Supplemental material for

Dynamic Retrospective Filtering of Physiological Noise in BOLD fMRI: DRIFTER

Simo Särkkä, Arno Solin, Aapo Nummenmaa,
Aki Vehtari, Toni Auranen, Simo Vanni, and Fa-Hsuan Lin

DRIFTER Toolbox for SPM

We have implemented the methods that were presented in the paper to an SPM toolbox, which is available for download on the toolbox web page:

`http://www.lce.hut.fi/research/mm/drifter/`

The toolbox has been tested with SPM8 (available at `http://www.fil.ion.ucl.ac.uk/spm/software/spm8/`) and offers both batch functionality and a graphical user interface trough the Batch Editor.

Resources on Kalman Filtering, RTS Smoothing and IMM

In the manuscript, there was no space to review the basic Kalman filtering and smoothing theory, and we had to resort to using citations instead of writing down the implementation retails of the underlying methods. However, the implementation details are very well documented in standard text books such as (Bar-Shalom et al., 2001; Grewal and Andrews, 2001) in the reference list of the article. Furthermore, implementations of all methods are included in the DRIFTER toolbox for SPM.

The MATLAB implementations of the filters and smoothers together with various other related methods can be found in the EKF/UKF toolbox for Matlab, whose link is given below. Links to documentation and course material using the same notation are also included.

- EKF/UKF toolbox for MATLAB can be found here:
  `http://www.lce.hut.fi/research/mm/ekfukf/`

- The documentation for the above toolbox provide a quite extensive overview of the methods:

  http://www.lce.hut.fi/research/mm/ekfukf/documentation.pdf

- Filtering and smoothing theory can be found, for example, in the following lecture notes:

  http://www.lce.hut.fi/~ssarkka/course_k2011/pdf/course_booklet_2011.pdf

Details on Benchmarking the Methods

In Section 3.3. *RMSE and SNR based benchmarking* in the article we presented a overview of how the methods were compared. A more detailed description of the benchmarking is presented here.

Root mean square error (RMSE)

The results were numerically benchmarked using two separate methods. The root mean squared error (RMSE) value was calculated for each estimate. The RMSE of an estimate $(\hat{x}_1, \ldots, \hat{x}_N)$ can be defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (x_i - \hat{x}_i)^2}, \tag{1}$$

where $(x_1, \ldots, x_N)$ is the actual known solution.

Signal-to-noise ratio (SNR)

The second benchmarking method was to compare the signal-to-noise ratios (SNR) of the estimates. Assume that the signal $x(t)$ is disturbed by a zero mean noise term $\epsilon(t)$, then the observation is

$$y(t) = x(t) + \epsilon(t). \tag{2}$$

The SNR of $y(t)$ can be defined as the ratio of the standard deviations of the signal and the noise:

$$\text{SNR} = \sqrt{\frac{\text{Var}[x(t)]}{\text{Var}[\epsilon(t)]}}. \tag{3}$$

Component standard deviation estimates

To be more precise, we assume that each voxel match a time series that can be split into components

$$y(t) = x(t) + c(t) + r(t) + \epsilon(t),$$

where $y(t)$ is the observed fMRI signal. As a result of running the methods, estimate values for each component are returned; $\hat{x}(t)$ is the cleaned BOLD signal estimate, $\hat{c}(t)$ and $\hat{r}(t)$ are the cardiac- and respiration-induced noise component estimates, and $\hat{\epsilon}(t)$ is the white measurement noise component estimate.

In RETROICOR and DRIFTER($x + \epsilon$) estimates of the physiological noise components are removed (i.e. an estimate of $x(t) + \epsilon(t)$ is returned), whereas in the DRIFTER(x) solution the estimated noise-free activation signal is separated (i.e. $\hat{x}(t)$ is returned explicitly). Utilizing these results, the following standard deviation values are studied:

- $\sigma_y^2 = \mathrm{Var}[y(t)]$ — The variance of the original observed fMRI signal.
- $\sigma_x^2 = \mathrm{Var}[x(t)]$ — The variance of the activation signal component. Real: From the actual underlying signal $x(t)$. Uncorrected: From the observed signal $y(t)$. RETROICOR / DRIFTER($x + \epsilon$): From the cleaned signal estimate component $\hat{x}(t) + \hat{\epsilon}(t)$. DRIFTER($x$): From the noise-free signal estimate $\hat{x}(t)$.
- $\sigma_c^2 = \mathrm{Var}[c(t)]$ — The variance of the cardiac noise component. Real: From the real cardiac component. Uncorrected: N/A. RETROICOR / DRIFTER: From the estimate of the cardiac component $\hat{c}(t)$.
- $\sigma_r^2 = \mathrm{Var}[r(t)]$ — The variance of the respiratory noise component. Real: Calculated from the real cardiac component. Uncorrected: N/A. RETROICOR / DRIFTER: From the estimate of the respiration-induced component $\hat{r}(t)$.
- $\sigma_\epsilon^2 = \mathrm{Var}[\epsilon(t)]$ — The variance of the white measurement noise component. Real: From the real added white noise $\epsilon(t)$. Uncorrected / RETROICOR / DRIFTER($x + \epsilon$): N/A. DRIFTER(x): Calculated by subtracting $y(t) - \hat{x}(t) - \hat{r}(t) - \hat{c}(t)$ and calculating the variance of the resulting noise estimate $\hat{\epsilon}(t)$.
- σ_n^2 — The variance of the remaining noise. Real: From the real added white noise $\epsilon(t)$. Uncorrected: From all noise in signal; $c(t)+r(t)+\epsilon(t)$ in the simulated case and $y(t)$ in-between activations (i.e. when $x(t) \approx 0$) in the empirical data case. RETROICOR / DRIFTER($x+\epsilon$): From the cleaned signal estimate component $y(t) - x(t) - \hat{c}(t) - \hat{r}(t)$ in the simulated case, and $y(t) - \hat{c}(t) - \hat{r}(t)$ in-between activations (i.e. when

$x(t) \approx 0$, see details later on) in the empirical data case. DRIFTER(x): From the cleaned signal estimate component. This is $\hat{x}(t) - x(t)$ in the simulated case, and $\hat{x}(t)$ in-between activations in the empirical data case.

As defined in equation (3), the SNR can now be calculated by

$$\mathrm{SNR} = \frac{\sigma_x}{\sigma_n}, \tag{4}$$

where both σ_x and σ_n differ for each five methods.

The above analysis holds on a per-voxel basis. The standard deviations vary spatially over the data, and therefore we normalize the values by dividing each standard deviation estimate by the signal standard deviation σ_y in the voxel. The arithmetic means of these normalized sigmas are then used in the analysis. This gives an insight into the effect of noise in the data. The same approach is used in calculating the signal-to-noise ratios; the mean of all the SNRs over all the voxels is used.

For fair comparison of the methods, four seconds of data was removed from the beginning and end of each time-series. This is due to both possible transients in the data and the peak detection method in RETROICOR.

Estimating the measurement noise in empirical data

In the case of analysis of empirical data, we choose voxels mostly from the high-order object-sensitive cortex for analysis. The voxels were chosen by studying activations in SPM that were contrasted using the uncorrected datasets. As there is no real signal to compare to, we estimated the noise variance σ_n^2 by studying parts of the data with minimal activations (i.e. $x(t) \approx 0$ at rest). These parts of the signal were extracted by using the stimulus timing information. Due to the post-stimulus effects of the hemodynamic responses, a two-second period of adaption was excluded after the end of each stimulus. However, the hemodynamic response will not be exactly zero after the two-second period, because before the response returns to zero, there is a delay and an undershoot effect, which typically last more than a couple of seconds (Handwerker et al., 2004).

Thus, as the neural activation induces a delayed hemodynamic response, the activation still has a small contribution to the variance estimates. To diminish this, we added a windowed de-trending smoothing filter to the estimation of the standard deviation of unexplained noise from the in-between activation data. The de-trending was implemented through a Savitzky-Golay smoothing filter (with polynomial degree one, and a window size of 5.1 s). The de-trending deals with some of the remaining hemodynamic effects and scanner drifting, but remaining post-stimulus undershoot effects still have a small contribution to the signal. With long TR, the filtering might cause slight underestimation of the noise, which may also have an effect to the results. The comparison is however fair, because the same benchmarking is used for all the methods.

An alternative way to study the remaining white measurement noise would be to use a linear model for modeling the HRF convolved response in the brain (as is done in SPM). This would however result in studying the model residual rather than the actual measurement noise.