

Neural Representations of Procedural Knowledge

Robert A. Mason and Marcel Adam Just

Center for Cognitive Brain Imaging, Carnegie Mellon University

Abstract

Although declarative concepts (e.g., *apple*) have been shown to be identifiable from their fMRI signatures, the correspondence has yet to be established for executing a complex procedure like tying a knot. Seven participants were trained to tie seven knots. Their neural representations of these seven procedures were assessed with fMRI as they imagined tying each knot. A subset of the trained participants physically tied each knot in a later fMRI session. We demonstrate that: (1) procedural knowledge of tying a particular knot is reliably identified from its fMRI signature; (2) procedural signatures were found in frontal, parietal, motor, and cerebellar regions; and (3) planning to tie knots was reliably identifiable prior to physical tying, using a classifier trained on mental tying signatures, indicating the similarity of the procedural signatures in the two cases. The findings indicate that fMRI activation patterns can illuminate the representation and organization of procedural knowledge.

Keywords

neural representations, procedural knowledge, skill learning, fMRI, knot-tying

Tying a knot is an ancient and useful human act that is the epitome of everyday procedural knowledge. Making an appropriate knot remains central to the work of sailors, fishermen, and surgeons, among others. Knot-tying requires the coordination of several different types of processing, such as planning, motor activity, and spatial processing, involving the use of a diverse set of neural resources. The goal of this research was to characterize the neural representations of the procedural knowledge of how to tie seven individual knots.

Procedural knowledge has been a focus of neuroimaging from the beginnings of PET imaging (Grafton et al., 1992) and the nature of the underlying neural representations continues to be addressed using fMRI (for a review see Grafton, 2010). More specifically, neuroimaging studies have investigated knot tying, generally focusing on which brain regions are involved or on the effect of learning (Cross, Hamilton, Cohen, & Grafton, 2017). Involved regions can be related to motor control (post-central gyrus), object manipulation and visualization (intra-parietal sulcus (IPS) and other parietal regions) and object encoding (fusiform gyrus) (Cross et al., 2012). The involvement of cortical regions can be modulated by the degree of knot learning: shallow learning associated with parietal cortex activation and deeper learning associated with medial structures like posterior cingulate and precuneus (Tracy et al., 2003). Although these prior studies are informative about which regions may be involved in knot-tying in general, they do not focus

on the neural representations of knot-tying procedures for individual knots. They do not tell us where or how individual procedures are neurally represented.

Multi-voxel pattern analysis (MVPA) has the potential to identify the neural signature of the tying of a specific knot. In contrast to univariate analyses that consider the activity of each voxel independently of any distal voxels, MVPA can identify a spatially-distributed activation pattern that is associated with a particular concept. In the case of a declarative concept such as *apple*, it is possible to specify that the concept is neurally represented by a pattern of activation levels in a set of voxels distributed over multiple brain systems (Carota, Kriegeskorte, Nili, & Pulvermüller, 2017; Huth, Nishimoto, Vu, & Gallant, 2012; Just, Cherkassky, Aryal, & Mitchell, 2010; Mitchell et al., 2008). This pattern is the concept's fMRI signature.

Unlike the representation of a declarative concept, the neural representation of a procedure should contain information about a mental process that unfolds over some time period while a sequence of actions is executed. The current study attempts to uncover the neural representation of the temporally-unfolding procedure for tying several different types of knots and to identify the brain structures involved.

To examine procedural knot-tying knowledge, participants with little previous knot-tying experience (besides shoelaces and square knots) received computer-based video training in tying knots. This training technique

has proven to be effective for learning how to tie knots (Brandt & Davies, 2006; Rogers, Regehr, Yeh, & Howdieshell, 1998; Schwan & Riempp, 2004) and imagined knot-tying is one of the training techniques surgeons use for teaching suturing (Torkington, Smith, Rees, & Darzi, 2000). Neural representations were assessed using a machine learning technique applied to the fMRI data collected while the participants imagined tying each knot after being prompted with the name and a picture of the knot.

A subset of participants also physically tied the knots. Although there are differences between motor imagery and movement execution there are some common underlying brain networks (see Hardwick, Caspers, Eickhoff, & Swinnen, 2018 for a meta-analysis). For example, simulation theory postulates a common network for motor imagery, execution and observation (Jeannerod, 2001). Other theories suggest the use of the same network for simulation and execution (Guillot & Collet, 2005), and still others suggest action simulation more heavily loads on executive functions (Glover & Baran, 2017). Our approach was to empirically test whether a classifier trained on the activation patterns associated with each knot during mental tying could identify the knots on the basis of the activation pattern evoked at the beginning of physical knot tying (just prior to movement initiation). The inclusion of physical tying trials enables a cross-task comparison of the procedural representations.

The goal of the current study was to determine the neural representations of the procedure of tying of individual knots using a machine-learning, cross-validation, prediction approach (Breiman, 2001). This type of analysis has an advantage over hypothesis testing and statistical regression modeling for complex tasks with a potentially large number of predictor variables. Here we are testing the ability to identify a neural representation of tying a knot based on the concomitant brain activation patterns.

The usefulness of MVPA for exploring procedural rules has recently been demonstrated in motor-action tasks that differ in level of abstraction (opening a bottle vs opening with particular tools vs opening different containers (Wurm & Lingnau, 2015) or when participants had to perform hand movements like aiming, squeezing, or extension (Pilgramm et al., 2016; Zabicki et al., 2016). Zabicki et al (2016) is particularly informative in the use of both Representational Similarity Analysis (RSA) and MVPA to test models of imagery, execution, and cross modality to examine the type of processing in various cortical regions. They note that “One conclusion for the similarity of executed and imagined actions therefore might be that the similarity of MI and ME is highest for higher levels of action processing like the planning of

a movement.” In this paper, we focus on the planning of the movements and experimentally isolate the planning.

MVPA has also been used in button pressing tasks involving sequenced behavior (Woolgar, Thompson, Bor, & Duncan, 2011; Yokoi, Arbuckle, & Diedrichsen, 2018). In rule-mapping button pressing (in which the stimulus information indicates which finger and which hand to respond with), activation patterns were observed in different brain regions corresponding to stimulus and task attributes, such as rules (e.g., inferior frontal sulcus (IFS), insula, and IPS), position of a cue (IFS and IPS), and background color (insula and IPS). In contrast, no main effects of these variables were found in a univariate analysis.

This project enhances the understanding of how sequences of actions (Woolgar et al., 2011; Yokoi et al., 2018) and imagined and executed movements are neurally represented (Pilgramm et al., 2016; Zabicki et al., 2016) in the context of a complex, sequential, and naturalistic procedural task. This project advances the prior research in three ways: (1) it identifies procedural representations in a naturalistic, complex task; (2) it investigates a procedure that involves a sequence of several hand movements rather than a single action; and (3) it characterizes the mental representations of procedural planning (before the execution of a task) based on the imagined performance of the task. The project centrally addresses the challenge of identifying the neural signatures of the tying of individual knots, and determining the role of those signatures in the planning stage that preceded tying.

Methods

Participants. Seven right-handed adults (6 females and 1 male, between the ages of 20 and 36) from the Carnegie Mellon community participated in several experimental sessions. They came to the lab a total of four times (Phase 1: training session 1 and testing session 1; Phase 2: training session 2 and testing session 2). For Phase 1, the first session was a training session in which the participants learned to tie a set of seven knots. The session ended when they reach an efficiency level of tying each knot successfully twice with instructions and once without instructions. This was followed by a testing session in the scanner within one day of the training session. In the testing session they imagined tying each knot six times (mental tying) as described below in the “Testing session” sections. Phase 2 was a replication of Phase 1. The only change was an increase in the criteria for ending the training session. The phase 2 training session ended when each knot was successfully tied five times without written instructions present. Phase 2 followed Phase 1 within a month (mean of 13 days, range 7-28). Three of these trained participants returned for a

Phase 3. Because several months had passed since Phase 2 testing, Phase 3 provided a re-training session in which the participants re-trained to tie 7 knots until they were as proficient as in Phase 2. These participants were then scanned for the Phase 3 testing session in which they physically tied all seven of the knots six times, providing 126 instances of physical knot tying in the scanner. The precise Phase 3 instructions are described below in the “Classification of the motor planning that precedes the physical tying of knots” section. All participants gave signed informed consent approved by the Carnegie Mellon Institutional Review Board. Six participants contributed data from a total of 12 mental tying test sessions (data from one participant was discarded due to excessive motion $> 1\text{mm}$ in the y plane and $> 3\text{mm}$ in the z plane). Our measurement approach was to use predictive modelling as opposed to hypothesis testing, where Breiman argues that power considerations are very different (Breiman, 2001). This technique is common in machine learning assessment, and represents a shift from a data modelling culture. The goal is “to find an algorithm $f(x)$ such that for future x in a test set, $f(x)$ will be a good predictor of y .” Having 6 subjects with 2 sessions and 6 repetitions of each knot tying trial enables a strong test of predictability using cross-validation in algorithmic modelling. The classification analysis is described in detail below. This approach provides an unbiased estimate of predictive accuracy rather than reliance on goodness of fit in regression models.

Training Session. The training session consisted of watching a series of instructional videos of how to tie each type of knot and then physically tying them. The video for each knot began with the name of the knot and then described and demonstrated, step-by-step, how that knot is tied (all knots consisted of five steps except the clove hitch (3 steps) and the taut line hitch (6 steps)). Participants were instructed to tie the knot step-by-step as they watched the video. At the end of the video for each knot describing how a given knot was tied, they were instructed to again attempt to tie the knot, this time while the full set of instructions remained on the screen. They were only allowed one attempt to tie the knot before they were moved on to the instructions for the next knot, regardless of the success or failure of that attempt (and of the attempt while watching the video). After cycling through all seven knots, the participants attempted to tie each knot when presented with only the name and a picture of a tied knot. For any knot that they failed to tie correctly, they restarted the training/test cycle for only that knot. When they correctly tied each knot they had previous mis-tied, the training session ended. The participant returned on the following day for the fMRI session. The second training session was the same except that the learning criterion was increased to 5 correct attempts in a row for each knot. The

extended set of training instructions, including the description of knot-tying terminology (e.g., quadrant, loop and intersection) and a link to a sample training video can be found in the Supplemental Materials.

Testing Session Task instructions. The participant’s task was to mentally go through the steps of tying a knot, as follows:

Please think of each step of the knot-tying separately. Do not move on to the next step before you have mentally completely performed the previous step. Please mentally work through every motion associated with each step in your mind to completion and do not short-cut around a particular step. Also please do not skip steps that are unclear but rather do your best to work through them as you did for previous steps.

Testing Session Experimental Paradigm. The stimuli were seven nautical knots: Square, Two Half Hitch, Timber Hitch, Taut line, Clove Hitch, Bowline, and Sheetbend. Each knot was presented six times (in six different random permutation orders of the 7 knots). On each trial, the name of the knot and a picture of the tied form of the knot were presented for 7s, during which the participant mentally tied the knot. (See Figure 1 for an example.) One end of the rope was black to allow the instructions to distinguish the two ends. This was followed by a 7s rest period, during which the participant fixated on an “X” displayed in the center of the screen. There were three additional presentations of a fixation “X”, 17s each, distributed across the session to provide a baseline measure.



Figure 1. One of the seven knot photos shown to participants. One end of the rope was colored black so the two ends could be distinguished in the instructions.

In the second phase of the fMRI acquisition, after participants had completed *mental* knot-tying task, they were asked to *physically* tie each of the knots only once (as prompted in a random order). The name and a picture of the knot were presented to indicate which knot to tie. Participants were instructed to quickly and accurately tie the knot as soon as they could after the name and picture appeared. The knot name and picture were displayed for 40s, followed by a prompt that required a button press to continue to the next knot. Before advancing to the next knot, the experimenter removed the tied knot and supplied the participant with a new rope.

fMRI Procedures. Functional images were acquired on a Siemens Verio (Erlangen, Germany) 3.0T scanner at the Scientific Imaging and Brain Research Center of Carnegie Mellon University using a gradient echo EPI pulse sequence with TR = 1000 ms, TE = 25 ms and a 60° flip angle. Twenty 5-mm thick AC-PC aligned slices were imaged with a gap of 1 mm between slices using a 32-channel head coil. The acquisition matrix was 64×64 with 3.125-mm×3.125-mm×5.0-mm in-plane resolution. Images were corrected for slice acquisition timing, motion, and linear trend, and were normalized to the Montreal Neurological Institute (MNI) template without changing voxel size (3.125 × 3.125 × 6 mm). The gray matter voxels were assigned to anatomical areas with reference to Automatic Anatomical Labeling (AAL) masks (Tzourio-Mazoyer et al., 2002). The original set of 90 AAL regions was modified in two cases: an IPS ROI was generated to parcel out sulcal activation from the superior and inferior parietal regions; and the pre-central, post-central and supplementary motor region were combined into a motor cortex ROI.

The percent signal change (PSC) relative to the fixation condition was computed at each gray matter voxel for each stimulus presentation (the PSC data was converted to z-scores). The main measure used in the subsequent analyses consisted of the voxel activation levels on each of the six brain images acquired within a 6 s window, offset 5s from the stimulus onset (i.e. images 5 to 10). In several previous machine learning studies of fMRI-assessed neural representations (e.g., Just, Cherkassky, Aryal, & Mitchell, 2010), the window and the offset were 4s to account for the delay in hemodynamic response. Here an additional second was added to the offset and 1s added to the window to account for the complexity of mentally tying a knot as compared to thinking about a well-known, concrete object. An exploration of data from two pilot participants suggested that this 5 to 10 s window provided maximal classification accuracies (see supplemental materials). In the analysis of the physical tying data, a smaller earlier 4 to 7s window was used to be consistent with the motor planning interval as described below; all other parameters remained the same.

Time course analysis and voxel selection. Prior to the selection of stable voxels, the activation time course for each knot in each presentation was normalized to a mean of zero and a standard deviation of 1 in the critical 6s window. This eliminated the effect of the amplitude of activation and only the pattern of activation over time for each knot remained.

Voxels in the training set were selected using the criteria that their activation levels for different knots were stable across repetitions and at the same time their activation patterns distinguished between the knots; selection was a three-step procedure. First, the mean pairwise correlation of the 6 activation levels (in the 6s window) *across presentations* was computed for each voxel (i.e., correlations were calculated for each knot at each voxel and then the correlations were averaged over knots). High correlations indicated that their activation profile (i.e., the tuning curve of the voxel) over the 7 knots was stable across the multiple presentations of the stimulus set. (The activation profile of a voxel refers to the vector of its 7 responses (activation levels) at 6 time points to the 7 knots during that presentation). Second, the mean pairwise correlation for the 6 s window *across knots* was calculated for each voxel (i.e., correlations were generated for each knot at each voxel and then averaged over presentations). Low correlations indicated distinctiveness between knots. Third, the measures from the first two steps were combined. The voxel selection criterion measure was the summed value of similarity on the first correlation (high correlation across presentations) and the distance (negation of the correlation) in the second correlation (low correlation across knots). For the whole brain analysis, the 120 voxels with the highest summed similarity-distance were selected, and for the individual anatomic region analyses, the 10 voxels with the highest summed similarity-distance scores within each anatomical region were selected. Previous research has indicated that 120-240 voxels provides maximal classification (see <http://www.ccbi.cmu.edu/publications.html#neurosemantics> for a variety of tasks that use similar numbers of features). The lower number of 120 voxels was selected here a-priori. Selection of voxels occurred only in the training set and then they were used to classify activation of that set of voxels in the test set. For purposes of visualization only to indicate the locations of the selected voxels (shown in Figure 2), the summed similarity-distance was calculated across the entire dataset (no separate test set was left out) in each participant and then a map of the union of those voxels was generated.

Classification of mental tying of individual knots. Gaussian Naive Bayes (GNB) classifiers applied in a cross-validation procedure (Just et al., 2010) were used to identify the 7 knots from their fMRI time-series signatures.

The classifier was trained using selected voxels (as described above) from only a subset of the data (the training set), and then tested on the remaining data (the test set). On each of the folds of the cross-validation procedure used for the within-participant classification, the training set on which the GNB classifier was trained consisted of the data for each of the 7 knots from 4 of the presentations (selecting the 120 most stable voxels in the training presentations), and the test set consisted of the mean of the remaining two presentations (averaging over two test presentations to simply reduce noise). There were 15 folds of cross-validation, defined by the 15 ways that 4 of the 6 presentations can be chosen to be in the training set. The data for this and the majority of the analyses of the mental knot-tying, unless otherwise stated, consisted of the fMRI time series (normalized to remove overall amplitude differences among the 7 knot signatures) consisting of 6 consecutive images obtained 1 s apart from the 5th to the 10th image after stimulus onset, reflecting the relative activation at successive time points in the mental knot-tying procedure.

Normalized rank accuracy (hereafter, simply accuracy) of the classification is the normalized rank of the correct label in the classifier's posterior-probability-ordered list of classes. If the classifier were operating at chance, the correct label would on average appear in the middle of the ranked list, producing a chance level normalized rank accuracy of .50. Accuracies are calculated for each item in each fold and then averaged across folds, and then across items. Statistical significance levels are obtained using random permutation testing (for the 7-class classification). The 7 knot labels associated with each of the 6 repetitions of imagined knot tying were randomly permuted, with 350 such iterations for each participant.

For the between participant classification, the classifier was trained on data from 11 sessions (5 participants in 2 sessions each, plus one participant from a single session) and tested on the left-out participant-session. In this analysis, each participant's data for the test set was averaged over the six presentations. The mean signal over the 6 second window corrected for the possibility that the rate of mental knot tying was different across participants. Additionally, the between participant data was conducted on the non-normalized data.

Classification of the motor planning that precedes the physical tying of knots. This analysis tested whether a classifier that was trained on the signature of the mental tying of knots could classify the signature of the motor planning that precedes physical tying. Three participants from the original set returned and physically tied all the knots multiple times in the scanner. The following procedure isolated the mental planning of the tying of a knot. On each trial, the participant was first handed a rope that they then held in the ready position,

and they then pressed a "ready" button. Following the button press, they were shown the name and image of the target knot for 4 sec. Their instructions were to "First, plan how to tie the knot and do not begin to tie the knot until after "START TYING NOW" appears on the screen." The experimenter observed the participants' hands to ensure that they followed these instructions precisely on every trial. The end of the planning period was indicated by the instruction "START TYING NOW!" printed in red above the knot name. The red prompt disappeared after two seconds.

The four second delay was motivated by two factors. First, the delay had to be long enough that the hemodynamic response from multiple images could be averaged. Second, it had to be the approximate length of time for a planning stage. This was determined by estimating the time prior to any motor action. For this purpose, two participants repeated the physical tying task in the scanner environment while video recordings of their hand movements were made. The analysis of the videos indicated that the participants never began to tie the knots until at least 3 s had passed (with a range of 3-7 s). On the basis of this testing, the temporal window used as the fMRI test data for classifying the planning stage were the 4 images from images 4-7 after stimulus onset (taking into account the hemodynamic delay of the fMRI signal), representing only the time to mentally plan the tying of the knot and never including activation due to physical motion.

The classification of the motor planning of knot-tyings was based on the neural signature of the mental tying of the knots. The classifier was trained on the mean mental knot-tying data for each knot, averaged across all presentations and all participants, and it was tested on the mean motor planning data averaged over the six presentations and the three Phrase 3 participants. The test data in this additional data set was averaged across all physical tying trials (six attempts for each knot) for the seven knots. The activation data for the time window of each knot was normalized so as to minimize the impact of knot difficulty on activation amplitude and to maximize the impact of the sequence pattern of activation levels. 120 voxels (as in the mental tying analysis) were used. The occipital cortex was excluded from this analysis due to the differences in the display between the training and test conditions (the mental tying display showed the name and image of the tied knot; the physical tying display additionally provided a view of the participant's hands through the use of a split mirror system).




Figure 2. Example of a normalized time course of activation for two knots from a single participant (data are averaged over the 120 voxels depicted in Figure 3). Error bars indicate standard deviations across voxels

Results

The Procedural knowledge of how to tie a particular knot can be identified from its fMRI signature. The main finding was that it was possible to accurately identify which of seven knots a participant was mentally tying from the accompanying fMRI activation patterns. Recall that the data consist of the normalized fMRI time series from 6 separate images, reflecting the relative activation at successive times in the mental knot-tying procedure. Figure 2 shows the mean time courses from two of the seven knots, indicating the type of information that the classifier was using. Despite the difficulty/amplitude difference between the two knots having been normalized out, their time courses are discriminable from each other.

The mean classification rank accuracy across two sessions was high (mean = 0.70; range = 0.59 to 0.79) and significantly above chance for all participants, indicating that the procedure of mentally tying each knot evokes an identifiable activation signature. (Every participant's accuracy was reliably above chance at $p < .05$, where the within-participant critical value was obtained by random permutation testing, range = 0.59 to 0.63). The classifier's features were the activation levels at the 6 time points for 120 most stable and discriminative voxels (720 features in total) in the cerebral cortex or cerebellum (using correlation across presentations and differences between knots of their activation time courses as measures of stability and discriminability, as described in the Methods. This feature selection approach was used for all analyses except where noted).

The between-subjects classification of the procedural neural signatures of the knot-tying procedures was not

reliable (mean rank accuracy = 0.49), indicating a lack of commonality of the signatures across participants. This finding stands in contrast to the commonality of the neural representations of declarative concepts (e.g. Just et al., 2010). The absence of commonality of procedural neural signatures across participants may have been due to individual differences in the rate at which participants mentally tied the knots. Consistent with this account is the finding reported below that there was reliable between-participant commonality in the knot-tying signatures when the data across the 6 time points in the time series were averaged, which would have minimized individual differences in the rate at which the knots were mentally tied.

The neural representation of this procedural knowledge (the 120 most stable, discriminative voxels across participants) can largely be localized to several brain regions (only those regions contain at least 1 stable voxel were further analyzed; regions were taken from a set of 90 automated anatomical labeling (AAL) regions) (Tzourio-Mazoyer et al., 2002). These regions include spatial processing regions: bilateral parietal (inferior, superior, and IPS); language and executive regions (bilateral pars triangularis, pars opercularis, superior temporal, middle and superior frontal); motor processing regions (pre- and post-central sulcus, and supplementary motor area); visual processing regions (occipital cortex); object processing regions (bilateral inferior temporal and fusiform); and the cerebellum. When the classification was performed using voxels from only one of these six groupings of regions at a time (selecting 10 stable voxels from each of the 22 anatomical regions involved, then allocating them into the six groupings of regions), the mean rank accuracy across these individual anatomically-defined ROIs was 0.61

(range = 0.53 to 0.66 of participant means for the 22 ROIs; the range of the accuracies in these ROIs for individual participants was .38 to .81). The highest accuracies were obtained in the left superior parietal ROI (.66), occipital lobe (.66), left IPS (.65), motor cortex (.65), right superior parietal (.63), bilateral superior frontal (.61) and cerebellum (.57). (Classification accuracies using voxels from several of these groupings are shown in Figure 3, overlaid on the locations of 120 stable voxels from all participants; the accuracies for the 22 individual ROIs are shown in Table 1.)

The knot-tying procedures were reliably identifiable even without recourse to information in occipital cortex (which may have contained a visual representation of the probe stimulus knot) or motor cortex. These findings indicate that multiple regions contain the procedural information, but their individual accuracy is typically and unsurprisingly less than their conjoint accuracy. This provides an indicator of the multi-locus, network representation of the procedural knowledge.

Table 1. The 22 ROIs significantly above chance rank accuracy grouped into functional networks. Italics indicate rank accuracies reported in Figure 3.

	rank accuracy
Spatial	
Left Superior Parietal Lobe	0.66
Left Intraparietal Sulcus	0.65
Right Superior Parietal Lobe	0.63
Right Intraparietal Sulcus	0.60
Right Inferior Parietal Lobe	0.58
Left Inferior Parietal Lobe	0.58
<i>Parietal (using 10 voxels from each of the above regions; voxels = 60)</i>	<i>0.64</i>
Language and Executive	
Right Superior Frontal Gyrus	0.61
Left Superior Frontal Gyrus	0.61
Left Middle Frontal Gyrus	0.58
Left Pars Opercularis	0.57
Right Middle Frontal Gyrus	0.57
Left Pars Triangularis	0.57
Right Pars Opercularis	0.56
Right Superior Temporal Sulcus (Posterior)	0.56
Left Superior Temporal Sulcus (Posterior)	0.56
Right Pars Triangularis	0.53
<i>Frontal (using 10 voxels from the above regions excluding Temporal; voxels = 80)</i>	<i>0.64</i>
Motor	
<i>Motor Cortex (Pre-central sulcus, Post-central sulcus, SMA; voxels = 10)</i>	<i>0.65</i>
Visual and Object Processing	
Occipital Lobe	0.67
Right Inferior Temporal Gyrus (Posterior)	0.58
Left Inferior Temporal Gyrus (Posterior)	0.56
Fusiform Gyrus	0.56
Cerebellum	
<i>Cerebellum (voxels = 10)</i>	<i>0.57</i>




Figure 3. The classification accuracies for each grouping of regions are overlaid on the locations of the 120 most stable voxels (as defined in the results). The ellipses schematically represent the regions from which the anatomical ROIs were taken. The occipital (visual processing) and inferior temporal (object processing) groupings are not depicted. Classification accuracies were calculated using 10 voxels from each of several anatomical ROIs within each region. [Accuracies correspond to bilateral groupings but are depicted only on the left hemisphere].

Signal intensity and amplitude-based classification. To determine whether there was a commonality of the signal intensity based neural representation of the knot-tying across participants, a classifier was trained on the data of all but one participant in both mental tying sessions, and then was tested on the data of the left out participant. The mean rank accuracy across participants was 0.87 (range .69 to 1.00) excluding occipital cortex voxels. This analysis excludes occipital cortex voxels to minimize the possibility that the classifier is simply using visual features of the knot and the picture to distinguish the knot trials (The mean was 0.88 (range .79 to .95) when occipital voxels were included). As noted in the methods section, this analysis used the mean activation over the critical 5 to 10 s window rather than the time series. Thus there is a great deal of commonality across participants in the signal intensity of the neural representation of tying a particular knot when the speed of tying is approximately equated by averaging over the time series. This technique of averaging over the time series was used in predicting knot tying execution from imagined knot tying (section 3 below) in which speed of tying might vary across individuals in both the mental tying and physical tying data.

Amplitude based classification could also be conducted in a within-participant analysis in which it would not be necessary to use the mean of the time series data. When the fMRI time series was not normalized and hence included signal amplitude information, the classification rank accuracy remained unchanged, using as

features the 120 most stable voxels from anywhere in the cortex. The seven knots were classified with a high mean rank accuracy of 0.70 (range = 0.59 to 0.80). This finding indicates an ability to identify knots from brain activation data based when difficulty differences are included in the signal.

The neural signature of the mental tying resembles the motor plan evoked for its physical tying. The neural signature of the mental tying of a knot, when used as training data for a classifier, allowed the classifier to identify individual knots by the neural signature of the motor plan for the physical tying procedure, indicating the similarity between the two types of signals. A classifier trained on the neural signature of the mental tying of the knots reliably identified the planning of individual knots with a classification accuracy of 0.791, very significantly above chance level, indicating the similarity between the two types of neural representations. The motor planning data, acquired before the participant started any hand movements (offset by the hemodynamic delay interval), did not contain any motor action information. Thus, the motor plan for tying a particular knot can be identified from its fMRI signature obtained during mental tying.

Discussion

This study reveals that the procedural knowledge of how to tie a particular knot has a distinct, identifiable neural signature. Although a few previous studies have

shown that motor movements (Bednark, Campbell, & Cunnington, 2015; Graydon, Friston, Thomas, Brooks, & Menon, 2005; Kim, Ogawa, Lv, Schweighofer, & Imamizu, 2015; Penhune, 2013; Pilgramm et al., 2016; Wiestler & Diedrichsen, 2013; Wurm & Lingnau, 2015), and stimulus-response rules (Apšvalka, Cross, & Ramsey, 2018; Woolgar et al., 2011) can be identified from their neural representations, there has not been a previous demonstration of the neural identifiability of complex procedures. This decoding of procedures uses the same approach as has been successfully applied to other non-procedural representations, such as representations of concrete object concepts (Just et al., 2010; Mitchell et al., 2008), scientific concepts (Mason & Just, 2015, 2016), emotion concepts (Kassam, Markey, Cherkassky, Loewenstein, & Just, 2013), and perceptual representations (Huth et al., 2012).

Despite the fact that the vocabulary of motor steps for tying the 7 knots was similar across the knots, it was still possible to identify the unique procedural knowledge signature across time for each knot, applying a machine learning algorithm to the multivoxel time-series activation patterns. Presumably, the order of the steps affects how each one is executed, imbuing the activation pattern with a unique knot signature.

Procedural representations organize a sequence of actions. The most general implication of this study is there are procedural representations that organize a sequence of actions, as proposed in the classic paper “The problem of serial order in behavior” (Lashley, 1951). Lashley argued that because of timing considerations, a sequence of steps in a procedure has to be organized within a mental structure that is larger than a simple linkage from one step to the next. He argued that serial order in behavior is supported by higher level mental structures that organize a sequence. The current study assesses this higher order structure in terms of its neural implementation and shows that it uniquely identifies each knot. This higher order structure manifested itself as an activation pattern distributed across several cortical regions and the cerebellum. This temporally-distributed activation pattern is evoked as a motor plan prior to the execution of motor actions. It was possible to demonstrate that the planning of physically tying a knot consists of retrieving a knot-tying motor plan that resembles the neural signature of mental knot-tying. Here the neurocognitive representation of a procedure consists of a sequence of motor actions and states. However procedural tasks can be conceptual as in the case of solving algebra problems or computer programming. In future research, neural representations of conceptual steps (and their planning) can be compared to motor steps to explore the possibility of neural representations of procedures at a more general or abstract level.

Brain locations of knot-tying procedural representations. It is notable that the neural signature of a knot-tying procedure was distributed across several brain regions. Likely roles include a mental visualization of the knot at a given state of its being tied (presumably supported by parietal regions specialized for spatial representations), and a motor plan for performing the next move (presumably supported by motor cortex and the cerebellum) but there may be overlap in the processes and the various cortical regions (Apšvalka et al., 2018; Graydon et al., 2005; Ito, 2008; Kim et al., 2015; Oosterhof, Tipper, & Downing, 2012; Wurm & Lingnau, 2015).

Identification of procedural representations for Brain-Computer-Interfaces (BCI). These findings suggest the possibility of neurally distinguishing complex sequences. Decoding of procedural representations can be useful in future developments of BCI for robotic control, having already been used to decode handedness of motion (Min, Marzelli, & Yoo, 2010) and to distinguish between mental versus auditory imagery associated with a movement (Yoo et al., 2007). Direction, types of hand movements, and tool use are also classifiable procedures (Handjaras et al., 2015) as was the intention to make a simple hand movement (Gallivan, McLean, Valyear, Pettypiece, & Culham, 2011). Here, the mental activity of thinking about tying a specific knot and the motor plan preceding physical tying trials were identified in our study and could potentially be used as input to a robotic device. Practical use is still limited due to the need use an fMRI scanner, but the findings indicate the possibility of using non-invasively assessed neural representations of procedures to control a robotic system.

Future extensions. This research represents an early stage of a programmatic investigation of the neural representations of complex procedural knowledge. The ability to identify a motor plan stage in the fMRI signal is related to the Lashley hypothesis that a motor plan is generated before movement rather than being serially constructed. This is a first step towards using fMRI to characterize the cognitive and neural components of procedural skills. The decomposition of procedural concepts should be enabled by studying various combinations of the subprocesses and individual movements involved.

Acknowledgements

The Office of Naval Research (Grant N00014-16-1-2694). Supported this research. We thank Zach Anderson and Robert Vargas for assistance in recording training videos, testing participants, running experiments and data processing.

Author Contributions

R.M. and M.J. contributed equally to all aspects of this project.

Corresponding Author

Robert A. Mason, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213
Email: rmason@andrew.cmu.edu

References

- Apšvalka, D., Cross, E. S., & Ramsey, R. (2018). Observing action sequences elicits sequence-specific neural representations in frontoparietal brain regions. *The Journal of Neuroscience*, 1597–18. <https://doi.org/10.1523/JNEUROSCI.1597-18.2018>
- Bednark, J. G., Campbell, M. E. J., & Cunnington, R. (2015). Basal ganglia and cortical networks for sequential ordering and rhythm of complex movements. *Frontiers in Human Neuroscience*, 9, 421. <https://doi.org/10.3389/fnhum.2015.00421>
- Brandt, M. G., & Davies, E. T. (2006). Visual-spatial ability, learning modality and surgical knot tying. *Canadian Journal of Surgery*, 49(6), 412–416. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3207546&tool=pmcentrez&rendertype=abstract>
- Breiman, L. (2001). Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3), 199–215. <https://doi.org/10.2307/2676681>
- Carota, F., Kriegeskorte, N., Nili, H., & Pulvermüller, F. (2017). Representational Similarity Mapping of Distributional Semantics in Left Inferior Frontal, Middle Temporal, and Motor Cortex. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhw379>
- Cross, E. S., Cohen, N. R., De, A. F., Hamilton, C., Ramsey, R., Wolford, G., & Grafton, S. T. (2012). Physical experience leads to enhanced object perception in parietal cortex: Insights from knot tying. *Neuropsychologia*, 50, 3207–3217. <https://doi.org/10.1016/j.neuropsychologia.2012.09.028>
- Cross, E. S., Hamilton, A. F. de C., Cohen, N. R., & Grafton, S. T. (2017). Learning to tie the knot: The acquisition of functional object representations by physical and observational experience. *PLOS ONE*, 12(10), e0185044. <https://doi.org/10.1371/journal.pone.0185044>
- Gallivan, J. P., McLean, D. A., Valyear, K. F., Petteypiece, C. E., & Culham, J. C. (2011). Decoding Action Intentions from Preparatory Brain Activity in Human Parieto-Frontal Networks. *Journal of Neuroscience*, 31(26), 9599–9610. <https://doi.org/10.1523/JNEUROSCI.0080-11.2011>
- Glover, S., & Baran, M. (2017). The motor-cognitive model of motor imagery: Evidence from timing errors in simulated reaching and grasping. *Journal of Experimental Psychology: Human Perception and Performance*, 43(7), 1359–1375. <https://doi.org/10.1037/xhp0000389>
- Grafton, S., Mazziotta, J., Presty, S., Friston, K., Frackowiak, R., & Phelps, M. (1992). Functional anatomy of human procedural learning determined with regional cerebral blood flow and PET. *The Journal of Neuroscience*, 12(7), 2542–2548. <https://doi.org/10.1523/JNEUROSCI.12-07-02542.1992>
- Grafton, S. T. (2010). The cognitive neuroscience of prehension: recent developments. *Experimental Brain Research*, 204(4), 475–491. <https://doi.org/10.1007/s00221-010-2315-2>
- Graydon, F. X., Friston, K. J., Thomas, C. G., Brooks, V. B., & Menon, R. S. (2005). Learning-related fMRI activation associated with a rotational visuo-motor transformation. *Cognitive Brain Research*, 22(3), 373–383. <https://doi.org/10.1016/j.cogbrainres.2004.09.007>
- Guillot, A., & Collet, C. (2005). Contribution from neurophysiological and psychological methods to the study of motor imagery. *Brain Research Reviews*, 50(2), 387–397. <https://doi.org/10.1016/J.BRAINRESREV.2005.09.004>
- Handjaras, G., Bernardi, G., Benuzzi, F., Nichelli, P. F., Pietrini, P., & Ricciardi, E. (2015). A topographical organization for action representation in the human brain. *Human Brain Mapping*, 36(10), 3832–3844. <https://doi.org/10.1002/hbm.22881>
- Hardwick, R. M., Caspers, S., Eickhoff, S. B., & Swinnen, S. P. (2018). Neural correlates of action: Comparing meta-analyses of imagery, observation, and execution. *Neuroscience & Biobehavioral Reviews*, 94, 31–44. <https://doi.org/10.1016/j.neubiorev.2018.08.003>
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, 76(6), 1210–1224. <https://doi.org/10.1016/j.neuron.2012.10.014>
- Ito, M. (2008). Control of mental activities by internal models in the cerebellum. *Nature Reviews Neuroscience*, 9(4), 304–313. <https://doi.org/10.1038/nrn2332>
- Jeannerod, M. (2001). Neural Simulation of Action: A Unifying Mechanism for Motor Cognition. *NeuroImage*, 14(1), S103–S109. <https://doi.org/10.1006/NIMG.2001.0832>
- Just, M. A., Cherkassky, V. L., Aryal, S., & Mitchell, T. M. (2010). A neurosemantic theory of concrete noun representation based on the underlying brain codes. *PloS One*, 5(1), e8622. <https://doi.org/10.1371/journal.pone.0008622>
- Kassam, K. S., Markey, A. R., Cherkassky, V. L., Loewenstein, G., & Just, M. A. (2013). Identifying emotions on the basis of neural activation. *PloS One*, 8(6), e66032. <https://doi.org/10.1371/journal.pone.0066032>
- Kim, S., Ogawa, K., Lv, J., Schweighofer, N., & Imamizu, H. (2015). Neural Substrates Related to Motor Memory with Multiple Timescales in Sensorimotor Adaptation. *PLOS Biology*, 13(12), e1002312. <https://doi.org/10.1371/journal.pbio.1002312>
- Lashley, K. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior: The Hixon Symposium* (pp. 112–146). Hoboken, NJ: Wiley. <https://doi.org/10.1093/rfs/hhq153>

- Mason, R. A., & Just, M. A. (2015). Physics instruction induces changes in neural knowledge representation during successive stages of learning. *NeuroImage*, *111*, 36–48. <https://doi.org/10.1016/j.neuroimage.2014.12.086>
- Mason, R. A., & Just, M. A. (2016). Neural Representations of Physics Concepts. *Psychological Science*, *27*(6), 904–913. <https://doi.org/10.1177/0956797616641941>
- Min, B.-K., Marzelli, M. J., & Yoo, S.-S. (2010). Neuroimaging-based approaches in the brain–computer interface. *Trends in Biotechnology*, *28*(11), 552–560. <https://doi.org/10.1016/j.tibtech.2010.08.002>
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting Human Brain Activity Associated with the Meanings of Nouns. *Science*, *320*(5880), 1191–1195. <https://doi.org/10.1126/science.1152876>
- Oosterhof, N. N., Tipper, S. P., & Downing, P. E. (2012). Visuo-motor imagery of specific manual actions: A multi-variate pattern analysis fMRI study. *NeuroImage*, *63*(1), 262–271. <https://doi.org/10.1016/j.neuroimage.2012.06.045>
- Penhune, V. B. (2013). Neural encoding of movement sequences in the human brain. *Trends in Cognitive Sciences*, *17*(10), 487–489. <https://doi.org/10.1016/j.tics.2013.08.008>
- Pilgramm, S., de Haas, B., Helm, F., Zentgraf, K., Stark, R., Munzert, J., & Krüger, B. (2016). Motor imagery of hand actions: Decoding the content of motor imagery from brain activity in frontal and parietal motor areas. *Human Brain Mapping*, *37*(1), 81–93. <https://doi.org/10.1002/hbm.23015>
- Rogers, D. A., Regehr, G., Yeh, K. A., & Howdieshell, T. R. (1998). Computer-assisted Learning versus a Lecture and Feedback Seminar for Teaching a Basic Surgical Technical Skill. *The American Journal of Surgery*, *175*(6), 508–510. [https://doi.org/10.1016/S0002-9610\(98\)00087-7](https://doi.org/10.1016/S0002-9610(98)00087-7)
- Schwan, S., & Riempp, R. (2004). The cognitive benefits of interactive videos: learning to tie nautical knots. *Learning and Instruction*, *14*(3), 293–305. <https://doi.org/10.1016/j.learninstruc.2004.06.005>
- Torkington, J., Smith, S. G., Rees, B. I., & Darzi, A. (2000). The role of simulation in surgical training. *Annals of the Royal College of Surgeons of England*, *82*(2), 88–94. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10743423>
- Tracy, J., Flanders, A., Madi, S., Laskas, J., Stoddard, E., Pyrros, A., ... DelVecchio, N. (2003). Regional brain activation associated with different performance patterns during learning of a complex motor skill. *Cerebral Cortex*, *13*(9), 904–910. <https://doi.org/10.1093/cercor/13.9.904>
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., ... Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, *15*(1), 273–289. <https://doi.org/10.1006/nimg.2001.0978>
- Wiestler, T., & Diedrichsen, J. (2013). Skill learning strengthens cortical representations of motor sequences. *ELife*, *2*, e00801. <https://doi.org/10.7554/eLife.00801>
- Woolgar, A., Thompson, R., Bor, D., & Duncan, J. (2011). Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. *NeuroImage*, *56*(2), 744–752. <https://doi.org/10.1016/j.neuroimage.2010.04.035>
- Wurm, F., & Lingnau, A. (2015). Decoding Actions at Different Levels of Abstraction. *The Journal of Neuroscience*, *35*(20), 7727–7735. <https://doi.org/10.1523/JNEUROSCI.0188-15>
- Yokoi, A., Arbuckle, S. A., & Diedrichsen, J. (2018). The Role of Human Primary Motor Cortex in the Production of Skilled Finger Sequences. *The Journal of Neuroscience*, *38*(6), 1430–1442. <https://doi.org/10.1523/JNEUROSCI.2798-17.2017>
- Yoo, S.-S., O'Leary, H. M., Lee, J.-H., Chen, N.-K., Panych, L. P., & Jolesz, F. A. (2007). Reproducibility of trial-based functional MRI on motor imagery. *The International Journal of Neuroscience*, *117*(2), 215–227. <https://doi.org/10.1080/00207450600582546>
- Zabicki, A., de Haas, B., Zentgraf, K., Stark, R., Munzert, J., & Krüger, B. (2016). Imagined and Executed Actions in the Human Motor System: Testing Neural Similarity Between Execution and Imagery of Actions with a Multivariate Approach. *Cerebral Cortex*, *27*(9), 4523–4536. <https://doi.org/10.1093/cercor/bhw257>