Temporal kernel CCA and its application in multimodal neuronal data analysis

**Felix Bießmann · Frank C. Meinecke · Arthur Gretton ·
Alexander Rauch · Gregor Rainer ·
Nikos K. Logothetis · Klaus-Robert Müller**

Abstract Data recorded from multiple sources sometimes exhibit non-instantaneous couplings. For simple data sets, cross-correlograms may reveal the coupling dynamics. But when dealing with high-dimensional multivariate data there is no such measure as the cross-correlogram. We propose a simple algorithm based on Kernel Canonical Correlation Analysis (kCCA) that computes a multivariate temporal filter which links one data modality to another one. The filters can be used to compute a multivariate extension of the cross-correlogram, the canonical correlogram, between data sources that have different dimensionalities and temporal resolutions. The canonical correlogram reflects the coupling dy-

---

F. Bießmann (✉) · F.C. Meinecke · K.-R. Müller
Machine Learning Group, TU Berlin, Franklinstr. 28/29, 10587 Berlin, Germany
e-mail: fbiessma@cs.tu-berlin.de

F.C. Meinecke
e-mail: meinecke@cs.tu-berlin.de

K.-R. Müller
e-mail: klaus-robert.mueller@tu-berlin.de

A. Gretton
Dept. Empirical Inference, MPI Biological Cybernetics, Spemannstr. 38, 72076 Tübingen, Germany
e-mail: arthur.gretton@tuebingen.mpg.de

A. Rauch · N.K. Logothetis
Dept. Physiology of Cognitive Processes, MPI Biological Cybernetics, Spemannstr. 38, 72076
Tübingen, Germany

A. Rauch
e-mail: arauch@tuebingen.mpg.de

N.K. Logothetis
e-mail: nikos.logothetis@tuebingen.mpg.de

G. Rainer
Visual Cognition Laboratory, University of Fribourg, Chemin du Musee 5, 1700 Fribourg, Switzerland
e-mail: gregor.rainer@unifr.ch

namics between the two sources. The temporal filter reveals which features in the data give rise to these couplings and when they do so. We present results from simulations and neuroscientific experiments showing that tkCCA yields easily interpretable temporal filters and correlograms. In the experiments, we simultaneously performed electrode recordings and functional magnetic resonance imaging (fMRI) in primary visual cortex of the non-human primate. While electrode recordings reflect brain activity directly, fMRI provides only an indirect view of neural activity via the Blood Oxygen Level Dependent (BOLD) response. Thus it is crucial for our understanding and the interpretation of fMRI signals in general to relate them to direct measures of neural activity acquired with electrodes. The results computed by tkCCA confirm recent models of the hemodynamic response to neural activity and allow for a more detailed analysis of neurovascular coupling dynamics.

Keywords Canonical correlation analysis · CCA · kCCA · tkCCA · Neurovascular coupling

1 Introduction

With the advent of modern brain imaging techniques, such as functional magnetic resonance imaging (fMRI), it became feasible to measure neural activity in the whole brain non-invasively (Ogawa et al. 1990). Despite its relatively young age, fMRI is nowadays one of the most often used techniques in clinical or experimental neuroscience studies. Combinations of fMRI with machine learning methods (Norman et al. 2006) improved the specificity of fMRI analysis methods to an extent that some studies set out to use fMRI and machine learning techniques in the context of lie detection (Langleben et al. 2005) or mind reading of hidden intentions (Haynes et al. 2007). However, the relationship between the blood-oxygen level dependent (BOLD) signal measured with fMRI and the underlying brain activity is still not fully understood (Logothetis 2008). The most straightforward—yet technically challenging—way to investigate this relationship, the so called neurovascular coupling, is to simultaneously measure the electrophysiological activity of the neurons and the fMRI response (Logothetis et al. 2001). Brain activity, that is the change in the electrical charge across the neurons' cell membranes, is reflected directly in the invasive electrophysiological measurements. Changes in neuronal activity correlate with the level of blood oxygenation which can be measured non-invasively using fMRI. The two methods, invasive electrophysiology and non-invasive fMRI, thus provide two views on brain activity, the former directly, the latter indirectly via the hemodynamic response. While the electrophysiological signal is measured at high temporal resolution, the fMRI signal has a low temporal resolution, due to both technical and physiological reasons. The spatial resolution of electrophysiological data can be high, but in these simultaneous recordings it is typically measured only at one electrode. In contrast, the fMRI signal measures the BOLD signal potentially in the whole brain, yielding data with some ten thousands of dimensions, i.e. voxels. When investigating the relationship between these two signals, there are two main problems: first the data is inherently multivariate (with different dimensionalities for each modality), and secondly the coupling is not instantaneous. Canonical Correlation Analysis (CCA) solves the first problem by computing a filter for each variable that maximizes the correlation between the two signals. CCA (Hotelling 1936) works by recursively computing maximally correlated projections of a pair of data sources, where each projection is uncorrelated with the previous projection. The projection directions then reveal the features of the two sources

that are most corelated. The second problem could be solved by looking at the full cross-correlogram instead of just the correlation coefficient between the sources. However, for multivariate data it is not clear how to compute such a correlogram, as the filters for each source might change as a function of the timelag τ. In this work we present an algorithm that tackles both problems at the same time. It can deal with delayed couplings between sources and is well suited for high dimensional data as it is based on kernel canonical correlation analysis. Kernel canonical correlation analysis (see Akaho 2001; Bach and Jordan 2002; Fukumizu et al. 2007, for a review on kCCA see Shawe-Taylor and Cristianini 2004, for a general review of the '*kernel trick*' see Schölkopf and Smola 2001 or Müller et al. 2001) has been widely applied in finding dependencies between random variables in high, or even infinite, dimensions. Using CCA in its kernelized form allows us to avoid dealing with co-variance matrices on very high dimensional data (as is the case with fMRI), which can be computationally prohibitive and numerically unstable. When using linear kernels the projections found by kCCA can be easily interpreted in the respective data space. kCCA with linear kernels has been successfully applied to analyze dependencies between electrophysiological data and visual stimuli (Macke et al. 2008) or between fMRI data and emotionally salient stimuli (Hardoon et al. 2007) to name just two examples. In both examples the correlation between the measured brain signal and an experimentally controlled stimulus is maximized. If the coupling between the data sources has some delay, it is advantageous to account for that when using CCA. Previous attempts to include temporal structure in a CCA framework are presented in (Friman et al. 2002) and (Macke et al. 2008). In the first case, only one data source is considered (timeseries of fMRI images) and its auto-correlation is maximized by applying linear kCCA to time shifted versions of the data. In (Macke et al. 2008), the authors employ CCA for mapping receptive fields of early visual neurons. Temporal structure of the correlations between the visual stimulus and the cell activity is accounted for by appending to each sample of neural activity its temporal context after stimulus onset.

While the aforementioned approaches do incorporate the temporal structure of the dependencies to some extent, the temporal dynamics of canonical correlations between multiple sources have not yet been accounted for in association with kCCA methods. Instead of estimating stationary projections as in standard kCCA, the proposed algorithm computes a multivariate temporal filter that maximizes the correlation between the two data sources over a certain time window. The temporal filter can be interpreted as a convolution on either data source and the filter coefficients can give valuable insights in the temporal coupling dynamics between the data sources.[1] Moreover, the filters can be used to compute a generalization of the cross-correlogram via CCA, which we denote the *canonical correlogram*. We begin our paper in Sect. 2 with an overview of linear kCCA. In Sect. 2.3, we generalize the original CCA algorithm to take into account time structure in a way that leads to easily interpretable results. Finally, in Sect. 3, we provide experimental results of our approach on both real and artificial datasets.

---

[1]In the case of neurovascular coupling the relationship is commonly referred to as the *hemodynamic response function* (HRF) (see for instance Logothetis and Wandell 2004 or Friston et al. 2000); while the neurons in early visual areas of the brain respond immediately after visual stimulation, the BOLD response to a stimulus peaks only 4–6 seconds afterwards; to correct for this time lag when correlating fMRI data with experimental stimuli, the stimulus timecourse is often delayed in time with respect to the fMRI signal so that stimuli and BOLD response appear roughly simultaneously; a popular alternative is to convolve the stimulus timeseries with a temporal filter modeling the dynamics of a generic HRF.

2 Temporal kernel canonical correlation analysis

2.1 Instantaneous canonical correlation analysis

For two multivariate variables $x \in \mathbb{R}^M$ and $y \in \mathbb{R}^N$, Canonical Correlation Analysis (CCA) estimates two normalized linear filters $w_x \in \mathbb{R}^M$ and $w_y \in \mathbb{R}^N$, called *canonical variates*, such that the correlation between the projections $w_x^\top x$ and $w_y^\top y$ is maximized (see e.g. (Anderson 1958) for a review):

$$
\begin{aligned}
\operatorname*{argmax}_{w_x, w_y} \quad & w_x^\top C_{xy} w_y \\
\text{s.t.} \quad & w_x^\top C_{xx} w_x = 1, \\
& w_y^\top C_{yy} w_y = 1,
\end{aligned}
\tag{1}
$$

where $C_{xx} \in \mathbb{R}^{M \times M}$ and $C_{yy} \in \mathbb{R}^{N \times N}$ denote the respective auto-covariance matrices and $C_{xy} \in \mathbb{R}^{M \times N}$ the crosscovariance matrix between x and y. If the number of dimensions is larger than the number of samples, it becomes more efficient and numerically more stable to optimize the filters w_x and w_y using kernel methods, i.e. kernel CCA (Akaho 2001; Bach and Jordan 2002; Fukumizu et al. 2007). For centered data, i.e. $\sum x = \sum y = 0$, the linear kernel matrices K_X and K_Y are given as the inner product of the data matrices $X = [x_1, x_2, x_3, \ldots, x_L] \in \mathbb{R}^{M \times L}$ and $Y = [y_1, y_2, y_3, \ldots, y_L] \in \mathbb{R}^{N \times L}$:

$$
\begin{aligned}
K_X &= X^\top X, \\
K_Y &= Y^\top Y.
\end{aligned}
\tag{2}
$$

The filters in the input space of each variable are then given as a linear expansion of the data points

$$
\begin{aligned}
w_x &= X\alpha, \\
w_y &= Y\beta,
\end{aligned}
\tag{3}
$$

where the vectors $\alpha \in \mathbb{R}^L$ and $\beta \in \mathbb{R}^L$ are the solution of the generalized eigenvalue problem in kernel space:

$$
\begin{bmatrix} 0 & K_X K_Y \\ K_Y K_X & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \rho \begin{bmatrix} K_X^2 & 0 \\ 0 & K_Y^2 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.
\tag{4}
$$

Note that even though in this paper we use only linear kernels, all technical considerations are equally valid for general (i.e. nonlinear) kernels; however the interpretation of the variates becomes more difficult. For example, in fMRI the linear filters w_y can easily be visualized as field pattern in voxel space. The interpretation of nonlinear CCA projections is not always that straightforward, but depending on the problem they can give valuable insights, too. Also, for the sake of simplicity, we considered only the largest eigenvalue ρ. However, the whole subsequent analysis can be performed with the following eigenvalues in the same manner. The eigenvalue ρ is the correlation coefficient between the projections $w_x^\top X$ and $w_y^\top Y$. As the eigenspectrum of (4) is symmetric, we just pick the largest positive value and neglect its negative counterpart.

In practice, to prevent kCCA from overfitting, it is often necessary to regularize the kernel matrices. A number of different types of regularization have been proposed (Shawe-Taylor and Cristianini 2004; Bach and Jordan 2002; Fukumizu et al. 2007); in this paper, the formulation of (Fukumizu et al. 2007) has been adopted, however in the experiments presented in Sect. 3.3 the diverse nature of the two data sources makes it necessary to have separate regularization parameters κ_X and κ_Y for the respective kernels; the regularized version of the generalized eigenvalue problem is then:

$$
\begin{bmatrix} 0 & K_X K_Y \\ K_Y K_X & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \rho \begin{bmatrix} K_X^2 + \kappa_X K_X & 0 \\ 0 & K_Y^2 + \kappa_Y K_Y \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.
\tag{5}
$$

The values for κ_X and κ_Y can be chosen using e.g. a resampling-based procedure that optimizes the distance to randomized surrogates of the data (see Sect. 2.5 for details).

2.2 Time-shifted kernel CCA

When the dependencies between signals exhibit temporal dynamics it is often more interesting to look at time-shifted instead of instantaneous canonical correlations. If there is some kind of delayed coupling between the signals or they both reflect the same underlying process, but with different delay τ, the canonical correlation between the two sources will be higher if the signals are shifted in time relative to each other. For *time-shifted kernel CCA*, one would simply replace the kernel of one signal, say K_X, by the kernel $K_\tau = X_\tau^\top X_\tau$, where X_τ denotes the respective signal, time-shifted by the delay τ. Since only *relative* time shifts between the signals matter, the other kernel stays untouched. The canonical variates are recovered analogously to (3):

$$
\begin{aligned}
w_x(\tau) &= X_\tau \alpha_\tau, \\
w_y(\tau) &= Y \beta_\tau.
\end{aligned}
\tag{6}
$$

Note that even though only one signal is shifted in time, both canonical variates do depend on τ, due to the τ-dependency of the expansion coefficients α_τ, β_τ. With this type of time-shifted kCCA it is possible to probe the signals and their canonical correlations for different time delays.

Note that for CCA the number of samples has to be equal for all data sources. When the data sources have different temporal resolutions, this implies that one source has to be downsampled. If X has the higher resolution, it can be shifted in small increments *before* downsampling, which means that the temporal resolution of the variates is upper bounded only by the sampling rate of X and does not depend on Y.[2]

In principle, calculating the time-shifted CCA sequentially for multiple time shifts $\tau \in \{\tau_1, \tau_2, \ldots, \tau_T\}$ would allow to compute an entire *canonical correlogram* between the data sources for a given time window, along with the corresponding τ-dependent canonical

---

[2]If X has a temporal resolution of 1 Hz (sample duration of 1 second) and Y a resolution of 0.25 Hz (sample duration of 4 seconds), we can shift X in increments of 1 s *before* downsampling it to 0.25 Hz. If, on the other hand, X had the lower sampling rate and Y the higher one, we could shift X only in increments of 4s, yielding larger gaps between possible time shifts τ. This means that the temporal resolution of the variates is upper bounded only by the sampling rate of X and does not depend on Y; the sliding window approach will not allow to resolve high frequency 'fine structure', but still allows to locate e.g. maxima or minima of the correlation with a slow signal with a good precision.

variates. We will refer to this approach in the following as *sequential time shifted kCCA* or sequential ts-kCCA. However, this approach has some inherent problems. In real applications like fMRI data analysis finding the correct regularization parameters κ_X and κ_Y turns out to be crucial. In general, parameter tuning by cross validation or resampling will demand different optimal parameters for different τ. Using different regularizers, however, would spoil the comparability of the canonical correlations and the canonical variates obtained for different time lags. In order to have an interpretable correlogram, the regularizers must be fixed for all τ. One might employ some heuristics for finding the right τ to optimize the parameters and use that parameter set for all other time lags, for instance if one has *a priori* knowledge about the nature of the data sources and their coupling. But in general it is not obvious what is the right choice for a common regularizer.[3]

While this problem might take a back seat if the range of optimal κ_X and κ_Y is not too large, there are two more severe problems that affect the temporal coherence of the canonical variates. The first reason for temporal incoherence is that kernel CCA always picks the largest positive eigenvalue from the symmetric eigenspectrum, so the canonical components (i.e. the projections of the data matrices onto the canonical variates) are always positively correlated. kCCA simply chooses the sign of the canonical variates accordingly, but it only takes into account relative signs (between variables). This means that if the variates $w_x(\tau_1)$ and $w_y(\tau_1)$ are a solution at time shift τ_1, then the variates $-w_x(\tau_1)$ and $-w_y(\tau_1)$ form a completely equivalent solution. Thus kCCA does not distinguish the absolute sign of the canonical variates but retrieves the relative sign between the two, rendering the absolute sign of a single variate completely arbitrary. If we start a new kCCA run on the next time shift τ_2, there is no guarantee that the choice of sign matches the one at the previous time step. In fact it turns out that kCCA sometimes produces sudden sign flips between neighboring τ-values. This problem could be tackled using heuristics such as enforing smoothness of the canonical variates over time. Starting from $w_x(\tau_1)$, one could fix the sign of each time-lagged variate $w_x(\tau_k)$ with $k \geq 2$ by setting

$$w_x(\tau_k) \leftarrow w_x(\tau_k) \, \text{sign} \left(w_x(\tau_{k-1})^\top w_x(\tau_k) \right). \tag{7}$$

This makes sure that the angle between adjacent variate vectors is smaller than 90 degrees. Note that this heuristic will fail when the estimates of the variates $w(\tau)$ are very noisy. In some cases however, this will lead to a sensible temporal structure of the canonical variates, as the results in Sect. 3.2.2 show. The second reason for the incoherent evolution of the variates in τ are indeterminacies that can occur at certain time shifts due to degenerate eigenvalues. If for some τ-values the first eigenvalue comes very close or even equal to the second eigenvalue, the very notion of 'largest eigenvalue' loses its meaning, since every vector in the space spanned by these eigenvectors is a solution with identical eigenvalue. In this case, kCCA picks an arbitrary vector from that eigenspace which will cause jumps in the canonical variates.

2.3 Temporal kernel CCA

In the previous section we have seen that in general a sequential kCCA will not yield solutions that are coherent for different τ with respect to the signs and the regularization across time lags. If a temporal interpretation of the results is intended, i.e. an interpretation of the

---

[3]A comparison with alternative regularization strategies, e.g. presented in (Hardoon et al. 2004), would be an interesting topic of future research but is beyond the scope of this work.

evolution of canonical correlation and canonical variates over τ, this poses severe problems. Even though it is sometimes possible to correct for these temporal incoherencies afterwards by some heuristics or visual inspection, this endeavor is tedious (and error prone) even for simple data sets; for data sources like fMRI it will quickly become unfeasible. To obtain temporally coherent filters from the start, we therefore propose a *temporal kernel CCA (tkCCA)* to solve the kCCA for all time shifts τ at once.

We have seen that CCA and linear kernel CCA estimate filters w_x and w_y such that the canonical components $w_x^\top X$ and $w_y^\top Y$ are maximally correlated. In application domains like the simultaneous measurement of electrophysiology and fMRI, the highest correlation is obtained at some time shift between the two modalities. In general we have to assume that there is not just a fixed delay, but a temporal filter that links one modality with the other one. In particular, the translation from neural activity to the BOLD signal is often modeled by a hemodynamic response filter (see e.g. Friston et al. 2000). We will therefore generalize the CCA optimization problem by allowing w_x to be an arbitrary spatio-temporal filter:

$$\max_{w_x(\tau), w_y} \text{Corr}\left(\sum_\tau w_x(\tau)^\top X_\tau, w_y^\top Y\right). \tag{8}$$

To maximize this correlation, we embed the data set X into a higher dimensional space where the different dimensions are occupied by copies of the original data, each shifted in time by a different τ from a set $\{\tau_1, \tau_2, \ldots, \tau_T\}$ of delays:

$$\tilde{X} = \begin{bmatrix} X_{\tau_1} \\ X_{\tau_2} \\ \vdots \\ X_{\tau_T} \end{bmatrix} \in \mathbb{R}^{MT \times L}. \tag{9}$$

By using the same time-delay embedding for the $w_x(\tau)$,

$$\tilde{w}_x = \begin{bmatrix} w_x(\tau_1) \\ w_x(\tau_2) \\ \vdots \\ w_x(\tau_T) \end{bmatrix} \tag{10}$$

we reduce the optimization problem in (8) to an ordinary CCA problem in the embedded space:

$$\phi = \max_{\tilde{w}_x, w_y} \text{Corr}\left(\tilde{w}_x^\top \tilde{X}, w_y^\top Y\right). \tag{11}$$

The kCCA can now be calculated for the kernels $\tilde{K}_X = \tilde{X}^\top \tilde{X}$ and $K_Y = Y^\top Y$, i.e. between the embedded version $\tilde{X}$ of one source and the unchanged version Y of the other one. The temporal resolution of the correlogram is solely determined by the spacing of the τ-values, in particular it can be higher than the temporal resolution of the signals in Y. Note that the embedded kernel $\tilde{K}_x$ is of the same size as the kernel K_x. Thus after the kernel $\tilde{K}_X$ has been computed tkCCA is computationally not more demanding than normal kCCA. Consequently the solution of embedded tkCCA is much faster to obtain than in the sequential time shifted kCCA. Another important advantage is that crossvalidation procedures or bootstrap resamplings can operate on the already computed kernel, which speeds up these

computations further. Since tkCCA just leads to a single eigenvalue equation of the type given by (5), neither the canonical correlation ρ nor the corresponding vectors α and β depend on τ. However, it is easy to recover the τ-dependent filter $w_x(\tau)$. In analogy to (3) we obtain

$$\tilde{w}_x = \tilde{X}\alpha,$$
$$w_y = Y\beta \tag{12}$$

and the delay-dependent canonical variates $w_x(\tau)$ are just given by the respective parts of $\tilde{w}_x$:

$$w_x(\tau) = X_\tau \alpha. \tag{13}$$

The canonical correlogram can now be calculated as the cross-correlation function between the canonical components (i.e. the data sets projected onto the filters $w_x(\tau)$ and w_y):

$$\rho(\tau) = \mathrm{Corr}\left(w_x(\tau)^\top X_\tau, w_y^\top Y\right)$$
$$= \frac{w_x(\tau)^\top X_\tau Y^\top w_y}{w_x(\tau)^\top X_\tau X_\tau^\top w_x(\tau) \cdot w_y^\top YY^\top w_y}$$
$$= \frac{\alpha^\top K_\tau K_Y \beta}{\alpha^\top K_\tau^2 \alpha \cdot \beta^\top K_Y^2 \beta}. \tag{14}$$

By design, the proposed tkCCA method does not suffer from the problems of temporal inconsistency sketched in Sect. 2.2. It is worth mentioning that the estimated canonical correlogram is different from a usual correlogram between two one-dimensional time signals in the sense that it is always positive.[4] This is due to the already mentioned fact that kCCA is ignorant with respect to the sign of the correlation; possible negative correlations between the signals are encoded in the relative sign of the canonical variates. So, to interpret the results of tkCCA, one always has to take into account both the canonical correlogram $\rho(\tau)$ and the corresponding canonical variates $w_x(\tau)$ and w_y. Just as an aside, while the great majority of work on CCA considers the solution of CCA simply as projection directions that maximize the correlation between the sources, an alternative perspective is presented in (Bach and Jordan 2006), where the authors propose a generative model for the solutions of CCA. It would be interesting to explore this generative model perspective for tkCCA as well.

2.4 Other possible approaches

One could think of alternatives to the presented approaches of temporal kCCA employing generalizations of kCCA to more than two variables, so called *multi-way CCA* (see Kettenring 1971 for CCA and e.g. Bach and Jordan 2002 for the kCCA version). Using an implementation based on multi-way CCA as in Friman et al. 2002 (see also Sect. 1) with slight modifications could also yield canonical variates that might reveal temporal structure

---

[4]Since the $\rho(\tau)$ are not directly calculated by kCCA, but by (14), this is not strictly true, i.e. only up to noise fluctuations.

in the couplings.[5] Essentially the modifications would involve computation of a multi-way kCCA where the first variable is one version of either of the two data sources and the other T variables are time shifted versions of the respective other data source, for instance one kernel K_F for the fMRI data and T kernels for the time shifted electrophysiology. The optimization could be formulated analogously to (Bach and Jordan 2002) and performing incomplete Cholesky decompositions on the kernel matrices can reduce the eigenvalue problem to a tractable size.

Unfortunately this multi-way formulation has a serious drawback: Time shifted versions of one data source are treated as separate variables, each having the same importance as the other data source. This means that if the auto-correlations of the time shifted signal are high, the canonical correlation can be maximized using mainly the auto-correlation from one data source and neglecting the cross correlation terms with the other data source. We performed all simulations presented in the following section with a multi-way temporal kCCA version in order to confirm these points. As we are not interested in the auto-correlation of one data source but rather the coupling over time between two data sources, multi-way kCCA with time lagged versions of one source as additional variables is not a favorable option.

2.5 Optimization of regularization parameters

In simulations as well as in applications on real data it is important to select good regularization parameters κ_Y and κ_X. If the regularizers are chosen too small, the tkCCA shows overfitting effects, resulting in high canonical correlations even for time series without correlations, if the regularizers are too large, the existing correlation structure in the data is suppressed. We therefore picked the regularizers by a surrogate data technique: For each set of regularization parameters (one for each variable) we performed tkCCA first on the true data and then on a surrogate data set, which contains the same samples as the true data but the time order of samples has been randomly permuted in one data source.[6] Since the correlation structure is destroyed by reshuffling, the canonical correlation $\hat{\phi}$ on the surrogate data will be different (i.e. smaller) than the canonical correlation ϕ on the true data. We also tried standard cross-validation for hyperparameter optimization, but the reshuffling is easier to implement, computationally cheaper and performs as well as cross-validation. The regularizers are now chosen such that this difference is as pronounced as possible:

$$(\kappa_X^*, \kappa_Y^*) = \underset{\kappa_X, \kappa_Y}{\operatorname{argmax}} \mathrm{E}[(\phi - \hat{\phi})^2]. \tag{15}$$

We approximated this expectation by an average over ten randomly shuffled surrogates and optimized the parameters over a range from 10^0–10^{-4}. Note that for this model selection procedure we do not need to compute the full canonical correlogram $\rho(\tau)$ but only the canonical correlation in the embedded space ϕ.

---

[5]The work presented in (Friman et al. 2002) is actually just a two-way CCA, as there is only one time shifted timeseries involved; however, it can be easily extended to more than one time lag in a framework such as the one presented in (Bach and Jordan 2002).

[6]Since only *relative* time dependencies between the two modalities matter, only one data set has to be shuffled.

3 Simulations and experiments

In the following, we apply the proposed tkCCA method to both simulated and real experimental data. After a short toy data example in Sect. 3.1 highlighting the main advantages of tkCCA, we proceed with more realistic simulations illustrating the effect of autocorrelations (Sect. 3.2.3) and different coupling parameters (Sect. 3.2.4 and 3.2.5) on the tkCCA result. Finally in Sect. 3.3 we apply tkCCA on simultaneous recordings of neural and hemodynamic activity and discuss how tkCCA reflects the neurovascular coupling mechanisms.

3.1 A simple illustration of the tkCCA method

Before we apply tkCCA to the more complex data from simulated and real experiments, we illustrate the algorithm and its results on the following simple example. From the univariate time series $s(t)$ (depicted in Fig. 1, panel A, top) we create two bivariate signals $x(t)$ and $y(t)$, one of them containing a time-lagged version of $s(t)$:

$$x(t) = \alpha\, s(t) + \epsilon \begin{bmatrix} n_1(t) \\ n_2(t) \end{bmatrix}, \tag{16}$$

$$y(t) = \beta\, s(t - \tau^*) + \epsilon \begin{bmatrix} n_3(t) \\ n_4(t) \end{bmatrix}, \tag{17}$$

where the vectors $\alpha, \beta \in \mathbb{R}^2$ mix $s(t)$ into the different dimensions of the signals $x(t)$ and $y(t)$ and $n_i(t) \sim \mathcal{N}(0, 1)$ are Gaussian i.i.d. sources. The signals $x(t)$ and $y(t)$ represent measurements obtained by different modalities. In the simulations, τ^* was set to 6 and $\alpha = \beta = [0.1\ 0.9]^\top$ and the noise level to $\epsilon = 0.15$. We then used sequential time-shifted kCCA and tkCCA in order to recover temporal filters that reflect the dependencies hidden in the data. The example illustrates that even for simple settings of low dimensionality, the canonical variates estimated by sequential kCCA are incoherent with respect to their sign and absolute value and thus difficult to interpret (Fig. 1, panel B, left column, middle and bottom plot). In contrast, the variates of tkCCA are close to zero for all values of τ except the variate at the delay specified in the simulations (Fig. 1, panel B, right column, middle plot). For both filters, tkCCA assigns the correct mixing coefficients α and β. As expected for this simple scenario, the canonical correlogram looks very similar for both methods— in both implementations, the mixing coefficients at $\tau = \tau^*$ are close to correct values and thus at a time lag of 6 samples, there is a pronounced peak found by both implementations. However, as it is crucial for the interpretation of this correlogram to take into account the variates, the solution found by tkCCA is preferable.

3.2 Application to simulated fMRI/electrophysiogical data

3.2.1 Data generation

To illustrate the behavior of the tkCCA method in detail, we perform controlled simulations that mimic the fMRI-experiments we will present later. In those experiments, a visual stimulus is presented to a subject while the neural activity in the primary visual cortex and the corresponding BOLD response is measured simultaneously. A detailed description will be given in Sect. 3.3. Just like the experimental data, the simulated data contains two different time series designed to match the main characteristics of the real data: One represents the
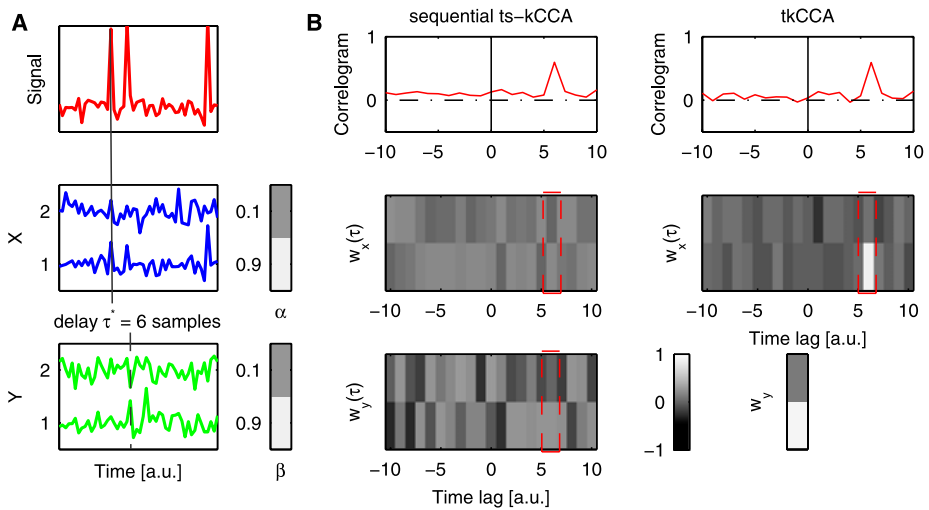
Fig. 1 Toy data example illustrating sequential time-shifted kCCA and tkCCA; **A**: signal $s(t)$ (*top*) and two two-dimensional measurement time series X (*blue*, no delay) and Y (*green*, 6 samples delay); time dimension is in arbitrary units [a.u.]; **B, left**: Results for sequential kCCA, canonical correlogram (*top*) and variates $w_x(\tau)$ (*middle*) and $w_y(\tau)$ (*bottom*); **B, right**: Results obtained using tkCCA; in contrast to sequential kCCA, only $w_x(\tau)$ has a temporal dimension (*middle plot*), whereas w_y is constant over time (*bottom*); the correct mixing is reliably recovered in both variates; importantly, tkCCA assigns high weights only at $\tau = 6$ in the time-dependent variate $w_x(\tau)$; as expected there are no significant differences for the correlograms in this simple example (*panel B, top row*)

neural activity in 8 different frequency bands[7] at a certain electrode position in the brain, the other one represents the BOLD activations in a 50×50-pixel image patch. To show how tkCCA can capture hidden structure in the data, we varied different model parameters in a controlled fashion. Data generation is performed in three steps (a schematic diagram of the data generation process is shown in Fig. 2):

1. *Electrophysiology:* First, a time series $E(f, t)$ is generated, representing the band power of the neural activity at frequency f and time t. At each frequency, the neurophysiolog-ical signal is modeled as a weighted sum of the visual stimulus $s(t)$ and spontaneous neural activity, represented as Gaussian i.i.d. noise $\epsilon(f, t) \sim \mathcal{N}(0, 1)$:

$$E(f, t) = \sqrt{\gamma \alpha(f)} \cdot s(t) + \sqrt{(1 - \gamma \alpha(f))} \cdot \epsilon(f, t). \tag{18}$$

   A frequency specific gain $0 \leq \alpha(f) \leq 1$ controls the relative stimulus susceptibility of each frequency band and the overall signal-to-noise ratio of the stimulus is controlled by the parameter $0 \leq \gamma \leq 1$. (The stimulus $s(t)$ is a boxcar time series with unit variance, the square roots ensure that $E(f, t)$ is properly normalized.)

2. *Neuro-vascular coupling:* The fMRI data is created from the electrophysiological data by first applying a hemodynamic response filter h (as implemented e.g. in the fMRI analysis toolbox *SPM5*, SPM5 2005) to each frequency band and then taking the weighted sum

---

[7]Different bands of electrophysiological bandpower are selective for independent stimulus features as re-ported in e.g. (Belitski et al. 2008) (see also Sect. 3.2.3).
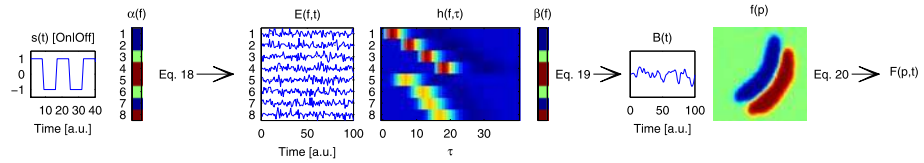
Fig. 2 Data generation: neural data $E(f, t)$ is simulated by a boxcar stimulus projected in 8 frequency bands with frequency specific gain, controlled by $\alpha(f)$ (18); the one dimensional BOLD time series $B(f)$ is simulated by convolving $E(f, t)$ with a time-frequency filter $h(f, \tau)$ and projecting the result on a frequency specific coupling filter $\beta(f)$ (19); the high dimensional BOLD time series $F(p, t)$ is then generated as the outer product with the spatial pattern $f(p)$ (see (20))

over all frequency bands:

$$B(t) = \frac{1}{Z} \sum_f \beta(f) \left(\sum_\tau h(f, \tau) E(f, t - \tau) \right), \tag{19}$$

where Z is a normalization parameter that sets $B(t)$ to unit variance. Note that we allow the hemodynamic response filter h to be frequency dependent. These filters together with the coupling weights $0 \leq \beta(f) \leq 1$ define the neuro-vascular coupling.

3. *BOLD signal:* Spatially resolved fMRI time series are obtained as the product of the time series representing the hemodynamic response $B(t)$ with a spatial pattern $f(p)$ and some additional noise:

$$F(p, t) = \text{sign}(\eta f(p)) \sqrt{|\eta f(p)|} \cdot B(t) + \sqrt{|(1 - \eta f(p))|} \cdot \epsilon(p, t). \tag{20}$$

The spatial pattern represents a correlational structure similar to those arising from correlated activity of neural ensembles. The signal-to-noise ratio of the spatial pattern is controlled by the parameter $0 \leq \eta \leq 1$.

In the following sections, we will show the influence of different model parameters on the tkCCA results. In the simulation model, we introduced three frequency dependent parameters, the relative stimulus susceptibility $\alpha(f)$, the neurovascular coupling strength $\beta(f)$ and the hemodynamic response filter $h(f, \tau)$. To separate the effects of these parameters, we performed three different sets of simulations. In each of them, we allowed only one of those three parameters to have an explicit frequency dependence, the other two were kept constant for different frequencies f. This approach allows us to uniquely identify the influence of each parameter on the tkCCA results. Within each of these settings, we varied both the overall stimulus susceptibility γ and the signal-to-noise ratio in the simulated BOLD, i.e. η. When varying η, we keep $\gamma = 0.1$ fixed (Panel B in Figs. 4–6), when varying γ, we keep $\eta = 0.05$ fixed (Panel C in Figs. 4–6). The length of each simulated time series is 400 data points, with a boxcar stimulus $s(t)$ of period length 16 to approximately match the properties of the real experimental data analyzed in Sect. 3.3. We calculated the tkCCA between the time series $x(t) = E(t, f)$ and $y(t) = F(p, t)$, i.e. we performed the time-delay embedding for the simulated electrophysiology while keeping the simulated BOLD unshifted. In all simulations reported in Sects. 3.2.3 to 3.2.5, the spatial BOLD field pattern $f(p)$ was recovered almost perfectly by the corresponding canonical variate w_y when $\eta > 0$. We therefore abstained from showing w_y in Figs. 4 to 6.

We designed the simulation data such that we can independently parametrize the coupling between the sources via $\beta(f)$ and $h(f)$ and the stimulus strength $\alpha(f)$. Since the

auto-correlations in each data source can distort the resulting variates and correlograms it is important to differentiate between correlations induced by the stimulus and correlations due to the actual coupling between the sources. We expect the proposed method to be primarily sensitive to the coupling parameters $\beta(f)$ and the temporal dynamics $h(f)$ and less susceptible to stimulus induced correlations, controlled by $\alpha(f)$. This ensures that the variates recovered by tkCCA reflect the coupling between the data sources rather than auto-correlations in either data source.

3.2.2 Comparison of tkCCA with alternative approaches

In Sects. 2.2 and 2.4 we argued from a theoretical perspective that alternative approaches like sequential ts-kCCA and multiway kCCA pose serious problems. To confirm these considerations, we compare the performance of the these methods and tkCCA on the simulated neuro-vascular coupling data described above.

In accordance with the expectations, panel A of Fig. 3 shows that tkCCA is faster than its competitors. In panels B and C, we analyze the performances on different data sets, each row corresponds to one of them. In the first row, the frequency dependence is confined to the coupling strength β, in the second row to the stimulus susceptibility α and in the third row only the hemodynamic response filter h depends on f. The other parameters are kept fixed with respect to the frequency. The colors (green, red, blue) indicate the different algorithms. The first three columns in panel C show the averaged[8] canonical variates (i.e. estimated filters) $w_x(\tau)$ and the rightmost column the *effective neuro-vascular coupling*

$$g_x(f, \tau) = \beta(f)h(f, \tau) \tag{21}$$

as defined by the data generating process. It is clearly visible that tkCCA (blue) captures the simulated neuro-vascular coupling better than ts-kCCA (green) or the multi-way kCCA
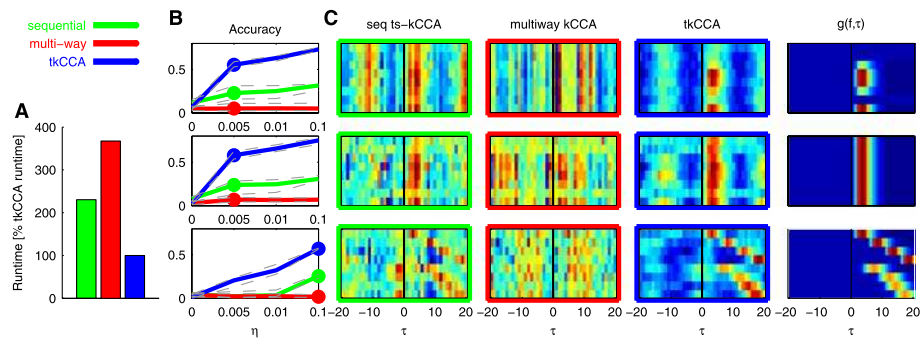

Fig. 3 Comparison of sequential time-shifted kCCA, multi-way kCCA and tkCCA; tkCCA is faster (*panel A*) and more accurate (*panel B*) than alternative approaches in all settings; **A**: runtime normalized to fastest method; **B**: Accuracy for different SNR's ($\eta = [0, 0.005, 0.01, 0.1]$) of correlational structure (stimulus SNR $\gamma = 0.1$); **C**: canonical variates $w_x(\tau)$, average over 10 simulations; the *dot* in the accuracy plots depicts the setting for the variates on the right; settings shown are frequency dependent coupling (*top row*, cf. Fig. 5), frequency dependent stimulus susceptibility (*middle row*, cf. Fig. 4) and frequency dependent temporal dynamics (*bottom row*, cf. Fig. 6)

---

[8]Prior to averaging we set $w_x(\tau) \leftarrow w_x(\tau) \operatorname{sign}(\langle g_x, w_x \rangle)$.

(red). To quantitatively compare the performances, we measure the accuracy as the absolute correlation coefficient between the simulated and the estimated filters:

$$|\mathrm{Corr}(g_x, w_x)| = \frac{|\langle g_x, w_x \rangle|}{\sqrt{\langle g_x, g_x \rangle}\sqrt{\langle w_x, w_x \rangle}}, \tag{22}$$

where the scalar product between two time-frequency filters is defined as $\langle a, b \rangle = \sum_{f,\tau} a(f, \tau) b(f, \tau)$. Panel B shows the accuracies on each data set for varying SNRs (i.e. different values of η) and constant overall stimulus susceptibility γ. Colored dots mark the η-values that correspond to the filters in panel C. In all settings, tkCCA clearly outperforms its competitors. With increasing η the accuracy of the estimated variates $w_x(\tau)$ becomes better both for the sequential time shifted kCCA[9] and tkCCA, for multi-way kCCA however, the accuracy stagnates. This is due to the low-rank approximations of the kernels computed on X_τ used in multi-way kCCA which are necessary in order to boil down the generalized eigenvalue problem to a tractable size (see also Bach and Jordan 2002). Since the first singular values are dominated by the auto-correlations in X, there is not much left in the kernels other than the stimulus induced auto-correlations after the incomplete Cholesky decomposition. In other words similar to a standard principal component analysis incomplete Cholesky decomposition just keeps the strongest spectral components for each time lag. These strongest components are not necessarily the ones that account for the cross-correlation which CCA tries to optimize. Another effect of the low-rank decomposition of K_X is the coarse discretization of the variates $w_x(\tau)$ when estimated using multi-way kCCA. When projecting the results back into input space (cf. (12)), the resulting $w_x(\tau)$ show striped patterns that do not change smoothly over time.

Overall the comparison shows that in terms of accuracy and computational cost, tkCCA clearly outperforms multi-way and sequential time shifted kCCA. Note that even though tkCCA is able to reproduce the simulated neuro-vascular coupling with a good accuracy, it estimates filters that deviate systematically from g. This is because of the periodic stimulus that manifests itself in the data. Any correlation at time delay τ will give rise to sidelobes in the variates at delay $\tau + n\theta$, where θ is the period length of the stimulus.

However, the strength of the sidelobes does not depend on the coupling encoded in β and h, but only on the stimulus susceptibility α. Consequently, the frequency dependence of β is only visible in the main maximum (panel C, upper row, tkCCA—blue box), not in the sidelobes. If the stimulus susceptibility α depends on f (panel C, middle row, tkCCA—blue box), it is the other way around: here, only the periodic sidelobes show the frequency dependence.[10] This allows to distinguish the true coupling between two sources from a spurious coupling due to auto-correlations. In the following sections we focus on tkCCA only and analyze these effects in more detail.

3.2.3 Influence of the stimulus susceptibility

In this section, we simulate a frequency independent neurovascular coupling by keeping both $\beta(f)$ and $h(f)$ fixed for all frequencies f. The only parameter that depends on the frequency is the stimulus susceptibility $\alpha(f)$ (see Fig. 4, Panel A). Such a frequency specific stimulus selectivity in cortical oscillations has been reported in e.g. (Belitski et al. 2008).

---

[9]We used the heuristic in (7) to repair the signs for sequential time-shifted kCCA solutions.

[10]The small frequency dependence of the main maximum is due to the superposition with one of the periodic maxima. After subtracting the periodic background, this dependence vanishes.
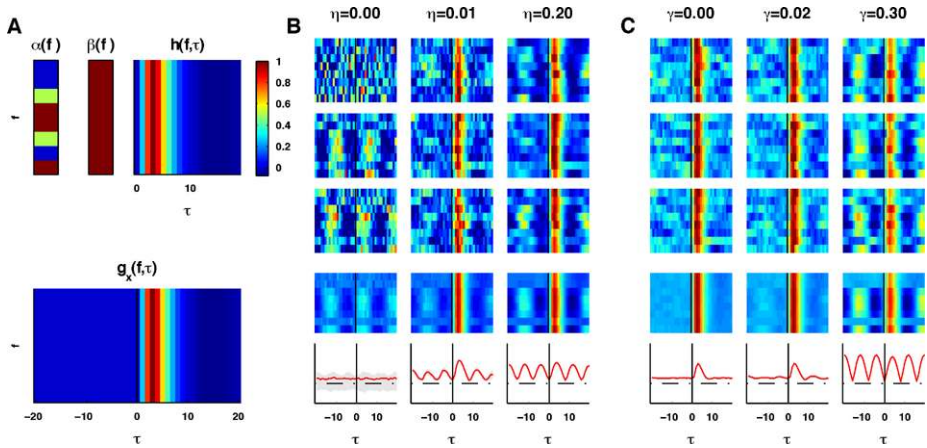
Fig. 4 A: Parameter settings for detection of stimulus driven dependencies; $\alpha(f)$ was frequency dependent, $\beta(f)$ and $h(f, \tau)$ were uniform across all values of f; $g_x(f, \tau)$ was computed as in (21); **B, C**: Time-dependent canonical variates $w_x(\tau)$ recovered by tkCCA on simulated data; single examples are shown in *rows 1–3*, a mean over 100 simulations in *row 4* and the corresponding correlograms (in *red*, 25*th* and 75*th* percentiles in *gray*) in *row 5*; **B**: Variates $w_x(\tau)$ for different values of η ($\gamma = 0.1$); at $\eta = 0$ (*no coupling*), stimulus driven auto-correlations dominate the variates (*Panel B, left column*); with increasing η oscillations appear in the band coefficients in bands with non-zero $\alpha(f)$; **C**: Variates $w_x(\tau)$ for different values of γ ($\eta = 0.05$); At $\gamma = 0$ (*no stimulus*), the correlation induced by spontaneous activity dominates the variates; with increasing γ, the stimulus induced dependencies are reflected as oscillations of filter weights in bands with non-zero $\alpha(f)$

The panels B and C of Fig. 4 show the variates, i.e. the estimated frequency-temporal filters $w_x(\tau)$ for three different repetitions of the simulation (first three rows) and the averaged filter over 100 repetitions (fourth row). The last row depicts the median of the canonical correlograms.[11] In all plots, $\tau = 0$ is indicated by a vertical line in the middle and the time lags were $-20 \leq \tau \leq 20$.

In panel B we see the evolution of the tkCCA results for changing BOLD noise levels $0 \leq \eta \leq 0.2$ while $\gamma = 0.1$ is kept constant. For $\eta = 0$ (leftmost column), the simulated BOLD is pure noise and does not contain any neurophysiological contribution. In this case the canonical correlogram is almost flat. The variates show weak oscillations at the stimulus frequency (16 samples). These oscillations reflect stimulus induced auto-correlations in the simulated neural data. It might seem odd that tkCCA finds a temporal structure in the variates, even when there is no correlational structure at all in the artificial BOLD data ($\eta = 0$). Due to the high dimensionality of the BOLD time series, it is very likely that tkCCA will find a spatial pattern, whose time course matches that of some dimensions of the other data source. Since the temporal structure is similar in our simulations (all simulated neural data shared the same stimulus), this overfitting to auto-correlations can show up in the average variates. But as the phase of the oscillation might vary, the effect on the average variates is not as strong as on single estimates. Importantly, the average canonical correlogram does not exhibit these overfitting effects (cf. Fig. 4 to 6, panel B, bottom left). However, as in this setting the stimulus was only picked up by certain frequency bands (in contrast to the results in Figs. 5 and 6), the oscillations in the variates appear only for nonzero values of $\alpha(f)$.

---

[11]The area between the 25%- and 75%-percentiles is shaded in gray, but the correlogram is so stable for different realizations of the data that this is not visible in most plots.
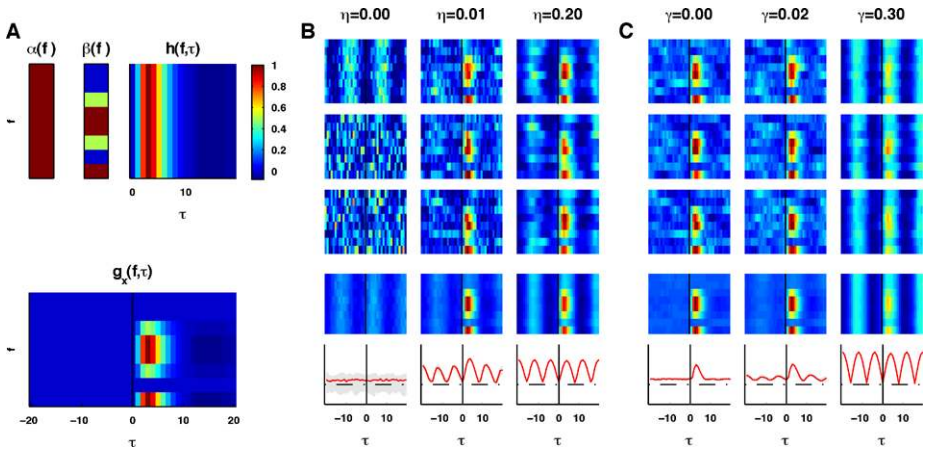
Fig. 5 **A**: Parameter settings for detection of frequency specific coupling; $\beta(f)$ was frequency dependent while $\alpha(f)$ and $h(f)$ were uniform across all values of f; **B, C**: Time-dependent canonical variates $w_x(\tau)$ recovered by tkCCA on simulated data; single examples are shown in *rows 1–3*, a mean over 100 simulations in *row 4* and the corresponding correlograms (in *red*, 25th and 75th percentiles in gray) in *row 5*; **B**: Variates $w_x(\tau)$ for η ($\gamma = 0.1$); if there is no coupling ($\eta = 0$), some variates reflect stimulus driven auto-correlations; for $\eta > 0$, the correct coupling coefficients $\beta(f)$ are found at the right time lags; **C**: Variates $w_x(\tau)$ for γ ($\eta = 0.05$); for all values of γ tkCCA finds the right $\beta(f)$ at the right time lag; at high values of γ, stimulus induced dependencies appear in the correlogram but $w_x(\tau)$ is still reveals the right $\beta(f)$

Already for very small η values, the coupling dynamics are well reproduced, which means that tkCCA is very robust even if the signal of interest has a variance that is much smaller than the noise variance.

In panel C, the stimulus gain γ is varied between 0 and 0.3, while the BOLD SNR is kept fixed at $\eta = 0.05$. For all γ settings, the estimated filters $w_x(\tau)$ capture the true coupling structure, as reflected by the pronounced maximum at $\tau = 4$. As the stimulus dependency of the neural time series increases with γ, sidelobes appear that reflect the periodicity of the visual stimulus. As expected, they are only visible in frequency bands where $\alpha(f) \neq 0$. While the stimulus-induced oscillations in the variates have the same frequency as the stimulus, the canonical correlogram contains oscillations at double that frequency since it shows only the *absolute* of the correlation.

3.2.4 Influence of the coupling strengths

In this section we demonstrate how tkCCA reacts to frequency-dependent strength of the neurovascular coupling. To this end, we choose the stimulus susceptibility $\alpha(f)$ and hemodynamic response $h(f, \tau)$ such that they are fixed for all frequencies f and allow only the parameter $\beta(f)$ to depend on f (see Fig. 5, Panel A).

Panels B and C in Fig. 5 again show variations of this experiment for different values of η (BOLD signal to noise level) and γ (stimulus susceptibility of the neural activity). Like in Sect. 3.2.3, the respective other parameter was kept fixed at $\gamma = 0.1$ and $\eta = 0.05$. From Panel B we see that as soon as the simulated BOLD contains small contributions from the electrophysiological signal (i.e. $\eta \geq 0.01$), tkCCA is able to recover the neurovascular coupling modeled by the coupling coefficients $\beta(f)$ and the time lag given by the coupling dynamics $h(f, \tau)$. The influence of the coupling coefficients $\beta(f)$ is clearly visible not only in the average canonical variates (fourth row), but even in each single trial (rows 1–3). Panel
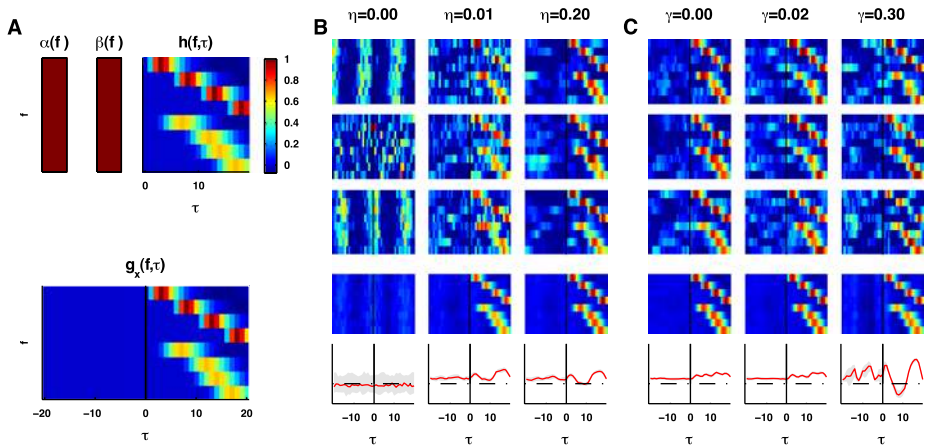
Fig. 6 A: Parameter settings for detection of frequency specific temporal coupling dynamics; $\alpha(f)$ and $\beta(f)$ was uniform, whereas $h(f, \tau)$ was frequency dependent; we chose four different delay parameters and two different widths for eight frequency specific hemodynamic response functions; **B**: Time-dependent canonical variates for different values of η ($\gamma = 0.1$); coupling parameters at the time lags specified in $h(f, \tau)$ are reliably captured at very low η; with increasing η, the estimated coupling dynamics become more pronounced; **C**: Time-dependent canonical variates for different values of γ ($\eta = 0.05$); at $\gamma = 0$ (*no stimulus*), coupling coefficients are recovered reliably; the correlograms exhibits oscillations at the stimulus frequency (16 samples) for high values of γ, but the interpretation of the correlogram is non-trivial, as it reflects many different temporal dynamics (i.e. the stimulus and different frequency band dynamics)

C shows that with increasing influence of the periodical stimulus, the canonical variates display increasing sidelobes. In contrast to Fig. 4, however, the sidelobes and the overall periodicity now extends to the whole range of frequencies due to the f-uniform stimulus susceptibility $\alpha(f)$. This allows to clearly separate the effects of $\alpha(f)$ and $\beta(f)$. As already seen in the last section, the canonical correlograms (panels B and C, last row) capture well the temporal correlation structure induced by spontaneous activity (main peak) and the periodical stimulus. As pointed out earlier in Sects. 3.2.2 and 3.2.4, when $\eta = 0$ (i.e. there is *no coupling*) the canonical variates are dominated by stimulus induced auto-correlations. As soon as there is any coupling between the neural and BOLD data ($\eta > 0$) however, tkCCA clearly picks the correct coupling coefficients.

3.2.5 Influence of temporal dynamics in the neurovascular coupling

In a third set of experiments, we keep both $\alpha(f)$ and $\beta(f)$ constant in f. The only frequency dependence is in the width and delay of the hemodynamic response function $h(f, \tau)$ (see Fig. 6, panel A). We chose four different delay parameters and two different widths to obtain a different hemodynamic response for each of the eight frequency bands. Of course this setting represents an extreme case and these differences in coupling dynamics are not to be expected from real data. Similar to the simulations presented before, panels B and C show realizations of this setting for different values of η and γ. Again, panel B shows that the correct temporal coupling dynamics is found at levels of η as small as 0.01, while for $\eta = 0$, the only temporal structure in the variates is due to auto-correlations in the simulated neural data. Since $h(f, \tau)$ mixes the stimulus with different delays into the simulated BOLD, the latter does not exhibit a clear periodicity any more. Consequently, the estimated variates/filter $w_x(\tau)$ do not show the sidelobes we know from the previous settings. While

the canonical correlogram does exhibit oscillations at the stimulus frequency for $\tau > 0$, the temporal structure in the correlogram is not as easy to interpret as in the previous examples. This is due to the superposition of different effects in the correlogram: The temporal dynamics of the stimulus and those of the coupling as defined in $h(f, \tau)$.

3.3 Application to experimental fMRI/electrophysiological recordings

We applied tkCCA to real data acquired by simultaneous recordings of neurophysiological activity and BOLD signal in response to visual stimulation. The data was recorded in the primary visual cortex of three male rhesus monkeys (Macaca mulatta). All surgical operations and anaesthesia procedures have been described in detail elsewhere (Logothetis et al. 1999, 2002). The experiments were approved by the local authorities (Regierungspräsidium, Tübingen, Germany) and were in full compliance with the guidelines of the European Community for the care and use of laboratory animals (EUVD 86/609/EEC).

3.3.1 Experimental design and data acquisition

Visual stimuli were presented binocularly using a gamma-corrected SVGA fiber-optic system (AVOTEC, Silent Vision, USA). The stimulation protocol consisted of 32 second blocks of visual stimulation using a rotating polar checkerboard stimulus at a size of $10° \times 10°$ of visual angle followed by a 32 second blank period of isoluminant gray. Checkerboard rotation direction was reversed every 8 s to minimize adaptation. The duration of the entire scan was 39.5 minutes corresponding to 37 blocks of visual stimulation.

The amplifiers for the electrophysiological recordings were custom made and have an analog compensation mechanism for the noise induced by gradient-switching of the scanner. The details of interference compensation and signal conditioning are given elsewhere (see Logothetis et al. 2001; Öltermann et al. 2007 and Goense and Logothetis 2008). Data have been recorded at a temporal resolution of 20 kHz. After performing additional software based artifact removal we downsampled the data to 7 kHz and performed spectral analysis on the time series.[12] Frequency bands cutoffs (taken from Goense and Logothetis 2008) were 0.1–8 Hz, 8–12 Hz, 12–24 Hz, 24–40 Hz, 40–60 Hz, 60–120 Hz, 120–1000 Hz and 1–3 kHz. Temporal bin size was set to 1 second. The spectral power in each band was centered and divided by its respective standard deviation.

fMRI data consisted of a timeseries of 3D images (256-by-128-by-7 voxels, each having a spatial extent of 0.25-by-0.25-by-1 mm), acquired at a temporal resolution of 0.25 Hz. Each session comprised 592 consecutive images of fMRI data (39.5 minutes of recording). The first preprocessing step was a spatial coregistration to the first image recorded using code from the fMRI analysis package *SPM5*, (SPM5 2005)). After manually segmenting the image for brain tissue (more specifically primary and secondary visual cortex), each fMRI image was spatially smoothed with a three-dimensional Parzen (or de la Vallé-Poussin) window (radius 1 mm or 4-by-4-by-1 voxels) as defined in (Harris 0978). We then chose a spherical region of interest (15 mm radius) around the recording electrode to restrict the number of voxels to approximately 4000. This restriction has been made merely for the sake of interpretability of the results. In principle kCCA could deal with a lot more voxels,

---

[12]Separation of electrophysiological activity into frequency sub-bands allows for distinct interpretations of neural activity; for instance, the lower parts of the spectrum (< 100 Hz) can be interpreted as the summed input to a population of neurons, whereas the high frequency content (> 1 kHz) can be interpreted as the (spiking) output from this population to other brain regions (Logothetis 2008).

but in order to relate the BOLD signal to the (spatially confined) neural activity recorded at the electrode it is reasonable to restrict the spatial extent of analysis to a physiologically plausible region. As a last preprocessing step, we removed linear trends from the fMRI data and centered the timecourses in each voxel.

3.3.2 Results for the experimental data

We applied the same analysis methods as in Sect. 3.2 to the data acquired in the experiments described above. Examples for canonical variates of fMRI, temporal canonical variates of electrophysiological data and the canonical correlograms are shown in Fig. 7. Panel A shows single experiment examples. Single experiment here refers to one 39.5 minute recording session. The top row are two representative fMRI variates w_y computed by tkCCA, a yellow arrow indicates the position of the recording site. Below are the corresponding time-dependent variates $w_x(\tau)$ for the electrophysiological data. Panel B shows the canonical correlograms in the top panel (plotted are median and 25th and 75th percentile, $n = 10$) and in the bottom figure the mean canonical variate for the electrophysiological data. Note that the temporal resolution of the canonical correlograms and variates $w_x(\tau)$ is 1 Hz, whereas the temporal resolution of the fMRI data is only 0.25 Hz. The two single experiment examples show that the canonical variate for the fMRI data source has distinct patches of high coefficients in primary visual cortex close to the recording site of the electrophysiological data. High values in the canonical variate of the fMRI data (Fig. 7, panel A, top row) indicate correlated activity in the neural tissue at those locations. In contrast to the simulation data, there is no ground truth available, as CCA is an unsupervised method. But considering that tkCCA yielded robust estimates of the spatial correlational structure at very high noise levels in simulations (at $\eta > 0.01$, the ground truth of the fMRI variate was always recovered by tkCCA, results not shown), we can assume that the spatial patterns revealed by tkCCA do reflect the correlational structure in primary visual cortex. This correlational structure can arise for instance from intra-cortical connections or common preferences for specific features in the visual
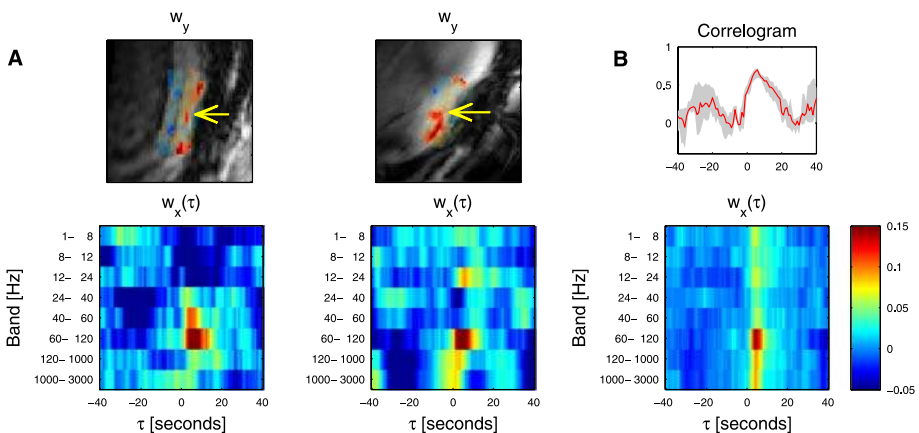

Fig. 7 A, top: single experiment examples of fMRI variates w_y, high positive values (*red*) indicate functional connectivity between the respective sites; position of the recording electrode is indicated by *yellow arrows*; **A, bottom**: time-dependent variates $w_x(\tau)$ for the electrophysiological data, corresponding to the variates above; **B, top**: canonical correlograms; median is plotted in *red*, in *gray* the 25th and 75th percentile ($n = 10$); note the pronounced peak at 5 ± 1.7 seconds (mean $\pm$ standard deviation) and the maximum at half the stimulus period (32 seconds); **B, bottom**: average canonical variate $w_x(\tau)$ for electrophysiological data ($n = 10$)

stimulus. Figure 7, panel A (bottom row) shows the time-dependent variates $w_x(\tau)$ for the electrophysiological data corresponding to the fMRI variates above. The coefficients of the filters indicate that frequency content in the band from 60 Hz to 120 Hz at a time lag of 5 seconds contributes most to the neurovascular coupling. In some experiments (cf. Fig. 7, A, bottom right) there is a weak oscillation in the filters with a period of approximately 64 seconds, the length of one stimulation period. In other experiments (cf. Fig. 7, panel A, bottom left) the time-dependent variate seems to be mainly capturing dependencies independent of the stimulus (there are only weak oscillations in the variate at the stimulus frequency).

It is important to note that the spatial filters for the BOLD data w_y have to be analyzed on a single experiment basis as the exact position of the head, the radio-frequency coils in the scanner and also the electrode position vary across sessions and subjects. Thus averaging across subjects is not possible, due to the lack of spatial normalization routines.[13] Electrophysiological variates and the canonical correlogram do not suffer from this problem. Examples for grand averages in Fig. 7 panel B, confirm the results from single experiments: It is mainly bandpower in the range between 60 Hz and 120 Hz at a time lag of 5 ± 1.7 seconds (mean $\pm$ one standard deviation, $n = 10$) that dominates the neurovascular coupling. At time lags close to the maximal canonical correlation, the percentiles around the median correlogram indicate a very low variance across experiments. However, with increasing time lags the variance also increases. This is a consequence of the fact that some recordings have a higher signal-to-noise ratio (with respect to the visual stimulus) than others. The activity in the fMRI signal and the electrophysiology does not reflect the visual stimulation equally well in all experiments and thus the stimulus induced sidelobes in the correlogram are less pronounced in some trials. The timepoint of the maximal hemodynamic response is in line with well established models of the HRF (Friston et al. 2000; Buxton et al. 2004; Logothetis and Wandell 2004), the correlation between the neural activity and BOLD response is maximized when the electrophysiological data is shifted by approximately 5 seconds relative to the fMRI time series. In addition to the confirmation of HRF models, the multivariate filters allow to draw conclusions about the correlational structure in the brain and the temporal dynamics of the coupling.

Although we shifted in time only the electrophysiology data, it is equivalent for the correlogram to shift in time the BOLD data. Note however that the number of dimensions in the fMRI data source is significantly larger which renders regularization more critical. Due to the lower temporal resolution of the raw fMRI data, the canonical variates and correlogram consequently also have a temporal resolution of only 0.25 Hz (instead of 1 Hz in the case of the embedding of the faster data source). Leaving aside these drawbacks, the embedding of the spatially resolved data source allows to recover spatio-temporal maps which can be visualized as movies showing how the correlational structure in the cortex changes in response to stimulation (see Fig. 8).

4 Conclusion

We presented a generic framework for analyzing data from multiple sources with non-instantaneous dependencies. The algorithm is based on kCCA and thus applicable to high dimensional data such as fMRI images. By embedding one data source into its temporal

---

[13]In human fMRI experiments, the single subject data is often mapped on a *template brain* in order to compare results across subjects; this procedure requires extensive spatial smoothing and a template brain; these standards have not yet been established for non-human primate brain scans.
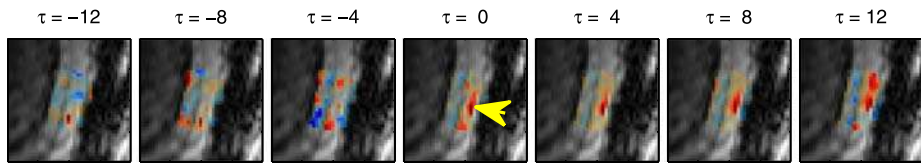
$\tau = -12$	$\tau = -8$	$\tau = -4$	$\tau = 0$	$\tau = 4$	$\tau = 8$	$\tau = 12$

Fig. 8 Spatio-temporal variates obtained using tkCCA and temporal embedding of the fMRI data; data was the same as in Fig. 7, panel A, *top left*; a *yellow arrow* indicates the position of the recording electrode; temporal resolution was 0.25 Hz, τ is increased in 4 second steps; as expected from the results in Fig. 7, the maximum of the correlogram was at $\tau = 4$; for $\tau > 0$, the maps exhibit a pronounced positive peak in the coefficients around the recording site

context before computing the kernel, tkCCA can find critical coupling features and temporal dynamics of the dependencies between the data sources. In tkCCA, one canonical variate has a temporal dimension that can be visualized. It can be interpreted as the contribution of each feature at each time lag to the canonical correlation. Instead of a canonical correlation coefficient, tkCCA computes a canonical correlogram that reflects the coupling dynamics between the two data sources. As in classical kCCA, the solutions can be efficiently formulated as a generalized eigenvalue problem. Importantly, tkCCA recovers correlograms which are coherent over time with respect to both, the sign and scale of correlations.

Based on theoretical considerations and results from extensive simulations we showed that tkCCA is able to reliably capture the coupling dynamics between two data sources even in the presence of strong auto-correlations in either data source. Applying tkCCA to brain activity measurements, simultaneously recorded direct invasive measurements and indirect non-invasive measurements, results in estimates of the canonical correlogram that are highly reproducible across sessions and subjects. The estimated time point of maximal neurovascular correlation confirms well established models of the HRF (Friston et al. 2000; Buxton et al. 2004; Logothetis and Wandell 2004). The filters computed by tkCCA are novel in that they allow to draw conclusions about the features in the data that give rise to these neurovascular coupling phenomena. This can give valuable insights in functional connectivity in or between cortical regions. Up to date, most fMRI connectivity studies have to rely entirely on unsupervised methods that use only fMRI data or supervised methods that correlate BOLD data with a stimulus time series (including some model assumptions about the HRF). The proposed method in combination with the simultaneous recordings offers for the first time a purely data-driven approach for the multivariate analysis of neurovascular coupling mechanisms.

Future directions of research could try to explore multi-dimensional correlograms or extensions to more than two data sources. Moreover, one could think of more sophisticated regularization techniques as proposed for instance in (Blaschko et al. 2008) or extending the tkCCA framework to tensor data. Other potential applications of temporal kCCA could be experiments including pharmacological interventions such as (Rauch et al. 2008a, 2008b) that interfere with neurovascular coupling mechanisms. If certain pharmacological agents have an effect on those mechanisms, the canonical correlations should reflect this change. Insights into how different experimental manipulations influence canonical correlations and variates might help to understand better neurovascular coupling mechanisms. Note that tkCCA is readily applicable also in other fields of science beyond brain imaging. A MATLAB implementation of tkCCA as presented in this work can be downloaded from http://user.cs.tu-berlin.de/~fbiessma/code.html.

References

Akaho, S. (2001). A kernel method for canonical correlation analysis. In *Proceedings of the international meeting of the psychometric society (IMPS)*. Berlin: Springer.

Anderson, T. W. (1958). *An introduction to multivariate statistical analysis*. New York: Wiley.

Bach, F. R., & Jordan, M. I. (2002). Kernel independent component analysis. *Journal of Machine Learning Research*.

Bach, F. R., & Jordan, M. I. (2006). *A probabilistic interpretation of canonical correlation analysis*. Technical Report.

Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M. A., Logothetis, N. K., & Panzeri, S. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience*.

Blaschko, M. B., Lampert, C. H., & Gretton, A. (2008). Semi-supervised Laplacian regularization of kernel canonical correlation analysis. In W. Daelemans, B. Goethals, & K. Morik (Eds.), *19th European conference on machine learning*. Antwerpen: Springer.

Buxton, R. B., Uludag, K., Dubowitz, D. J., & Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *Neuroimage*.

Friman, O., Borga, M., Lundberg, P., & Knutsson, H. (2002). Exploratory fMRI analysis by autocorrelation maximization. *Neuroimage*.

Friston, K. J., Mechelli, A., Turner, R., & Price, C. J. (2000). Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. *Neuroimage*.

Fukumizu, K., Bach, F. R., & Gretton, A. (2007). Statistical consistency of kernel CCA. *Journal of Machine Learning Research*.

Goense, J. B. M., & Logothetis, N. K. (2008). Neurophysiology of the BOLD fMRI signal in awake monkeys. *Current Biology*.

Hardoon, D. R., Szedmak, S., & Shawe-Taylor, J. (2004). Canonical correlation analysis: an overview with application to learning methods. *Neural Computation*.

Hardoon, D. R., Mourao-Miranda, J., Brammer, M., & Shawe-Taylor, J. (2007). Unsupervised analysis of fMRI data using kernel canonical correlation. *Neuroimage*.

Harris, F. J. (1978). On then use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*.

Haynes, J. D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Current Biology*.

Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*.

Langleben, D. D., Loughead, J. W., Bilker, W. B., Ruparel, K., Childress, A. R., Busch, S. I., & Gur, R. C. (2005). Telling truth from lie in individual subjects with fast event-related fMRI. *Human Brain Mapping*.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*.

Logothetis, N. K., & Wandell, B. A. (2004). Interpreting the bold signal. *Annual Reviews of Physiology*.

Logothetis, N. K., Guggenberger, H., Peled, S., & Pauls, J. (1999). Functional imaging of the monkey brain. *Nature Neuroscience*.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Öltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*.

Logothetis, N. K., Merkle, H., Augath, M., Trinath, T., & Ugurbil, K. (2002). Ultra high-resolution fMRI in monkeys with implanted RF coils. *Neuron*.

Macke, J. H., Zeck, G., & Bethge, M. (2008). Receptive fields without spike-triggering. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *21th neural information processing systems conference*. Cambridge: MIT Press.

Müller, K. R., Mika, S., Rätsch, G., Tsuda, K., & Schölkopf, B. (2001). An introduction to kernel-based learning algorithms. *IEEE Transactions on Neural Networks*.

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*.

Öltermann, A., Augath, M. A., & Logothetis, N. K. (2007). Simultaneous recording of neuronal signals and functional NMR imaging. *Magnetic Resonance Imaging*.

Ogawa, S., Lee, T. M., Nayak, A. S., & Glynn, P. (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magnetic Resonance in Medicine*.

Kettenring, J. R. (1971). Canonical analysis of several sets of variables. *Biometrika*.

Rauch, A., Rainer, G., Augath, M., Öltermann, A., & Logothetis, N. K. (2008a). Pharmacological MRI combined with electrophysiology in non-human primates: effects of lidocaine on primary visual cortex. *Neuroimage*.

Rauch, A., Rainer, G., & Logothetis, N. K. (2008b). The effect of a serotonin-induced dissociation between spiking and perisynaptic activity on bold functional MRI. *Proceedings of the National Academy of Sciences*.

Schölkopf, B., & Smola, A. J. (2001). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. *Adaptive computation and machine learning*. Cambridge: MIT Press.

Shawe-Taylor, J., & Cristianini, N. (2004). *Kernel methods for pattern analysis*. Cambridge: Cambridge University Press.

SPM5. (2005). *Statistical parametrical mapping toolbox*. http://www.fil.ion.ucl.ac.uk/spm/.