Published in final edited form as:

J Cogn Neurosci. 2011 June ; 23(6): 1346–1357. doi:10.1162/jocn.2010.21436.

Neural correlates of sequence learning with stochastic feedback

Bruno B. Averbeck¹, James Kilner², and Christopher D. Frith^{2,3}

¹Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, London, WC1N 3BG, UK

²Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, WC1N 3BG, UK

³Niels Bohr Interacting Minds Project, Århus University Hospital, Århus C, Denmark

Abstract

Although much is known about decision making under uncertainty when only a single step is required in the decision process, less is known about sequential decision making. We carried out a stochastic sequence learning task in which subjects had to use noisy feedback to learn sequences of button presses. We compared flat and hierarchical behavioral models and found that while both models predicted the choices of the group of subjects equally well, only the hierarchical model correlated significantly with learning related changes in the MEG response. The significant modulations in the MEG signal occurred 83 ms before button press and 67 ms after button press. We also localized the sources of these effects and found that the early effect localized to the insula, whereas the late effect localized to the premotor cortex.

Introduction

Coherent sequences of decisions are an important feature of organized behavior and real world decisions are often made with uncertainty. Given the importance of coherent behavior and the fact that it is often disrupted by brain disorders (Andreasen, 1979a, 1979b; McKenna & Oh, 2005; Penfield & Evans, 1935; Schwartz, Reed, Montgomery, Palmer, & Mayer, 1991), there is a need to understand the neural processes involved in the learning and orchestration of sequences of decisions. However, most work on decision making has focused on single step decisions (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Fellows & Farah, 2007; Glimcher & Rustichini, 2004; Hsu, Bhatt, Adolphs, Tranel, & Camerer, 2005; Samejima, Ueda, Doya, & Kimura, 2005; Wittmann, Daw, Seymour, & Dolan, 2008).

In the current study we addressed several aspects of sequential decision making. The first questions we addressed were when and where do the cognitive processes relevant to learning sequence from stochastic feedback take place? We used magneto-encephalographic (MEG) imaging, which has high temporal precision, to record brain activity during the task. This allowed us to define in time when cognitive processes became active and allowed us to analyze data on a movement-by-movement basis, without making movements artificially slow. Furthermore, we used source localization tools to estimate where in space these processes were taking place. As such, we examined how the signals related to learning in the task evolve over time and space.

Additionally, we were interested in the extent to which subjects used optimal strategies to learn in our task. In an interesting contradiction, two dominant fields of enquiry separately describe behavioral processes as either optimal (Knill & Saunders, 2003) or heuristic, which implies sub-optimality (Kahneman, Slovic, & Tversky, 1982). To a large extent these perspectives come from studying different classes of behaviors (Trommershauser, Maloney, & Landy, 2008). Many who study sensory-motor integration have found that subject performance can have features of optimality (Ernst & Banks, 2002; Jacobs, 1999; Kersten, 1999; Knill & Richards, 1996; Knill & Saunders, 2003; Kording & Wolpert, 2006; Poggio, Torre, & Koch, 1985; Todorov & Jordan, 2002; Trommershauser, Maloney, & Landy, 2003a, 2003b), whereas those who study decision under risk have consistently shown that subject performance is not optimal and instead heuristic (Gilovich, Griffin, & Kahneman, 2002; Kahneman et al., 1982; Payne, Bettman, & Johnson, 1992; Tversky & Kahneman, 1986), the neural correlates of which are currently being explored (Tobler, Christopoulos, O'Doherty, Dolan, & Schultz, 2008). Thus, the difference may have to do with the class of behavioral process being studied as well as other features, including whether or not models are updated with learning (Trommershauser et al., 2008).

The stochastic sequence learning task we used has elements of both sensory-motor integration and decision making, as the uncertainty in our task is external and not due to noise in the sensory-motor system. Furthermore, by bringing decision making into a sequential framework, we examined aspects of the cognitive processes that underlie decision making behaviors that do not exist in single-step decision making paradigms. Specifically, we asked whether or not subjects were able to learn optimal sequences of decisions when information at other points in the sequence affected the current choice. This relationship among parts of a sequence can be modeled using a hierarchical model. Many important human behaviors, including speech, contain hierarchical structure between sequence elements, but it is not clear if these processes are solved by a domain general system, or by a domain specific language system (Fiebach & Schubotz, 2006).

Methods

Subjects and task

Fourteen subjects (7 male) carried out a stochastic sequence learning task while being scanned. Informed consent was obtained from each subject in accordance with procedures approved by the Joint Ethics Committee of the National Hospital for Neurology and Neurosurgery and the Institute of Neurology, London. Each trial began with the presentation of a green outline circle at the center of the screen which cued the subjects to execute a movement (Fig. 1A). They responded by executing a button press with either their left or right thumb. After each response, they were given feedback about whether or not they had pressed the correct button for that movement of the sequence. If they were correct, the outline circle filled green, if they were incorrect the outline circle filled red. After the feedback had been given for 200 ms, they were again presented with a green outline circle which cued the next movement of the sequence. They again pressed one of two buttons and were given feedback. This was repeated 4 times, such that each trial was composed of 4 button presses, with each button press followed by feedback. In 15% of the cases the wrong feedback was given. In other words, if they had pressed the correct button, they were given red feedback, and if they had pressed the incorrect button, they were given green feedback. Thus, the feedback from any individual button press did not necessarily allow the subjects to correct their mistakes and execute the correct sequence in the subsequent trial. Integrated over trials, however, the subjects could infer the correct button press sequence.

Exact task timing depended on the reaction times of the subjects and in the case of visual feedback, on where the screen refresh cycle was when we initiated draw commands. The

mean reaction time following the pacing cue was 385 ms (std 284 ms). Visual feedback followed the button press by 40 ms (std 6 ms). The mean time between feedback and the presentation of the subsequent pacing cue was 226 ms (std 9 ms), and the mean time between button presses was 602 ms (std 275 ms).

We used 6 of the 16 possible sequences (i.e. 4 button presses with either the left or right thumb gives 16 possible sequences), which were balanced for first order button press probabilities and contained at least two left and two right button presses. The sequences used were LLRR, RRLl, LRLR, RLRL, LRRL and RLLR. Each subject executed 8 sets of 6 blocks, where each block was one sequence (Fig. 1B). A single set consisted of all 6 sequence blocks where the order of the blocks was chosen pseudo-randomly. For example, a block of sequence 4 followed by a block of sequence 1, etc. In each block, subjects had to determine the correct sequence using the stochastic feedback and execute it correctly 8 times before they advanced to the next block. Subjects were informed when the block switched, and thus they knew when they had to start learning a new sequence. If they failed to complete the sequence correctly 8 times by 20 trials they were advanced to the next sequence. If the subjects completed all blocks of trials each sequence was executed correctly 64 times within a session, resulting in a total of 384 correct sequences and 1536 correct button presses for each subject. In addition, subjects received at least one set of training on all six sequences before entering the MEG and were instructed that they only had to determine which of the six sequences was correct in the current block. The training familiarized them with the six sequences, the stochastic feedback and other aspects of the task. The subjects were not told specifically that all sequences would require two left and two right button presses. Thus, within the experiment the subjects were familiar with the mechanics of the task and the sequences and their task was to use the stochastic feedback for each sequence to estimate which sequence was correct, and then execute that sequence.

The task was challenging, and we found post-hoc that in about 9% of the blocks the 12 subjects retained for analysis failed to learn the sequence. Specifically, they went through 20 trials without learning to criterion. With respect to the MEG analysis described below, we carried out the analysis both with and without the blocks of data in which they failed to reach criterion. While including or not including this data affected the exact F values of the MEG analysis, the rank order of the fit of the 3 models was the same (see Results) and the plots showing the spatial distribution of the effects were similar. We report results for the analyses in which we excluded the blocks in which subjects did not learn.

Data analysis

Analyses proceeded in a series of steps (Fig. 1C). First, we fit behavioral models to each subject's learning data. Then, we extracted the movement-by-movement estimates of learning from each subject's behavioral model. This learning estimate (quantified as the probability that the subject knew which button they should press at each point in time) was then regressed on the data in sensor space, to see if the button press related activity was modulated by learning. We then used source localization on a subject by subject basis to localize the significant effects that we found in sensor space. After localization was done for each subject, the results were used to carry out SPM statistics on whether or not the localizations were significant, across subjects.

Behavioral model

We fit Bayesian statistical models to the subjects' behavior. The models allowed us to quantify trial-by-trial how much the subjects had learned about which sequence or which button was correct in the current block. The subjects could press either the left or right button at each point in the sequence and therefore they had a binary decision. The model

assumes that the subjects were trying to learn the sequence, and, therefore, that they were trying to optimize the number of times green feedback was received. Statistically, this can be accomplished by remembering how often green feedback was given for the left (or right) button at each point in the sequence. For example, if green feedback was given more often for the left button, then the left button should be pressed. Thus, the model integrates information about red vs. green feedback given for left and right button presses individually for each of the four button presses in the sequence.

The model began with a binomial likelihood function for each movement of the sequence, given by:

$$p(D_T|\theta_{i,j}) = \theta_{i,j}^{r_{i,j}} (1 - \theta_{i,j})^{N_j - r_{i,j}} \quad (1)$$

Where $\theta_{i,j}$ is the probability that pressing button i ($i \in \{left, right\}$) on movement j ($j \in \{1, \dots, 4\}$) would be followed by green feedback. The variable $r_{i,j}$, defined below, is the number of times reward (green feedback) was given when button i was pressed on movement j (or red feedback was given when the other button was pressed), and N_j is the number of trials. The vector D_T represents all the data collected up to trial T for the current block, which in this case are the values of r and N . This was the only data relevant to inferring the correct sequence of button presses. Importantly, the model does not contain any information about previous sequences from the current set. Subjects were not told that all sequences were given in each set, and therefore it was unlikely that they would be able to infer set boundaries and use this information to improve learning.

The probability that the left button should be pressed for movement j after T trials (i.e. that it is more likely to be the correct button) is given by:

$$\Delta_{j,T}^F \equiv P_{j,T}^F(B=L) = p_T(p(\theta_{left,j}|D_T) > p(\theta_{right,j}|D_T)) = \int_0^1 p(\theta_{left,j}|D_T) \int_0^{\theta_{left,j}} p(\theta_{right,j}|D_T) d\theta_{right,j} d\theta_{left,j} \quad (2)$$

We have written the posterior here (i.e. $p(\theta_{right,j}|D_T)$). Button probabilities were equally likely in the experiment so the prior was flat and the posterior is just the normalized likelihood for this estimate. The superscript F on Δ indicates that this is the probability used for button presses in the flat model.

For the hierarchical model, a few additional computational steps were necessary to compute the button probabilities. In effect, the hierarchical model uses feedback about all four button presses in the sequence to estimate the probability that either the left or the right button is correct at each individual point in the sequence. This is the feature which distinguishes it from the flat model and this is possible because not all button press combinations were used in the experiment and, therefore, certain possibilities could be eliminated. We eliminated these possibilities computationally as follows. Given that we had an estimate of the probability that either the left or the right button should be pressed at each point in the sequence, we could calculate the probability that each of the sequences was correct. This probability was given by:

$$P(D_T|s_k) = \prod_{j=1}^4 \left(\Delta_{j,T}^F B_{j,k} + (1 - \Delta_{j,T}^F)(1 - B_{j,k}) \right), \quad (3)$$

where $B_{j,k}$ is 1 if the left button was correct for movement j , sequence k (s_k), and 0 if the right button was correct. Formally, B is the conditional probability of the left button being correct given the sequence, but because it is a delta function, we have not used probability notation. The posterior probability of each sequence was then given by Bayes rule:

$$P(S) \equiv P(s_k|D_T) = \frac{P(D_T|s_k) p(s_k)}{P(D_T)}. \quad (4)$$

Again, we have used a flat prior on sequences, as this was in fact the prior in the experiment, so $p(s_k) = 1/6$ for all sequences.

In the hierarchical model, these sequence probabilities can then be used to infer the button probabilities. As mentioned above, the hierarchical model uses information about all four button presses to infer the probabilities of each individual button press (eq. 3). In equation 4 the individual button probabilities were used to infer the sequence probabilities. These sequence probabilities can then be used to infer the button probabilities under the hierarchical model. This was calculated as:

$$\Delta_{j,T}^H \equiv P_{j,T}^H(B=L) = \sum_{k=1}^6 B_{j,k} P(s_k|D_T), \quad (5)$$

Where again $B_{j,k}$ was 1 if the left button was correct for movement j , sequence k and 0 if the right button was correct. Incidentally, although the second order transitions between buttons were not balanced (i.e. L->R vs. L->L) the hierarchical model subsumes and models this fact by modeling all the correlations between button probabilities.

Equations 1-5 describe an ideal observer model, but it is likely that the subjects' behavior will deviate from this model. In order to better predict their behaviour we added 2 parameters to the basic model that allowed for differential weighting of positive and negative feedback. These parameters affect both the flat and the hierarchical model. When the models were fit to the subject behavior, different parameters were fit under each model to individual subjects, as shown in Results.

The differential weighting was implemented by using the following equation for the feedback:

$$\begin{aligned} fb_{positive} &= 0.5 + \alpha \\ fb_{negative} &= 0.5 - \beta \end{aligned} \quad (6)$$

The subscripts *positive* and *negative* indicate whether the feedback was positive (green) or negative (red). The total reward (feedback), in equation 1 was then given by:

$$r_{i,j} = \sum_{t=1}^T f b_{positive,i,j} u(t) + f b_{negative,i,j} (1 - u(t)). \quad (7)$$

The parameter $u(t)$ is one if green feedback was given and zero if red feedback was given on trial t . Thus, α and β scale the amount that is learned from positive and negative feedback. For an ideal observer both parameters would be 0.5. The parameters, α and β were fit to individual subject decision data by maximizing the likelihood of the subject's sequence of decisions, given the model parameters. Thus, we maximized:

$$p^M(D^*|\alpha, \beta) = \prod_{t=1}^N (\Delta_t^M C_t + (1 - \Delta_t^M)(1 - C_t)), \quad (8)$$

where C is the choice that the subject made for each movement ($C=1$ for left, $C=0$ for right) and the superscript on Δ and p indicates the probability calculated under either the flat or the hierarchical model (i.e. $M = F$ or $M = H$) and D^* indicates the vector of decision data, C_t , or the sequence of button presses. This function was maximized using nonlinear function maximization techniques in Matlab. Examination of the likelihood function around the maximum likelihood estimates showed that likelihood values dropped off smoothly and in a relatively Gaussian way as the parameter values changed. We also tried multiple starting values for the parameters (0, 0.1, 0.3 and 0.5) and found that the algorithm always converged to the same answer. Thus, our estimates were not likely local minima. In part this could be due to the large amount of data that was collected for each subject.

Two of the subjects performed very poorly in the task (subjects 5 and 8) and their parameter values reflected this as they had values near zero for both α and β , implying no integration of the feedback. Once the parameters were found that maximized the likelihood, either equation 2 ($P^F(B)$) or equation 5 ($P^H(B)$) was used to generate the probability, on a trial by trial basis, that the subject knew which button to press at each point in the sequence, and equation 4 was used to estimate the sequence probability ($P(S)$), given the feedback from the current block.

Finally, for model comparison on the behavioral data the log-likelihood ratio for the two models was calculated as:

$$llr = \log p^F(D^*|\alpha, \beta) - \log p^H(D^*|\alpha, \beta). \quad (9)$$

Because the models had the same number of parameters, no corrections for degrees of freedom had to be done. We could simply compare the fit of the two models on a subject-by-subject basis and look at the distribution of the llr across subjects to see if it was significantly positive or negative.

We also examined three other models which were extensions of the basic model. None of them fit the subject behavior better than the flat model and therefore we will not go into the formal description of the models. However, it is worth describing them briefly as we did explore various possibilities, beyond the flat and hierarchical models, for how subjects might have been performing the task. We were trying to model the possibility that subjects were integrating information over the first few trials of the block, but at some point they decided which sequence they thought was correct in the block. In that case, the subjects' probability estimate would jump from its current value to 1. Effectively, we were trying to

model the possibility that when the probability estimate crossed a threshold the subjects were sufficiently confident of which sequence they were executing and their belief estimate went to 1 or some very high value. We did this in three ways. First, we passed the button press probabilities (equation 2) generated by the basic model through a soft-max function (Bishop, 1995). This function is used to convert value estimates in reinforcement learning models into probabilities. It contains a temperature parameter which controls how much probabilities are amplified. In effect, it causes the probabilities during learning to go much more quickly to 1. We estimated the temperature as a free parameter, but found that it did not improve the fit. The second approach was to implement a threshold. In this case, when the belief value crossed that threshold it was set to 1. In the third case we switched from the button probability generated by equation 2 to the sequence button probability generated by equation 5 when the probability passed a threshold. The threshold was always allowed to vary as a free parameter. As stated, however, none of these 3 approaches worked better than the basic button model.

MEG data acquisition and pre-processing

Magneto-encephalographic (MEG) data were recorded using 275 3rd order axial gradiometers with the Omega275 CTF MEG system (VSM medtech, Vancouver, Canada) located in a magnetically shielded room. The signals were recorded at a sampling rate of 480 Hz. Visual stimulus lag was estimated using a photo-diode to measure the onset time at the screen relative to the signal sent from the task control computer to the data acquisition computer and was found to be 25 ms. This delay was used in calculating all of our timing values. Data analysis was carried out using SPM5 (Wellcome Department of Imaging Neuroscience) and custom written Matlab routines to implement the behavioral model.

We began the analyses by low-pass filtering the MEG signal at 50 Hz and downsampling to 120 Hz. The data were then epoched into a 400 ms window centered on the button press. Trials with a response which had an absolute value greater than 3000 Femto-Tesla were discarded as outliers. Statistical analysis and source localization was carried out using SPM. Details of the GLM model that was fit are given below. Gaussian random field theory was used to control for multiple comparisons in either 2-D space x 1-D time (sensor space) or 3-D space (source space) (Kiebel, Kilner, & Friston, 2007; Kilner, Kiebel, & Friston, 2005). Sensors were converted into a 2-D space using Gaussian interpolation. Smoothing was never done across both time and space. Thus, the signals were first filtered in time and then filtered in space at a single time point.

When relevant sensor space effects were identified, we estimated the sources of these effects using reconstruction algorithms in SPM5. For each subject, we constructed a forward model describing the transformation between distributed dipole sources and the magnetic field distribution measured by the MEG sensors. Sources were modeled using the 7204 vertex template cortical mesh available in SPM5, defined in Talairach and Tournoux coordinates. It was coregistered to the sensor locations via three fiducial marker positions (Mattout, Henson, & Friston, 2007) and the gain matrix of the lead-field model was then computed using a spherical head model and source estimates were computed using Restricted Maximum Likelihood estimation to invert the forward model within a parametric empirical Bayes framework (Mattout, Phillips, Daunizeau, & Friston, 2007; Mattout, Phillips, Penny, Rugg, & Friston, 2006). This inversion proceeded using multiple sparse empirical priors for covariance components (Friston, Harrison et al., 2008). The greedy search algorithm (Friston, Chu et al., 2008) provided an optimal mixture of sparse prior components. This produced source reconstructions for each experimental condition and for each subject.

We then compared activation levels on the mesh across subjects using a random effects model (i.e. a second level model in SPM). Voxels were accepted at an uncorrected p value

of 0.01 and all significance values are reported at cluster level corrected for whole brain. Sources were estimated in a 2x2 design, probability x button, where probability was low or high (i.e. early or late learning). Significance of the probability factor was calculated by first computing main effect contrasts subject by subject and then doing univariate t-tests.

For the GLM regressions which looked for an effect of the learning related probabilities on the MEG sensor response we carried out analyses in two ways. First, we examined the probability of the button which was actually pressed, under both the flat and hierarchical models. Thus, the first level regressor was given as:

$$n(t) = \Delta_t^M C_t + (1 - \Delta_t^M)(1 - C_t). \quad (10)$$

Interestingly, this regressor was not significant under either model (data not shown). We then used a regressor which had only the probability of the correct button press or sequence at the current point in time. This was given by:

$$m(t) = \Delta_t^M B_{i,j} + (1 - \Delta_t^M)(1 - B_{i,j}). \quad (11)$$

This regressor is what is plotted in Fig. 4. It tracks learning more closely and is the same whether or not the most probable or the least probable button was pressed. The sequence probability given by equation 4 was also used, in which case the currently correct sequence was plugged in as the r.h.s. of equation 11.

When carrying out these analyses, we correlated the regressor, $m(t)$, on a movement by movement basis with the time-varying MEG responses at each sensor and each time point in the 400 ms window around button press, similar to what is done in fMRI data with voxels. In our case, we treated the 2-D sensor space by 1-D time space as a 3-D volume, as described above. Thus, we fit the following GLM to each point in the 3-D volume

$$q_{x,y,l} = a_0 + b m(t) + c g(t) + d m(t) g(t) \quad (12)$$

Where $q_{x,y,l}$ indicates the response or MEG equivalent voxel in the 3-D volume (i.e. x and y = position and l = time around button press) and $g(t)$ is the button press for that movement, which is a dummy variable with values -1 and 1 . The last parameter is an interaction term. This model was fit to individual subjects and then the parameter estimates for the individual subjects were taken to the second level in an SPM random effects approach.

Results

Behavior

Fourteen subjects carried out the stochastic sequence learning task (Fig. 1), in which they had to learn sequences of four left/right button presses (e.g. LLRR, LRLR, etc.). Explicit feedback was given after each button press, but 15% of the time inaccurate feedback was given. Two of the subjects failed to learn the sequences to criterion, so their results will not be discussed further. It took the remaining 12 subjects, on average, 3.2 trials of learning with each new sequence before they executed a complete trial correctly, where a correct trial was defined as pressing all four buttons in the sequence correctly (Fig. 2A). The subjects' performance reached a plateau by about 3 trials correct (about 6 total, i.e. combined correct and error trials) in each block and they executed the sequence in the remaining trials with few errors. When the performance was examined as the serial position of the movement,

results were similar. It could be seen, however, that the early movements of the sequence were learned slightly faster (Fig. 2B). There was also a small bias to perform better on the early and last movements of the sequence between correct trials 3 and 4 (Fig. 2B), consistent with the primacy and recency gradients seen in most sequence tasks (Averbeck, Chafee, Crowe, & Georgopoulos, 2002). However, later in the block the performance reached a ceiling and this effect could not be seen.

In the task, the subjects made a sequence of four left or right button presses. Thus, there were 16 possible sequences (2^4). However, we only used 6 of the 16 possible sequences in the experiment. Because we only used a subset of the possible sequences, there were correlations between the correct button presses and feedback about button presses at other points in the sequence could be used to better infer the correct button at the current point in the sequence. For example, if in a particular block the subject was certain that the first two button presses were LL, they would predict that the subsequent button presses were RR even if the evidence for RR was equivocal, because LLRR was the only sequence we used that started with LL. These correlations between button presses can be represented with a hierarchical structure.

We explored the hypothesis that subjects took this hierarchical or correlational structure into account when they were learning the sequences. To do this, we predicted the subjects' behavior with 2 different models, one which did not take the hierarchical structure into account (Fig. 3A – Flat model), and one which did (Fig. 3B – Hierarchical model). Behavioral parameters were optimized for each subject under each model. We found that subjects learned more from positive feedback than negative feedback under both models (Fig. 3C). The parameter values were also lower for the hierarchical model than the flat model. This is due to the fact that the hierarchical model is more efficient with the data, because it correctly models the actual stochastic process used in the experiment. Because parameters were optimized under each model, learning rates for the two models were similar, although the hierarchical model tended to learn more smoothly across trials, because it integrated information across buttons, as can be seen in a single example block from a single subject (Fig. 4).

We next examined whether we could find differences in the ability of the two models to predict the behavior of the subjects. We did this by computing a t-test on the log-likelihood ratio of the two models across subjects. A negative log-likelihood ratio for a single subject favored the hierarchical model. There was, however, no significant difference between the models, across our subjects ($t_{11} = 1.8$, $p = 0.09$). We did find that the behavior of the subjects who learned better was better described by the hierarchical model. Thus, the total number of trials it took subjects to finish the task, which is a measure of how efficiently they learned, was negatively correlated with the log-likelihood ratio ($r = -0.76$, $p < 0.01$, $n = 12$). In other words, the hierarchical model better predicted the behavior of subjects who learned more efficiently.

MEG responses – Parametric effect of probability

We measured magneto-encephalographic (MEG) responses while subjects learned and executed the sequences. We began by examining effects in sensor space, i.e. by looking at task effects on the temporal response in the interpolated scalp map around the time of button press (200 ms before to 200 ms after button press). All analyses in sensor space were carried out by converting the 2-D sensor x 1-D time data into a 3-D volume, which allowed us to carry out corrections for multiple comparisons using the tools developed for fMRI data (see methods). Reported statistical results are based on clusters of samples which exceeded a threshold.

In our analyses of the MEG data we compared the relative ability of three different learning related variables, taken directly from the behavioral models discussed above, to predict the change in the MEG response with learning (Fig. 4). We examined the button probabilities under the flat ($P^F(B)$) and hierarchical models ($P^H(B)$), and we also considered the probability of the sequence ($P(S)$). In all cases these probabilities were derived on a movement-by-movement basis with the models fit to the behavioral data of the individual subjects, including separate modeling of positive and negative feedback, as described above. These probabilities represent how much has been learned about the sequence, prior to the current button press. Thus, they represent the subjects' knowledge of which button should be pressed, or which sequence is correct, in the current block. We also included the button which was pressed and the interaction between the probability and the button which was pressed in the analysis (see equation 12 in methods for first-level model).

As we were interested in studying the learning effect, we only carried out our analysis on blocks in which the subjects completed 8 sequences correctly, although results in sensor space were similar when all blocks were used (data not shown). There were no significant clusters for the button probability under the flat model ($P^F(B)$; $p = 0.24$, t-test, cluster level, $df = 11$) and there was one cluster which just missed significance under the button probability for the hierarchical model ($P^H(B)$; $p = 0.052$, t-test, cluster level, $df = 11$). The sequence effect, however, had two significant clusters ($P(S)$; $p < 0.05$, t-test, cluster level, $df = 11$), one had a maximum at 83 ms before button press (Fig. 5) and extended above threshold from 108 ms to 67 ms before button press. The other had a maximum at 67 ms after button press (Fig. 6) and extended above threshold from 58 ms after until 150 ms after button press. There was also a significant interaction between sequence probability and the button which was pressed in the cluster which followed the button press ($p < 0.05$, t-test, cluster level, $df = 11$, max at 67 ms, above threshold from 58 to 100 ms after button press), but not for the cluster which preceded the button press.

These modulations of the MEG response manifested as changes in the temporal evolution of the signal just before and after button press (Fig. 5B), i.e. the temporal response was different depending on how well the sequence had been learned in the current block. In all cases there was an increase in the signal early in the block, when the sequence was not well learned ($P(S) = 0.33$) leading to a decrease later in the block when the sequence was well learned ($P(S) = 1.0$). This could also be seen in sensor space, as the parametric regressor was negative in the region with significant effects (Fig. 5C). Thus, as the probability increased across the block (Fig. 4) the MEG signal decreased.

As there were differences in how well the hierarchical and non-hierarchical models predicted the behavior of individual subjects, we were interested in whether there would be correlations between the relative fit of the models, and how well the sequence model fit the behavior of individual subjects. To test this we examined the correlation of the relative fit of the models, measured with the log-likelihood ratio, and the contrast estimates for individual subjects. We did not, however, find that there were significant correlations, after corrections for multiple comparisons, in the fit of the models and the contrast estimates for the sequence model.

Given that these effects were only significant under the sequence model and not the flat model, we next assessed whether or not the contrast estimates (i.e. the parametric regressors from the GLM) were significantly larger for the sequence model than the flat model by comparing the distribution of contrast values across subjects between the two models at the peak location, 83 ms before and 67 ms after button press. Neither of these distributions were significantly different ($p = 0.55$, unequal variance t-test, $n = 24$). We did, however, find that there was a significant difference in the variance of the second level contrast distribution at

83 ms before ($p < 0.05$, F test, $df = 11, 11$) and an almost significant difference at 67 ms after ($p = 0.056$, F test, $df = 11, 11$) between the sequence model and the flat model. Thus, the increased significance under the sequence model is due to lower variance in the second level contrasts, as opposed to a larger mean of the parametric modulator.

Next, we used source localization to identify the possible locations of the sequence probability effects. Significant effects in source space were assessed by estimating distributed activation levels on a cortical mesh for individual subjects, and then carrying out second level statistics in SPM on these activation levels. We first assessed the source of the probability effect seen at -83 ms, using a window from -117 ms to -67 ms. We used a window slightly larger than the window over which the sensor effect was significant as the algorithm rarely converged for small windows. We found a significant source ($p < 0.05$) bilaterally in the insula, just lateral to the striatum (Fig. 5D).

Next, we carried out source localization for the significant effect 67 ms after button press, using a window from 50 to 150 ms after button press. We found a significant source for the probability effect at this time in premotor cortex, (Fig. 6D; $p < 0.05$). There was an additional significant source ($p < 0.05$) in early visual cortex bilaterally (left side $x = -28$, $y = -82$, $z = -14$), perhaps reflecting an attentional effect on the assessment of the feedback and a source at the frontal pole ($x = 12$, $y = 64$, $z = -12$).

Discussion

We examined the behavioral and neural correlates of learning in a stochastic sequence learning paradigm. Comparing flat and hierarchical behavioral models suggested that both predicted the subjects' decisions equally well. However, only the sequence probability from the hierarchical model resulted in significant correlations with the MEG signal and this significance was due to less variance in the contrast estimates across subjects. When we localized the learning related signals we found a cluster of activity in the insula that preceded the button press and a cluster in the premotor cortex that followed the press.

Behavior

Our sequence task had important features that allowed us to assess how well subjects were learning and whether or not they were optimal. The learning coefficients show that subjects were not optimal, as positive and negative feedback should have been weighted similarly, whereas in fact subjects relied more on positive feedback as has been seen previously in sequence learning (Averbeck, Sohn, & Lee, 2006). In the task, a hierarchical structure is optimal, whereas a flat structure is not. The behavioral data suggested that subjects that learned better tended to learn in a more hierarchical manner. Perhaps additional training on the sequences would have benefitted the subjects who did not learn as well. Future experiments could clarify this point.

Imaging results

Two different learning related signals emerged in the MEG sensor data, one just before button press and one just after button press. The early signal was near the mid-line, whereas the later signal was lateralized over the right side, although the left side signal may have been just below significance. When we carried out source localization on these two signals we found a source in the insula for the early signal and a source in premotor cortex for the later signal. Some caution in interpreting these results is necessary, however, as it is difficult to know how precisely the MEG sensor signals can be localized.

Previous fMRI studies have shown activation in the insula during motor learning (Floyer-Lea & Matthews, 2004) and this area has a direct projection to the striatum (Chikama,

McFarland, Amaral, & Haber, 1997) and as such it likely takes part in a network of areas related to updating actions on the basis of feedback. Previous work has also shown activity in this area during outcome anticipation that is either negatively or positively valenced (Critchley, Mathias, & Dolan, 2001; Knutson & Greer, 2008; Volz, Schubotz, & von Cramon, 2004). Although the MEG signal which was localized to the insula preceded the button press by 83 ms, it is likely that the press has already been initiated at the cortical level at this time. Therefore, this signal may represent anticipation of either a red or green outcome, where the anticipation is scaled by how much the subject has learned in the block. Once the sequence is well learned green feedback is highly likely. It is also interesting that, unlike the signal which follows the button press, this signal was not modulated by the button which was pressed as we did not find an interaction effect in the sensors. This makes it unlikely that this signal was directly involved in learning, as there was no information about the action. This signal may be more related to one's subjective sense of progression through the block, as many studies implicate the anterior insula in subjective interoception (Craig, 2009).

Activation in premotor cortex has been seen in tasks with hierarchical structure (Koechlin & Jubault, 2006; Schubotz & von Cramon, 2003, 2004), and the sequence probability is hierarchical as it represents the entire series of button presses in the order that they unfold. As there is an interaction between sequence probability and the button which was pressed, this signal may have a more direct role in updating the probability information based on the feedback.

A series of studies by Koechlin and colleagues have also suggested that, when tasks have explicit hierarchical structure, task factors that map to different levels of the cognitive hierarchy map to different locations in frontal cortex (Badre & D'Esposito, 2007; Koechlin & Jubault, 2006; Koechlin, Ody, & Kouneiher, 2003; Koechlin & Summerfield, 2007). The behaviour being studied in their work, however, differs in important ways from the behaviour we have studied. Specifically, the studies by Koechlin and colleagues did not examine decision making in a framework where subjects had to deal with uncertainty about the relationship between actions and outcomes. Rather, the previous studies used rule based cognitive tasks where the link between stimulus/action/feedback was deterministic given the behavioural rule. Stochasticity in these tasks was implied by the frequency with which the rule changed across blocks. However, all of the information was always provided by the task, and the mapping between actions and outcomes was deterministic. This is very different from the task we have utilized, which required subjects to deal with uncertainty in an effort to infer the sequence (rule) which was in operation. In our task, even if one knew which sequence was correct in a particular block, one would not be able to predict the feedback on an individual button press. As such, different cognitive processes are likely required to solve our task. The advantage of our approach is that we were able to test directly whether or not subjects were using hierarchical or flat statistical models when learning the sequences. Thus, we have provided imaging evidence for hierarchical control in a task which could have been solved using a flat model, although the behavioural data was more equivocal. It is not clear what the alternative model would be in the tasks used by Koechlin and colleagues. In their experiments, however, it is less of an issue as they were not studying learning but rather performance of a complex cognitive task.

One Important caveat to the model comparison approach we have taken, with respect to the imaging data, is that we have examined significance by linearly correlating the model prediction with the MEG signal (Behrens, Hunt, Woolrich, & Rushworth, 2008; Behrens, Woolrich, Walton, & Rushworth, 2007; Daw et al., 2006; Wittmann et al., 2008). Linear correlations, however, do not allow one to infer conclusively that the underlying neural responses are necessarily favoring one model over the other. More specifically, the sequence

probability is a nonlinear function of the button probability under the flat model. As such, there cannot be more information in the underlying neural response about the sequence variable than there is about the button probability due to the data processing inequality (Cover & Thomas, 1991). Thus, our inference relates only to the specific functional form that we have examined, the linear relationship, and does not tell us about the detailed neural representation of this probability. For this, single unit studies can be more valuable, partly for practical reasons. It is in many cases possible to examine the relationship between single unit firing rate responses and various behavioral variables graphically, and fit models accordingly. Also, there are often more trials available for fitting more complex models. The high-dimensionality of MEG data makes examining the relationship between time varying signals and task variables highly complex. Interestingly, the data from the single unit studies has consistently shown, in many brain areas and in many similar tasks, that sequence information is explicitly represented in the brain (Averbeck et al., 2006; Nakamura, Sakai, & Hikosaka, 1998; Tanji, 2001).

Comparison with single unit studies

One of the goals of the present study was to approach a question we have already examined in macaques at the single cell level (Averbeck et al., 2006) in humans using an imaging approach. We had originally intended to use a task that was as similar as possible to the task used in the macaque study. However, unlike macaques which require several trials to learn a 3 movement sequence with explicit feedback, human participants given a 4 movement sequence learn it in about 1 trial (unpublished data). This rapid learning makes studying the learning process difficult and this is why we made the task more difficult by introducing the stochastic feedback.

The second difference has to do with the nature of the information that can be extracted from single unit data vs. MEG imaging data. Specifically, in the macaque study, we were able to track learning by following the emergence of a signal in single neurons that explicitly represented the sequence which was correct in the current block. However, we were not able to extract sequence specific information from the MEG signal (unpublished data). Thus, we had to use a different approach to examine the learning related changes in neural activity. Given this difference, however, the premotor signal which follows the button press is in many respects comparable to the location we studied in the macaque, as the activity in the macaque was just anterior to the frontal eye fields and we used eye movements as our behavioral output in the macaque. Thus premotor cortex and caudal area 46 may have similar functions for different effectors.

Conclusion

In conclusion, we found that, when subjects learned efficiently they learned hierarchically. Furthermore, the learning related variable which most strongly correlated with the imaging data was the probability of the sequence, a parameter which is present in the hierarchical model, but not in the flat model. Thus, the imaging data and to some extent the behavioural data suggest that when efficient subjects were faced with learning a sequence which had hierarchical structure, they were able to take advantage of that structure.

References

- Andreasen NC. Thought, language, and communication disorders. I. Clinical assessment, definition of terms, and evaluation of their reliability. *Arch Gen Psychiatry*. 1979a; 36(12):1315–1321. [PubMed: 496551]
- Andreasen NC. Thought, language, and communication disorders. II. Diagnostic significance. *Arch Gen Psychiatry*. 1979b; 36(12):1325–1330. [PubMed: 496552]

- Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP. Parallel processing of serial movements in prefrontal cortex. *Proc Natl Acad Sci U S A*. 2002; 99(20):13172–13177. [PubMed: 12242330]
- Averbeck BB, Sohn JW, Lee D. Activity in prefrontal cortex during dynamic selection of action sequences. *Nat Neurosci*. 2006; 9(2):276–282. [PubMed: 16429134]
- Badre D, D'Esposito M. Functional Magnetic Resonance Imaging Evidence for a Hierarchical Organization of the Prefrontal Cortex. *J Cogn Neurosci*. 2007
- Behrens TE, Hunt LT, Woolrich MW, Rushworth MF. Associative learning of social value. *Nature*. 2008; 456(7219):245–249. [PubMed: 19005555]
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci*. 2007; 10(9):1214–1221. [PubMed: 17676057]
- Bishop, CM. *Neural Networks for Pattern Recognition*. 1 ed. Oxford University Press; Oxford: 1995.
- Chikama M, McFarland NR, Amaral DG, Haber SN. Insular cortical projections to functional regions of the striatum correlate with cortical cytoarchitectonic organization in the primate. *J Neurosci*. 1997; 17(24):9686–9705. [PubMed: 9391023]
- Cover, TM.; Thomas, JA. *Elements of Information Theory*. Jon Wiley and Sons, Inc.; New York: 1991.
- Craig AD. How do you feel -- now? The anterior insula and human awareness. *Nat Rev Neurosci*. 2009; 10:59–70. [PubMed: 19096369]
- Critchley HD, Mathias CJ, Dolan RJ. Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron*. 2001; 29(2):537–545. [PubMed: 11239442]
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature*. 2006; 441(7095):876–879. [PubMed: 16778890]
- Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*. 2002; 415(6870):429–433. [PubMed: 11807554]
- Fellows LK, Farah MJ. The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment per se? *Cereb Cortex*. 2007; 17(11):2669–2674. [PubMed: 17259643]
- Fiebach CJ, Schubotz RI. Dynamic anticipatory processing of hierarchical sequential events: a common role for Broca's area and ventral premotor cortex across domains? *Cortex*. 2006; 42(4):499–502. [PubMed: 16881258]
- Floyer-Lea A, Matthews PM. Changing brain networks for visuomotor control with increased movement automaticity. *J Neurophysiol*. 2004; 92(4):2405–2412. [PubMed: 15381748]
- Friston K, Chu C, Mourao-Miranda J, Hulme O, Rees G, Penny W, et al. Bayesian decoding of brain images. *Neuroimage*. 2008; 39(1):181–205. [PubMed: 17919928]
- Friston K, Harrison L, Daunizeau J, Kiebel S, Phillips C, Trujillo-Barreto N, et al. Multiple sparse priors for the M/EEG inverse problem. *Neuroimage*. 2008; 39(3):1104–1120. [PubMed: 17997111]
- Gilovich, T.; Griffin, D.; Kahneman, D., editors. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press; Cambridge, UK: 2002.
- Glimcher PW, Rustichini A. Neuroeconomics: the consilience of brain and decision. *Science*. 2004; 306(5695):447–452. [PubMed: 15486291]
- Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF. Neural systems responding to degrees of uncertainty in human decision-making. *Science*. 2005; 310(5754):1680–1683. [PubMed: 16339445]
- Jacobs RA. Optimal integration of texture and motion cues to depth. *Vision Research*. 1999; 39(21):3621–3629. [PubMed: 10746132]
- Kahneman, D.; Slovic, P.; Tversky, A. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press; Cambridge: 1982.
- Kersten, D. High level vision as statistical inference. In: Gazzaniga, MS., editor. *The New Cognitive Neurosciences*. MIT Press; Cambridge: 1999.
- Kiebel, SJ.; Kilner, JM.; Friston, KJ. Hierarchical models for EEG and MEG. In: Friston, KJ.; Ashburner, JT.; Kiebel, SJ.; Nischols, TE.; Penny, WD., editors. *Statistical parametric mapping: The analysis of functional brain images*. Academic Press; London: 2007. p. 211–222.

- Kilner JM, Kiebel SJ, Friston KJ. Applications of random field theory to electrophysiology. *Neurosci Lett*. 2005; 374(3):174–178. [PubMed: 15663957]
- Knill, DC.; Richards, W. Perception as Bayesian Inference. Cambridge University Press; Cambridge: 1996.
- Knill DC, Saunders JA. Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research*. 2003; 43(24):2539–2558. [PubMed: 13129541]
- Knutson B, Greer SM. Anticipatory affect: neural correlates and consequences for choice. *Philos Trans R Soc Lond B Biol Sci*. 2008; 363(1511):3771–3786. [PubMed: 18829428]
- Koechlin E, Jubault T. Broca's area and the hierarchical organization of human behavior. *Neuron*. 2006; 50(6):963–974. [PubMed: 16772176]
- Koechlin E, Ody C, Kouneiher F. The architecture of cognitive control in the human prefrontal cortex. *Science*. 2003; 302(5648):1181–1185. [PubMed: 14615530]
- Koechlin E, Summerfield C. An information theoretical approach to prefrontal executive function. *Trends Cogn Sci*. 2007; 11(6):229–235. [PubMed: 17475536]
- Kording KP, Wolpert DM. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Science*. 2006; 10(7):319–326.
- Mattout J, Henson RN, Friston KJ. Canonical source reconstruction for MEG. *Computational Intelligence and Neuroscience*. 2007
- Mattout, J.; Phillips, C.; Daunizeau, J.; Friston, KJ. Bayesian inversion of EEG models. In: Friston, KJ.; Ashburner, JT.; Kiebel, SJ.; Nichols, TE.; Penny, WD., editors. Statistical parametric mapping: The analysis of functional brain images. Academic Press; London: 2007.
- Mattout J, Phillips C, Penny WD, Rugg MD, Friston KJ. MEG source localization under multiple constraints: an extended Bayesian framework. *Neuroimage*. 2006; 30(3):753–767. [PubMed: 16368248]
- McKenna, P.; Oh, T. Schizophrenic Speech: Making Sense of Bathrooms and Ponds that Fall in Doorways. Cambridge University Press; Cambridge: 2005.
- Nakamura K, Sakai K, Hikosaka O. Neuronal activity in medial frontal cortex during learning of sequential procedures. *J Neurophysiol*. 1998; 80(5):2671–2687. [PubMed: 9819272]
- Payne JW, Bettman JR, Johnson EJ. Behavioral decision research: a constructive processing perspective. *Annual Review of Psychology*. 1992; 43:87–131.
- Penfield W, Evans J. The frontal lobe in man: a clinical study of maximum removals. *Brain*. 1935; 58:115–133.
- Poggio T, Torre V, Koch C. Computational vision and regularization theory. *Nature*. 1985; 317(6035):314–319. [PubMed: 2413361]
- Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science*. 2005; 310(5752):1337–1340. [PubMed: 16311337]
- Schubotz RI, von Cramon DY. Functional-anatomical concepts of human premotor cortex: evidence from fMRI and PET studies. *Neuroimage*. 2003; 20(Suppl 1):S120–131. [PubMed: 14597305]
- Schubotz RI, von Cramon DY. Sequences of abstract nonbiological stimuli share ventral premotor cortex with action observation and imagery. *J Neurosci*. 2004; 24(24):5467–5474. [PubMed: 15201318]
- Schwartz MF, Reed ES, Montgomery M, Palmer C, Mayer NH. The quantitative description of action disorganization after brain damage: a case study. *Cognitive Neuropsychology*. 1991; 8:381–414.
- Tanji J. Sequential organization of multiple movements: involvement of cortical motor areas. *Annu Rev Neurosci*. 2001; 24:631–651. [PubMed: 11520914]
- Tobler PN, Christopoulos GI, O'Doherty JP, Dolan RJ, Schultz W. Neuronal distortions of reward probability without choice. *J Neurosci*. 2008; 28(45):11703–11711. [PubMed: 18987206]
- Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*. 2002; 5(11):1226–1235.
- Trommershauser J, Maloney LT, Landy MS. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A Opt Image Sci Vis*. 2003a; 20(7):1419–1433. [PubMed: 12868646]

- Trommershauser J, Maloney LT, Landy MS. Statistical decision theory and trade-offs in the control of motor response. *Spat Vis.* 2003b; 16(3-4):255–275. [PubMed: 12858951]
- Trommershauser J, Maloney LT, Landy MS. Decision making, movement planning and statistical decision theory. *Trends Cogn Sci.* 2008; 12(8):291–297. [PubMed: 18614390]
- Tversky A, Kahneman D. Rational choice and the framing of decisions. *Journal of Business.* 1986; 59(4):S251–S278.
- Volz KG, Schubotz RI, von Cramon DY. Why am I unsure? Internal and external attributions of uncertainty dissociated by fMRI. *Neuroimage.* 2004; 21(3):848–857. [PubMed: 15006651]
- Wittmann BC, Daw ND, Seymour B, Dolan RJ. Striatal activity underlies novelty-based choice in humans. *Neuron.* 2008; 58(6):967–973. [PubMed: 18579085]

Fig. 1.

Task. A. Top shows sequence of images presented in a single trial. Bottom shows timing of events. Each cue, button press, feedback series is presented 4 times in a single trial. B. Sequence of sets, blocks and trials. Each set is composed of 6 sequence blocks. Subjects progress through a single sequence block by learning and executing the sequence correctly 8 times. C. Sequence of analyses presented in results.

Fig. 2.

Learning rate. A. Fraction correct complete sequences (i.e. 4 correct button presses) and number of trials it takes to get one trial correct. Both are plotted as a function of the number of correct trials in the current block. B. Fraction correct individual button presses for first through fourth button press in each sequence.

Fig. 3.

Probability model of learning. A. Flat architecture, in which feedback about individual movement affects only that movement. B. Hierarchical architecture, in which feedback about individual movements is used to infer the sequence, which is then used to infer the movements. C. Mean and SEM of parameter values across the 12 subjects for positive (α) and negative (β) feedback under the two models.

Fig. 4. Probabilities for an example block under flat and hierarchical models. Probability that the feedback up to the current trial supported the correct button press under the flat model (orange line), the correct button press under the hierarchical model (blue line) or the correct sequence (green line) at each point in the sequence, as it evolves within a single example block of trials. Red lines under x-axis indicate the movements on which the subject pressed the incorrect button and asterisks indicate the movements in which the incorrect feedback was given.

Fig. 5. Significant sensor modulation with probability 83 ms before button press based on sequence model. A. Interpolated map of p-values in sensor space. Only p-values below 0.001 are shown, and all values plotted are the negative log of the p-value (i.e. a value of 2 is 0.01). B. Temporal response at a single sensor (indicated in panel A with white circle). Each line represents predicted response averaged across subjects for one probability, with probabilities indicated on right of panel. C. Contrast of parametric regressor representing probability main effect in the general-linear model. D. Estimated source of effect. Plot shows a significant cluster of voxels (peak activation: $x = -30$, $y = 2$, $z = 10$) in the insula. Clusters were significant bilaterally.

Fig. 6. Significant sensor modulation with probability 67 ms after button press. A. Interpolated map of p-values in sensor space. Only p-values below 0.001 are shown, and all values plotted are the negative log of the p-value (i.e. a value of 2 is 0.01). B. Temporal response at a single sensor (indicated in panel A). C. Contrast of probability main effect in the general-linear model. D. Estimated source of effect. Plot shows a significant cluster of voxels (peak activation: $x = -50$, $y = 6$, $z = 24$) in Premotor cortex.