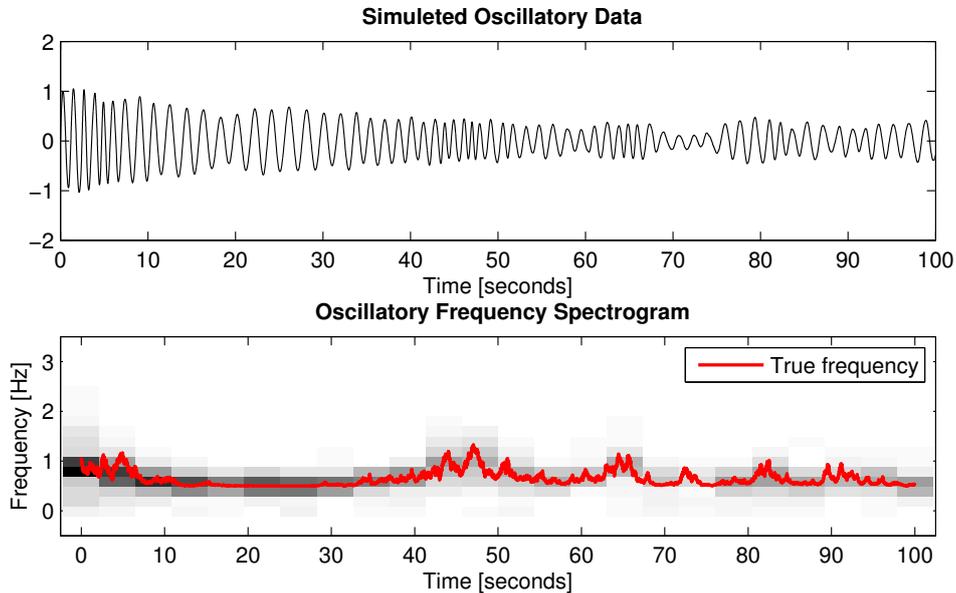


Figure 4: A simulated trajectory of a stochastic oscillator. The oscillating signal is shown in the upper figure and the signal spectrogram and the actual frequency trajectory in the lower figure.

into this category, and they can be identified if their frequencies are known beforehand. Obviously, the quality of the estimates decreases drastically as the Nyquist limit is reached, and some frequencies might cause problems, such as multiples of the Nyquist frequency. The reference signals, which are used for determining the frequencies of cardiac and respiratory signals, need to have at least one harmonic below the Nyquist frequency or otherwise it is difficult to estimate its frequency unambiguously.

We use simulated data, where we consider only one stochastic oscillator with no harmonics. The data is simulated using the same models that was presented in the previous section. The frequency trajectories are drawn randomly such that the random walk (Winener process) $w_k = w_{k-1} + \mathcal{N}(0, t_k - t_{k-1})$ defines a trajectory that is transformed into a frequency time series $f(t_k) = \frac{1}{2} + \frac{1}{1 + \exp(0.1w_k)}$ (in Hertz). Now the frequency $f(t)$ is constrained to the interval 0.5–1.5 Hz.

Figure 4 shows an example signal that has been simulated using the randomly drawn frequency trajectory. The stochastic oscillator signal shows changes in amplitudes and the effect of changing frequency. The lower figure shows the actual frequency trajectory together with the spectrogram of the

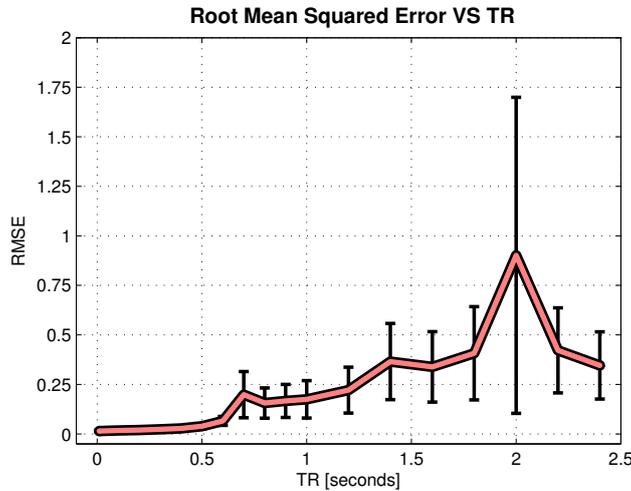


Figure 5: The average RMSE results per TR. The standard deviation is visualized with error bars.

simulated data.

We use the following parameters in our simulation: Each simulated trajectory is of length 100s, with discrete time interval lengths (i.e. TR) of $\Delta t = \{0.01, 0.05, 0.1, 0.2, \dots, 1, 1.2, \dots, 2.4\}$. We define $Q_c = 0.01$, and zero-mean Gaussian measurement noise with variance $\sigma^2 = 0.01^2$ is added to the simulated data. We simulate $N = 100$ periodic signals. The DRIFTER method is applied to the data and the estimated signal is compared against the true signal. The root mean squared error (RMSE) is reported for different TRs. The code implementation used in this study can be found attached as Appendix A at the end of this report.

The results in Figure 5 show that the method works best on fast-sampled data, and gradually the error increases as the TR grows. Rather unsurprisingly the error is larger around those sampling frequencies that are multiples of the Nyquist frequency of the average oscillator.

3.2 Estimation of Frequencies Without References

Even if no reference signals are available, it is often possible to estimate the frequencies of the cardiac and respiratory signals from the fMRI data. In practice, this is only possible when the sampling frequency is high enough such that at least the fundamental frequency is above the Nyquist limit. Thus the absolute minimum for sampling frequency would be a bit below

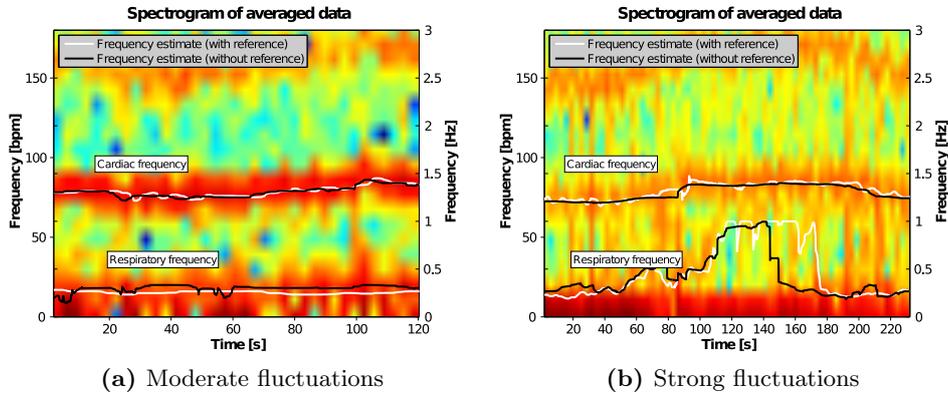


Figure 6: Example spectrograms of spatially averaged fMRI data. IMM based frequency trajectories are shown both for estimates based on the mean time series (white) and external data (black).

3 Hz (TR \sim 300 ms), assuming a cardiac frequency around 60–80 bpm. This limit can, however, be pushed by taking the timings of individual slices into account (e.g., as in Frank *et al.*, 2001).

With short-enough TR, the average over all the voxels in the fMRI slices can be used to construct an artificial reference signal. This time series should contain oscillating components representing both the cardiac and respiration-induced noises, because the cardiac and respiratory effects tends to be quite coherent throughout the brain (Glover *et al.*, 2000). By using data from only a few well-chosen voxels, the effect of the oscillatory components can be made even clearer. Regions around arterial voxels tend to show the cardiac pulsation more clearly, whereas for example voxels in the eyes show the respiration-related oscillations.

As a brief example, Figure 6 presents two spectrograms of such spatially averaged time series data corresponding to moderate and strong fluctuations in physiological noise. Here the results are based on actual fMRI data (datasets labelled 1 and 12) from Särkkä *et al.* (2012). The visualizations feature frequency trajectories estimated both with and without references which are in black and white, respectively.

4 Discussion and Conclusions

In this study we have gone through the formulation of the image-based Bayesian method for retrospective elimination of periodical physiological noise in fMRI, DRIFTER. A stochastic state-space model for modeling the signal components was constructed as a composition of stochastic oscillators and a slowly moving Wiener velocity model. An interacting multiple model (IMM), an adaptive Kalman filtering approach, tracked the oscillation frequencies in reference signals of physiological noise. These frequency estimates of respiration- and cardiac-induced noises were passed to Kalman filter and Rauch–Tung–Striebel smoother algorithms which carried out the Bayesian inference step.

DRIFTER is a robust and efficient algorithm for estimation and removal of physiological noise in fMRI data. Yet there are questions that may hinder the spreading to a wider audience. Most notably: (i) the method is model-based, (ii) the method requires pre-defined parameters to run, (iii) the method requires reference signals, and (iv) the proof-of-concept demonstrations has focused on fast-sampled fMRI data. All of these characteristic features were discussed in Särkkä *et al.* (2012) in a rather theoretical fashion.

However, to provide a more user-friendly point-of-view the following observations can be made: As for (i), the modeling can be restricted to the oscillatory structure of the noises alone, as the measurement noise estimate can be added back to the slowly-moving cleaned BOLD estimate. This restrains the risk of destroying relevant data. Yet, the study design has to be planned so that the stimulus responses do not mix up with the physiological oscillations (as discussed, e.g., by Krüger and Glover, 2001; Huettel *et al.*, 2004).

To address point (ii), the predefined parameter values are rather robust. The model is not sensitive to the parameter values as long as larger values are favoured in order to avoid too ‘stiff’ models (see discussion in Särkkä *et al.*, 2012), which makes it easy to try out the method without having to concentrate on choosing suitable parameters. Fortunately the parameters have clear physical interpretations, and it is rather straight-forward to combine the method with some parameter estimation methods (see, e.g., Särkkä and Nummenmaa, 2009).

In this study, the points (iii) and (iv) were looked into by providing two brief examples. Avoiding the use of external reference signals is possible if the TR is fast enough, as was demonstrated in Figure 6 with both moderately and strongly fluctuating physiological effects in the data. In frequency identification without references, image averaging seems efficient especially in the case of respiratory signals, because of its spatially homogeneous nature (Glover *et al.*, 2000), but for the cardiac it might be beneficial to only use the areas that are known to have stronger cardiac signal contribution (see, Dagli *et al.*, 1999; Glover *et al.*, 2000). However, tracking a strongly fluctuating signal can be tricky and subject to uncertainty, as was seen in Figure 6b. One alternative to increase the temporal resolution of this approach is to use the slice-timing information in multi-slice EPI data, which provides fast-sampled data even in the case of long TR. As shown by Frank *et al.* (2001), it is possible reconstruct the physiological signals by treating individual slices as separate time-ordered observations.

In most current fMRI studies the sampling rate is restricted to approximately two seconds. This restriction is enforced by the slice acquisition techniques and the need for around 30 slices for spatial whole-brain coverage in the data. Even with this slow sampling rates the variabilities associated with physiological sources of motion are present but become distributed throughout the fMRI time series. Provided we have good estimates for the frequency trajectories of the quasi-periodic noise components, it is possible to estimate their contribution to the signal using DRIFTER. The demonstrative results in Figure 5 suggest that DRIFTER can be exploited even with slow-TR data. However, care must be taken in choosing the TR, as aliasing becomes especially troublesome at multiples of the base frequency of the signal.

To conclude, the DRIFTER method shows prominent possibilities and it will hopefully gain users worldwide. Admittedly it has its weaknesses which stem from its strengths: robustness, adaptivity and modeling signal structure. In this work it has been shown that the most obvious weaknesses can be circumvented without rocket science, by careful study design (repetition times and reference signal recordings) and checking the estimation outcome with common sense.

References

- Bar-Shalom, Y., Li, X.-R., and Kirubarajan, T. (2001). *Estimation with Applications to Tracking and Navigation*. Wiley Interscience.
- Biswal, B., DeYoe, E., and Hyde, J. (1996). Reduction of physiological fluctuations in fMRI using digital filters. *Magnetic Resonance in Medicine*, 35(1):107–113.
- Buxton, R. (2009). *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques*. Cambridge University Press, second edition.
- Candès, E. J. and Wakin, M. B. (2008). An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30.
- Cheng, H. and Li, Y. (2010). Respiratory noise correction using phase information. *Magnetic Resonance Imaging*, 28(4):574–582.
- Chuang, K.-H. and Chen, J.-H. (2001). IMPACT: Image-based physiological artifacts estimation and correction technique for functional MRI. *Magnetic resonance in medicine*, 46(2):344–353.
- Dagli, M. S., Ingeholm, J. E., and Haxby, J. V. (1999). Localization of cardiac-induced signal change in fMRI. *NeuroImage*, 9(4):407–415.
- Deckers, R. H., van Gelderen, P., Ries, M., Barret, O., Duyn, J. H., Ikonomidou, V. N., Fukunaga, M., Glover, G. H., and de Zwart, J. A. (2006). An adaptive filter for suppression of cardiac and respiratory noise in MRI time series data. *NeuroImage*, 33(4):1072–1081.
- Frank, L. R., Buxton, R. B., and Wong, E. C. (2001). Estimation of respiration-induced noise fluctuations from undersampled multislice fMRI data. *Magnetic Resonance in Medicine*, 45(4):635–644.
- Glover, G. H., Li, T.-Q., and Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*, 44(1):162–167.
- Grewal, M. S. and Andrews, A. P. (2001). *Kalman Filtering, Theory and Practice Using MATLAB*. Wiley Interscience.

- Hu, X., Le, T. H., Parrish, T., and Erhard, P. (1995). Retrospective estimation and correction of physiological fluctuation in functional MRI. *Magnetic Resonance in Medicine*, 34(2):201–212.
- Huettel, S. A., Song, A. W., and McCarthy, G. (2004). *Functional Magnetic Resonance Imaging*. Sinauer Associates, Sunderland, MA.
- Hutton, C., Josephs, O., Stadler, J., Featherstone, E., Reid, A., Speck, O., Bernarding, J., and Weiskopf, N. (2011). The impact of physiological noise correction on fMRI at 7 T. *NeuroImage*, 57(1):101–112.
- Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45.
- Krüger, G. and Glover, G. H. (2001). Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magnetic Resonance in Medicine*, 46(4):631–637.
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D. N., Hoppel, B. E., Cohen, M. S., and Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences of the United States of America*, 89(12):5675–5679.
- Le, T. H. and Hu, X. (1996). Retrospective estimation and correction of physiological artifacts in fMRI by direct extraction of physiological activity from MR data. *Magnetic Resonance in Medicine*, 35(3):290–298.
- Ljung, L. and Glad, T. (1994). *Modeling of Dynamic Systems*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W.-L., and Nichols, T. E. (2006). Non-white noise in fMRI: Does modelling have an impact? *NeuroImage*, 29(1):54–66.
- Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872.
- Øksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications*. Springer-Verlag, sixth edition.

-
- Oppenheim, A. V., Schaffer, R. W., and Buck, J. R. (1999). *Discrete-Time Signal Processing*. Prentice Hall, second edition.
- Rauch, H., Tung, F., and Striebel, C. (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA journal*, 3(8):1445–1450.
- Särkkä, S. (2006). *Recursive Bayesian Inference on Stochastic Differential Equations*. Doctoral dissertation, Helsinki University of Technology.
- Särkkä, S. and Nummenmaa, A. (2009). Recursive noise adaptive Kalman filtering by variational Bayesian approximations. *IEEE Transactions on Automatic Control*, 54(3):596–600.
- Särkkä, S., Solin, A., Nummenmaa, A., Vehtari, A., Auranen, T., Vanni, S., and Lin, F.-H. (2012). Dynamical retrospective filtering of physiological noise in BOLD fMRI: DRIFTER. *NeuroImage*, 60(2):1517–1527.
- Solin, A. (2010). *Cubature Integration Methods in Non-Linear Kalman Filtering and Smoothing*. Bachelor’s thesis, Faculty of Information and Natural Sciences, Aalto University, Finland.
- Thomas, C. G., Harshman, R. A., and Menon, R. S. (2002). Noise reduction in BOLD-based fMRI using component analysis. *NeuroImage*, 17(3):1521–1537.
- Triantafyllou, C., Hoge, R., Krueger, G., Wiggins, C., Potthast, A., Wiggins, G., and Wald, L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *NeuroImage*, 26(1):243–250.

Appendices

Appendix A: Simulation Code

The following code requires the DRIFTER toolbox (version 2012-04-25 used in this work) which is available for download on the web page <http://becs.aalto.fi/en/research/bayes/drifter/> and distributed under the GNU General Public License.

Listing 1: The full simulation ‘runSimulation.m’ code for Mathworks Matlab (tested in Matlab 2010b).

```

1  function runSimulation
   %% runSimulation – Test the performance of DRIFTER
   %
   % Description:
5  %   This code runs the simulation study in the report section
   %   '3.1 Sampling Rates and Aliasing'. We use simulated data, where we
   %   consider only one stochastic oscillator with no harmonics. The
   %   data is simulated using the stochastic oscillator setup presented
   %   in the study.
10 %
   %   The frequency trajectories are drawn randomly such that the
   %   random walk (Winener process) defines a trajectory that is
   %   transformed into a frequency time series f (in Hertz). The drifter
   %   method is applied to the data after the simulated signal has been
15 %   transformed to downsampled observations (defined by dts). The mean
   %   squared error results are then captured and visualized.
   %
   %   Running this code requires the DRIFTER toolbox that is available
   %   for download on http://becs.aalto.fi/en/research/bayes/drifter/
20 %
   %   Copyright 2012 Arno Solin

   %% Simulate M draws

25  % Add path to the DRIFTER toolbox
   addpath /path/to/the/DRIFTER/toolbox/

   % Number of random draws
   M = 100;
30
   % Different time discretizations (TR) to consider
   dts = [0.01 0.05 0.1:0.1:1 1.2:.2:2.4];

   % Allocate space for results
35  MSE = zeros(M,numel(dts));
   C = zeros(M,numel(dts));

   % Open parallel pools to run the code in
   matlabpool open
40
   % Run in parallel

```

```

    parfor i=1:M
        [MSE(i,:),C(i,:)] = simulate_and_estimate(dts,100);
    end
45

    %% Visualize results

    figure(1); clf; hold on
50
    % Plot error bars and data points
    errorbar(dts,mean(sqrt(MSE),1),std(sqrt(MSE),[],1), ...
        'Color','k','LineWidth',2)
    plot(dts,mean(sqrt(MSE),1),'sk')
55
    %bplot(dts,mean(sqrt(MSE),1),3,'-',[0 0 0;1 .5 .5])

    % Set labels and title
    title('\bf Root Mean Squared Error VS TR','FontSize',12);
    xlabel('TR [seconds]'); ylabel('RMSE')
60

    % Modify appearance
    xlim([-1 2.5]); ylim([-0.05 2]); set(gca,'YTick',0:.25:4)
    box on; grid on; set(gcf,'Color','w')

65
    % Set paper size
    set(gcf,'PaperUnits','centimeters');
    set(gcf,'PaperSize',[14 10])
    set(gcf,'PaperPosition',[0.25 2.5 14 10])

70
    % Save figure
    saveas(gcf,'results.eps','eps')

end

75
function [MSE,C] = simulate_and_estimate(dts,Tend)
%% simulate_and_estimate - Estimate results for one draw
%
% Syntax:
80 % [MSE,C] = simulate_and_estimate(dts,Tend)
%
% In:
% dts - Time discretizations (in seconds)
% Tend - End time (in seconds)
85 %
% Out:
% MSE - Mean squared error for each dt
% C - Standard deviation estimate from the filtering estimate
%
90 % Description:
% A realization of a frequency trajectory in the band 0.5–1.5 Hz is
% simulated and then the function 'simulate_periodic_data' is called
% to simulate realizations of stochastic oscillatory signals with
% given parameters. Noise is added to simulate noisy observations
95 % and then the DRIFTER method is run with the simulated data. The MSE
% for each dt/TR is captured and returned.
%
% See also:
% drifter, simulate_periodic_data
100 %
% Copyright 2012 Arno Solin

%% Simulate periodic data

```

```

105 % Parameters
    N = Tend/dts(1);
    T = dts(1)*(0:N-1);
    Qc = 0.01;
    x0 = [0;1];
110
    % Set up frequency trajectory
    f = 0.1*cumsum(randn(1,N)); % random walk
    f = 1./(1+exp(f));         % transformed
    f = 0.5+1*f;               % in range (Hz)
115
    % Simulate full periodic data
    x = simulate_periodic_data(N,dts(1),f,Qc,x0);

120 %% Run analysis

    % Allocate space for results
    MSE = zeros(1,numel(dts));
    C = zeros(1,numel(dts));
125
    % Figure
    figure(1); clf

    % Loop
130 for i=1:numel(dts)

        % Indices
        ind = 1:int16(dts(i)/dts(1)):N;

135        % Set up DRIFTER
        clear data models
        data.data = x(1,ind)+0.01*randn(1,numel(ind));
        data.dt = dts(i);
        data.qr = 1e-9;
140        data.sd = 0.01;
        data.scalefactor = 1;
        data.mean = 0;
        models{1}.frequency = f(ind)*60;
        models{1}.dt = dts(i);
145        models{1}.qr = Qc;

        % Run DRIFTER
        [data,models,~,SPP] = drifter(data,models);

150        % Calculate MSE
        MSE(i) = mean((x(1,ind)'-squeeze(models{1}.estimate)).^2);

        % Take care of the covariance
        C(i) = mean(sqrt(squeeze(SPP(1,1,:))));
155
    end

end

160 function [x] = simulate_periodic_data(N,dt,f,Qc,x0)
    %% simulate and estimate - Estimate results for one draw
    %
    % Syntax:
165 % simulate_periodic_data(N,dt,f,Qc,x0)

```

```
%  
% In:  
% N - Number of steps to simulate  
% dt - Time discrteization step length (in seconds)  
170 % f - A vector of requencies in Hz  
% Qc - Spectral denisty of the dynamic noise term  
% x0 - Initial state  
%  
% Out:  
175 % x - 2xN-vector of simulated states  
%  
% Description:  
% Simulate a stochastic oscillator with a given frequency trajectory.  
% The model is set up as a continuous-time state space model, or a  
180 % stochastic differential equation. Runnig this code requires the  
% 'lti_disc' function that is available i.e. in the DRIFTER toolbox.  
%  
% See also:  
% lti_disc  
185 %  
% Copyright 2012 Arno Solin  
  
%% Simulate a stochastic oscillator  
  
190 % Allocate space for results  
x = zeros(2,N);  
  
% Initial state  
if nargin < 5 || isempty(x0)  
195 x(:,1) = randn(2,1);  
else  
x(:,1) = x0;  
end  
  
200 % The rest of the states  
for k=2:N  
  
% Dynamic model  
F = [ 0 2*pi*f(k);  
205 -2*pi*f(k) 0];  
  
L = [0;1];  
  
% Discretize  
210 [A,Q] = lti_disc(F,L,Qc,dt);  
  
% Determine if stochastic oscillator or deterministic  
if Qc<eps  
x(:,k) = A*x(:,k-1);  
215 else  
x(:,k) = A*x(:,k-1) + chol(Q)*randn(2,1);  
end  
end  
  
220 end  
end
```