Modeling state-related fMRI activity using change-point theory

Martin A. Lindquist[1*], Christian Waugh[2] and Tor D. Wager[3]

1. Department of Statistics, Columbia University, New York, NY, 10027

2. Department of Psychology, University of Michigan, Ann Arbor, MI, 48109

3. Department of Psychology, Columbia University, New York, NY, 10027

ADDRESS:

* Corresponding author:

Martin Lindquist
Department of Statistics
1255 Amsterdam Ave, 10th Floor, MC 4409
New York, NY 10027
Phone: (212) 851-2148
Fax: (212) 851-2164
E-Mail: martin@stat.columbia.edu

Abstract

The general linear model (GLM) approach has arguably become the dominant way to analyze functional magnetic resonance imaging (fMRI) data. It tests whether activity in a brain region is systematically related to some known input function. However, it becomes impractical when the precise timing and duration of psychological events cannot be specified *a priori*. In this work, we introduce a new analysis approach that allows the predicted signal to depend non-linearly on the transition time. It uses ideas from statistical control theory and change-point theory to model slowly varying processes with uncertain onset times and durations of underlying psychological activity. Our approach is exploratory in nature, while retaining the inferential nature of the more rigid modeling approach. It is a multi-subject extension of the exponentially weighted moving average (EWMA) method used in change-point analysis. We extend existing EWMA models for individual subjects (a single time series) so that they are applicable to fMRI data, and develop a group analysis using a hierarchical model, which we term HEWMA (Hierarchical EWMA). The HEWMA method can be used to analyze fMRI data voxel-wise throughout the brain, data from regions of interest, or temporal components extracted using ICA or similar methods. We validate the false-positive rate control of the method and provide power estimates using simulations based on real fMRI data. We further apply this method to an fMRI study (n = 24) of state anxiety. A toolbox implementing all functions in Matlab is freely available from the authors.

Keywords: fMRI, Change-point, EWMA

Introduction

The voxel-wise general linear model (GLM) approach has arguably become the dominant way to analyze functional magnetic resonance imaging (fMRI) data. This model is well-suited for testing whether variability in a voxel's time course can be explained by a set of *a priori* defined regressors that model predicted responses to psychological events of interest. The GLM approach has proven particularly powerful for dealing with event-related designs, as a sequence of sparse events occurring at random intervals affords a relatively specific predicted response, and a good fit to the data is often interpreted in terms of signal evoked by a particular psychological event (Wager, Hernandez, Jonides, & Lindquist, in press; Worsley & Friston, 1995). Mixed block/event-related designs have been used to investigate state-effects in fMRI (Visscher, Miezin, Kelly, Buckner, Donaldson, McAvoy, Bhalodia, & Petersen, 2003), but inferences on states are subject to the limitations in interpretability of block designs, and only states that are under relatively precise experimental control (i.e., they can be turned on and off repeatedly by experimental manipulations) can be studied. New kinds of statistical models are needed to capture state-related changes in activity, particularly when psychological states have uncertain onset times, temporal intensity profiles, and durations.

As an example, consider an fMRI study planned to assess the effects in the brain of a new short-acting drug (Fig. 1A)—an event that is not easily repeated many times in a block design. Researchers might be interested in the kinetics of the drug in different brain areas. Some areas might respond to the initial administration, and others might show sustained responses consistent with plasma concentrations. Yet other areas may taper off or ramp up over the period of administration, as the brain coordinates an orchestrated response to different components of the experience. State-related activity is critical in many other domains, including drug pharmacology

(Breiter et al., 1997; Wise, Williams, & Tracey, 2004), memory (Donaldson, Petersen, Ollinger, & Buckner, 2001; Otten, Henson, & Rugg, 2002), motivation and emotion (Gray, Braver, & Raichle, 2002). For example, a participant watching a funny film might experience amusement that builds gradually over time. Without moment-by-moment reports, which may alter the experience (Taylor, Phan, Decker, & Liberzon, 2003), the progression of amusement over time will be difficult to specify. This is true in general for emotional states, which are difficult to turn on and off rapidly (required for blocked or ER designs) and may last longer than or not as long as the putative eliciting event.

The GLM's effectiveness in modeling such state-related activity is limited in several ways, and developing new approaches could lead to fruitful alternatives. One issue concerns the match between the desired inference and the inference provided by the model. For example, researchers may be interested in making inferences on the onset and duration of brain activity in a number of research settings. Suppose a GLM is specified that incorporates a sustained (10-s long) response to a stimulus, as is common with epoch-related designs. The GLM provides inferences on the magnitude of the hemodynamic responses. These inferences will be appropriate for testing hypotheses regarding the magnitude, but not for hypotheses concerning the duration of activation. For example, suppose that there is only a brief hemodyamic response to the stimulus. The regressor will fit partially, and may well reach statistical significance as the GLM tests whether the magnitude of the regressor is non-zero. However, it is clear that inferring a sustained response from this fit would be inappropriate. Thus, given a hypothesis about the duration of activity, it would clearly be more informative to perform inferences directly on the duration of activation. The GLM does not provide for this type of inference.

Two other related issues with the GLM involve the potential for mis-modeling (Neter,

Kutner, Nachtsheim, & Wasserman, 1996), which can reduce sensitivity and lead to incorrect inferences, and limitations in reproducibility (Liou, Su, Lee, Aston, Tsai, & Cheng, 2006; Genovese, Noll, & Eddy, 1997) for some paradigms. First, in the example above, even a partial fit to the regressor may lead to a significant result. Therefore, model significance alone is not enough to imply that the specified model is a better choice than other plausible alternatives. More flexible models reduce the risk of mis-modeling and increase sensitivity when the precise shape of the response cannot be accurately specified *a priori*, and this is an advantage of the flexible approach to modeling state-related activity we propose in this paper. Second, suppose that a brain region reproducibly responds to a particular emotion (e.g., anxiety), though the onset and time course varies across participants and studies depending on the particular task demands and individual propensities. In this case, there is no GLM model that will give reproducible activation in the region of interest. However, a more flexible model may capture consistencies in activation magnitude while allowing for variations in timing.

For each of the situations outlined above, model flexibility is crucial. Models that permit more data-driven estimates of activation, and inferences on when and for how long activation occur, may be advantageous for discovering state-related activations whose timing can only be specified loosely *a priori*. Rather than treating psychological activity as a zero-error fixed effect specified by the analyst (as in GLM analysis) and testing for brain changes that fit the specified model, data-driven approaches attempt to characterize reliable patterns in the data, and relate those patterns to psychological activity *post hoc*. One particularly popular data-driven approach in the fMRI community is independent components analysis (ICA), a variant on a family of analyses that also includes principal components and factor analysis (Beckmann & Smith, 2004; Calhoun, Adali, & Pekar, 2004; McKeown & Sejnowski, 1998). Recent extensions of these

methods can identify brain activity patterns (components) that are reliable across participants, treating participant as a random effect (Beckmann & Smith, 2005; Calhoun et al., 2004), and can identify state-related changes in activity that can subsequently be related to psychological processes. However, these methods do not provide statistics for inferences about whether a component varies over time and when changes occur in the time series. In addition, because they do not contain any model information, they capture regularities whatever the source; thus, they are highly susceptible to noise, and components can be dominated by artifacts.

Thus, the capacity for statistical inference is a strength of the GLM, whereas the ability to discover activation patterns that are related to a relatively unconstrained psychological model (i.e., the approximate onset and offset of a mental state under loose experimental control) is a strength of data-driven methods. In this paper, we present a model-driven approach for identifying changes in fMRI time series in individual and group data that allows for valid population inference. What sets this work apart from the GLM approach is that the predicted signal depends non-linearly on the model parameter (transition time). In essence, the suggested approach can be seen as an extension of the GLM framework to allow for unknown transition times. The approach uses ideas from statistical control theory and change-point theory to model slowly varying processes with uncertain onset times and durations of underlying psychological activity. Thus, like data-driven methods, the model uses the data to come up with estimates of whether, when, and for how long activation occurred with only minimal specification of *a priori* constraints. Concretely, in our model, one need only specify the length of a "no activation" baseline period and, for some applications, that a brain region's activity level alternates between two states (e.g., stimulated and rest). The estimates can then be compared with psychological or physiological parameters of interest—such as subjective ratings, behavior, or physiological

responses—to constrain interpretation. Thus, a benefit of the methods are that they are semi-model-free methods of detecting activation and are therefore insensitive to variations in the phase lag and shape of the hemodynamic response across the brain. Though the methods presented here share the attractive features of data-driven analysis methods, it is still in its core a model-driven approach and it retains the inferential nature of the more rigid modeling approach: population statistical inferences are made on activation patterns.

The change-point analysis that we develop is a multi-subject extension of the exponentially weighted moving average (EWMA) change-point analysis (also called statistical quality control charts (Neubauer, 1997; Roberts, 1959; Shehab & Schlegel, 2000)). Activity during a baseline period is used to estimate noise characteristics in the fMRI signal response. This activity is used to make inferences on *whether, when*, and *for how long* subsequent activity deviates from the baseline level. We extend existing EWMA models for individual subjects (a single time series) to include AR(p) and ARMA(1,1) noise processes, then develop a group 'random effects' analysis using a hierarchical model, which we term HEWMA (Hierarchical EWMA).

The HEWMA method may be used to analyze fMRI data voxel-wise throughout the brain, data from regions of interest, or temporal components extracted using ICA or similar methods. Here, we provide power and false-positive rate analyses based on simulations, and we apply HEWMA to voxel-wise analysis of an anxiety-producing speech preparation task. We demonstrate how the method detects deviations from a pre-task-instruction baseline and can be used to characterize differences between groups of individuals in both evoked fMRI activity and changes in states.

The method also provides inferences on the number and timing of changes in the state of

activity, denoted change-points (CPs). Knowledge of the CPs may provide a basis for discriminating anticipatory activity from responses to a challenge (e.g., activity that begins in anticipation of pain from that elicited by painful stimuli; (Koyama, McHaffie, Laurienti, & Coghill, 2005; Wager, 2005)). CP maps may also be used to identify brain regions that become active at different times during a challenge (e.g., the early, mid, or late phases of a tonic painful stimulus). These maps may provide meaningful characterizations of differences among individuals: For example, the onset time of brain responses to anxiety may provide clinically relevant markers of anxiety disorders. In studies of emotion, the speed of recovery from adverse events is thought to be an important predictor of emotional resilience (Fredrickson, Tugade, Waugh, & Larkin, 2003; Tugade & Fredrickson, 2004), and CPs could provide direct brain measures of recovery time. In cognitive psychology, brain CPs in problem solving and insight tasks may provide a direct neural correlate of traditional time-to-solution measures in cognitive studies (Cheng & Holyoak, 1985; Christoff et al., 2001). Inferences on the duration of activation could have similar advantages, including tests on the duration of pharmacological or emotion-related activity and the duration of cognitive task-set related activity in cognitive state-item designs.

Methods

There are a large variety of change-point detection problems that present themselves in the analysis of time series and dynamical systems. In this section, we first develop a method for detecting changes in activation patterns for a single time series using exponentially weighted moving averages (EWMA), and then develop a hierarchical extension (HEWMA) appropriate for

multisubject fMRI studies. Further, we introduce methods for estimating the exact timing and duration of the detected change.

I. The exponentially weighted moving-average (EWMA) model

Given a process that produces a sequence of observations $\vec{x} = (x_1, x_2, \ldots x_n)^T$ (e.g., an fMRI time series), we first consider a two-state model where the data is modeled as the combination of two normal distributions, one with mean θ_0 and covariance matrix Σ, and the second with mean θ_1 and the same covariance Σ. During a baseline acquisition period, the process generates a distribution of data with mean θ_0, and while in this state, the process is considered to be *in-control*. The observations follow this distribution up to some unknown time τ, the change-point, when the process changes (i.e., a new psychological state results in increased or decreased neural activity), resulting in the generation of fMRI observations from the second distribution with mean θ_1 (see Fig. 1A). While in this second state, the process is deemed to be *out-of-control*, or in the out-of-control (OOC) state. The statistical model for this framework can be written as follows:

$$x_t = s_t + \varepsilon_t \text{ for } t = 1, \ldots n \tag{1}$$

where s_t denotes the signal and ε_t the noise at time t. The signal is specified by

$$s_t = \begin{cases} \theta_0 & \text{for } t = 1, \ldots \tau \\ \theta_1 & \text{for } t = \tau + 1, \ldots n \end{cases} \tag{2}$$

and $\vec{\varepsilon} = (\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n)^T$ follows a mean-zero normal distribution with covariance matrix Σ, i.e.

$$\vec{\varepsilon} \sim N(0, \Sigma). \tag{3}$$

9

The diagonals of the n x n matrix Σ are the noise variance estimates for each observation, and the off-diagonals are the covariance among observations induced by autocorrelation in the fMRI time series. At a later stage we will relax the constraint of a single change-point and allow s_t to move between the two states.

The problem of detecting and estimating change-points has been studied extensively in the change-point and statistical quality control literature. In statistical control the exponentially weighted moving average (EWMA) control chart has become a popular and flexible approach for detecting deviations from some baseline mean. It is based on the EWMA statistic, z_t. The statistic is a temporally smoothed version of the data, and is defined as follows:

$$z_t = \lambda x_t + (1 - \lambda) z_{t-1} \qquad\qquad \text{for } t = 1, \dots n \qquad\qquad (4)$$

or, equivalently,

$$z_t = \lambda \sum_{j=0}^{t-1} (1 - \lambda)^j x_{t-j} + (1 - \lambda)^t \theta_0 \qquad\qquad \text{for } t = 1, \dots n \qquad\qquad (5)$$

where $0 < \lambda < 1$ is a constant smoothing parameter chosen by the analyst, and the starting value z_0 is set equal to the baseline mean, θ_0. Thus, each value of z_t is a weighted average of the current observation x_t and the previous value of the EWMA statistic. Notably, since the EWMA statistic is a weighted average of the current and all past observations, it is relatively insensitive to violations to the normality assumption. Smoothing the data (i.e., decreasing λ) can increase power to detect deviations from the null model by regularizing the data; however, the optimal choice of smoothing parameter depends on the nature of the deviations. A general rule of thumb is to choose λ to be small (more smoothing) if one is interested in detecting small but sustained shifts in the process, and larger (less smoothing) if the shifts are expected to be large but brief.

The optimal choice of λ is discussed in greater detail in Lucas & Saccucci (Lucas & Saccucci, 1990) .

Without loss of generality we will assume that $z_0 = 0$, and therefore we can rewrite the sequence of EWMA statistics in matrix notation as

$$\vec{z} = \Lambda \vec{x} \tag{6}$$

where $\vec{z} = (z_1, z_2, \ldots z_n)^T$ and

$$
\Lambda = \lambda
\begin{vmatrix}
1 & & & & & 0 \\
(1-\lambda) & 1 & & & & \\
& (1-\lambda) & 1 & & & \\
\vdots & & & \ddots & & \\
(1-\lambda)^{n-2} & & & & 1 & \\
(1-\lambda)^{n-1} & (1-\lambda)^{n-2} & \ldots & & (1-\lambda) & 1
\end{vmatrix}
\tag{7}
$$

is a lower triangular smoothing matrix.

II. Statistical inference on the presence of activation in EWMA

In general, we are interested in making inference on two features of the model stated in Eqs. 1-3. First, we seek to develop statistical tests to determine *whether* a change in distribution (i.e., a departure from the baseline state of fMRI activity) has indeed taken place, i.e., whether to reject the null hypothesis $\theta_0 = \theta_1$. If a change is detected, we would also like to estimate *when* exactly the change took place, i.e. estimate the unknown parameter τ.

For detecting activation, the null hypothesis is that there is a single, baseline state where the mean (denoted µ) is constant. The alternative is that there are two or more states with different mean activity levels, though here we consider only the simplest two-state (baseline and activation) alternative. Thus,

$$H_0 : \mu = \theta_0 \text{ for } t = 1,\ldots n$$

$$H_a : \mu = \theta_0 \text{ for } t = 1,\ldots \tau \text{ and } \mu = \theta_1 \text{ for } t = \tau + 1,\ldots n \tag{8}$$

Our aim is to assess the probability of observing the data under the null hypothesis $\bar{x} \sim N(\theta_0, \Sigma)$.

In our analysis the parameters of the covariance matrix are estimated using data acquired during a baseline period in which the subject remains in a resting state (Fig. 1A and 1B). For each EWMA statistic z_t following the baseline period, we compute a test statistic T_t:

$$T_t = \frac{z_t - \theta_0}{\sqrt{Var(z_t)}}. \tag{9}$$

where $Var(z_t)$ is the variance of the EWMA statistic at time t. The statistic T follows a t-distribution with df degrees of freedom (discussed below), providing p-values according to classical inference. Further, control limits (analogous to confidence intervals) for detecting values of z_t that vary significantly from baseline can be calculated as follows:

$$\theta_0 \pm t^* \sqrt{Var(z_t)} \tag{10}$$

where t^* is a critical value from the t-distribution corresponding to the desired false positive-rate (see control limits in Fig. 1B). If z_t at any time exceeds the control limits, the process is deemed to have changed states and the null hypothesis is rejected. The control limits incorporate correction for multiple tests across time, as discussed below.

EWMA statistic variance

Under white noise, it can be shown (Montgomery, 2000) that

$$Var(z_t) = \sigma^2 \frac{\lambda}{2 - \lambda} \left(1 - (1 - \lambda)^{2t}\right) \tag{11}$$

For AR(1) noise, the minimum autocorrelation model appropriate for fMRI data, $Var(z_t)$ has previously been derived (Schmid, 1997). Explicit derivations of $Var(z_t)$ for a variety of

different underlying noise models (e.g. AR(p) and ARMA(1,1)) can be found in Lindquist and

Wager (Lindquist & Wager, In press). To summarize, for an AR(p) process, the autocorrelation

function is given by

$$\gamma(k) = \sum_{m=1}^{p} A_m G_m^{-h} \tag{12}$$

where G_m are the roots of

$$1 - \phi_1 z - \phi_2 z^2 + \ldots - \phi_p z^p = 0 \tag{13}$$

for $m = 1, \ldots p$ and A_m are constants (Brockwell & Davis, 2002). In this situation,

$$Var(z_t) = \frac{\lambda\gamma(0)}{(2-\lambda)}(1-(1-\lambda)^{2t}) + 2\lambda^2 \times$$

$$\sum_{m=1}^{p} \frac{A_m G_m}{G_m - (1-\lambda)} \left[\frac{(1-\lambda)}{G_m \lambda(2-\lambda)}(1-(1-\lambda)^{2(t-1)}) - \left(\frac{1-\lambda}{G_m}\right)^t \sum_{i=0}^{t-2}[(1-\lambda)G_m]^i \right]. \tag{14}$$

For an ARMA(1,1) process, the autocorrelation function can be calculated (Brockwell & Davis,

2002) as follows:

$$\gamma(h) = \begin{cases} \dfrac{1+\theta^2+2\phi\theta}{1-\phi^2}\sigma^2 & \text{if } h = 0 \\ \dfrac{(1-\phi\theta)(\phi-\theta)}{1-\phi^2}\sigma^2\phi^{h-1} & \text{if } h \geq 1 \end{cases}$$

Here,

$$Var(Z_t) = \frac{\lambda\gamma(0)}{(2-\lambda)}(1-(1-\lambda)^{2t}) + \frac{\lambda}{(2-\lambda)}\frac{2(1-\lambda)\gamma(1)}{(1-\phi(1-\lambda))}\left(1-(1-\lambda)^{2(t-1)}\right)$$

$$-\frac{2\lambda^2\gamma(1)(\phi(1-\lambda))^t}{\phi(1-\phi(1-\lambda))}\sum_{i=0}^{t-2}\left(\frac{1-\lambda}{\phi}\right)^i$$

13

In general, for a process with fixed length and covariance matrix Σ, the covariance matrix of the time series of EWMA statistics ($\vec{z}$), is given by $\Sigma^* = \Lambda\Sigma\Lambda^T$. Hence the variance of $\vec{z}$ is in turn given by the diagonal elements of Σ^*.

Correction for search across time and for spatial correlation

The p-values that are calculated at each time point must be controlled for search over the time series. Bonferroni correction is overly conservative because of positive dependence across time. A more sensitive procedure, which we adopt, uses Monte Carlo integration. Under the null hypothesis, the sequence of test statistics $\vec{T} = \{T_1, T_2, \ldots, T_n\}$ follows a multivariate t-distribution with covariance matrix Σ^* and *df* degrees of freedom calculated using Satterthwaite's approximation. Familywise error rate (FWER) control across the time series is provided by randomly generating vectors of n-length t-values from the $t(\Sigma^*, df)$ distribution and using their maxima to estimate a distribution of maximum null-hypothesis t-values (Nichols & Holmes, 2002):

$$T_{max} = \max_{k \in \{1,\ldots n\}} \{|T_k|\} \tag{15}$$

We use the multivariate T random number generator provided in Matlab (Mathworks, Natick, MA) with 10,000 samples at each voxel, which runs in less than 0.3*s* for *n=200* on a personal computer with a 2.53 GHz Pentium 4 processor and 1.00 GB of RAM. The 95[th] percentile of the distribution of T_{max} provides a critical t-value, T^*, for two-tailed FWER control[1]. If any $|T| > T^*$, the voxel is significantly activated (or deactivated) relative to the baseline period.

It should be noted that the null distribution of a likelihood ratio test for the hypotheses in Eq. 8 has been worked out analytically in Hawkins (1977) and Worsley (1979) for the case when

---

[1] The test is two-tailed because T_{max} is defined based on the absolute value of *T*.

14

ε_t are independent and identically distributed (*iid*) normal random variables. However, the distributions of the test statistics are rather complicated and numerical methods are used to study them. Siegmund (1985) provides analytic approximations for these distributions. However, all the results are for the *iid* normal case, while we are dealing with data that has significant autocorrelation. Using the permutation methods described above allows us to expand the application to comfortably handle temporally autocorrelated noise models.

Finally, we use false detection rate (FDR) (Genovese, Lazar & Nichols (2003)) to correct for correlation over space. As FDR works on the p-values, it is directly applicable after applying EWMA.

III. Estimation and inference on change-points (CPs)

There are several methods for estimating τ, the time point at which the state shift takes place. In the 'zero-crossing' method (Nishina, 1992), the last time point at which the process crosses θ_0 before $|T| > T^*$ is the estimate of τ. More formally, let us assume that t^A is the first time the EWMA statistic exceeds the control limits:

$$t^A = \min\left\{t \mid |T_t| > T^*\right\}. \tag{16}$$

The change-point estimator of τ for an *increase* in the process mean is given by the latest time point prior to t^A in which z_t lies below $\hat{\theta}_0$, which is the estimated baseline mean. That is, we define

$$\hat{\tau} = \max\left\{t \mid t \leq t^A, z_t \leq \hat{\theta}_0\right\}. \tag{17}$$

The change-point estimator of τ for a *decrease* is defined in an analogous manner. The method's main advantages are that it is conceptually straightforward and computationally efficient.

Activation duration and multiple change-points

A limitation of the zero-crossing method is that only a single, initial CP is assessed, and there is no clear provision for assessing return to the baseline state or the presence of multiple CPs. One possible approach would be to assume a return to baseline once the EWMA statistic crosses back across the control limit. Hence, the total number of OOC points can give a rough indication of the activation duration, though durations are likely to be biased toward zero.

As an alternative, we can use a Gaussian mixture model to classify each observation as either belonging to the baseline (in control) or activated (OOC) distribution. This procedure also allows us to estimate the length of time spent in the activated state. Though, the mixture model is philosophically different from change-point estimation using the zero-crossing method (it is an estimation rather than an inferential technique), we present it here as a flexible approach towards studying state changes in an fMRI time series, which can be extended to multiple activation states. In the mixture model approach, we assume that the data has been pre-whitened using the covariance estimates obtained from the EWMA stage of the analysis. We then model the fMRI time course as a mixture of two normal distributions, with different means, as follows:

$X_0 \sim N(\theta_0, \sigma_1^2)$ for the baseline state, and $X_1 \sim N(\theta_1, \sigma_2^2)$ for the activated state (see Fig. 1C for an example).

We can write this mixture as

$$X = (1 - \Delta)X_0 + \Delta X_1 \qquad (18)$$

where the random variable Δ is equal to one with probability *p* and equal to zero with probability *1-p*. The density function of X can be written

$$f_X(x) = (1 - p)f_{X_0}(x) + pf_{X_1}(x). \qquad (19)$$

where $f_{X_i}(x)$ is the normal probability density function with mean θ_i and variance σ_i^2.

We can fit this model to the data using maximum likelihood methods. The unknown parameters of the model are $(\theta_0, \theta_1, \sigma_1^2, \sigma_2^2, p)$ and the log-likelihood can be written:

$$l(\theta_0, \theta_1, \sigma_2^2, \sigma_2^2, p \mid \mathbf{x}) = \sum_{i=1}^{n} \log\left(f_X(x_i)\right) \tag{20}$$

The parameters that maximize this term can be found using the EM-algorithm (see Appendix A).

Once we have determined the maximum likelihood estimates of the parameters in this model, we need to classify each data point according to which state they belong to. This can be done using Bayes' formula. The probability that a data point belongs in the active (OOC) state is given by:

$$P(active \mid x_i) = \frac{p f_{X_1}(x_i)}{f_X(x_i)}. \tag{21}$$

If $P(active \mid x_i) > 0.5$, then the time point is classified as belonging to the active state, otherwise it is classified as belonging to the original baseline state. The mixture model approach is attractive as it does not require one to specify the number of change-points present in the time course *a priori*.

It is important to note that in the EWMA framework the noise covariance was not assumed to vary between the baseline and OOC states. In this section we have relaxed this assumption to allow the covariance to vary (up to a scaling term) between the states. This was done to increase the flexibility of the mixture model. However, we return to the assumption of non-varying noise covariance in the next section where the HEWMA framework is presented.

IV. Hierarchical EWMA (HEWMA) for population inference

The EWMA procedure outlined above is suitable for studying a single time series for an individual subject. In fMRI analysis we are typically most interested in detecting whether an

effect is present over an entire group of subjects. Group changes are of primary interest in the neurosciences for making population inferences. For this reason, we developed a group analysis using a hierarchical (i.e., mixed effects) extension of EWMA.

It should be noted that there does exist a multivariate extension of the EWMA framework - the multivariate exponentially weighted moving average (MEWMA) (Lowry, Woodall, Champ & Rigdon (1992)). However, direct application of MEWMA is not appropriate for our purposes. For one, the MEWMA approach is only useful for performing a fixed-effects analysis, while we are primarily interested in mixed effects analysis and the ability to make population wide inference. Secondly, the MEWMA approach does not differentiate between the direction of the change in mean, while we are interested in studying activations and deactivations separately.

To circumvent these two shortcomings of the MEWMA approach, we instead chose to develop the hierarchical exponentially weighted moving-average (HEWMA) model that allows us to perform a mixed-effects analysis on fMRI group data using the same type of analysis that the EWMA method allows for single-subject data. We use the EWMA statistic and covariance matrix defined previously, together with a between-subjects covariance term, to obtain the HEWMA statistic, which is a weighted population average, and its corresponding covariance matrix. Thereafter, the Monte-Carlo procedure used in the EWMA framework is applied to get p-values and test the hypothesis of consistent activation within the group, or alternatively differences in activation between groups.

The data consist of a time course from one voxel, an ROI or a component for each of m subjects. Independent analysis is performed for each voxel/ROI/component time series (i.e., the massive univariate approach). Let x_t^i denote the data for subject i at time t, where $i = 1, \ldots m$ and $t = 1, \ldots n$. Our hierarchical mixed-effects model takes the following form:

$$x_t^i = s_t^i + \varepsilon_t^i \tag{22}$$

$$s_t^i = s_t^{pop} + \eta_t \tag{23}$$

where η_t is the subject-level noise term at time t and s_t^i, the underlying signal for subject i, is considered to be either in a baseline or activated state at each time point. Finally, the population signal

$$s_t^{pop} = \begin{cases} \theta_0 & \text{for } t = 1, \dots \tau \\ \theta_1 & \text{for } t = \tau + 1, \dots n \end{cases}, \tag{24}$$

shows how the subject mean state-values are drawn from a larger population. In matrix format we can write Eq. 22 as

$$\vec{x}^i = \vec{s}^i + \vec{\varepsilon}^i \tag{25}$$

where $\vec{x}^i = \left(x_1^i, x_2^i, \dots x_n^i\right)^T$, $\vec{s}^i = \left(s_1^i, s_2^i, \dots s_n^i\right)^T$ and $\vec{\varepsilon}^i = \left(\varepsilon_1^i, \varepsilon_2^i, \dots \varepsilon_n^i\right)^T$. Here $\vec{\varepsilon}^i \sim N(0, \Sigma_i)$ and the vector $\vec{s}^i$ can be considered the equivalent of beta weights, one weight for each time point. Similarly, we can write Eq. 23 as

$$\vec{s}^i = \vec{s}_{pop} + \vec{\eta} \tag{26}$$

where $\vec{s}_{pop} = \left(s_1^{pop}, s_2^{pop}, \dots s_n^{pop}\right)^T$ and the noise vector $\vec{\eta}$ is $N(0, \Sigma_B)$. Here Σ_B represents the noise covariance matrix between subjects, and the off-diagonals are zero for independent subjects; thus, $\vec{\eta}$ is an n x 1 vector of iid $N(0, \sigma_B)$ random variables. This implies that the true subject-level effect varies over time and may differ among subjects.

The multi-level model can be written in single-level format as:

$$\vec{x}^i = \vec{s}_{pop} + \vec{\eta} + \vec{\varepsilon}^i = \vec{s}_{pop} + \vec{\xi}^i \tag{27}$$

where $\vec{\xi}^i \sim N(0, \Sigma_i + \Sigma_B)$; that is, the overall noise term is distributed with a variance equal to the sum of between-subjects and within-subjects variances. In the continuation we write $V^i = \Sigma_i + \Sigma_B$. Thus, in the HEWMA model the unknown parameters are θ_0, θ_1, τ, and the variance components Σ_i and Σ_B. Note that if we have first performed a single-level analysis (EWMA) on each subject (see Eq. 4) we can use the EWMA statistics (z_t^i for subject i at time t) in the second level. In this case, we can assume that the within-subject variance components Σ_i are known, and brought forward, from the first level of analysis.

Estimation of the HEWMA statistic and its variance components

When calculating the total variation in subject i's EWMA statistic we write,

$$V_i^* = \Lambda V_i \Lambda^T$$

$$= \Lambda \Sigma_i \Lambda^T + \Lambda \Sigma_B \Lambda^T$$

$$= \Sigma_i^* + \Sigma_B^*, \tag{28}$$

where Σ_i^* is the variance brought forward from the single subject analysis and Σ_B^* can be calculated using the fact that $\Sigma_B = I_n \sigma_B^2$. The only unknown term that needs to be estimated within the HEWMA framework is therefore the parameter σ_B. Hence, we can write $\Sigma_B^* = \alpha \Lambda^T \Lambda$ where $\Lambda^T \Lambda$ denotes the known portion and α the unknown portion of the covariance matrix.

In our approach we estimate the unknown variance component using restricted maximum likelihood (ReML). Our approach is equivalent to that used in SPM2 (Friston et al., 2002), which uses an EM-algorithm to estimate the parameters of interest. To simplify notation we begin by rewriting the problem in matrix form. Let,

$$\vec{z}_G = \left[\vec{z}^{1\mathrm{T}} \ \vec{z}^{2\mathrm{T}} \ \cdots \ \vec{z}^{mT} \right]^T \tag{29}$$

be the combined vector of EWMA-statistics for all m subjects (G denotes group). Recall, that in the hierarchical model we can write the covariance matrix for each individual subject as

$V_i^* = \Sigma_i^* + \Sigma_B^*$. Hence, it follows that the covariance matrix for $\vec{z}_G$ can be written:

$$V_G^* = \begin{bmatrix} V_1^* & & & 0 \\ & V_2^* & & \\ & & \ddots & \\ 0 & & & V_m^* \end{bmatrix} \tag{30}$$

Further, let $G = \begin{bmatrix} I_n & I_n & \cdots & I_n \end{bmatrix}^T$ be an $mn \times n$ matrix where I_n is the $n \times n$ identity matrix. Using this notation we can define the HEWMA-statistic, $\vec{z}_{pop}$, as

$$\vec{z}_G = G\,\vec{z}_{pop} + \vec{\xi}_{pop} \tag{31}$$

where the covariance matrix $\vec{\xi}_{pop}$ is $N\big(0, V_G^*\big)$. We can then estimate the HEWMA statistic using generalized least squares regression:

$$\hat{\vec{z}}_{pop} = \big(G^T V_G^{*-1} G\big)^{-1} G^T V_G^{*-1} \vec{z}_G$$

$$= \left[\sum_{i=1}^m V_i^{*-1}\right]^{-1} \sum_{i=1}^m V_i^{*-1} \vec{z}^{\,i} \ . \tag{32}$$

Hence the estimate of $\vec{z}_{pop}$ will be a weighted average of the individual subject's EWMA statistics. The covariance matrix for the HEWMA statistic can be written as:

$$V_{pop}^* = \big(G^T V_G^{*-1} G\big)^{-1} = \left[\sum_{i=1}^m V_i^{*-1}\right]^{-1} \tag{33}$$

In order to estimate the HEWMA statistic and its covariance matrix, we first need an estimate of V_i^*. In our approach the within-subject component Σ_i^* is assumed known from the first level and the general form of Σ_B^* is also assumed known up to a scaling term. Hence we can write the total

21

variance for subject i as $V_i^* = \alpha \Lambda^T \Lambda + \Sigma_i^*$, where α is the unknown and $\Lambda^T \Lambda$ the known part of Σ_B^*.

The problem of estimating the parameter α, is similar to the variance component estimation procedure performed in SPM2 (Friston et al., 2002). There they find the so-called hyper-parameter using an EM-algorithm and we will follow the same general outline here. Let, $V_G^* = \alpha Q_G + \Sigma_G$ where

$$Q_G = I_m \otimes \Lambda^T \Lambda \qquad \text{and} \qquad \Sigma_G = \begin{bmatrix} \Sigma_1^* & & & 0 \\ & \Sigma_2^* & & \\ & & \ddots & \\ 0 & & & \Sigma_m^* \end{bmatrix}. \tag{34}$$

The parameters α and Z_{pop} are estimated iteratively using the EM-algorithm in Appendix B.

Corrected p-values and population change-point estimates

The final step in the HEWMA framework is performing a Monte Carlo simulation to get corrected p-values (searching over time). This is done as described in the EWMA section, except here we use Z_{pop} and its covariance matrix to calculate the relevant test-statistics and define the multivariate t-distribution. Fig. 6 shows an example of the whole procedure where the HEWMA group activation, Z_{pop}, is estimated from the time courses of 24 subjects. Monte Carlo simulations are used to find corrected p-values. Fig. 6B shows the observed max T-statistic as a black line overlaid on the distribution of the max T-statistic under the null hypothesis. Finally, FDR is used to correct for spatial correlations.

Change-point estimation can either be performed directly on the HEWMA statistic using the methods described for the single subject case, or performed on the individual subjects' time courses. Though specific inferential techniques have not been developed for CP estimation in the

group (HEWMA) analysis, the latter approach can be taken with a few minor alterations of the single subject framework. When using the zero-crossing method, population inference is straightforward using a sign permutation test. Here the change-point is defined to be the first time point, prior to the HEWMA statistic being in the active state, in which a significant number of the individual time courses have crossed the baseline mean. Alternatively, it is possible to estimate the change-point for each individual subject and apply bootstrap methods to construct tests and confidence intervals for the population change-point. In addition, we have developed theory for using MLE methods for estimating the number of change-points, as well as their timing. This material will be published in a separate paper dealing solely with change-point estimation.

V. Simulations

In order to test the EWMA and HEWMA methodology and the efficiency of the change-point estimation procedure we performed a set of three simulation studies outlined in detail below. In Simulation 1, we simulate a single-subject dataset with four active regions, each with a different activation onset time, and estimate both the likelihood of activation and CPs using EWMA with the zero-crossing method. In Simulation 2, we assess the false positive rate (FPR) and power for the HEWMA method across values of the smoothing parameter λ. Simulation 3 assessed power and FPR across varying durations of the baseline period used for estimating variance components.

Note that in each of our simulations the noise is considered spatially independent. This was done to show the accuracy of the p-values in the case when they are considered uncorrected across space. In an application to real data, FDR (Genovese, Lazar & Nichols, 2002) can be used

to correct these p-values in he same manner that they are used to correct GLM-based p-values (Worsley & Friston, 1995).

Simulation 1. As shown in Fig. 2, we constructed a 64×64 phantom image containing a square region of size 48×48 representing a human brain. The image intensities are assigned values of 1 or 0 for the points inside or outside of the square, respectively. Four smaller squares, with dimensions 8×8, are placed inside the larger square to simulate ROIs with static contrast to the larger square. To simulate a dynamic image series, this base image is recreated 250 times according to a boxcar paradigm consisting of a prolonged period of activation of length 50 time points during which the signal within the four squares increases to 2. The onset time of activation varies between the four regions and takes values of 60, 80, 100 and 120 time points. Hence, each region has activation of similar length and intensity, but with varying onset times. Noise simulated using an AR(2) model with standard deviation equal to 1 is added to each voxel's time course. We analyzed the simulated data set using EWMA with $\lambda = 0.2$ and an AR(2) noise model. We further estimated onset times (change-points) and p-values for all active voxels.

Simulation 2. The next simulation sought to study the false positive rate (FPR) and perform power calculations for the HEWMA method. Actual fMRI noise was extracted from non-significant voxels of the brain (chosen because their HEWMA statistics gave rise to p-values above .95), obtained from the experimental data described in the next section. In total 50,350 noise time courses of length 215 time points were included in the study. The simulation mimicked a group analysis consisting of 20 subjects. For the FPR study, null hypothesis data with no activation was created by randomly sampling time series from the collection of noise time courses. This was done for each of 20 "subjects" and a random between-subject variation

with a standard deviation of size one third of the within-subject variation was added to each

subject's time course. A significance level of $\alpha = 0.05$ was used to determine "active" voxels.

For the power calculations the same procedure was repeated, with the difference that an active

period of length 50 time points was added to the noise data, with intensity equivalent to a

Cohen's d of 0.5 (Cohen, 1988). This coincides with values observed in experimental data

(Wager, Vazquez, Hernandez, & Noll, 2005).

The HEWMA method was performed on 5000 replications of each of these two data types

for λ values set equal to 0.1, 0.3, 0.5, 0.7 and 0.9. The analysis was further performed using

noise models in the HEWMA framework corresponding to white noise (WN), AR(1), AR(2) and

ARMA(1,1) noise. For each simulation the first 60 time points were used as a baseline period.

Simulation 3. The procedure was identical to that for Simulation 2, except that in this

simulation, the baseline length was either 20, 40, 60 or 80 time points. The analysis was again

performed using noise models in the HEWMA framework corresponding to white noise (WN),

AR(1), AR(2) and ARMA(1,1) noise. For each simulation λ was set equal to 0.2 and a

significance level of $\alpha = 0.05$ was used.

VI. Experimental fMRI data collection and analysis

Participants. We applied the HEWMA method to data from 30 participants scanned with

BOLD fMRI at 3T (GE, Milwaukee, WI). The experiment was conducted in accordance with the

Declaration of Helsinki and was approved by the University of Michigan institutional review

board. Six participants were excluded because of motion or nonlinear normalization artifacts,

leaving 24 participants. They were discarded prior to analysis due to substandard spatial

normalization and/or excessive head motion. It should be noted that these subjects would also have been excluded from a standard GLM analysis.

    Task design. The task used was a variant of a well-studied laboratory paradigm for eliciting anxiety (Dickerson & Kemeny, 2004; Gruenewald, Kemeny, Aziz, & Fahey, 2004; Roy, Kirschbaum, & Steptoe, 2001), shown in Fig. 5. The design was an off-on-off design, with an anxiety-provoking speech preparation task occurring between lower-anxiety resting periods. Participants were informed that they were to be given two minutes to prepare a seven-minute speech, and that the topic would be revealed to them during scanning. They were told that after the scanning session, they would deliver the speech to a panel of expert judges, though there was "a small chance" that they would be randomly selected not to give the speech.

    After the start of fMRI acquisition, participants viewed a fixation cross for 2 min (resting baseline). At the end of this period, participants viewed an instruction slide for 15 s that described the speech topic, which was to speak about "why you are a good friend." The slide instructed participants to be sure to prepare enough for the entire 7 min period. After 2 min of silent preparation, another instruction screen appeared (a 'relief' instruction, 15 s duration) that informed participants that they would not have to give the speech. An additional 2 min period of resting baseline followed, which completed the functional run.

    Heart rate was monitored continuously, and heart rate increased after the topic presentation, remained high during preparation, and decreased after the relief instruction (data will be presented elsewhere). Because this task involves a single change in state, as in some previous fMRI experiments (Breiter & Rosen, 1999; Eisenberger, Lieberman, & Williams, 2003), and the precise onset time and time course of subjective anxiety are unknown, this design is a good candidate for the HEWMA analysis.

Image acquisition. A series of 215 images were acquired using a T2*-weighted, single-shot reverse spiral acquisition (gradient echo, TR = 2000, TE = 30, flip angle = 90) with 40 sequential axial slices (FOV = 20, 3.12 x 3.12 x 3 mm, skip 0, 64 x 64 matrix). This sequence was designed to enable good signal recovery in areas of high susceptibility artifact, e. g. orbitofrontal cortex. High-resolution T1 spoiled gradient recall (SPGR) images were acquired for anatomical localization and warping to standard space.

Image analysis. Offline image reconstruction included correction for distortions caused by magnetic field inhomogeneity. Images were corrected for slice acquisition timing differences using a custom 4-point sync interpolation and realigned (motion corrected) to the first image using Automated Image Registration (AIR; (Woods, Grafton, Holmes, Cherry, & Mazziotta, 1998)). SPGR images were coregistered to the first functional image using a mutual information metric (SPM2). When necessary, the starting point for the automated registration was manually adjusted and re-run until a satisfactory result was obtained. The SPGR images were normalized to the Montreal Neurological Institute (MNI) single-subject T1 template using SPM2 (with the default basis set). The warping parameters were applied to functional images, which were then smoothed with a 9 mm isotropic Gaussian kernel.

Individual-subject data were subjected to linear detrending across the entire session (215 images) and analyzed with EWMA. An AR(2) model was used to calculate the EWMA statistic (z_t) and its variance, and $\vec{z}$ and the variance estimates were carried forward to the group level HEWMA analysis. We used custom software (see author note for download information and appendices) to calculate statistical maps throughout the brain, including HEWMA (group) *t* and p-values for activations (increases from baseline) and deactivations (decreases from baseline); individual and group CPs (calculated on the group HEWMA t-time series) using the zero-

crossing method, activation duration as estimated by the number of OOC points, and CP and run-length estimates using the Gaussian mixture model described above.

Significant voxels were classified into sets of voxels showing similar behavior using k-means clustering. We considered activations and deactivations separately, and used k-means clustering on the group CP and longest activated run length (from the mixture model) to assign voxels into classes with similar behavior. To do this, we used the k-means algorithm implemented in Matlab 7.4 (Mathworks, Natick, MA) with the v x 2 matrix of values for the v significant activated (or deactivated) voxels as input. This choice was arbitrary and is primarily for data visualization; other clustering algorithms may also be used effectively. The number of classes was determined by visual inspection of the joint histogram of CP and duration values. Twelve classes were used for the analyses reported here. Sets of contiguous supra-threshold voxels ('regions') of the same class were the unit of analysis for interpretation.

Examining the systematic features in the time courses of regions of interest permits us to make inferences about the role the region plays in the speech preparation task. Rather than being limited to testing whether a voxel is activated during the preparation interval on average compared with baseline periods, the HEWMA method can detect a number of different types of interesting systematic features, including sustained activation during the preparation task, transient activation during instruction presentation, fluctuations in baseline activity that may be related to the start of scanning, and others. Of particular interest are voxels whose activation onset is near the time of task onset (at 60 TRs or 120 s) and whose activity is sustained throughout the task (at least 60 TRs/120 s), which may reflect sustained anxiety.

Results

I. Simulation results

Simulation 1. We used EWMA to create a significance map and a change-point map that accurately depicts the difference in onset time between regions. Setting the smoothing parameter $\lambda = 0.2$ and the significance level $\alpha = 0.05$, we analyzed each voxel using the EWMA procedure outlined in the Methods section. Fig. 2A depicts the theoretical significance map, with equal amount of activation present within each of the four active regions. Fig. 2C depicts the actual significance map obtained using EWMA. This indicates that we were able to accurately detect a large number of active voxels within the region of activity, with a minimal number of false positives outside of the region. Fig. 2B depicts the theoretical change-point map, where the intensity varies depending on the onset time. Fig. 2D depicts the actual change-point map obtained by calculating the zero-crossing change-point estimate for each voxel deemed active in the prior analysis. The CP map provided accurate estimates of the onset times for the four regions, as indicated by the similar intensity values for the true values (Fig. 2C) and estimates (Fig. 2D). Examination of the error in CP estimation showed a distribution that was centered at zero with a slight left–skew (mean = -2.0, St. Dev. = 6.3, median = 0 and IQR = 5).

Simulation 2. Figs. 3A and 3B shows the FPR and power calculations for each noise type as a function of the smoothing parameter λ. Fig. 3A shows that the number of false positives increases for each noise type as a function of λ (i.e., with less smoothing). This is natural, as low values of λ entail a greater amount of smoothing, and minimize the risk of the null hypothesis data venturing too far from the baseline mean. The nominal alpha level of 0.05 is shown by the horizontal dashed line. As λ increases, the amount of smoothing decreases and the FPR exceeds $\alpha = 0.05$. The ARMA model performs worse than the other models in FPR and

power, while the other models appear to behave in a similar manner (though are somewhat conservative, with FPRs near 0.01 with low λ). All models control the FPR appropriately with $\lambda \leq .4$.

Studying Fig. 3B, it appears that though the power increases slightly for each noise type as a function of λ, it does not vary in a significant manner. All noise models gave roughly equivalent results. However, we suggest the use of the AR(2) model as it has the flexibility to model periodic noise oscillations that are often produced in fMRI as a result of physiological changes. Fig. 5A shows an ROC curve for this noise model corresponding to each of the five smoothing parameters.

In summary, our simulation studies indicate that a low value of λ will give a test with reasonable power and low false positive rates. Increasing the value of λ will lead to a slight increase in power, but at the cost of an increase in FPR. In the continuation we use $\lambda = 0.2$.

Simulation 3. Fig. 3C and 3D shows the FPR and power calculations for each noise type as a function of the baseline period length. Studying Fig. 3C, it is clear that the number of false positives decreases for each noise type as the baseline period increases. This is natural as a long baseline period allows for more data to accurately estimate the parameters of the model. The ARMA model performs significantly worse than the other models, and gives rise to a dramatically inflated FPR for baseline lengths less than 60 time points. This performance may be due to the fact that the baseline period is too short to get an accurate estimate of the variance components for this model type, as the ARMA parameter estimation is more complex than for the other models (i.e. MLE vs. method of moments). Studying Fig. 3D, it appears that though the power increases slightly for each noise type as a function of baseline length, it does not vary substantially across baseline durations except for the ARMA model which performed worse than

the other models. We again suggest the use of the AR(2) model, and Fig. 5B shows an ROC

curve for this noise model corresponding to each of the four baseline lengths.

In summary, our simulation studies indicate that an increased baseline leads to increased

power as well as decreased FPR, which is advantageous. However, it is interesting to note that

the WN, AR(1) and AR(2) models are all robust enough to handle a baseline period as short as

20 time points. Naturally, these values depend on the noise characteristics, so collecting more

baseline data (e.g., 60 time points) is recommended. The ARMA model, on the other hand,

requires a baseline of at least 60 time points.

II. FMRI results

The HEWMA analysis on the experimental data revealed task-related changes consistent

with previous literature on neuroimaging of emotion (Phan, Wager, Taylor, & Liberzon, 2004; T.

D. Wager, K. L. Phan, I. Liberzon, & S. F. Taylor, 2003) including activations in dorsolateral

and rostral medial prefrontal cortices, middle temporal gyrus, and occipital cortex (Figs. 6 and

7). Deactivations were found in ventral striatum and ventral anterior insula. The activation is

consistent with what might be expected in a cognitively complex task, which involved visual

cues at two periods during the task, the mental effort and subvocal rehearsal required to prepare a

speech, and the anxiety elicited by the task context. In particular, the rostral PFC has been

strongly implicated in processing of self-relevant information and the representation and

regulation of aversive emotional states (Ochsner et al., 2004; Phan, Wager, Taylor, & Liberzon,

2002; Quirk & Gehlert, 2003; Quirk, Russo, Barron, & Lebron, 2000; Ray et al., 2005; T.

Wager, K. L. Phan, I. Liberzon, & S. F. Taylor, 2003). Information about the onset and duration

of activation provided by the HEWMA analysis can constrain interpretation of the roles of these

regions in task performance.

Estimates were made of the time of onset and duration of activity for activated regions. The range of estimated onset times, from around 40 TRs (80 s) to around 180 TRs (360 s) from the start of scanning indicates that different regions were activated at different times during scanning. Most significant voxels showed activation onsets around the time of task onset (60 TRs), when the visual cue to begin speech preparation was presented (Fig. 7). Likewise, activation duration estimates from the Gaussian mixture model ranged from transient increases (approx. 10 TRs or 20 s) to sustained increases (~80 TRs, 160 s). The 2-D histogram of significant voxels by CP and estimated duration is shown in Fig. 7B. K-means classification was used to provide a way to identify classes of activated voxels with similar onset and duration estimates. Class membership is indicated by color in the histogram in Fig. 7B. The diversity of onset times and durations suggests that a variety of different GLM models would be required to detect these activations.

Examining the group time courses of voxels in these regions corroborates this view and provides additional evidence. Time courses for two patterns of responses are shown in Fig. 8, including a medial prefrontal region showing sustained activity throughout the anxiogenic task and an occipital region showing transient responses to the instruction periods (when visual stimuli were presented). The baseline period in Fig. 8 is indicated by the shaded gray box in each panel, and the HEWMA-statistic time course is shown by the thick black line (+/- one standard error across participants, shown by gray shading). Control limits are shown by dashed lines; thus, the region is significantly activated if the HEWMA-statistic exceeds the control limit at any time point. Importantly, the time course of activation in rMPFC parallels the structure of the task in that activation begins around the onset of speech instructions and is sustained throughout the

duration of the preparation interval. This region has been related in many neuroimaging studies to subjective anxiety and self-referential processing (Breiter & Rosen, 1999; Dougherty et al., 2004; Eisenberger et al., 2003; Ray et al., 2005; Wang et al., 2005), and is thus a neurophysiologically plausible candidate to show sustained activation related to the anxiogenic task.

Discussion

Typically statistical methods in fMRI can be categorized into two broad categories: hypothesis and data driven approaches. Hypothesis-driven approaches test whether activity in a brain region is systematically related to some known input function. In this approach, typically, the general linear model (GLM) is used to test for differences in activity among psychological conditions or groups of participants. However, for many psychological processes, the precise timing and duration of psychological activity can be difficult to specify in advance. In this situation the GLM approach becomes impractical, as the psychological activity cannot be specified *a priori*. Data-driven methods, such as independent components analysis (ICA), give an account of the data using few *a priori* assumptions. Instead, they attempt to characterize reliable patterns in the data, and relate those patterns to psychological activity post hoc. The main drawback of these methods is that they do not provide statistics for making inferences about whether a component varies over time and when changes occur in the time series.

The purpose of the HEWMA method is to allow for the detection of systematic changes in activity with a variety of onset times and activation durations. The differences in onset and duration are likely to reflect differences in functional anatomy, i.e., the way in which each activated region participates in the task. Both the EWMA and HEWMA methods for fMRI data analysis are designed to detect regions of the brain where the signal deviates reliably from a

baseline state. The methods make no *a priori* assumption about the behavior of these changes, and the method will detect activation and deactivation, as well as regions with both short and prolonged activation duration. In this sense, both EWMA and HEWMA can be thought of as searches for activity differences across time, correcting for the multiple comparisons tested and accounting for the correlation among observations.

Once a systematic deviation from baseline has been detected, the second step in the analysis entails estimating *when* exactly the change took place, as well as the recovery time (if any). This estimation procedure can be performed using the zero-crossing method or a Gaussian mixture model. Other methods are under development as well (e.g. an MLE approach) and we present the methods above as a starting point for further work. Once these estimates are obtained, we can cluster the active voxels into groups whose estimates behave in a similar manner. This allows us to classify regions and even discard of regions for which the activation was triggered by effects such as drift or movement. As a final step, after estimating the change-point for each active voxel in the brain, we summarize the results in a Change-point Map (CPM). A CPM is an image of the brain with a color-coded change-point mask superimposed, whose intensity varies depending on the estimated onset time of activation.

In this paper we have assumed that the change points were fixed in time across groups of subjects. We feel this provides a useful first step towards estimating the latency of processes in the brain. There may be situations where the change-point, corresponding to a certain stimuli, differs across subjects. In this case it may prove beneficial to instead allow the change-points to vary randomly across subjects. This topic is often referred to as multi-path change-point problems (Asgharian & Wolfson, 2001), and will be the focus of future work.

In addition, we have chosen to detrend the fMRI time courses prior to the EWMA/HEWMA analysis in order to remove nuisance parameters (e.g. drift). However, it would be relatively straightforward to extend the EWMA framework to simultaneously estimate the trend regressors while performing the change-point detection. For example, there are methods are testing for change-points in simple linear regression (Kim & Siegmund 1989), which would directly allow for the modeling of drift components in the model. While we have not further explored these techniques at this time, this is an interesting direction for future research.

Recommended choices for analysis parameters

The main decisions that need to be made before applying these methods to experimental data are the choice of the smoothing parameter λ and the length of the baseline period. According to our power and false-positive rate analyses, a relatively low value of λ gives high power, with a strong control of the FPR. In our analysis of experimental data we typically choose a value of 0.2, as this appears to give rise to an adequate amount of smoothing for the analysis of fMRI data. The optimal value of λ will vary depending on whether brief or sustained changes are of greater interest, with lower values being more appropriate for more sustained activity. The length of the baseline period is another issue. The data within this period is used to estimate the baseline mean, as well as the within-subject variation, and our simulations show that, unsurprisingly, longer baseline periods produce more accurate estimates. However, our analysis indicates that the method is relatively stable for even very short baseline periods for white noise, AR(1) and AR(2) models. The ARMA model showed increased sensitivity to short baseline periods with drastically increased FPR and decreased power. Though some models give reasonable results with relatively short baseline periods, a longer baseline period of at least 60

time points is recommended to ensure stable estimates of the baseline activation variance. The

AR(2) model may be most appropriate for fMRI data, as it has the flexibility to model periodic

noise oscillations that are often produced in fMRI as a result of physiological changes (e.g.,

pulsatile motion of the brain due to breathing and cardiac activity).

Potential applications

HEWMA appears to be an appropriate analysis for group fMRI data, particularly when it

is not possible to replicate experimental manipulations within subjects (e.g., a state anxiety

induction that cannot be repeated without changing the psychological nature of the state).

Emotional responses are one prime candidate for applications of the method. But there are a

number of other domains in which it may be useful as well, and the method applies to any

longitudinal data with enough observations so that repeated measures ANOVA (for example) is

impractical. HEWMA may be particularly useful for arterial spin labeling and perfusion MRI

studies, which measure brain activity over time without the complicating factors of signal drift

and highly colored noise in fMRI (Liu, Wong, Frank, & Buxton, 2002; Wang et al., 2005).

Another potential use is in identifying voxels of interest and characterizing brain

responses in 'ecologically valid' tasks, such as free viewing of films. In a recent paper, for

example, participants watched a 60 min segment of an action movie (Hasson, Nir, Levy,

Fuhrmann, & Malach, 2004). The investigators examined the time-course of activity throughout

the brain and assessed whether increases in particular regions were systematically related to

features of the film (e.g., presentation of scenes, hands, faces). HEWMA could be used in this

situation to identify voxels that respond consistently across participants during viewing, which

would reduce false positives by providing a reduced set of voxels of interest and providing some

quantitative tools for characterizing the duration and number of activated periods.

HEWMA could also be applied to changes in state-related activity evoked by learning, e.g., tonic increases brain activity as a function of expertise, or to studies of tonic increases following solutions to 'insight' problem-solving tasks. Expertise results in functional and structural reorganization of cortex (Kilgard & Merzenich, 1998; Kourtzi, Betts, Sarkheil, & Welchman, 2005), and HEWMA could be used to more precisely characterize the time-course of both types of changes (i.e., do shifts occur gradually or suddenly?) In an 'insight' task, participants are presented with a problem that requires a novel combination of elements (Bowden, Jung-Beeman, Fleck, & Kounios, 2005; MacGregor, Ormerod, & Chronicle, 2001). Once participants solve the task, there is a qualitative shift in their understanding of how the elements of the problem relate that cannot be reversed (the solution is 'obvious' once one knows it). This profound shift is poorly understood, in part because appropriate methods have not been devised to study its brain mechanisms in healthy participants.

Another potential use is in longitudinal studies of brain function or structure, and how they change with development or with the progression of a neurological or psychiatric disorder. For example, Mayberg and colleagues have conducted several longitudinal studies of resting FDG PET activity in depressed patients over the course of treatment (Goldapple et al., 2004; Mayberg et al., 2002). The time course corresponding to brain activity changes throughout the treatment process is unknown, and HEWMA could be used to locate regions that respond to treatment and identify the time at which they do so.

Conclusions

In this paper we developed a new approach, HEWMA, that can be used to make inferences about individual or group fMRI activity, even when conditions are not replicated (e.g.

a single experimental induction of emotion). The HEWMA method is an extension of EWMA, a time series analysis method in statistical process-control theory and change-point theory, to multisubject data. It permits population inference, and can be used to analyze fMRI data voxel-wise throughout the brain, data from regions of interest, or temporal components extracted using ICA or similar methods. Simulations show that the method has acceptable false-positive rate control, and application to an fMRI study of anxiety shows that it produces reasonable and novel results with empirical data. A toolbox implementing all functions in Matlab is freely available from the authors (see Author Note).

The HEWMA approach can complement GLM-based and purely data-driven methods (such as ICA) by providing inferences about whether, when, and for how long systematic state-related activation occurs in a particular brain region. Although it is developed here for fMRI data analysis, the method could be useful in detection of deviation from a baseline state in any type of time series data, including ASL, longitudinal studies of brain structure or PET activity, and others.

Author Note

References

Asgharian, M. & Wolfson, D.B. (2001). Covariates in multi-path changepoint problems: modeling and consistency of the MLE. *The Canadian Journal of Statistics*, 29(4), 515-528.

Beckmann, C. F., & Smith, S. M. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging, 23*(2), 137-152.

Beckmann, C. F., & Smith, S. M. (2005). Tensorial extensions of independent component analysis for multisubject FMRI analysis. *Neuroimage, 25*(1), 294-311.

Bowden, E. M., Jung-Beeman, M., Fleck, J., & Kounios, J. (2005). New approaches to demystifying insight. *Trends Cogn Sci, 9*(7), 322-328.

Breiter, H. C., Gollub, R. L., Weisskoff, R. M., Kennedy, D. N., Makris, N., Berke, J. D., et al. (1997). Acute effects of cocaine on human brain activity and emotion. *Neuron, 19*(3), 591-611.

Breiter, H. C., & Rosen, B. R. (1999). Functional magnetic resonance imaging of brain reward circuitry in the human. *Ann N Y Acad Sci, 877*, 523-547.

Brockwell, P. J., & Davis, R. A. (2002). *Introduction to Time Series and Forecasting* (Second ed.). New York: Springer-Verlag.

Calhoun, V. D., Adali, T., & Pekar, J. J. (2004). A method for comparing group fMRI data using independent component analysis: application to visual, motor and visuomotor tasks. *Magn Reson Imaging, 22*(9), 1181-1191.

Calhoun, V. D., Adali, T., Stevens, M. C., Kiehl, K. A., & Pekar, J. J. (2005). Semi-blind ICA of fMRI: A method for utilizing hypothesis-derived time courses in a spatial ICA analysis. *Neuroimage, 25*(2), 527-538.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognit Psychol, 17*(4), 391-416.

Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., et al. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *Neuroimage, 14*(5), 1136-1149.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Dickerson, S. S., & Kemeny, M. E. (2004). Acute stressors and cortisol responses: a theoretical integration and synthesis of laboratory research. *Psychol Bull, 130*(3), 355-391.

Donaldson, D. I., Petersen, S. E., Ollinger, J. M., & Buckner, R. L. (2001). Dissociating state and item components of recognition memory using fMRI. *Neuroimage, 13*(1), 129-142.

Dougherty, D. D., Rauch, S. L., Deckersbach, T., Marci, C., Loh, R., Shin, L. M., et al. (2004). Ventromedial prefrontal cortex and amygdala dysfunction during an anger induction positron emission tomography study in patients with major depressive disorder with anger attacks. *Arch Gen Psychiatry, 61*(8), 795-804.

Eisenberger, N. I., Lieberman, M. D., & Williams, K. D. (2003). Does rejection hurt? An FMRI study of social exclusion. *Science, 302*(5643), 290-292.

Fredrickson, B. L., Tugade, M. M., Waugh, C. E., & Larkin, G. R. (2003). What good are positive emotions in crises? A prospective study of resilience and emotions following the terrorist attacks on the United States on September 11th, 2001. *J Pers Soc Psychol, 84*(2), 365-376.

Friston, K. J., Glaser, D. E., Henson, R. N., Kiebel, S., Phillips, C., & Ashburner, J. (2002). Classical and Bayesian inference in neuroimaging: applications. *Neuroimage, 16*(2), 484-512.

Genovese, C. R., Noll, D. C. and Eddy, W. F. (1997). Estimating Test-Retest Reliability in fMRI I: Statistical Methodology, *Magnetic Resonance in Medicine*, 38, 497--507.

Genovese, C. R., Lazar, N. A., & Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage, 15*(4), 870-878.

Goldapple, K., Segal, Z., Garson, C., Lau, M., Bieling, P., Kennedy, S., et al. (2004). Modulation of cortical-limbic pathways in major depression: treatment-specific effects of cognitive behavior therapy. *Arch Gen Psychiatry, 61*(1), 34-41.

Gray, J. R., Braver, T. S., & Raichle, M. E. (2002). Integration of emotion and cognition in the lateral prefrontal cortex. *Proc Natl Academy Science U S A, 99*(6), 4115-4120.

Gruenewald, T. L., Kemeny, M. E., Aziz, N., & Fahey, J. L. (2004). Acute threat to the social self: shame, social self-esteem, and cortisol activity. *Psychosom Med, 66*(6), 915-924.

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science, 303*(5664), 1634-1640.

Hawkins, D.M. (1977). Testing a sequencce of Observations for a Shift in Location. *Journal of the American Statistical Association*, 72(357), 180-186.

Kilgard, M. P., & Merzenich, M. M. (1998). Cortical map reorganization enabled by nucleus basalis activity. *Science, 279*(5357), 1714-1718.

Kim, H.J. & Siegmund (1989). The Likelihood Ratio test for a Change-point in Simple Linear regression. *Biometrika*, 76(3), 409-423.

Kourtzi, Z., Betts, L. R., Sarkheil, P., & Welchman, A. E. (2005). Distributed neural plasticity for shape learning in the human visual cortex. *PLoS Biol, 3*(7), e204.

Koyama, T., McHaffie, J. G., Laurienti, P. J., & Coghill, R. C. (2005). The subjective experience of pain: Where expectations become reality. *Proc Natl Acad Sci U S A, 102*(36), 12950-12955.

Lindquist, M. A., & Wager, T. D. (In press). Application of change-point theory to modeling state-related activity in fMRI. In P. Cohen (Ed.), *Applied Data Analytic Techniques for "Turning Points" Research*.

Liou, M., Su, H.R., Lee, J-D, Aston, J.A.D., Tsai, A.C. & Cheng, P.E.. (2006) A Method for Generating Reproducibility Evidence in fMRI studies. *Neuroimage*, 29:383-395.

Liu, T. T., Wong, E. C., Frank, L. R., & Buxton, R. B. (2002). Analysis and design of perfusion-based event-related fMRI experiments. *Neuroimage, 16*(1), 269-282.

Lowry, C.A., Woodall, W.H., Champ, C.W. & Rigdon, S.E. (1992). A Multivariate Exponentially Weighted Moving Average Chart. *Technometrics*, 34, 46-53

Lucas, J. M., & Saccucci, M. S. (1990). Exponentially weighted moving average control schemes: Properties and enhancements. *Technometrics, 32*, 1-29.

MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2001). Information processing and insight: a process model of performance on the nine-dot and related problems. *J Exp Psychol Learn Mem Cogn, 27*(1), 176-201.

Mayberg, H. S., Silva, J. A., Brannan, S. K., Tekell, J. L., Mahurin, R. K., McGinnis, S., et al. (2002). The functional neuroanatomy of the placebo effect. *Am J Psychiatry, 159*(5), 728-737.

McIntosh, A. R., Chau, W. K., & Protzner, A. B. (2004). Spatiotemporal analysis of event-related fMRI data using partial least squares. *Neuroimage, 23*(2), 764-775.

McKeown, M. J., & Sejnowski, T. J. (1998). Independent component analysis of fMRI data: examining the assumptions. *Hum Brain Mapp, 6*(5-6), 368-372.

Montgomery, D. C. (2000). *Introduction to Statistical Quality Control* (Fourth Ed.). New York: Wiley.

Neter, J., Kutner, M.H., Nachtsheim, C.J., & Wasserman, W. (1996). *Applied Linear Statistical Models* (Fourth Ed.) Irwin Publishers.

Neubauer, A. S. (1997). The EWMA control chart: properties and comparison with other quality-control procedures by computer simulation. *Clin Chem, 43*(4), 594-601.

Nichols, T. E., & Holmes, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp, 15*(1), 1-25.

Nishina, K. (1992). A Comparison of Control Charts from the Viewpoint of Change-point Estimation. *Quality and Reliability Engineering International, 8*, 537-541.

Ochsner, K. N., Ray, R. D., Cooper, J. C., Robertson, E. R., Chopra, S., Gabrieli, J. D., et al. (2004). For better or for worse: neural systems supporting the cognitive down- and up-regulation of negative emotion. *Neuroimage, 23*(2), 483-499.

Otten, L. J., Henson, R. N., & Rugg, M. D. (2002). State-related and item-related neural correlates of successful memory encoding. *Nat Neurosci, 5*(12), 1339-1344.

Phan, K. L., Wager, T., Taylor, S. F., & Liberzon, I. (2002). Functional neuroanatomy of emotion: a meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage, 16*(2), 331-348.

Phan, K. L., Wager, T. D., Taylor, S. F., & Liberzon, I. (2004). Functional neuroimaging studies of human emotions. *CNS Spectr, 9*(4), 258-266.

Quirk, G. J., & Gehlert, D. R. (2003). Inhibition of the amygdala: key to pathological states? *Ann N Y Acad Sci, 985*, 263-272.

Quirk, G. J., Russo, G. K., Barron, J. L., & Lebron, K. (2000). The role of ventromedial prefrontal cortex in the recovery of extinguished fear. *J Neurosci, 20*(16), 6225-6231.

Ray, R. D., Ochsner, K. N., Cooper, J. C., Robertson, E. R., Gabrieli, J. D., & Gross, J. J. (2005). Individual differences in trait rumination and the neural systems supporting cognitive reappraisal. *Cogn Affect Behav Neurosci, 5*(2), 156-168.

Roberts, S. W. (1959). Control Chart Tests Based on Geometric Moving Averages. *Technometrics, 1*, 239-250.

Roy, M. P., Kirschbaum, C., & Steptoe, A. (2001). Psychological, cardiovascular, and metabolic correlates of individual differences in cortisol stress recovery in young men. *Psychoneuroendocrinology, 26*(4), 375-391.

Schmid, W. (1997). *On EWMA charts for time series* (Vol. 5). Heidelberg: Physica-Verlag.

Shehab, R. L., & Schlegel, R. E. (2000). Applying quality control charts to the analysis of single-subject data sequences. *Hum Factors, 42*(4), 604-616.

Siegmund, D. (1985). *Sequential Analysis: Test and Confidence Intervals.* New York, Heidelberg, Berlin: Springer-Verlag.

Taylor, S. F., Phan, K. L., Decker, L. R., & Liberzon, I. (2003). Subjective rating of emotionally salient stimuli modulates neural activity. *Neuroimage, 18*(3), 650-659.

Tugade, M. M., & Fredrickson, B. L. (2004). Resilient individuals use positive emotions to bounce back from negative emotional experiences. *J Pers Soc Psychol, 86*(2), 320-333.

Visscher, K.M., Miezin, F.M., Kelly, J.E., Buckner, R.L., Donaldson, D.I., McAvoy, M.P., Bhalodia, V.M., Petersen, S.E. (2003). Mixed block/event-related designs can correctly separate transient and sustained activity in fMRI. *NeuroImage*, 19, 1694-1708.

Wager, T., Phan, K. L., Liberzon, I., & Taylor, S. F. (2003). Valence, gender, and lateralization of functional brain anatomy in emotion: A meta-analysis of findings from neuroimaging. *Neuroimage, 19*, 513-531.

Wager, T. D. (2005). The neural bases of placebo effects in pain. *Current Directions in Psychological Science, 14*(4), 175-179.

Wager, T. D., Hernandez, L., Jonides, J., & Lindquist, M. (in press). Elements of functional neuroimaging. In J. Cacioppo & R. J. Davidson (Eds.), *Handbook of Psychophysiology* (4th ed.). Cambridge: Cambridge University Press.

Wager, T. D., Phan, K. L., Liberzon, I., & Taylor, S. F. (2003). Valence, gender, and lateralization of functional brain anatomy in emotion: a meta-analysis of findings from neuroimaging. *Neuroimage, 19*(3), 513-531.

Wager, T. D., Vazquez, A., Hernandez, L., & Noll, D. C. (2005). Accounting for nonlinear BOLD effects in fMRI: parameter estimates and a model for prediction in rapid event-related studies. *Neuroimage, 25*(1), 206-218.

Wang, J., Rao, H., Wetmore, G. S., Furlan, P. M., Korczykowski, M., Dinges, D. F., et al. (2005). Perfusion functional MRI reveals cerebral blood flow pattern under psychological stress. *Proc Natl Acad Sci U S A, 102*(49), 17804-17809.

Wise, R. G., Williams, P., & Tracey, I. (2004). Using fMRI to quantify the time dependence of remifentanil analgesia in the human brain. *Neuropsychopharmacology, 29*(3), 626-635.

Woods, R. P., Grafton, S. T., Holmes, C. J., Cherry, S. R., & Mazziotta, J. C. (1998). Automated image registration: I. General methods and intrasubject, intramodality validation. *J Comput Assist Tomogr, 22*(1), 139-152.

Worsley, K.J. (1979). On the Likelihood Ratio test for a Shift in Location of Normal Populations. *Journal of the American Statistical Association*, 74(366), 365-367.

Worsley, K. J., & Friston, K. J. (1995). Analysis of fMRI time-series revisited--again. *Neuroimage, 2*(3), 173-181.

Figure Captions

Figure 1. A schematic overview of: (A) the model of true activation, (B) the EWMA statistic and its control bounds, and (C) the Gaussian mixture model (illustrated using a separate simulation). The parameter τ represents the true change-point, and $\hat{\tau}$ its estimate. N_0 is the first time point after the baseline period, while N is the total number of time points in the potential-activation period. N^{ooc} is the first time point that is positively identified as activated, and it is used to calculate $\hat{\tau}$ in the zero-crossing method. Control limits are critical values for the EWMA statistic, z_t, correcting for multiple dependent comparisons across time. In the mixture model, $\hat{\tau}_{1-4}$ are estimates of when there is an active-nonactive or nonactive-active state change and $\hat{\omega}_{1-2}$ are estimates of the durations of two example activation periods.

Figure 2. Results from a simulated single run experiment (Simulation 1) using an effect size equal to 1 and an AR(2) noise model. (A-B) The true significance and change-point maps, (C) the significance map obtained from the EWMA analysis of the data (using $\lambda = 0.2$ and the AR(2) control bounds) and (D) the change-point map estimated using the zero-crossing method. To the right are examples of both active and non-active time courses plotted together with their control limits. The data is represented by the light gray line and the EWMA statistic by the dark black line.

Figure 3. (A-B) Simulation 2. Simulated power and false positive rates for HEWMA with varying smoothness parameter λ, and a fixed baseline length of 60 TRs. (C-D) Simulation 3. Same plots with fixed $\lambda = 0.2$ and varying baseline length.

Figure 4. (A) Simulation 2. Receiver operating characteristic (ROC) curves show the fraction of true positives (true positive rate, TPR) vs. the fraction of false positives (false positive rate, FPR)

for HEWMA using our recommended noise model (AR(2)) with varying smoothness parameter λ, and a fixed baseline length of 60 TRs. The optimal primary threshold for discriminating active and inactive voxels is p < .08. In these simulations, higher lambdas (less smoothing) produced more sensitive results (but with false positive rates that exceed the nominal threshold; see Fig. 3). Sensitivity with real datasets will vary depending on the smoothness of the true underlying signal. (B) Simulation 3. ROC plots with fixed $\lambda = 0.2$ and varying baseline length.

Figure 5. A schematic of the experimental task design for the fMRI study. See text for a detailed explanation.

Figure 6. (A) HEWMA group activation, $\vec{z}_{pop}$, in the medial frontal cortex and the estimated change-point for onset of activation (CP, green line). (B) Results of Monte Carlo simulations for finding corrected p-values. Black line: observed max T; distribution: null hypothesis max T. (C) Case weights calculated by taking the inverse of Eq. (28). Weights are based on variability during the baseline interval. Higher variance will result in a lower weight for that subject. (D) The individual time courses for the 24 subjects.

Figure 7. (A) Left: Regions with significant activations in HEWMA (corrected over time and false discovery rate (FDR) corrected at alpha = .05 over space; (Genovese, Lazar, & Nichols, 2002)). Right: Regions color-coded according to K-means classification (7 classes). (B) Histogram of number of voxels by CP and activation duration, color-coded by estimated class. (C) Axial slices of regions shown in (A).

Figure 8. The brain surface is shown in lateral oblique and axial views. Significant voxels are colored according to significance. Increases are shown in red-yellow, and decreases are shown in light-dark blue. (A) Time course from a region showing sustained activity in rostral medial PFC (single representative voxel). The baseline period is indicated by the shaded gray box, and the HEWMA-statistic is shown by the thick black line (+/- one standard error across participants, shown by gray shading). The control limits are shown by dashed lines. (B) A similar plot showing transient responses to presentation of task instructions in visual cortex.
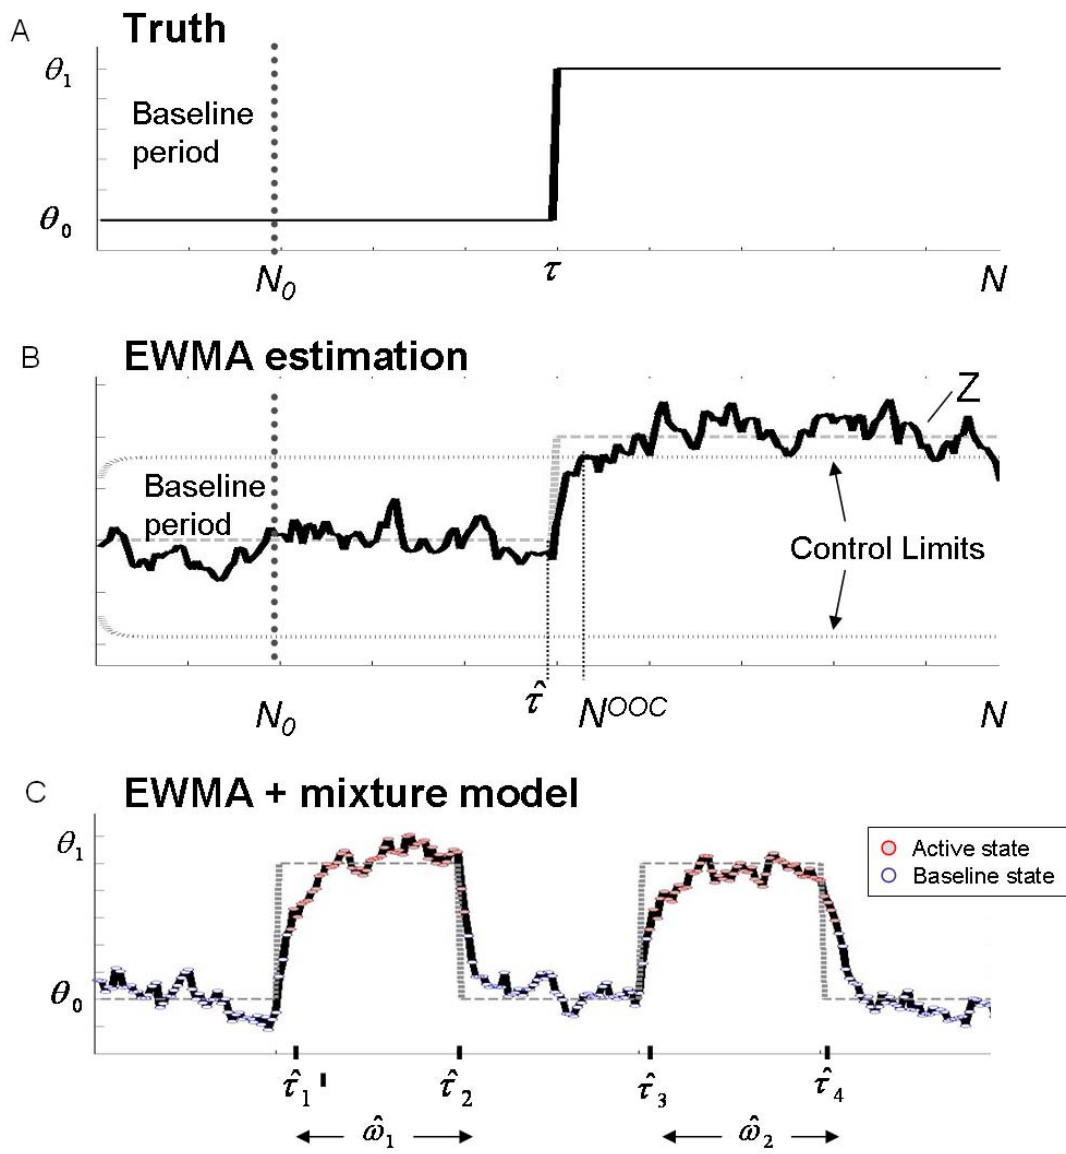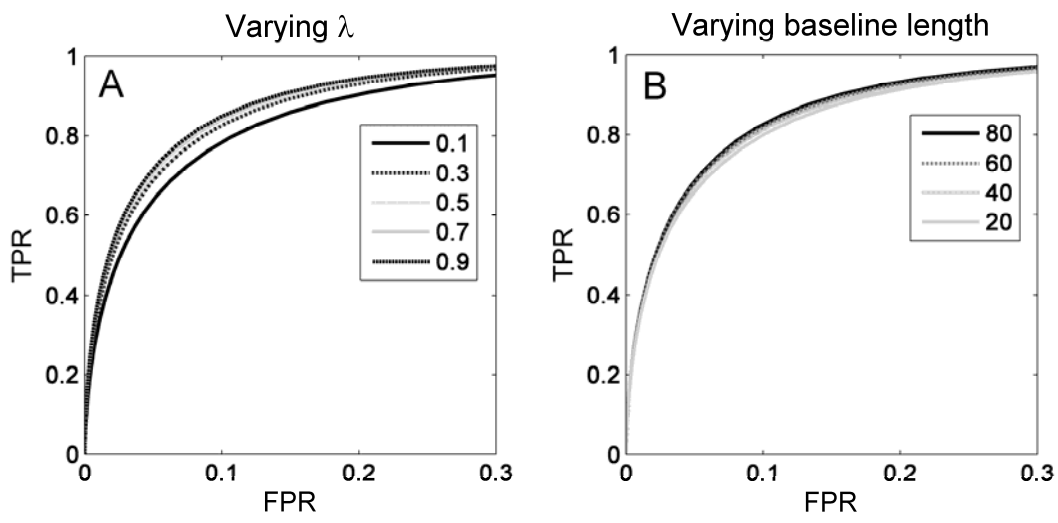
Figure1
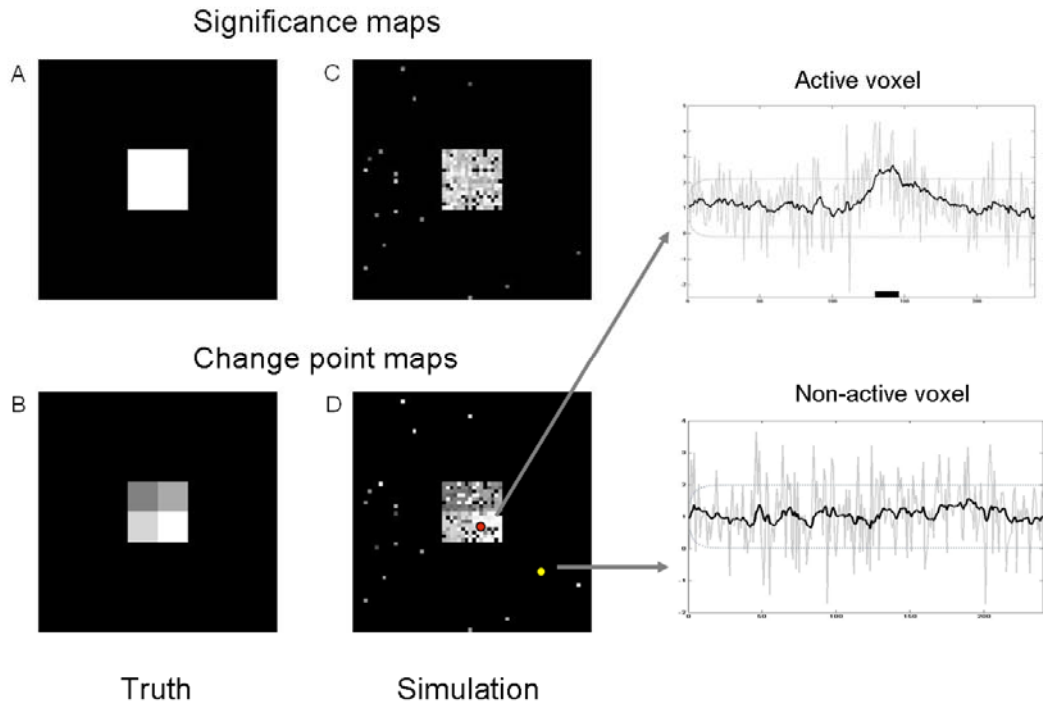
Figure2
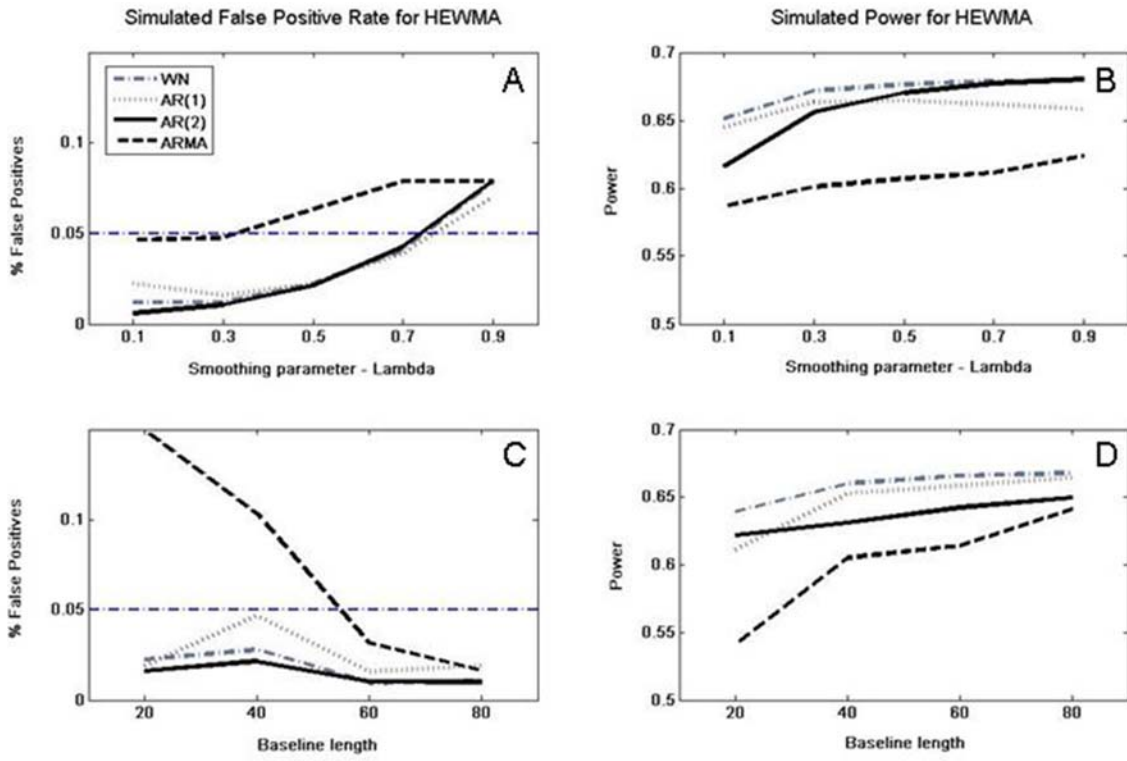
Figure 3

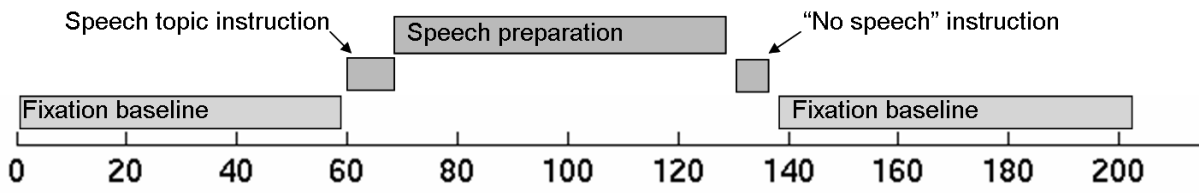**Figure 4**

Task design: Anxiogenic speech preparation task
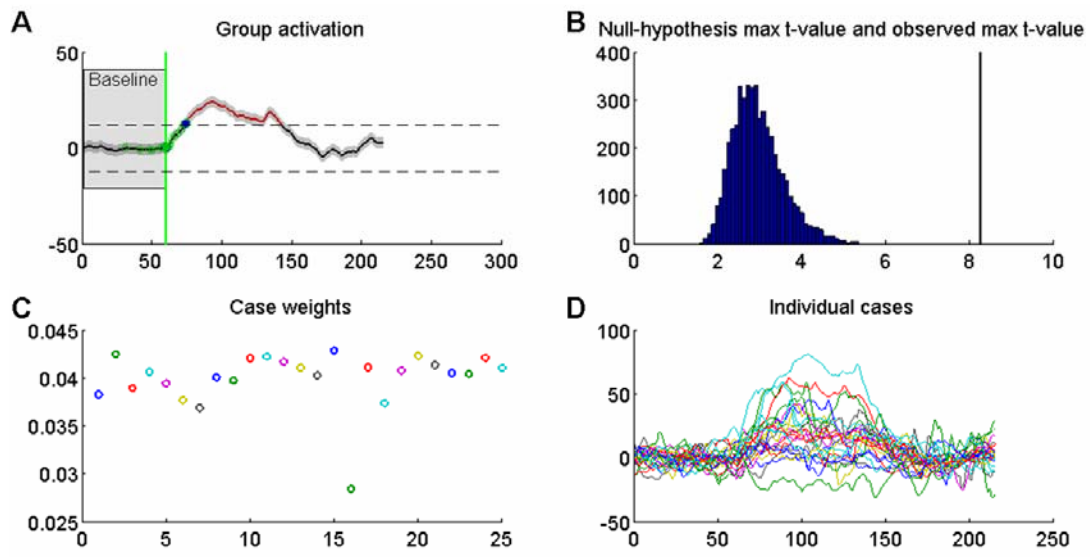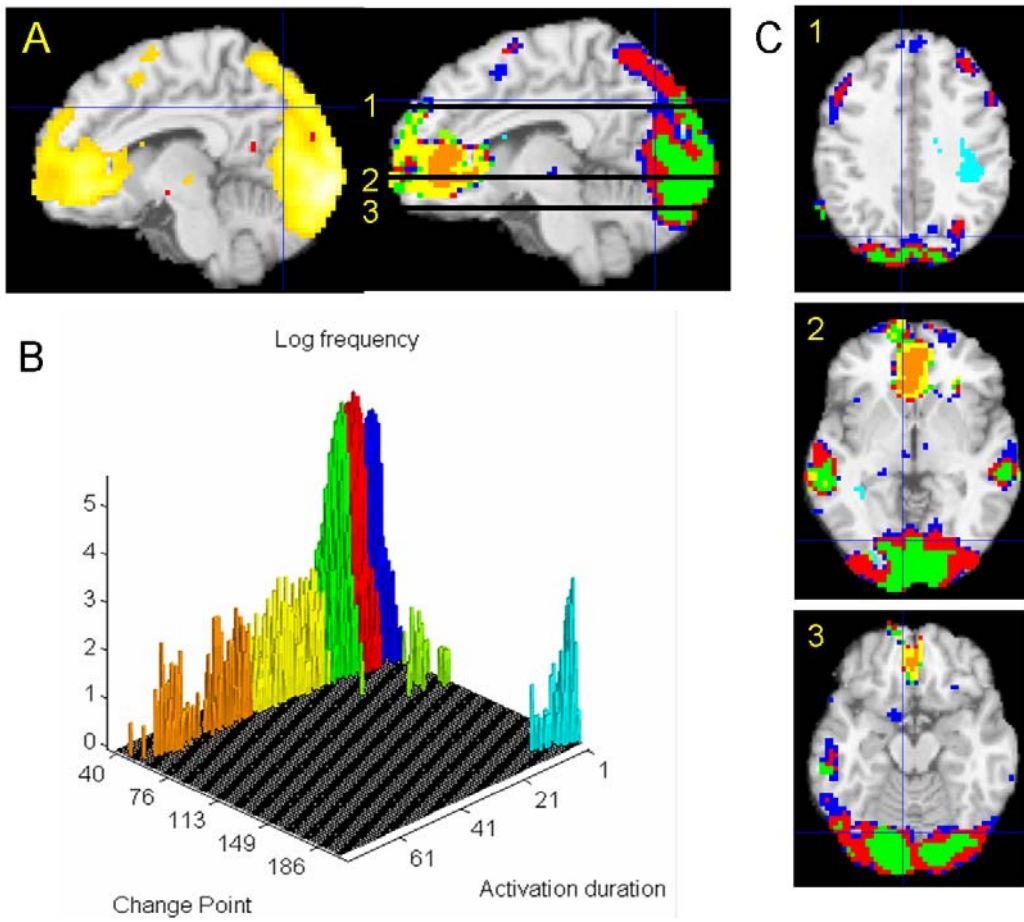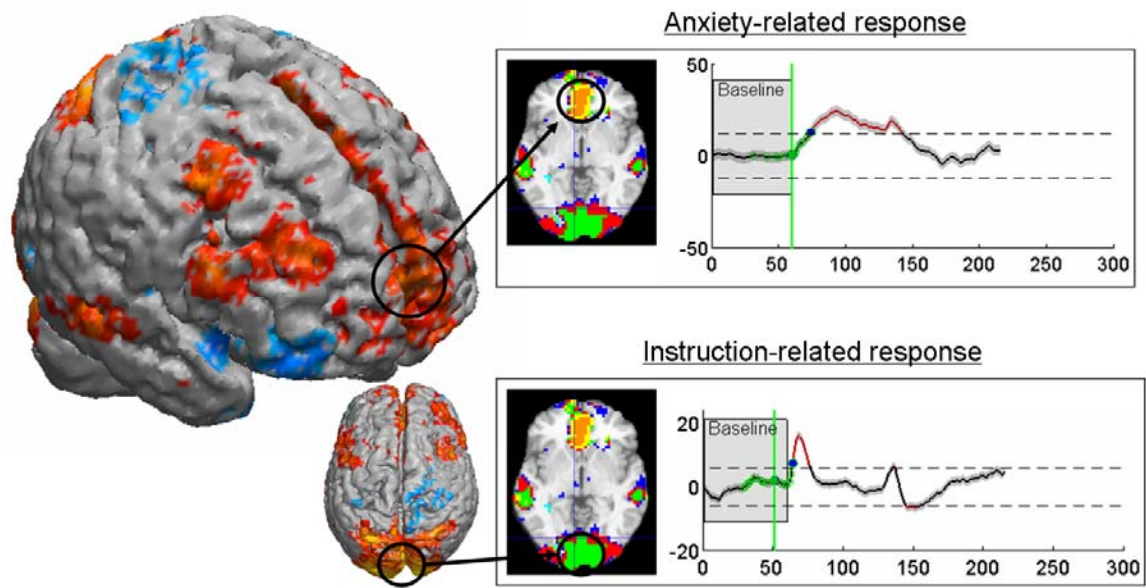

Figure 5

Figure 6

Figure 7

Figure8

Appendix A

EM-algorithm for Gaussian Mixture Model

The fMRI time course is modeled as a mixture of two normal distributions, with different means and variances: $X_0 \sim N(\theta_0, \sigma_0^2)$ and $X_1 \sim N(\theta_1, \sigma_1^2)$. We can write this as $X = (1 - \Delta)X_0 + \Delta X_1$, where the random variable Δ is equal to one with probability p and equal to zero with probability *1-p*. The density function of X can be written

$$f_X(x) = (1 - p)f_{X_0}(x) + pf_{X_1}(x).$$

where $f_{X_i}(x)$ is the normal probability density function with mean θ_i and variance σ_i^2. The log-likelihood can be written:

$$l(\theta_0, \theta_1, \sigma_0^2, \sigma_1^2, p \mid \mathbf{x}) = \sum_{i=1}^{n} \log\left(f_X(x_i)\right)$$

The unknown parameters $(\theta_0, \theta_1, \sigma_0^2, \sigma_1^2, p)$ that maximize this term can be found using the EM-algorithm.

Perform the following two steps until convergence:

E-step: $\quad \hat{\gamma}_i = \dfrac{\hat{p}f_{X_2}(x_i \mid \hat{\theta}_j, \hat{\Sigma}_j)}{(1 - \hat{p})f_{X_1}(x_i \mid \hat{\theta}_j, \hat{\Sigma}_j) + \hat{p}f_{X_2}(x_i \mid \hat{\theta}_j, \hat{\Sigma}_j)} \qquad$ for $i = 1, \ldots T$

M-step: $\quad \hat{p} = \dfrac{1}{T}\sum_{i=1}^{T}\hat{\gamma}_i, \quad \hat{\theta}_j = \dfrac{\sum_{i=1}^{T}\hat{\gamma}_i x_i}{\sum_{i=1}^{T}\hat{\gamma}_i}, \quad \hat{\sigma}_j = \dfrac{\sum_{i=1}^{T}\hat{\gamma}_i(x_i - \hat{\theta}_j)(x_i - \hat{\theta}_j)^T}{\sum_{i=1}^{T}\hat{\gamma}_i} \qquad$ for $j = 1, 2$

Appendix B

EM-algorithm for HEWMA

Let Λ be the lower triangular smoothing matrix (defined in Eq. 7) and $G = [I_n \; I_n \; \cdots \; I_n]^T$

an $mn \times n$ matrix where I_n is the $n \times n$ identity matrix. Further, let $\vec{z}_G$ be the combined vector of

EWMA-statistics for all m subjects and $V_G^* = \alpha Q_G + \Sigma_G$ its covariance matrix, where

$$Q_G = I_m \otimes \Lambda^T \Lambda \qquad \text{and} \qquad \Sigma_G = \begin{bmatrix} \Sigma_1^* & & & 0 \\ & \Sigma_2^* & & \\ & & \ddots & \\ 0 & & & \Sigma_m^* \end{bmatrix}.$$

In our implementation the within-subject component Σ_i^* is assumed known from the first level,

and Q_G is also assumed known. The unknown parameter α is estimated iteratively using the

EM-algorithm below. This is equivalent to the variance component estimation performed in

SPM2 (Friston et al., 2002).

Perform the following two steps until convergence:

E-step: $\qquad V_G^* = \alpha Q_G + \Sigma_G$

$\qquad\qquad C = \left(G^T V_G^{*-1} G \right)^{-1}$

M-step: $\qquad P = V_G^{*-1} - V_G^{*-1} G C G^T V_G^{*-1}$

$\qquad\qquad g = -\frac{1}{2} tr(PQ_G) + \frac{1}{2} tr\left(P^T Q_G P \vec{z}_G \vec{z}_G^T \right)$

$\qquad\qquad H = \frac{1}{2} tr\left(P^T Q_G P Q_G \right)$

$\qquad\qquad \alpha = \alpha + H^{-1} g$