Published in final edited form as:

Neuroimage. 2007 May 1; 35(4): 1480–1494.

Continuous carry-over designs for fMRI

Geoffrey Karl Aguirre

Department of Neurology, University of Pennsylvania

Abstract

This paper describes *continuous carry-over* fMRI experiments. In these studies, stimuli are presented in an unbroken, sequential manner, and can be used to estimate simultaneously the mean difference in neural activity between stimuli as well as the effect of one stimulus upon another (carry-over effects). Neural adaptation, which has been the basis of many recent fMRI studies, is shown to be a specific form of carry-over effect. With this approach, the adapting effects of stimuli may be studied in a continuous sequence, as opposed to within isolated pairs or blocks. Additionally, the average, direct effect of a stimulus upon neural response can form the basis of a simultaneously obtained distributed pattern analysis, allowing comparison of neural population coding on focal (within voxel) and distributed (across voxel) spatial scales. These studies are ideally conducted with serially balanced sequences, in which every stimulus precedes and follows every other stimulus. While m-sequences can provide this stimulus order, the *type 1 index 1* sequence of Finney and Outhwaite may be used in fMRI studies for those experimental designs for which an m-sequence solution does not exist. Continuous carry-over designs with serially balanced sequences are argued to be particularly well suited to the characterization of “similarity spaces,” in which the perceptual similarity of stimuli is related to the structure of neural representation both within and across voxels. These concepts are illustrated with a worked example involving the neural representation of color. It is shown that data from a single scanning session is sufficient to detect direct and carry-over effects, as well as demonstrate the correspondence of the similarity structure of distributed patterns of neural firing and the perceptual similarity of a set of colors.

Introduction

Many neuroscience experiments examine the relationship between variation in a stimulus and the corresponding neural response. For example, an experiment might measure the stimulus response function that relates the brightness of a flash of light to the magnitude of neural activity. More complex variation in a set of stimuli is possible, and may be expressed as changes along component axes. Thus, a set of rectangles may be defined by their length and width, or musical notes may vary in tone, duration and amplitude. The perceptual similarity of pairs of stimuli will be related, sometimes in a complex manner, to changes along these different stimulus axes. A general question is if stimuli that are perceived to be similar are represented by neural states that are themselves similar, and what form that these neural responses take (Edelman, 1998). Naturally, different types of neural codes, expressed in different cortical regions, at different spatial scales, might be used to represent aspects of stimulus variation. Blood oxygen level dependent (BOLD) functional magnetic resonance imaging (fMRI) has the potential to measure several types of neural representation, and thus provides a method of examining the mapping between neural and perceptual similarity spaces.

Correspondence: Geoffrey Karl Aguirre, Department of Neurology, University of Pennsylvania, 3 West Gates, 3400 Spruce Street, Philadelphia PA 19104, Tel: 215-662-3390, Fax: 215-349-5579, aguirreg@mail.med.upenn.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

In traditional fMRI studies, the average BOLD response to multiple presentations of a given stimulus is measured to determine the direct effect of stimulus variation upon the amplitude of neural response. This provides a stimulus response function that relates modulation of a stimulus to the average response across a population of neurons within a voxel. In contrast, adaptation fMRI (Grill-Spector and Malach, 2001, Henson and Rugg, 2003) is used to support inferences regarding neural population codes within voxels. Based upon the habituation of neural responses seen in electrophysiology studies (Desimone, 1996), an adaptation fMRI experiment presents pairs or blocks of stimuli that are the same or different along a dimension of theoretical interest. A reduction in the magnitude of response to the stimuli that share a feature is evidence that the identified voxel contains a population of neurons that code the stimulus dimension. Finally, distributed pattern analysis (Haxby et al., 2001), examines the structure of stimulus representation across voxels. Although composed of intermixed and complementary populations of neurons, voxels may have a slight bias towards greater responses for particular stimuli. While the signal from a single voxel may be insufficient (in the face of noise or dimensionality of the response) to differentiate between behavioral states, the aggregate information across a population of voxels may allow the classification of neural responses according to category of stimulus.

These different modes of fMRI inference yield complementary information regarding the neural representation of stimuli and have traditionally been obtained as part of separate experiments. It may be advantageous, however, to measure these effects simultaneously in an fMRI experiment both for the sake of efficiency as well as for the opportunity to examine the relative contribution of different forms of neural coding to the representation of stimulus variation. In this paper I describe a class of fMRI experiments, termed *continuous carry-over* designs, that facilitate efficient, simultaneous measurement of these different types of neural representation. This type of design is particularly well suited to the study of complex stimulus similarity spaces.

A carry-over effect is the modulation of the neural response to the current stimulus by the previously presented stimulus. This is in contrast to the direct-effect—the average amplitude of neural activity to a stimulus independent of its context—that is the basis of a stimulus response function within a voxel or distributed pattern analysis across voxels. Studies of neural adaptation are studies of a form of carry-over effect, often constrained to examine the effect in isolated pairs of stimuli (Kourtzi and Kanwisher, 2000). The chief insight of the current work is that the neural adaptation effects that have been examined in these fMRI studies can be observed not only in the paired presentation of stimuli, but in the continuous modulation of response to stimuli presented in an unbroken stream. In this fashion, fMRI studies of stimulus perception are brought closer to approaches used in the study of single-unit electrophysiology (Schwartz et al., 2002).

A second insight, not unique to the current study (Buracas and Boynton, 2002), is the importance of the particular order in which the stimuli are presented. Efficient and unbiased estimation of both the direct effect of a stimulus upon the amplitude of neural response and the carry-over effect induced by the previous response is best achieved using *serially balanced sequences*, which are constructed so that the presentation of every stimulus follows every other stimulus.

In the theory section of this paper, we will consider a simple model for both the direct effect of a stimulus upon neural response, and the effect of context that the preceding stimulus imposes on the current stimulus. Neural adaptation is defined as a specific form of more general neural carry-over effects. Other forms of carry-over effect might be sought, and in this paper we will examine a simple model of gain-control, termed here *bias*. We will then consider types of serially balanced sequences that are suitable for use in a continuous carry-over design. While

m-sequences (Buracas and Boynton, 2002) are found to be appropriate, we will consider an alternative balanced sequence for fMRI (the Finney and Outhwaite *type 1 index 1* sequence) that shares several important properties with m-sequences, but can be used for stimulus spaces defined by a (nearly) arbitrary number of stimuli.

The feasibility of this approach is then demonstrated with an example BOLD fMRI study of the neural representation of color. The perception of color by human observers can be described within a three-dimensional space, such as the red-green-blue space of color digital displays and the cyan-magenta-yellow space for color reproduction in print. Here we make use of the three dimensional, OSA-UCS color space which has well defined perceptual properties, further considered below. The neural response to a set of 27 color stimuli was studied using a continuous carry-over approach with a serially balanced sequence. The perceptual similarity of different colors was found to be reflected by both within-voxel adaptation effects and the distributed pattern of responses across the cortex. Additional influences of the stimuli were found, including a direct effect of color, in which cortical regions were found to respond more strongly to certain locations in color space, and a non-symmetric bias effect (orthogonal to adaptation), in which the response to the lightness of a color was influenced by the lightness of the preceding color.

Theory

Carry-over effects

The response of a system to a stimulus may reflect both the current stimulus and the past history of stimuli and responses (Finney and Outhwaite, 1956). In modeling the spike response of a neuron to trains of (stochastically arranged) stimuli, one can measure both the linear “excitatory” effects that stimuli have upon neural responses, as well as the non-linear “gain control” effects, expressed predominantly as a tendency for stimuli to suppress the likelihood of firing over short time scales (Schwartz et al., 2002). With thousands of observations at high temporal frequency, and (in most cases) a binary input and response measure, spike-triggered covariance approaches can model high-order relationships between input and response in time. Beyond characterizing the effect of prior stimuli upon the response to the current stimulus, recurrent neural network models are used to characterize the response to the current stimulus as a function of the previous activation state of the network, which is influenced both by previous stimuli and the interactions between nodes of the network (Botvinick and Plaut, 2004). With sufficient model accuracy, one could predict the precise activation state of a neural system after any arbitrary pattern of stimuli, as well as the expected neural response to the next stimulus.

For measurements obtained using BOLD fMRI, model simplification is necessary. These limitations relate to the integration of the BOLD fMRI signal over the responses of millions of neurons within a voxel and over several seconds of time. Here, we consider models analogous to those of behavioral studies of sequential magnitude estimation. Such studies attempt to discern the *direct effect* of the current stimulus upon a subject’s response, as well as the *carry-over effect* of previous stimuli, where the effect of the previous stimulus is manifest in its relationship to the current stimulus. We can formalize this relationship using the notation of Ferris and colleagues (Ferris et al., 2001), and consider the voxel-wise neural response Y_{it} to stimulus i (of n possible stimuli) at sequence position t , and restrict the sequential dependence of the response to influences from the previous $t-1$ stimulus, j :

$$Y_{it} = \mu + \tau_i + \phi_{ji} \quad (1)$$

where μ is the overall mean response, τ_i is the direct effect of the i th stimulus, and ϕ_{ji} is the carry-over effect of stimulus j on stimulus i . Note that τ represents the stimulus effect that has

traditionally been studied in fMRI experiments. If the estimated τ varies across stimuli, then one could create a stimulus-response function that relates stimulus variation to amplitude of neural response. If the differences in τ between stimuli are quite subtle, it still may be possible to distinguish the neural response to stimuli by combining the direct effect measurements across voxels, and performing a distributed pattern analysis.

Two other aspects of this simplified model are worth noting initially. First, the carry-over effect is modeled as proportional to the previous stimulus (j) as opposed to the previous response (τ_j). While it is theoretically possible to construct models of fMRI data that relate the instantaneous state of brain activation to the response to the next stimulus, these are not suitable for modeling the within-voxel population code for stimulus representation as is sought by neural adaptation experiments. This is because the fMRI signal from a voxel is integrated over the firing of millions of neurons, making the underlying population code inaccessible to direct measurement. This may lead to the circumstance where τ for a given voxel is the same for all stimuli, yet carry-over effects are still present. Consequently, the model uses the stimuli themselves (or the behavioral responses to them) as a proxy for the activation state of the system prior to the presentation of the next stimulus.

Second, the carry-over effect is modeled only for the previous stimulus, as opposed to the previous two, three, or more stimuli. This is not a necessary property of the model, but is a reasonable simplification for the experimental parameters and stimuli considered in this paper, for which behavioral carry-over effects are much larger for the immediately preceding stimulus compared to more distant stimuli (Gescheider, 1988). If a small number of stimuli are presented rapidly and with appropriate higher-order counterbalancing (discussed below), then the model could be expanded to represent the effect of earlier stimuli.

We now consider the form that describes the carry-over effect, ϕ . The carry-over effect is fully described by an $n \times n$ matrix that indicates the modulatory effect upon neural response of the transition from any stimulus to any other stimulus. Figure 1 presents an example ϕ matrix for a set of 16 stimuli (perhaps 16 notes of increasing pitch, or 16 tactile stimuli of increasing temperature). To estimate the ϕ matrix, an experiment might simply measure the effect of every stimulus preceding every other stimulus. In practice, model parameterization can increase statistical and explanatory power. Specifically, the ϕ matrix may be modeled by a set of basis matrices that have useful interpretations. I offer here one approach in which the primary carry-over effects of interest are proportional to the difference in intensity of sequential stimuli, although other parameterizations are certainly possible.

First, there may be some component of the ϕ matrix that is symmetric about the diagonal axis. Such symmetry implies that this component of the carry-over effect is the same regardless of the direction of the transition between stimuli. For example, the effect upon neural activity is the same whether stimulus i_3 precedes stimulus i_5 or i_5 precedes i_3 . We may further model this effect as being linearly proportional to the perceptual difference between sequential stimuli. We will term this linear effect λ_1 , and Figure 1 presents this and subsequent elements of the basis set of matrices. In a manner analogous to the modeling of the direct effect of parametric stimulus variation upon magnitude of neural response (Buchel et al., 1998), non-linear relationships between sequential stimulus differences and symmetric carry-over effect can be modeled using an orthogonalized polynomial expansion. Thus, second-order (quadratic) and third-order (cubic) relationships are modeled by λ_2 and λ_3 . These effects can also be expressed as:

$$\phi_{ji} \approx \lambda_1 |j - i| + \lambda_2 |j - i|^2 + \lambda_3 |j - i|^3 \dots + \kappa_{j=i} \quad (2)$$

The element κ represents the additional refinement that identical pairs of stimuli are modeled separately from sequential stimuli that are different, even if only slightly. The motivation for this decomposition is that different mental operations may attend the perception of perfect stimulus repetition (e.g., as observed in Gomez et al., 2000).

Symmetric carry-over effects of this kind can be related to the notion of neural adaptation that has been studied in some fMRI experiments. Consider a voxel that represents the identity of stimuli by the pattern of activity within an ensemble of neurons. If the present stimulus is represented by a set of neurons that overlaps the pattern evoked by the previous stimulus, the neural response may be reduced in the sequentially activated neurons. The degree of habituation might be maximal for identical stimuli, middling for similar stimuli, and minimal for completely different stimuli. The presence of graded habituation for a particular dimension of stimulus variation is evidence that the identified cortical area represents that stimulus dimension by population code.

We may test the idea that the degree of symmetric adaptation observed is proportional to the perceptual (or conceptual) distance between sequential stimuli. This perceptual distance may be calculated from a fundamental stimulus property (such as the angle of rotation in degrees for stimuli that vary in orientation). For some stimulus sets, however, the perceptual distance between pairs of stimuli will be a complex function of changes in one or more parameters that define the stimulus space. In such cases, the perceptual effect of a transition from one stimulus to another is measured behaviorally, and expressed within an $n \times n$ similarity matrix, which relates the perceptual similarity of any one stimulus to any other. Figure 2 presents perceptual similarity matrices for several types of stimulus spaces. The correspondence between the similarity matrices shown in Figure 2 and the symmetric carry-over effect λ_I is clear: Changes in neural activity within a voxel that are well modeled by the carry-over effect λ_I implies the presence of a neural population code in which the degree of overlap in response pattern to one stimulus and another is linearly related to the perceptual similarity of those two stimuli.

There may be other forms of carry-over effect that are not symmetric across stimulus transitions. In behavioral studies of semantic priming, the priming effect can either be symmetric (*salt* $\rightarrow$ *pepper* is the same as *pepper* $\rightarrow$ *salt*), or asymmetric (*stork* $\rightarrow$ *baby* is not the same as *baby* $\rightarrow$ *stork*), and can be biased by the larger semantic context (Thompson-Schill et al., 1998). The carry-over effect of such stimuli could be examined as above, but additional matrices that do not impose symmetry across stimulus transitions would be needed.

Within the realm of early visual perception there are other, non-symmetric forms of carry-over effect that may be of interest and are proportional to the intensity differences between stimuli. Studies of magnitude estimation have identified the phenomenon of perceptual bias, which is closely related to gain-control. While performing magnitude estimation, subject responses to the current stimulus tend to be biased towards the previous response (Gescheider, 1988). That is, if the previous stimulus was more intense than the current stimulus, then the subject tends to score the current stimulus as more intense. This effect is asymmetric, in that the transition from stimulus i_3 to stimulus i_5 is not the same as for i_5 to i_3 . Various models for this carry-over effect have been proposed (DeCarlo and Cross, 1990), but we will consider here a model in which the effect is proportional to the magnitude difference between the prior and current stimulus:

$$\phi_{ji} \approx \rho(j - i) \quad (3)$$

When $\rho > 0$, the bias is assimilative (the response to stimulus i is larger when the previous stimulus was more intense), and when $\rho < 0$, then the carry-over bias is contrastive.

This asymmetric bias effect, and the symmetric effects considered above, are only a small subset of the possible orthogonal basis matrices that might be used to model carry-over effects. This particular set was selected because of its ready interpretation in the context of an example study of color perception that is detailed below. A further expansion of these elements, or a different structuring of the space entirely, may be appropriate for other experimental questions (e.g., the use of Volterra kernel expansions; Friston et al., 2000).

It should be noted that non-linearities in the transformation of neural activity to BOLD signal may be mistaken for some neural carry-over effects. Non-linear BOLD fMRI signal responses to changes in the timing and duration of stimulus input have been observed and quantified (Friston et al., 2000). The bias effect described above would be sensitive to, for example, a period of decreased vascular reactivity that follows a large amplitude response. This vascular non-linearity would be mistaken for contrastive neural bias. With appropriate model construction (and the use of serially balanced sequences), it is possible to dissociate the fast effects of potential neural non-linearities from the slower hemodynamic response (Kellman et al., 2003). Notably, some forms of carry-over effect are not sensitive to non-linearities in the hemodynamic response. Specifically, a non-linear vascular response cannot explain symmetric adaptation effects.

The presence of these direct and carry-over effects in the neural response to stimuli can be assessed within the context of the General Linear Model (Worsley and Friston, 1995):

$$Y_t = KG_{mt}\beta_m + e \quad (4)$$

where Y_t is a BOLD fMRI signal over time, G is an $m \times t$ matrix of covariates, β are the m unknown neural response parameters, K is a convolution matrix representing the hemodynamic response function, and e is the residual error of the model. Covariates of interest are created by modeling the expected neural response as a function of the direct and carry-over effects of the stimuli, as expressed in the restricted basis set of matrices. Through appropriate construction of the G matrix, the β values are rendered proportional to the effects of μ , τ , ρ , λ , and κ in the neural responses. We will presently consider the construction of G for an example experiment, but turn now to the ordering of the i stimuli over time.

Serially balanced sequences

In a continuous carry-over design, each stimulus follows the prior at a steady rate. Our goal is to measure both the average response to each stimulus type (the *direct effects*) as well as the influence of a stimulus on the following stimulus (*carry-over effects*, such as adaptation). When carry-over and direct effects are rendered orthogonal, both may be assessed efficiently and without bias. Direct and carry-over effects are orthogonal when the order of presentation of stimuli is serially first-order balanced, meaning that each stimulus is preceded equally often by every other stimulus (including self-adjacencies). If an experiment has n stimuli (labeled from $0 \dots n-1$), a first-order, serially balanced sequence n times, and have a minimum length will present each stimulus of n^2 (although one additional, discarded element is needed to precede the first stimulus of the sequence). Note that a baseline or null-trial condition is also treated as a stimulus, and cannot simply be added at random as the baseline trials themselves constitute a context that may alter the response to the stimuli that follow.

Maximum length sequences (*m-sequences*) are (nearly) serially-balanced. An m -sequence is a pseudo-random ordering of integers that is generated using maximal linear feedback shift registers (Buracas and Boynton, 2002). An m -sequence of length n^2-1 provides first-order counterbalancing of the stimuli, although not perfectly as it excludes an instance of the stimulus label zero and the adjacency of stimulus label zero with zero. Higher-order counterbalancing is possible, as more generally an m -sequence of length n^r-1 provides r -order counterbalancing

(although again excluding the sub-sequence in which label zero appears r times in a row). M-sequences have (essentially) no temporal structure that a subject might use to anticipate the identity of the next stimulus. A limitation, however, is that m-sequences are not available for all experiments with n stimuli. M-sequences only exist for n equal to the power of a prime integer (and in practice only when that integer is equal to 2, 3 or 5; Buracas and Boynton, 2002), yielding m-sequences for $n=2, 3, 4, 5, 8, 9, 16, 25, 27, 32, \dots$. For designs with a different number of stimuli, some other approach is needed.

An alternate serially balanced sequence is the *type 1 index 1* sequence in the classification of Finney and Outhwaite (Finney and Outhwaite, 1956). Although described several decades ago, recent work by Nonyane and Theobald (In press) makes creation and selection of these sequences feasible. For these sequences, all n treatments are presented in n “blocks” of different permutations of the ordering of the treatments, with the stimulus repeating at the termination of one block and the beginning of the next. No *type 1 index 1* sequences exist for n equal to 3, 4 or 5, but sequences have been found for n between 6 and 34 (Nonyane and Theobald, In Press). An example sequence provided by Nonyane and Theobald for $n=6$ is

0 012345 502143 304152 240351 132054 425310 (7)

This sequence might indicate the order in which 6 different stimuli (or 5 stimuli and a null-condition) are to be presented during an fMRI study. Note that the initial treatment is duplicated (leading to a sequence that is n^2+1 in length) so that all self-adjacencies are present, and that the first and last treatments are the same. Importantly, while “blocks” of stimuli may be discerned within the stimulus sequence, there is no requirement that the presentation of the stimuli be visibly divided into different blocked periods. Also, the assignment of labels to stimuli is arbitrary, thus the initial ordering of labels (0 1 2 3...) does not imply that the stimuli themselves are presented in an apparently ordered fashion.

For BOLD fMRI experiments with relatively few stimulus types (i.e., $n < 10$), m-sequences are generally preferred when they are available. Assuming a stimulus presentation time on the order of 1 second, it is possible to present m-sequences with second-order and higher counterbalancing within a reasonable period of data collection (i.e., less than an hour).

For many experiments with 10 or more stimuli there are few m-sequences available and even if available may not be preferable to *type 1 index 1* sequences. In most cases, only first-order counterbalancing of the stimuli will be achievable during the study period in a single subject, which both m and *type 1 index 1* sequences provide. Additionally, because of the “blocking” of the stimuli into permuted sets, *type 1 index 1* sequences are insensitive to order effects across the entire period of stimulus presentation. In other words, relatively less experimental variance will be present below the fundamental frequency of repetition of stimuli, unlike m-sequences that distribute power over all frequencies. This is desirable for BOLD fMRI studies as temporally autocorrelated noise renders low-frequency changes in the signal difficult to detect (Zarahn et al., 1997). Finally, *type 1 index 1* sequences provide perfect serial balance, whereas m-sequences provide very close but imperfect balance. Despite this, an m-sequence (if available) may nonetheless be preferable if it is necessary to minimize any ability of the subject to anticipate the structure of the stimulus order. Although the order of stimulus presentation within “blocks” is highly stochastic, there is some minimal regularity in a *type 1 index 1* sequence that might be detectable by the subject (each stimulus is guaranteed to appear twice within $2n-2$ trials, and there is a stimulus repetition every n trials).

Nonyane and Theobald have recently described a method (Nonyane and Theobald, In Press) to systematically generate all *type 1 index 1* sequences for a given n . Many *type 1 index 1* sequences exist for a given n : there are 624 for $n=6$; 175,588 for $n=7$ (Nonyane and Theobald,

In Press); and the number for larger sets of stimuli is likely vast. Different sequences vary in their suitability for use in an fMRI experiment, and in Appendix A I consider some metrics by which *type 1 index 1* sequences might be chosen for use in a BOLD fMRI study.

We turn now to an example of a continuous carry-over BOLD fMRI experiment, conducted with a serially balanced sequence. The experiment examines the neural representation of color, as manifest in 27 different colored stimuli (plus a null-trial). As no *m*-sequences exist for $n=28$, a *type 1 index 1* sequence was selected that would be maximally sensitive to carry-over adaptation effects that are linearly related to the perceptual similarity of the color stimuli, as expressed in Figure 2D. The specific sequence used, which is appropriate for the study of other 3-dimensional stimulus spaces composed of 27 stimuli (and a null-trial), is presented in Appendix B.

Methods

Subject and scanning parameters

A 28 year-old, right-handed woman participated in the study. Structural and functional data were collected on a 3.0 Tesla Siemens Trio scanner using an 8-channel head coil. High-resolution T1-weighted structural images were collected in 160 axial slices and near isotropic voxels ($0.9766 \text{ mm} \times 0.9766 \text{ mm} \times 1.0000 \text{ mm}$; TR = 1620 msec, TE = 3 msec, TI = 950 msec). Functional, blood-oxygenation-level-dependent (BOLD), echoplanar data were acquired in 3 mm isotropic voxels (TR = 3000, TE = 30). BOLD data were acquired in 42 axial slices, in an interleaved fashion with 64×64 in plane resolution. The functional data were collected in 6 runs of 136 TRs each. The first 6 seconds of each run consisted of “dummy” gradient and radio frequency pulses to allow for steady state magnetization during which no stimuli were presented and no fMRI data collected.

Stimulus sequence and accounting for overlap at scan boundaries

During scanning the subject was presented with stimuli in the order specified by the $n=28$, *type 1 index 1* sequence shown in Appendix B. Because the scans were shorter than the total duration of stimuli to be presented, it was necessary to break the complete sequence into sections corresponding to 7 minute scans. The scan boundaries, however, hamper the accurate measurement of the hemodynamic response to the stimulus sequence. At the start of a scan, several seconds are required for the BOLD response to build to a “steady-state”, and at the end of a scan, the hemodynamic response to the last presented stimuli continues after data collection has ceased. These effects are eliminated by re-presenting 10 elements of the sequence (15 seconds worth) at the initiation of each scan, and then trimming those periods from the data in subsequent processing. In the case of the first scan, stimuli from the end of the sequence are presented first. Because this trimming applies to the start and end of the scanning session, the initial, extra duplication stimulus can be dropped, and the final fMRI data represent the entire sequence of stimuli, with the delayed and dispersed BOLD fMRI response to the neural sequence “wrapping around” from end to start.

Stimuli and behavioral task

The stimuli consisted of a full screen presentation of one of the 27 colors depicted in Figure 2D, or a black (null-trial) screen. The stimuli were back-projected onto a screen viewed by the subject through a mirror mounted on the head coil, and were $27^\circ \times 18^\circ$ in size. Each stimulus was presented for 1450 msec, with a 50 msec ISI consisting of a black screen. During scanning, the subject was instructed to monitor a small fixation dot in the center of the screen. The dot was randomly assigned to appear either black or white on each trial, and the subject was to indicate the state of the fixation dot by button press. On null trials, for which the screen was black, the fixation dot was always white, and the subject was instructed to indicate the state of

the fixation dot as usual. The fixation task was designed to assist and monitor subject alertness, and to reduce the explicit comparison of one stimulus to the next. As the state of the fixation dot was random with respect to the color present on the screen (with the exception of null-trials), the fixation task was essentially orthogonal to the main manipulations of interest. We presented this cycle of stimuli twice to the subject, for a total of 40 minutes of scan data.

Image pre-processing

Off-line data analysis was performed using VoxBo (www.voxbo.org) and SPM2 (<http://www.fil.ion.ucl.ac.uk/>) software. Data were sinc interpolated in time to correct for the slice acquisition sequence, motion corrected with a six parameter, least squares, rigid body realignment routine using the first functional image as a reference, and normalized in SPM2 to a standard template in Montreal Neurological Institute (MNI) space. Normalization maintained 3mm isotropic voxels and used 4th degree B-spline interpolation. Two analyses of the data were conducted. In the first, which examined “focal” effects, the fMRI data were smoothed in space with a 1.5 voxel isotropic, Gaussian kernel. In the second, which examined “distributed” patterns of neural activity, the data were left unsmoothed. For each data set (the spatially smoothed and unsmoothed), the average power-spectrum across voxels and across scans was obtained, and the (square root) of the power spectrum fit with a 1/frequency function (Zarahn et al., 1997). This model of intrinsic noise was used during regression analyses with the Modified General Linear Model (Worsley and Friston, 1995) to inform the estimation of intrinsic temporal autocorrelation.

Anatomical data from the subject were processed using the FMRIB Software Library (FSL) toolkit (<http://www.fmrib.ox.ac.uk/fsl/>) to correct for spatial inhomogeneity and to perform non-linear noise-reduction. The resulting anatomical image was then processed within the BrainVoyager package (BrainInnovation;<http://www.brainvoyager.de/>) to identify the gray-white cortical boundary and produce inflated cortical surfaces for data presentation.

Statistical Analysis of “focal” effects

The purpose of the first analysis was to identify focal, voxel-wise signal changes that reflect direct and carry-over effects of the color stimuli. The design matrix used for this purpose is shown in Figure 3. Each covariate was constructed initially as a model of neural activity that was then convolved with a standard HRF (Aguirre et al., 1998). The following covariates were included:

- μ Models the mean response to any color stimulus versus the null trials.
- τ_G, τ_M, τ_B Models the direct, linear effect of stimulus variation along the three cardinal axes of the color space. For example, if the neurons within a voxel respond more strongly to green stimuli than red, a positive loading upon the τ_G covariate would be expected.
- λ_0 Models the carry-over effect of a stimulus following a blank trial. Because the similarity of any color to a blank screen was undefined for our stimulus set, these trials were modeled separately.
- κ Models the difference between any color stimulus that follows a different color, and color stimuli that follow an identical color. This allows identical trials to be modeled separately from different trials, as the similarity metric may have a discontinuous value for identical stimuli relative to different.
- $\lambda_1, \lambda_2, \lambda_3$ The first, second, and third-order carry-over adaptation effects of a color following another color. Those trials in which a color followed another, nonidentical color were identified, and the value of the covariate for that trial assigned the magnitude of the Euclidean distance between the two colors in the stimulus space (Figure 2D). These

values were then mean-centered, and all other time-points in the covariate set to zero, resulting in a covariate that models the linear, symmetric carry-over effect for colors following other, different colors. Quadratic and cubic covariates were created by raising the linear covariate to the second and third power, and orthogonalizing each covariate with respect to the others in the adaptation set (Buchel et al., 1998). An alternative formulation of the adaptation covariates, not used here, would decompose the effects along three cardinal axes, yielding first and second-order models of adaptation in the G, J and L directions.

- ρ** Linear, carry-over bias effect for changes in Lightness. For each trial in which one color followed another, the signed difference in Lightness was obtained. The L direction was selectively modeled as previous studies have shown carry-over bias effects for this stimulus property (Beck, 1966). A positive loading upon ρ indicates assimilative bias, in that the response to a stimulus is stronger when preceded by a lighter stimulus.

Not shown in Figure 3 are covariates of no-interest that modeled the intercept, scan effects, and the fixation task.

Functional regions of interest (ROIs) were defined using a “main effect” contrast, which consisted of the μ and λ_0 covariates evaluated jointly with an F-test. Tables 1 and 2 present the collinearity of the covariates of interest with each other before and after convolution with an HRF. In the initial model of neural response, there is essentially no correlation between the covariates used to define the main effect result for region of interest analysis (μ and λ_0) and the other covariates of interest. This is necessary for the validity of the analysis approach taken here, in which a main effect is used to define functional regions of interest within which interaction terms are assessed. There is substantial correlation (0.71) between the neural models of the direct and carry-over bias effect in the L color-space direction (τ_L and ρ). This correlation exists in the high frequency alternation between stimuli, as stimuli with a large Lightness value will also tend to have values larger than the preceding stimulus. This correlation would be reduced for stimulus spaces with more magnitude levels along a dimension. After convolution with the HRF, the correlation between these two covariates is substantially decreased (0.28), as a consequence of attenuation of this high frequency alternation. Convolution also acts to render the mean response (μ) and carry-over from blank trials (λ_0) highly collinear (-0.78). This is of no practical consequence as these two covariates were evaluated jointly with an F test in region definition.

Further expansion of the covariates included here could be considered. For example, higher-order representations of direct effects (i.e., quadratic forms for τ_G , τ_J , and τ_L), interactions between these color directions, and carry-over bias models (ρ) for other color-space axes. Not all possible effects were included in the model for two reasons. First, collinearity of the covariates (particularly after convolution) would be increased and, second, a restricted set of covariates were selected to allow focus upon key topics in this example dataset.

The “main effect” contrast was evaluated at a map-wise level, and thresholded to allow identification of separable foci of activity. Discrete patches of supra-threshold response were identified on the cortical surface representation and labeled by their presumed correspondence to visual areas by anatomical location. The voxels corresponding to these patches of interest were identified, the average imaging signal from each region obtained, and within-region statistical tests of (essentially orthogonal) covariates of interest conducted. For ease of presentation, only the data from the left hemisphere are considered for this example experiment.

Statistical Analysis of “distributed” effects

In this analysis, the similarity of patterns of neural activity evoked by different stimuli was assessed. The functional data were not smoothed in space. A design matrix (Figure 3, right side) that modeled the mean response to each color across all 54 presentations versus the null-trials (blank screen) was used. These covariates are simply delta functions placed at the time of occurrence of the presentation of a particular stimulus, and then convolved with the HRF.

A set of 1000 voxels were then selected from within the posterior visual areas identified by the focal, “main effect” analysis. For each of the 27 stimuli, the beta value for the corresponding covariate was obtained for each of the selected voxels. The average vector of stimulus beta values across voxels was obtained, and subtracted from each voxel vector. A neural similarity matrix was then constructed by calculating the correlation between the vector of beta values across voxels for each stimulus and every other stimulus. As the matrix is symmetric about the diagonal, only the lower triangle was retained, and the diagonal elements were excluded (as these have an obligatory value of unity and are thus uninformative). The correlation between the resulting neural similarity matrix and a perceptual similarity matrix (Figure 2D), and its decomposed elements, was then assessed.

Results and Discussion

We consider now the results of an example, continuous carry-over, fMRI study of color perception conducted in a single subject. The subject viewed 27 different colors (plus a black screen null-condition) presented in a serially balanced order. The colors were drawn from the OSA-UCS color space, in which the axes correspond to greenness-redness opponency (G), yellowness-blueness (J, from the French *jaune*), and lightness (L) (Figure 2D). By design, the perceptual salience of a single unit step along any of the axes in the OSA-UCS space is equivalent, and the perceptual salience of combined movement along the axes is nearly proportional to the linear, Euclidean distance between points.¹

As the purpose here is to illustrate principles of experimental design and opportunities for inference, as opposed to rigorous testing of neuroscientific hypotheses regarding color perception, results are considered for some effects that did not achieve standard significance levels (e.g., $p < 0.07$). Confirmation and extension of these findings, as well as a more comprehensive interpretation with regard to models of the neural representation of color, awaits additional data and analysis. Our primary goal here is to demonstrate that the representation of stimulus information at multiple levels of neural coding might be measured in a single experiment.

Focal effects – Main effect of stimulus presentation

The main effect of the presentation of any color stimulus versus the null-trials (as well as the added effect of presentation of stimuli following null-trials) was assessed at the whole brain level. At an $F(2,653)$ threshold of 14.9 (corresponding to a Bonferroni corrected, map-wise $\alpha = 0.05$), extensive, confluent areas of significant responses were seen throughout occipital, temporal and parietal cortex. To allow the identification of the centers of visual area responses, the threshold was raised to $F > 30$ (corresponding to a map-wise $\alpha = 1e-08$), and the resulting activations displayed upon the inflated cortical surface (Figure 4A). Clearly, sufficient power is present in the design to easily detect main effects within an individual subject. While retinotopic mapping was not performed to confirm these assignments, the locations of activity

¹In practice, it is not possible to produce a three-dimensional, isotropic, Euclidean space for color perception (Brainard, 2003). The similarity matrix could be adjusted to reflect the non-Euclidean properties of the stimulus space, but in this example study we tolerate the small deviation from perceptual accuracy to allow the generalizability of the method to other 3- dimensional perceptual spaces.

accord well with the known locations of several retinotopically organized visual areas, including (nomenclature after Wandell et al., 2005) V1, V2/V3v, hV4, V3A/B, MT+, as well as an area within the lateral occipital sulcus (LOS). There appeared to be two foci of activation along the calcarine sulcus. As the subject was performing a task at fixation during scanning, the V1 area of activation was divided into a posterior (foveal) and anterior (peripheral) region so that a response from the peripheral V1 could be obtained that was not affected by the fixation task. The average signals within these regions of interest were then interrogated for orthogonal, focal responses to the color stimuli. It is also within these regions of interest that voxels were identified for subsequent analyses of distributed responses to color (see below).

Focal effects – Carry-over adaptation

Stimuli that are processed by a common neural population may be expected to demonstrate partial adaptation to successively presented, similar stimuli, proportional to the degree of overlap in their representation. For the continuously presented stimuli, we asked what the response was to a color as a function of how similar (Figure 2D) it was to the immediately preceding color. Here, we focus upon the linear relationship between adapted neural response and BOLD fMRI signal. Within the defined regions, linear carry-over adaptation for color was seen in V1 and V2/V3v. Figure 4B presents the relationship between neural adaptation and the magnitude of difference in color space between adjacent stimuli for the fitted covariates (λ_1 , λ_2 , λ_3). Plotted separately is the degree of adaptation observed for identical stimuli (the κ effect). The model did not assume that the adaptation to an identical stimulus can be predicted from the adaptation to different stimuli of varying similarities. Nonetheless, the adaptation response seen in V1 and V2/V3v to identical stimuli is positioned roughly as an extension of the response to similar stimuli.

The presence of continuous adaptation to color space similarity within these regions suggests the presence of voxels that contain populations of neurons with overlapping responses to different color stimuli. This analysis did not decompose the adaptation response by direction within color space, but instead assessed the integrated movement along all three axes. It is therefore possible that the adaptation effect observed was driven to a greater extent by one particular direction in color space. It may be of interest in future analyses of this data, and in experiments of this kind, to determine if carry-over adaptation varies by direction of change in stimulus space, and if this differs by cortical region.

A whole-brain analysis was also conducted to identify cortical areas which demonstrate partial adaptation to color space similarity, but may be outside of the defined visual area regions of interest. An F-test was used to evaluate the conjoint explanatory power of the first, second, and third-order carry-over adaptation covariates. The response of an area identified in dorsolateral pre-frontal cortex is shown in Figure 4B. This area showed a strong modulation of response as a function of carry-over adaptation, although the form of the response is curious: identical stimuli produced the largest neural response, as did maximally different stimuli. Such complex response forms to continuous changes in stimulus features may be a property of higher-order association cortex (similar responses were seen in a midline, parietal region and the inferior pre-frontal cortex bilaterally), and are a ready topic for future investigation.

Focal effects – Carry-over bias

Behavioral studies of the perceived brightness of sequentially presented stimuli have shown a tendency towards an assimilation bias effect, in which stimuli are perceived as brighter if preceded by a bright stimulus and darker if preceded by a darker stimulus (Gescheider, 1988). In this continuous carry-over experiment we can test if carry-over bias is present in the neural response to stimuli of different Lightness. Within the visual area regions of interest, I measured the magnitude and significance of the scaling of a covariate that modeled carry-over

bias for changes in the Lightness of the color stimuli from one trial to the next. Figure 4C presents the results. Strong contrastive effects were observed within area V1, such that the response to a light stimulus was stronger when preceded by a darker stimulus. Conversely, in area MT+ an opposite, assimilation bias was observed.

Just as carry-over adaptation can be used to infer the representational properties of neuronal ensembles, carry-over bias can be used as an index of sensitivity of a neuronal population to a particular stimulus dimension. As defined here, carry-over bias can be considered a (rather constrained) model of gain-control, which has been studied as the non-linear modulation of neural response to changes in luminance or contrast. Systems that exhibit gain-control are capable of maximizing the information content of the dynamic range of responses (Yu and Lee, 2005). If a cortical region demonstrates carry-over bias for a particular stimulus dimension, this might be taken as evidence that neurons in the region process information regarding that stimulus dimension.

Interestingly, opposite bias effects were observed in different visual areas. Behavioral bias effects have been shown to be sensitive to the task context. For example, assimilation is enhanced by providing subjects with feedback when they are engaged in an absolute magnitude estimation task (Gescheider, 1988). It may be the case that sequential differences in stimuli are represented differently in different cortical areas, and manifest a different behavioral expression dependent upon the experimental context. As was noted earlier, we must be alert to the possibility that carry-over bias effects may be the result of a non-linear transformation of neural activity to BOLD fMRI signal, as opposed to a non-linear neural response. The presence, however, of bias effects of opposite sign in adjacent neural areas would tend to mitigate this concern.

The study of carry-over bias may be further extended in continuous carry-over studies. In the context of this particular experiment, the carry-over bias present for color stimuli that vary along other axes (e.g., Greenness and Yellowness) may be of interest. Further, non-linear relationships between intensity differences in the stimuli and neural response could be examined. The current data set is, however, limited in the degree to which carry-over effects along a single dimension might be measured, as there were only three levels to each of these factors. Experiments with multiple levels of a stimulus parameter will be better able to distinguish between more subtle models of carry-over effects. Finally, carry-over effects were modeled in this study only for the immediately preceding stimulus. In behavioral studies, assimilative bias is generally found with regard to the immediately preceding stimulus, and contrastive bias for earlier trials (Gescheider, 1988). A design with fewer total stimuli would allow the presentation order to be determined by an m-sequence with higher-order counterbalancing. In this way, carry-over effects from $t-2$ and earlier trials could be readily examined. With fewer stimuli allowing sufficient counterbalancing, more sophisticated models of carry-over effect (Schwartz et al., 2002) might be implemented.

Focal effects – Direct modulation by color axis

Because of first-order counterbalancing, unbiased measurements of the direct and carry-over effects may be obtained. The design matrix included covariates that modeled the magnitude of the evoked response to a stimulus as a function of its position along the three color axes (independent of the preceding stimulus). An F-test was used to identify regions in which there was a significant modulation of neural response by color direction. Figure 4D presents the magnitude of modulation of neural response along the three color axes for visual areas in which there was significant combined effect across all three color directions. In all five areas, the strongest modulation was observed for the Lightness dimension, with darker colors producing a greater level of neural response. Modulation of responses was much smaller for the other color axes, although consistent across regions. All five areas showed a greater response to green

compared to red stimuli, and greater response to yellow as compared to blue. These effects were most significant within area hV4, where the magnitude of signal modulation was large in proportion to the main effect of stimulation.

This finding indicates that the average response of neurons within some cortical regions is greater for certain points in the color space than others. In particular, we see that the integrated amplitude of neural response most strongly codes the relative Lightness of a stimulus, despite the perceptual equivalence of changes along any of the three axes. This suggests that different aspects of the perceived color of a stimulus are represented by different neural codes.

Distributed effects – neural pattern and perceptual similarity

The previous analyses have examined the effect of manipulations of color upon focal neural responses. These results indicate that collections of neurons, on the order of voxels or regions, represent continuous variation in color space by the absolute magnitude of their response and the pattern of their firing. We turn now to representation of color information across a larger spatial scale. As has been observed for other stimulus types (Haxby et al., 2001, Kamitani and Tong, 2005) the small bias of individual voxels (reflecting the bias of the population of neurons within) to respond differentially to stimuli, aggregated across voxels, carries substantial information regarding the perceptual state of the subject. For example, classification methods can be used to identify which of 8 angles of orientation of a bar of light (Kamitani and Tong, 2005), or which of 8 classes of visual object (Haxby et al., 2001), a subject is viewing. These findings provide fundamental information regarding the organization of neural codes for stimulus representation. Further insight may be provided by comparing the quantity and type of classification information carried by subsets of the available voxels, either by cortical region or based upon other selection criteria.

While the data obtained here could be analyzed using support vector machine (Kamitani and Tong, 2005) or non-linear techniques (Polyn et al., 2005), I examined the representation of the color similarity structure by correlation. The presentation of a color stimulus (as compared to the null-trials) evokes signal changes across multiple visual areas. It may be the case that stimuli that differ from one another may evoke distributed patterns of neural activity that themselves differ, perhaps in a manner that reflects their perceptual similarity. Such a result is not a foregone conclusion. For example, the relevant stimulus dimension may be encoded entirely by neural population codes that are amplitude-homogeneous on the spatial scale of voxels. (Although such a model would not have predicted the presence of focal, direct signal change effects represented by the τ_G , τ_J , and τ_L covariates, such as was seen above.) Alternatively, neurons might code the stimulus dimension by the magnitude of their population response. If, however, the response function is the same across voxels, or present only within a limited cortical area, the resulting pattern of activity would be the same for different stimuli, differing only in magnitude scaling. In such a circumstance, focal, direct effects of the stimulus change would be detectable while the distributed pattern analysis performed here would fail. Of course the representation of different properties of a stimulus (e.g., Lightness vs. Hue) may itself differ on these spatial scales and population codes.

For this analysis, sub-populations of 1000 voxels within the posterior visual areas were identified. These voxels were selected based upon the magnitude of the “main effect” used to define the regions in Figure 4A. Figure 5A shows the voxels used for this test, identified in the data without spatial smoothing. The matrix shown on the right in Figure 3 was then used to obtain the average response (β value) for each color stimulus across its multiple presentations for each of the 1000 voxels. Importantly, the calculated β value for each color combines stimulus presentations across all first-order contexts, and so represents an average response unbiased by carry-over effects. The vector of voxel values for each color was then compared to every other color with a correlation test, and entered into a neural similarity matrix (Figure

5B). Finally, a correlation was calculated for the relation between a neural similarity matrix and the perceptual, color similarity matrix and its color direction sub-components (Figure 5C).

For voxels identified as having the most significant response to any color versus a blank screen, there was a significant correlation ($r=0.21$) between the distributed patterns of response to different colors and the perceptual similarity matrix. This correlation was carried almost entirely by changes along the L color axis ($r=0.43$), with essentially no variability in the neural pattern explained by differences in color opponency. This result suggests that while stimulus lightness may be represented by distributed firing over many voxels, hue is represented by within voxel population codes.

Clearly, further analyses could be performed. In particular, it might be of interest to contrast the pattern-similarity performance of sets of voxels selected by visual area, as opposed to statistical test. Further, one might ask about the performance of voxels selected for carry-over adaptation effects in representing the similarity structure.² For the current purpose, however, it is sufficient to demonstrate that continuous carry-over designs using serially balanced sequences can identify distributed patterns of responses in addition to within voxel adaptation and direct effects. More powerfully, insights gained in an analysis of focal effects can be used to constrain and inform analysis of distributed responses in the same data set.

Limitations of continuous carry-over designs

While suitable for testing a wide range of neuroscientific hypotheses, continuous carry-over designs have some specific limitations. Continuous carry-over designs are not as powerful as a “block” design for the detection of direct effects. Thus, if the only purpose of an experiment is to identify direct effects for their own sake, then a continuous carry-over design is not optimal. It should be noted, however, that block-designs are subject to bias by unbalanced carry-over effects, and first-order counterbalancing in a block design may be difficult to achieve with a reasonable duration of stimulus presentation. Note also that ample statistical power to identify main effect responses was observed in our example experiment. Further improvements in power for detection of direct effects can be provided by adjustment of experimental parameters (as explored in Appendix A).

Because the continuous carry-over design considered here has the same duration of presentation for all stimuli, it is most similar to the “immediate adaptation” approach employed in many studies (e.g., Kourtzi and Kanwisher, 2000, Winston et al., 2004). Some studies of neural adaptation, however, present a prolonged adapting stimulus, followed by a brief test stimulus (e.g., Engel, 2005). Such an approach may be accommodated within a continuous carry-over design by treating the duration of presentation as another stimulus dimension, resulting in carry-over effects from long-to-short duration presentations (as well as the other factorial combinations). This solution is cumbersome, however, if only the long-to-short transition is of interest. Other adaptation studies examine the effect of repetition over longer time periods with intervening stimuli (Henson et al., 2002). Such designs are not easily accommodated within a continuous carry-over approach.

Conclusions

This paper describes a class of BOLD fMRI experiments that I have termed “continuous carry-over”. It was shown that serially-balanced sequences are well suited to determine the order of presentation of stimuli in such an experiment, and I introduced the Finney and Outhwaite *type*

²It should be noted, however, that it is not valid to perform a distributed analysis using voxels selected based upon the focal, direct effects (τ_G , τ_J , and τ_L). Voxels selected in such a manner necessarily have differential magnitude responses that reflect the similarity space. As a result, a distributed pattern that matches the perceptual space may be obtained even in null-hypothesis data.

1 index 1 sequence to be used in those cases in which *m*-sequences are not available. With proper covariate construction, an experiment can be used to detect simultaneously:

- 1) The main effect of any stimulus versus null-trials
- 2) Focal “direct effects”, which are the differences in the mean neural responses between stimuli
- 3) Focal “carry-over effects”, which are the responses to a stimulus as a function of its relation to preceding stimuli. Two forms of carry-over response were studied, a symmetric neural adaptation effect, and an asymmetric effect that reflects neural bias.
- 4) Distributed patterns of response unbiased by the context of their presentation

This type of design is particularly well suited to the study of perceptual similarity spaces and their neural representation. Example data were presented for the perception of color, which is an instance of a space defined by three axes and a (roughly) Euclidean distance metric (for the particular stimuli used). This approach can be readily applied to perceptual spaces that have different forms. Examples of these were considered in Figure 2. Particularly intriguing is the use of such designs to study conceptual similarity spaces where different similarity assessments may exist simultaneously, and be evoked by different instructions or experimental contexts.

Acknowledgements

The author would like to thank B A S Nonyane and C M Theobald for guidance in the implementation of their algorithm; D Hu for writing C code used to generate *type 1 index 1* sequences from random initial start points; A Thomas for assistance in devising a method to account for sequence overlap at scan boundaries; DH Brainard for the color stimuli used in the example experiment; and A Harris for editing suggestions. This work was supported by K08 MH 7 2926-01 and a Burroughs-Wellcome Career Development Award.

Appendix A: Selection of type 1 index 1 sequences

Different *type 1 index 1* sequences vary in their suitability for use in BOLD fMRI experiments, and we consider here two general metrics to guide their selection. First, sequences vary in the degree to which they balance the ordering of stimuli within blocks (described further below). Second, sequences vary in their Efficiency (Friston et al., 1999), which is the degree to which the temporal pattern of neural activity that they are designed to produce is represented in the BOLD fMRI signal following convolution with the hemodynamic response function (HRF).

In this Appendix, we will consider some computations and simulations that may be used to guide the selection of better *type 1 index 1* sequences for use in a BOLD fMRI study. This will be illustrated by selecting the optimal $n=28$ (27 stimuli plus a null-condition) sequence for study of the similarity space shown in Figure 2D. The sequence that is ultimately identified is presented in Appendix B.

Maximize block balance

As was observed earlier, *type 1 index 1* sequences are essentially immune to order effects across the entire duration of stimulus presentation, as all n stimuli are presented before the next set of stimuli. However, sequences may differ in the degree to which the ordering of the stimuli within blocks is evenly distributed—e.g., stimulus i_6 occurs equally often in all possible positions within a block—and any trends in the ordering of stimuli are minimized. While these criteria cannot be perfectly satisfied, different sequences depart from this ideal to different degrees. Nonyane and Theobald (Nonyane and Theobald, In Press) offer metrics for these departures, and sequences may be selected based upon minimization of these scores.

While it is unlikely that imperfect block balance would play a substantial role in the results of an imaging experiment (particularly with as many as 27 different stimuli), there is little cost to optimizing this metric. To do so, we first generate a set of sequences using a modification of the Nonyane and Theobald algorithm. Because this algorithm is systematically enumerative, and because successively generated sequences are very similar, only a small portion of the vast space of possible sequences can be generated in this manner in a reasonable period of time. A modification of the algorithm allows sequence searches to begin from random initial starting positions in the solution space, allowing us to sample the space of possible sequences more broadly (sample code is available from <http://www.bioss.ac.uk/staff/cmt.html/designseq.html>). Of the set of sequences generated (1500 in this application), I selected the sequence that had the smallest departure from orthogonality of treatment effects and within-block ordering. The selected sequence was then further optimized as described below.

Maximize Efficiency

We now consider two design choices for a serially balanced, continuous carry-over experiment: the effect of null-trials and trial duration. We will evaluate these choices with respect to experimental effects that might be induced in a continuous carry-over design: an average direct effect of the stimuli (in essence, the average τ for all n stimuli), and a first order (linear) carry-over adaptation effect (λ_1). The desirability of different designs can be compared by calculating the Efficiency ratio (E), which reflects the ability of a particular experimental design to create temporal patterns of neural activity that can pass through the HRF to be represented in the BOLD fMRI signal (Friston et al., 1999). The maximum E is unity, and for the sake of comparison, a two-condition, blocked fMRI design with optimal alternation between conditions provides an E of approximately 0.9.

Null-trials are treated as any other stimulus, and increase the n (from 27 to 28) of the to-be-generated sequence. The Efficiency of the direct effect may be improved by increasing the time spent in the null-trial condition. Appendix Figure 1A and 1B demonstrate that this is best done by maintaining the sequence as $n=28$ and selectively increasing the duration of the null trials (as opposed to increasing the n to 29 or higher, and assigning more than one label to the null-condition).

Next, we consider the duration of presentation of the stimuli. As the HRF attenuates high-temporal frequency changes in neural activity, we might expect that briefer stimulus presentations will have reduced Efficiency. Appendix Figure 1C shows that longer stimulus durations improve the detectability of both the direct and linear adaptation effects. In practice long stimulus presentations may be less efficient than predicted in the face of neural responses that habituate and temporally autocorrelated noise in BOLD fMRI data (Zarahn et al., 1997). Nonetheless, the simulation does suggest that very brief (< 1 second) stimulus presentations using balanced sequences have relatively poor efficiency.

Label permutation

At this stage we have identified an $n=28$ sequence that minimizes block ordering effects, and have made decisions regarding the duration of presentation of the stimuli, including the observation that it may be preferable to selectively lengthen the presentation of the null-trials. It is still possible to derive an even more efficient order of stimulus presentation from this particular sequence. Observe that the mapping between the numbers assigned to the treatments in the sequence and the numbers assigned to the stimuli is arbitrary: the labeling of the 28 treatments (27 stimuli plus a null-condition) can be arbitrarily permuted. This permutation does not alter the properties of the sequence with regard to the block ordering effects considered above, but different permutations may differ in the expected Efficiency of the design for use

in BOLD fMRI. Of the possible $28!$ (greater than $3e29$) re-orderings of the 28 stimulus treatments, we examined 1,000,000 and identified the sequences that maximized E for the linear adaptation carry-over (λ_1) and direct effect (τ) covariates.

Appendix Figure 1D presents the average observed improvement in the maximum sequence efficiency as a function of thousands of re-orderings searched. A modest improvement in the efficiency of the carry-over adaptation was obtained (~9% absolute, 38% relative), although no improvement was found for the efficiency of the average direct-effect contrast. This is sensible, as the power of the direct-effect contrast is determined by the frequency structure of alternations between stimulus and null-trial periods, which by design is fixed to occur on average every 28 stimuli. As a consequence, it is necessary only to search for sequences that are optimal for detecting carry-over effects, as any selected sequence will be as good as any other for detection of direct effects. The best performing sequence yielded $E(\lambda_1)=0.314$, and $E(\tau)=0.3664$, and is provided in Appendix B.

Appendix B: Selected type 1 index 1, n=28 sequence for study of 3-dimensional, Euclidean similarity space

This sequence was derived as described in the Appendix A. Values of “0” represent the null-trials, and have double the duration of the other events. Note that the initial “2” can be omitted if the implementation of the sequence uses the “scan overlap” technique described in the methods, as the last label of the sequence is also “2”.

```

2
2 23 6 12 19 4 20 26 3 25 1 5 18 10 24 27 16 9 22 17 13 15 7 21 0 0 14 8 11
11 24 15 16 7 25 12 17 21 4 8 1 14 19 20 6 27 9 5 13 18 22 0 0 2 3 23 26 10
10 9 8 17 16 14 12 15 24 22 2 6 11 18 7 0 0 13 4 23 27 21 25 20 3 19 5 26 1
1 12 27 5 19 10 20 15 0 0 9 7 18 21 26 17 4 13 25 22 24 3 14 16 8 6 2 11 23
23 25 17 2 14 9 11 6 8 0 0 10 12 4 19 13 20 1 7 26 5 22 18 3 24 16 27 15 21
21 23 14 11 10 13 12 2 5 17 9 24 0 0 15 22 7 16 6 19 3 4 25 26 20 8 27 1 18
18 1 6 23 10 27 22 3 12 16 24 19 0 0 7 17 15 9 21 14 13 5 11 4 2 26 8 20 25
25 13 22 8 12 7 14 1 21 6 4 27 19 17 20 5 9 2 18 23 3 26 24 11 15 10 16 0 0
0 0 21 27 7 19 18 24 13 6 20 11 14 22 25 23 1 4 12 8 26 16 10 3 9 15 2 17 5
5 15 17 7 22 11 2 8 21 3 16 23 18 19 9 12 6 0 0 4 26 25 14 24 20 10 1 13 27
27 0 0 8 14 17 26 11 20 23 12 3 5 2 25 18 16 13 7 4 10 6 24 1 9 19 22 21 15
15 20 18 25 5 4 3 10 17 22 27 26 9 23 16 1 11 13 24 21 2 12 0 0 19 6 14 7 8
8 2 24 26 6 3 0 0 12 13 23 19 11 25 15 14 4 5 1 20 27 18 9 17 10 22 16 21 7
7 11 1 0 0 22 23 13 26 12 20 9 4 6 16 17 3 2 19 21 18 5 14 10 25 8 15 27 24
24 10 5 7 9 13 16 3 8 25 21 22 15 11 12 14 18 0 0 17 1 27 6 26 2 20 19 23 4
4 15 13 17 25 7 1 8 3 20 2 16 11 5 10 0 0 23 24 14 6 18 26 22 19 27 12 21 9
9 18 2 15 1 10 4 14 21 5 0 0 11 27 20 17 19 12 22 13 8 24 23 7 3 6 25 16 26
26 18 27 3 1 24 9 20 7 5 16 12 11 8 22 14 2 21 17 6 13 10 23 15 4 0 0 25 19
19 25 4 11 26 13 2 10 14 27 17 18 15 5 21 1 3 22 9 0 0 16 20 12 23 8 7 24 6
6 1 25 3 13 9 10 26 15 8 4 18 17 14 23 21 12 24 2 7 20 0 0 5 27 11 19 16 22
22 1 26 27 8 19 24 5 20 16 4 9 14 15 23 17 11 3 21 10 2 0 0 18 12 25 6 7 13
13 14 5 8 23 2 4 24 25 27 10 19 26 7 6 9 3 18 11 17 0 0 1 22 20 21 16 15 12
12 26 21 20 24 17 8 18 4 1 15 19 7 2 13 3 27 14 25 0 0 6 5 23 22 10 11 9 16
16 19 2 1 23 9 26 0 0 27 25 10 7 15 6 17 24 12 18 14 20 4 21 8 13 11 22 5 3
3 15 25 2 27 23 5 6 22 26 4 16 18 20 13 19 14 0 0 24 8 10 21 11 7 12 9 1 17
17 27 4 22 12 5 24 18 13 0 0 3 7 10 8 16 25 11 21 19 1 2 9 6 15 26 23 20 14
14 26 19 15 3 17 12 1 16 2 22 6 10 18 8 5 25 9 27 13 21 24 4 7 23 11 0 0 20
20 22 4 17 23 0 0 26 14 3 11 16 5 12 10 15 18 6 21 13 1 19 8 9 25 24 7 27 2

```

References

- Aguirre GK, Zarahn E, D'Esposito M. The variability of human, BOLD hemodynamic responses. *Neuroimage* 1998;8:360–369. [PubMed: 9811554]
- Attneave F. Dimensions of similarity. *American Journal of Psychology* 1950;63:516–556. [PubMed: 14790020]
- Beck J. Contrast and assimilation in brightness judgements. *Psychonomic Science* 1966;1:342–344.
- Botvinick M, Plaut DC. Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychol Rev* 2004;111:395–429. [PubMed: 15065915]
- Brainard, DH. Color appearance and color difference specification. In: Shevell, SK., editor. *The Science of Color*. Optical Society of America; Washington, DC: 2003. p. 191-216.
- Buchel C, Holmes AP, Rees G, Friston KJ. Characterizing stimulus-response functions using nonlinear regressors in parametric fMRI experiments. *Neuroimage* 1998;8:140–148. [PubMed: 9740757]

- Buracas GT, Boynton GM. Efficient design of event-related fMRI experiments using M-sequences. *Neuroimage* 2002;16:801–813. [PubMed: 12169264]
- DeCarlo LT, Cross DV. Sequential Effects in Magnitude Scaling: Models and Theory. *Journal of Experimental Psychology: General* 1990;119:375–396.
- Desimone R. Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci U S A* 1996;93:13494–13499. [PubMed: 8942962]
- Edelman S. Representation is representation of similarities. *Behavioral and Brain Sciences* 1998;21:449–498. [PubMed: 10097019]
- Engel SA. Adaptation of oriented and unoriented color-selective neurons in human visual areas. *Neuron* 2005;45:613–623. [PubMed: 15721246]
- Ferris SJ, Kempton RA, Deary IJ, Austin EJ, Shotter MV. carry-over bias in visual assessment. *Perception* 2001;30:1363–1373. [PubMed: 11768489]
- Finney DJ, Outhwaite AD. Serially balanced sequences in bioassay. *Proc R Soc Lond B Biol Sci* 1956;145:493–507. [PubMed: 13359400]
- Friston KJ, Mechelli A, Turner R, Price CJ. Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. *Neuroimage* 2000;12:466–477. [PubMed: 10988040]
- Friston KJ, Zarahn E, Josephs O, Henson RN, Dale AM. Stochastic designs in event-related fMRI. *Neuroimage* 1999;10:607–619. [PubMed: 10547338]
- Gescheider GA. Psychophysical scaling. *Annu Rev Psychol* 1988;39:169–200. [PubMed: 3278675]
- Gomez RL, Gerken L, Schvaneveldt RW. The basis of transfer in artificial grammar learning. *Mem Cognit* 2000;28:253–263.
- Grill-Spector K, Malach R. fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol (Amst)* 2001;107:293–321. [PubMed: 11388140]
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 2001;293:2425–2430. [PubMed: 11577229]
- Henson RN, Rugg MD. Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia* 2003;41:263–270. [PubMed: 12457752]
- Henson RN, Shallice T, Gorno-Tempini ML, Dolan RJ. Face repetition effects in implicit and explicit memory tests as measured by fMRI. *Cereb Cortex* 2002;12:178–186. [PubMed: 11739265]
- Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 2005;8:679–685. [PubMed: 15852014]
- Kellman P, Gelderen P, de Zwart JA, Duyn JH. Method for functional MRI mapping of nonlinear response. *Neuroimage* 2003;19:190–199. [PubMed: 12781738]
- Kourtzi Z, Kanwisher N. Cortical regions involved in perceiving object shape. *J Neurosci* 2000;20:3310–3318. [PubMed: 10777794]
- Nonyane BAS, Theobald CM. Design Sequences for Sensory Studies: Achieving Balance for Carry-over and Position Effects. *British Journal of Mathematical and Statistical Psychology*. In Press
- Polyn SM, Natu VS, Cohen JD, Norman KA. Category-specific cortical activity precedes retrieval during memory search. *Science* 2005;310:1963–1966. [PubMed: 16373577]
- Schwartz, O.; Chichilnisky, EJ.; Simoncelli, EP. Characterizing neural gain control using spike-triggered covariance. In: Dietterich, TG., et al., editors. *Advances in Neural Information Processing Systems*. MIT Press; Cambridge, MA: 2002.
- Steyvers M. Multidimensional Scaling. *Encyclopedia of Cognitive Science*. 2002
- Thompson-Schill SL, Kurtz KJ, Gabrieli JDE. Effects of semantic and associative relatedness on automatic priming. *Journal of Memory and Language* 1998;38:440–458.
- Wandell BA, Brewer AA, Dougherty RF. Visual field map clusters in human cortex. *Philos Trans R Soc Lond B Biol Sci* 2005;360:693–707. [PubMed: 15937008]
- Winston JS, Henson RN, Fine-Goulden MR, Dolan RJ. fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *J Neurophysiol* 2004;92:1830–1839. [PubMed: 15115795]
- Worsley KJ, Friston KJ. Analysis of fMRI time-series revisited—again. *Neuroimage* 1995;2:173–182. [PubMed: 9343600]

- Yu Y, Lee TS. Adaptive contrast gain control and information maximization. *Neurocomputing* 2005;65(66):111–116.
- Zarahn E, Aguirre GK, D’Esposito M. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage* 1997;5:179–197. [PubMed: 9345548]

Figure 1. An example carry-over matrix (ϕ), and a restricted basis-set of matrices

The grayscale value within each cell of the matrix reflects the proportional size of the carry-over effect for the sequential transition between two stimuli. The labels given to the matrices correspond to those used in the text, with λ modeling a symmetric carry-over effect (for linear, quadratic and cubic components), ρ representing an asymmetric effect (termed *bias* in the text), and κ modeling the carry-over effect of seeing the same stimulus twice in succession. The ellipsis indicates that many other matrices, modeling other types of carry-over effects, could be added to this set.

Figure 2. Examples of stimulus spaces and their implied similarity matrices

A) On the left, an example set of 8 stimuli consisting of a bar of light rotated in 22.5° increments. On the right is the similarity matrix for these stimuli. **B)** A contrived semantic similarity space for birds. If subjects were asked to provide ratings of the similarity of pairs of different birds, a representational space such as depicted on the left of the figure might emerge, in which some groups of birds are found to be more similar than others. This might represent, for example, the 2-dimensional projection of the results of a multi-dimensional scaling analysis conducted upon the similarity ratings (Steyvers, 2002). **C)** A stimulus space that consists of X and Y position within a city-block grid. Unlike a Euclidean space, the distance between two points is defined by $|x_1 - x_2| + |y_1 - y_2|$, termed Manhattan or city-block geometry. Manhattan

geometry is also encountered in similarity spaces for stimuli with very separable dimensions (e.g., size and brightness; Attneave, 1950). **D**) The Optical Society of America (OSA), committee for Uniform Color Scales (UCS) color space. The similarity matrix assumes equal perceptual saliency of steps along any of the three axes and a Euclidean relationship between the axes.

Figure 3. Design matrices used in analysis of focal and distributed neural responses

Time runs along the vertical axis from top to bottom. On the left is shown the serially balanced sequence used in the experiment (and presented in Appendix B), with each label replaced with the corresponding stimulus color. Each column of the design matrices represents a particular covariate over time. The “focal” design matrix was used in the data analysis shown in Figure 4. The “distributed” design matrix, which modeled each presentation of a given color independent of context, was used in the analysis of distributed patterns of neural response shown in Figure 5B. In the example study, the sequence was presented twice for a total 40 minute duration, but only one half of the identical, replicated design matrix is shown.

Figure 4. Focal effects

A) The main effect of stimuli vs. null trials (and the added effect of stimuli following null-trials), presented on the inflated cortical surface. Points of activity correspond to F values greater than 30. The regions of activity seen here provided the signal used for the regional analyses presented in the subsequent panels. **B)** Carry-over adaptation effects within the regions defined in panel A. For each region, the t-value for the linear carry-over adaptation effect across all colors was obtained, and the fitted adaptation response plotted. The x-axis is the distance between adjacent stimuli in the color space, and the y-axis is the % fMRI signal change, relative to the average response to any color that followed a different color. Also shown is the response to a color that followed an identical color (step size of zero), modeled with a separate, orthogonal covariate (κ). In blue is the fitted response of a region in the left, dorsolateral pre-frontal cortex identified by a whole-brain F-test of the linear, quadratic and cubic symmetric carry-over (adaptation) covariates. **C)** For each visual area, the loading upon the linear, L-direction, asymmetric carry-over (bias) covariate is shown, expressed as % fMRI signal change. Negative values indicate a contrastive effect, in that (e.g.) the response to a darker stimulus was enhanced when it was preceded by a lighter stimulus. Contrast effects were found in primary visual cortex, and a trend towards the opposite, assimilative effect, was seen in area MT+. **D)** Region direct effects. For each visual area, the modulation of neural response as a function of the color of the stimulus was modeled for the three cardinal color

directions, and this set of covariates assessed with an F-test. Five visual areas showed a significant modulation of response by direct effects. For each area, the tri-plot indicates the magnitude and direction of modulation of response as a function of position along the color axes. Each tri-plot expresses the modulation as a proportion of a 0.2% signal change, with the positive direction being increases in signal as one moves in color space away from the lightest, yellowest, reddest stimulus (i_1)

Figure 5. Distributed neural similarity

A) Displayed on T1 axial slices, voxels located within posterior visual areas with the highest 1000 F-values for the main effect of stimulus presentation (μ and λ_0). **B)** Calculation of the neural similarity matrix. The vector of β values obtained across voxels for each color were compared by correlation. The resulting r value provides one cell of the neural similarity matrix. Only the lower triangle (omitting the diagonal) of this diagonally symmetric matrix is shown. **C)** Shown is the perceptual similarity matrix for color space (derived from Figure 2D), as well as decompositions of this matrix along the three color axes. The correlation between the neural similarity matrix shown in panel B with the perceptual similarity matrices is indicated. The G and J color directions were evaluated conjointly.

Appendix Figure 1. The efficiency ratio (obtained through simulation) of two covariates derived from a serially-balanced experimental design

In all panels, solid indicates $E(\tau)$ and dashed indicates $E(\lambda_1)$, as defined by the similarity matrix from Figure 2D. **A)** The effect of null-trial duration upon relative power. Selectively increasing the duration of the null-trials improves efficiency of covariates that model direct effects. **B)** Increasing the number of labels assigned to null-trials does not improve efficiency. **C)** Increasing the duration of stimulus presentation theoretically enhances the efficiency of covariates for both effects in a serially balanced design. **D)** Searches across permuted labeling of stimuli within a particular serially balanced sequence leads to a modest improvement in efficiency for detection of adaptation effects, but no change in the efficiency of covariates modeling direct effects.

Calculated correlation of individual covariates in the “focal effects” design matrix (Figure 3) with other covariates. The correlation provided here is for the neural model of the covariates, prior to convolution with the HRF. Values that are discussed in the text are indicated in bold.

	μ	λ_0	κ	λ_1	λ_2	λ_3	τ_G	τ_J	τ_L
λ_0	0.00								
κ	0.00	0.00							
λ_1	0.00	0.00	0.00						
λ_2	-0.01	0.01	-0.01	-0.05					
λ_3	0.00	0.00	0.00	-0.05	-0.08				
τ_G	0.00	0.00	-0.05	0.04	-0.04	0.02			
τ_J	0.00	-0.02	0.02	0.01	0.04	0.00	-0.05		
τ_L	-0.01	0.06	-0.02	-0.03	0.00	-0.08	0.03	-0.06	
ρ	-0.01	0.01	-0.01	-0.05	-0.04	-0.09	0.03	-0.05	0.71

Calculated correlation of individual covariates in the “focal effects” design matrix (Figure 3) with other covariates. The correlation provided here is for the BOLD fMRI model of the covariates, after convolution with the HRF. Values that are discussed in the text are indicated in bold

Table 2

	μ	λ_0	κ	λ_1	λ_2	λ_3	τ_G	τ_J	τ_L
λ_0	-0.74								
κ	-0.03	0.01							
λ_1	-0.01	0.00	0.01						
λ_2	0.07	-0.05	-0.08	0.09					
λ_3	-0.09	0.03	-0.06	0.00	-0.03				
τ_G	-0.01	-0.04	0.02	-0.03	0.06	0.00			
τ_J	0.00	0.02	-0.05	0.03	0.11	-0.01	0.01		
τ_L	-0.03	0.09	-0.07	-0.11	0.02	-0.04	-0.02	-0.05	
ρ	-0.06	0.08	0.04	0.00	-0.01	-0.08	0.02	0.11	0.22