Article

Partially Mixed Selectivity in Human Posterior Parietal Association Cortex

Highlights

- A small patch of AIP encodes body parts, body sides, and cognitive strategies

- Multiple encoded effectors argues against strict anatomical segregation

- Orthogonal coding of different effectors may enable functional segregation

- Other variables are organized by body part (partially mixed coding)

Authors

Carey Y. Zhang, Tyson Aflalo, Boris Revechkis, Emily R. Rosario, Debra Ouellette, Nader Pouratian, Richard A. Andersen

Correspondence

cyzhang@caltech.edu (C.Y.Z.), taflalo@caltech.edu (T.A.)

In Brief

Zhang, Aflalo, et al. study how body part, body side, and cognitive motor strategy are represented in human anterior intraparietal area (AIP). They find hierarchical coding determined by body part, in contrast to pure mixed selectivity, identifying a novel functional network organization.

CellPress

CellPress

Partially Mixed Selectivity in Human Posterior Parietal Association Cortex

Carey Y. Zhang,[1,4,6,]* Tyson Aflalo,[1,4,]* Boris Revechkis,[1,5] Emily R. Rosario,[2] Debra Ouellette,[2] Nader Pouratian,[3] and Richard A. Andersen[1]

[1]Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA 91125, USA
[2]Casa Colina Hospital and Centers for Healthcare, Pomona, CA 91767, USA
[3]Department of Neurosurgery, Interdepartmental Program in Neuroscience, and Brain Research Institute, David Geffen School of Medicine at UCLA, Los Angeles, CA 90095, USA
[4]These authors contributed equally
[5]Present address: Helynx Inc., Pasadena, CA 91103, USA
[6]Lead Contact
*Correspondence: cyzhang@caltech.edu (C.Y.Z.), taflalo@caltech.edu (T.A.)
http://dx.doi.org/10.1016/j.neuron.2017.06.040

SUMMARY

To clarify the organization of motor representations in posterior parietal cortex, we test how three motor variables (body side, body part, cognitive strategy) are coded in the human anterior intraparietal cortex. All tested movements were encoded, arguing against strict anatomical segregation of effectors. Single units coded for diverse conjunctions of variables, with different dimensions anatomically overlapping. Consistent with recent studies, neurons encoding body parts exhibited mixed selectivity. This mixed selectivity resulted in largely orthogonal coding of body parts, which "functionally segregate" the effector responses despite the high degree of anatomical overlap. Body side and strategy were not coded in a mixed manner as effector determined their organization. Mixed coding of some variables over others, what we term "partially mixed coding," argues that the type of functional encoding depends on the compared dimensions. This structure is advantageous for neuroprosthetics, allowing a single array to decode movements of a large extent of the body.

INTRODUCTION

The posterior parietal cortex (PPC) of humans has historically been viewed as an association area that receives diverse inputs from sensory cortex, "associates" these inputs for processing more cognitive functions such as spatial awareness, attention, and action planning, and delivers the outcomes of the associative process to more motor regions of the frontal cortex (Balint, 1909; Holmes, 1918; Mountcastle et al., 1975; Ungerleider and Mishkin, 1982). However, subsequent single-neuron recording experiments with behaving non-human primates (NHPs) point to a systematic organization of functions in PPC (Andersen and

Buneo, 2002). Of particular interest to the current investigation, separate cortical areas around the intraparietal sulcus (IPS) have concentrations of neurons selective for saccades (lateral intraparietal area [LIP]) (Andersen et al., 1987), reach (parietal reach region [PRR]) (Snyder et al., 1997), and grasping (anterior intraparietal area [AIP]) (Murata et al., 2000). These data suggest that this part of the PPC, rather than being one large association region, is rather composed of a number of anatomically separated cortical fields that are specialized for intended movements that are effector specific (eye, arm, hand).

More recent functional magnetic resonance imaging (fMRI) studies in humans have presented a mixed picture with some studies finding similar segregation for the types of intended movement in areas around the IPS (Astafiev et al., 2003; Connolly et al., 2003; Culham et al., 2003; Gallivan et al., 2011; Prado et al., 2005) and other studies finding largely an intermixing of effectors (Beurze et al., 2009; Heed et al., 2011a; Hinkley et al., 2009; Levy et al., 2007) as well as bimanual representation (Gallivan et al., 2013). These findings provide evidence for a degree of distributed and overlapping representation of effectors on both sides of the body within PPC.

With the first chronic single-neuron recordings of PPC in humans, we found similarities with the NHP studies. Neurons in human AIP are highly selective for different imagined grasp shapes, while neurons in nearby Brodmann area (BA) 5 are not (Klaes et al., 2015). However, the human neural recordings also pointed to some degree of distributed representation, with AIP neurons also selective for reach direction and with AIP and BA5 neurons being selective for reaches with either the left or the right limb or both (Aflalo et al., 2015). While we have found evidence that multiple effectors are encoded in the same anatomical region of cortex, these studies were carried out in separate sessions and thus the functional organization of multiple effectors within the same population of neurons remains unclear.

Pertinent to how different effectors are coded within PPC are recent results that address encoding strategies for multiple dimensions of representations and their computational advantages in association cortices more generally. Neurons in prefrontal cortex and PPC (Raposo et al., 2014; Rigotti et al., 2013) exhibit what has been termed mixed selectivity (Fusi et al.,

CrossMark

2016), a neural encoding scheme in which different task variables and behavioral choices are combined indiscriminately in a non-linear fashion within the same population of neurons. This scheme generates a high-dimensional non-linear representational code that allows for a simple linear readout of multiple variables from the same network of neurons (Fusi et al., 2016). A basic question is whether such an organization of functional variables is universal or, rather, is in part due to the types of functional variables that were compared or the cortical subregions selected for study.

In the current study, we examine the anatomical and functional organization of different types of motor variables within a 4 × 4 mm patch of human AIP. We varied movements along three dimensions: the body part used to perform the movement (hand versus shoulder), the body side (ipsilateral/left versus contralateral/right), and the cognitive strategy (attempted versus imagined movements). Each of these variables has been shown to modulate PPC activity (Andersen and Cui, 2009; Gallivan et al., 2013; Gerardin et al., 2000; Heed et al., 2011a). Thus, we are able to look at how different dimensions of motor variables are encoded, and whether different variable types are treated in an equivalent manner (e.g., all variables exhibiting mixed selectivity) or whether different functional organizations are found for different types of variables. Finally, we compare the hand and shoulder movements to speech movements, a very different type of motor behavior.

We find that movements of the hand and shoulder are well represented in human AIP, whether they are imagined or attempted, or performed with the right or left side. Single units were heterogeneous and coded for diverse conjunctions of different variables: there was no evidence for specialized subpopulations of cells that selectively coded one movement type. However, the different motor dimensions were not indiscriminately mixed, as body side and cognitive strategy were fundamentally different from body part at the level of neural coding. There was a high degree of correlation between movement representations of the right and left side, within, but not between body parts. The same was true for cognitive strategy. Thus, body part acted as a superordinate variable that determined the structure of how the other variables were encoded. Mixed coding of some movement variables, but not others, argues in favor of PPC having a partially mixed encoding strategy. Finally, while AIP lacks anatomical segregation of body parts, the mixed coding between body parts leads to what we call *functional* segregation of body parts. Such segregation is hypothesized to enable multiple body parts to be coded in the same population with minimal interference.

RESULTS

Recording from AIP of a female, C3/C4 tetraplegic participant 7 years post-injury (N.S.), we compared neural responses of attempted and imagined actions of the hand or shoulder on the right and left side of the body. Hand movements involved squeezing the hand into a fist and shoulder movements involved shrugging the shoulder. Shoulder shrugs are a staple of the participant's behavioral repertoire being a primary method to operate her motorized wheelchair. For imagined movements,

we instructed N.S. to visualize her limb performing the instructed action, while for attempted movements, she was instructed to send the appropriate motor command to move the instructed limb. In the case of shoulder movements, attempted movement resulted in overt motor execution, while for the hand, there was no resulting movement because of paralysis. For the shoulder, we confirmed behavioral compliance by measuring the presence of trapezius EMG activity during attempted but not imagined movement.
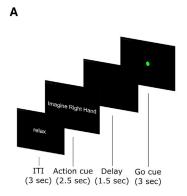
We used a delayed movement paradigm (Figure 1A). Following an inter-trial interval (ITI), N.S. was instructed to attempt or imagine movement of the left or right hand or shoulder. This instruction was extinguished during a delay period. A generic "Go" cue, visually identical across trial conditions, prompted movement. From initial pilot data, we knew that all hand and shoulder movements evoked activity in the population; however, we were unsure how the different conditions mapped onto individual neurons. Shoulder and hand movements are frequently performed together, opening the possibility that hand and shoulder movements would frequently be localized to the same neural population. We therefore introduced speech as a fundamentally different action that could provide an additional movement for comparison. During the speech conditions, N.S. simply said "left" or "right" as instructed. Eight repetitions of each trial type were pseudorandomly inter-leaved such that one repetition of each condition was performed before repeating a condition.
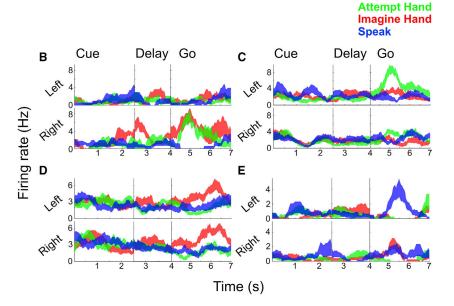
Figures 1B–1E show several well-tuned example units that highlight how neurons commonly coded for a complex assortment of different condition types. For instance, Example B codes for movements of the right hand, whether or not the movement was imagined or attempted. Example C codes exclusively for attempted movements of the left hand. Example D responds similarly for imagined actions of the left or right hand, but not attempted actions. Example E codes for when N.S. spoke "left."

To better understand the strength of tuning in the population to each condition, we fit a linear model to each neuron that explained firing rate relative to baseline (taken as the firing rate during the ITI) as a function of each task condition for both the Go and Delay phases. All eight primary movement conditions as well as speaking were represented in the neural population (Figure 2A). We also examined the magnitude of information content for recorded units by computing the area under the receiver operating characteristic curve (AUC) generated when comparing the Go/Delay period activity to ITI activity for each condition separately (Figure 2B). While there were significant differences between specific pairwise comparisons (Figure S3), results across these measures were comparable overall.

We performed a second type of linear analysis, fitting a linear model that explained the firing rate relative to baseline as a function of the three motor variables of strategy, body side, and body part (Figure 2C). This analysis revealed an asymmetry in how body part is represented compared to body side or strategy. In particular, relatively larger non-linear interactions of strategy and body side with body part indicate that body part may in some way structure the functional responses to the other variables, a point we directly address below.

We found significant differences in mean firing rates between hand and shoulder movements (hand greater than shoulder,

A

ITI   Action cue   Delay   Go cue
(3 sec)   (2.5 sec)   (1.5 sec) (3 sec)

Attempt Hand
Imagine Hand
Speak

Firing rate (Hz)

Time (s)

Figure 1. Neurons in PPC Exhibit Mixed Selectivity to Movement Variables
(A) Delayed movement paradigm. N.S. was cued as to what kind of movement to perform (e.g., imagine/attempt left/right hand/shoulder) and then cued to perform the movement after a brief delay. See STAR Methods for more details.
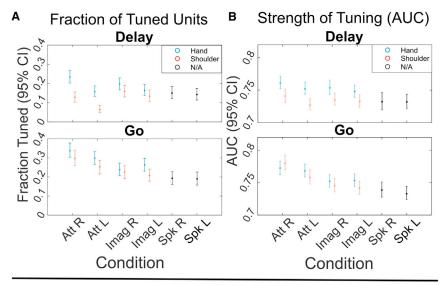(B–E) Single unit example responses over time (mean ± SEM) demonstrating diverse coding to the different conditions.
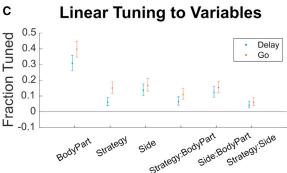
randomly mixed while others are organized with more structure (partially mixed, Figure 3E).

We first performed a degree of specificity analysis (Figure 4) to determine (1) whether highly specialized sub-populations of neurons are dedicated to each movement type, and (2) whether some variables exist as subsets or suppressed versions of other variables. A specificity index was computed as the normalized difference in beta values between motor variables for each neuron (taken from the linear models described above). Values near zero indicate equivalent neural responses to the two conditions being compared, while values near 1 (or −1) indicate exclusive neural responses for one condition. By proposition (1), we would expect values to be clustered near 1 (or −1) as, e.g., either a neuron is tuned to the right side or the left side. By proposition two, we would expect strong biases such that values would be clustered on one side of the range (between 0 and 1 *or* 0 and −1) as, e.g., a neuron tuned to imagine movement should be better (or equivalently) tuned to attempted movement. Inconsistent with these proposals, we found that specificity values were distributed over the full range (Figures 4A–4F). For instance, despite a small population bias for attempted movements, a sizable proportion of neurons were exclusively or more strongly activated for imagined movements (Figures 4A and 4B; see Figure 1D). The neural representation of motor imagery is thus not a subset, or less strongly represented version, of motor execution. Likewise, many neurons showed preferential coding for the left hand (Figures 4C and 4D) even with a population bias for the right hand. There was a strong specificity bias toward the movement conditions (imagined or attempted movements of the hand or shoulder) over the speech conditions (Figures 4G and 4H). This is expected given that speech tuning is found in a smaller proportion of neurons in a weaker fashion (Figure 2). Of special note, the results here are very similar for both movements of the shoulder (above the level of injury) and movements of the hand (below the level of injury). Thus, movements below and above the level of injury are coded in a similar manner.

MANOVA p = 5.614e-8) and attempted and imagined movements (attempt greater than imagine, p = 0.0020), and no significant differences in mean firing rates between left- and right-sided movements (p = 0.2951). Comparing the firing rates of all eight movement conditions (pooled together) with the firing rates of the speech conditions, we found a significant bias toward hand and shoulder movements over speaking (t test p = 4.7048e-8).

How are these different motor representations coded with respect to each other in the same region of cortex? Figure 3 shows five possibilities: (1) the eight movement condition representations could be anatomically segregated from each other, with a highly specialized sub-population of neurons dedicated to each (sparse mixed selectivity, Figure 3A); (2) an organization similar to (1), save that some variables are subordinate to others, for instance, imagined movements may be a subset or suppressed version of attempted movements (Figure 3B); (3) highly specialized sub-populations are tuned to each motor variable class exclusively (body part, body side, strategy) (pure selectivity, Figure 3C); (4) each motor variable class could be randomly mixed together (Churchland and Cunningham, 2014; Fusi et al., 2016) (all mixed, Figure 3D); and (5) some variables may be

Linear Tuning to Conditions

A Fraction of Tuned Units

B Strength of Tuning (AUC)

C Linear Tuning to Variables

Figure 2. Significant Tuning to Each Movement Condition

(A) The fraction of units in the population tuned for each condition in the Delay and Go phases, separated by body part and body side (95% confidence interval). A unit was considered tuned to a condition if the beta value of the linear fit for the condition (Linear analyses 1, STAR Methods) was statistically significant ($p < 0.05$, uncorrected). See also Figure S3 for pairwise comparisons between conditions and Figure S5 for results of individual sessions.

(B) The magnitudes of the units' tuning to each condition in the Delay and Go phases, as defined by the area under the receiver operating characteristic curve (AUC) between Delay/Go and ITI activity, separated by body parts (95% confidence interval). Only significant AUC values were included in analyses (shuffle test, $p < 0.05$ uncorrected). See also Figure S4 for the AUC values of excitatory (positively tuned) and inhibitory (negatively tuned) units presented separately, as well as Figure S3 for pairwise comparisons between conditions (Att R, attempt right; Att L, attempt left; Imag R, imagine right; Imag L, imagine left; Spk R, speak right; Spk L, speak left).

(C) Fraction of units with significant tuning to each motor variable and the interaction terms for both the Delay (blue) and Go (red) phases, as opposed to the eight movement conditions in (A) ($p < 0.05$, uncorrected 95% confidence intervals, see also Linear analysis 2 in STAR Methods).

We failed to find complete specialization of function across the population for single units, and the distributed and overlapping nature of responses makes it difficult to find structure in the responses of individual neurons. We therefore turned to population-based analyses to more readily identify how the different conditions are encoded with respect to each other. We measured all pairwise correlations between population responses for each of the eight movement conditions and looked for systematic structure in how the different motor variables (body part, body side, cognitive strategy) were coded (Figure 5). Correlation was used as a measure of similarity over other distance measures such as Euclidean or Mahalanobis distance because the sign of the correlation is potentially informative of the underlying structure. For example, two conditions represented by distinct neural populations (e.g., sparse mixed selectivity) would manifest as a negative correlation between the two conditions, while a positive correlation would indicate a degree of overlap between the populations. Asymmetric relationships between the different variables were immediately apparent.

Correlations between conditions that differed in body side or cognitive strategy were high if the comparisons were made within a body part. In stark contrast, correlations between conditions that differed in body part were low even if cognitive strategy

and body side were held constant (Figures 5A and 6A). Low correlation between body parts was also apparent when comparing speech with shoulder or hand (Figure 5B). Such low correlations despite activating overlapping neural populations are a signature of network responses that occupy distinct neural subspaces, thus minimizing crosstalk during planning and execution epochs (Churchland and Cunningham, 2014; Kaufman et al., 2014). Here the same principal may be at play for cortical representations of different effectors in an overlapping neural population. We term this "functional segregation" of body parts. That the functional organization is based around effector is especially apparent when the distances between conditions were hierarchically clustered (Figure 5C), with body part being the primary differentiating variable. Further, for a given body part, movements with more shared traits are coded more similarly than movements with fewer shared traits (Figure 6B). For instance, a neuron tuned to *imagined* left-hand movements was more likely tuned to *imagined* right-hand movements (but not *attempted* right-hand movements). Likewise, a neuron tuned to *right*-hand imagined movements was likely to be tuned to *right*-hand attempted movements (but not *left*-hand attempted movements). This functional segregation likely accounts for the non-linear interaction terms of Figure 2C.

Neural differences between hand and shoulder movements may be driven by the fact that the hand is below the level of injury while the shoulder is above the level of injury: in this case,
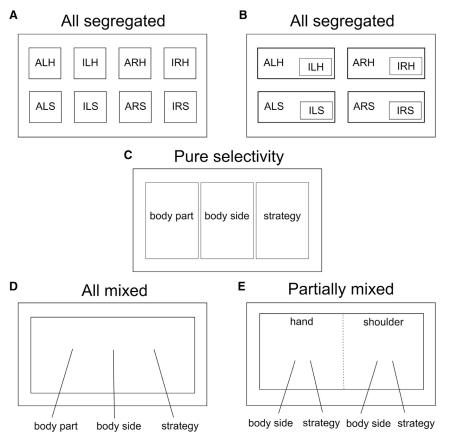
Figure 3. Possible Organizational Models of Neural Representations

(A) The neurons coding for each condition are anatomically segregated, i.e., distinct, non-overlapping networks (ALH, attempt left hand; ILH, imagine left hand; ARH, attempt right hand; IRH, imagine right hand; ALS, attempt left shoulder; ILS, imagine left shoulder; ARS, attempt right shoulder; IRS, imagine right shoulder).

(B) Conditions can be overlapping such that the responses to some conditions are subsets or weak versions of others, e.g., imagined movements being subsets of attempted movements.

(C) Neurons coding each of the motor variables (body part, body side, and strategy) are anatomically segregated.

(D) The neural population exhibits mixed selectivity, with individual neurons showing tuning to various conjunctions of variables.

(E) The neural population exhibits partially mixed selectivity, with the mixing of representations being dependent on the variables under investigation. Here, hand and shoulder are mixed leading to orthogonal coding of effectors (functional segregation); however, the other variables (body side and strategy) are mixed only within, but not between, effectors. This model is consistent with the results observed in this study. Note that solid lines in this diagram indicate anatomical boundaries of neural populations, while dotted line indicates functional boundaries/segregation.

proprioceptive feedback or long-term effects from the injury might be the primary difference. To address this issue, we replaced shoulder shrugging movements with shoulder abduction movements (shoulder abduction resulted in no overt movement) and repeated the correlation analyses. The results are similar when both body parts are chosen to be below the level of injury (Figures 6C and 6D). In particular, the largest degree of separation exists between body parts.

Functional segregation of body parts should lead to minimal shared information about other motor variables when compared across body parts. The motor dimensions can be thought of as categorical variables with two levels (e.g., body part has the levels of shoulder and hand). Given functional segregation, a classifier trained on one level ("Level A") should fail to generalize to the other category ("Level B") and vice versa (Figure 7A). Alternatively, for highly overlapping representations, a classifier trained on Level A should generalize to Level B and vice versa (Figure 7B). For example, given functional segregation between hand and shoulder, the neural signature that differentiates right- from left-sided movements for the hand should fail to generalize to the shoulder. We tested for this possibility by looking at patterns of generalization across trained classifiers. The results of such an analysis are shown in Figures 7C–7H. For Figure 7C, we trained a linear discriminant classifier on all shoulder movement trials to differentiate between left- and right-sided movements, regardless of strategy. The decoder performed well within its own training data as expected (leave-one-out cross-

validation, Figure 7C, left blue bar) but performed at chance differentiating left- from right-sided movements for hand trials (Figure 7C, right blue bar). The reverse was true when applying a classifier trained on hand trials to shoulder trials (Figure 7C, orange bars). Likewise, Figure 7D shows that a decoder trained to differentiate strategy using shoulder trials failed to generalize to hand trials, and vice versa. In contrast, decoders trained to differentiate strategy or body part *were* able to generalize and perform well across different body sides (Figures 7E and 7F) and different strategies (Figures 7G and 7H). Body part differences exhibit functional segregation, while cognitive strategy and body side do not.

Given that some motor variables are similar in their neural encoding, is it possible to decode the body part, body side, and cognitive motor strategy from the neural population? We constructed a neural classifier to differentiate all conditions (Figure 8). Cross-validated classification performance was high. However, as expected, misclassification tended to occur between conditions with more variables in common. This is especially true between attempted and imagined movements as predictable from the high degree of similarity in the neural responses (Figure 6A).

DISCUSSION

We tested how a variety of motor variables were coded at the level of single neurons in human AIP. This allowed us to address
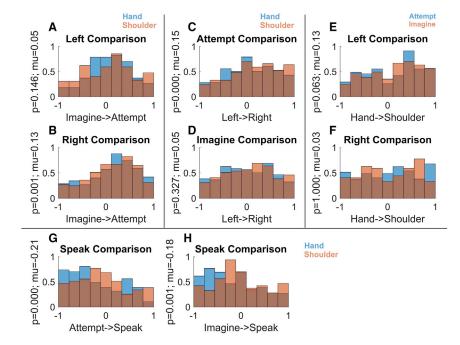
Figure 4. Specificity of Coding for Motor Variables

Each panel (A)–(F) shows the degree to which neurons code one variable exclusively, its opposite, or respond similarly for both. Only units with significant modulation for at least one condition in the comparison are included in the analyses (p < 0.05, Bonferroni corrected).

(A and B) Distribution of the degree of specificity to the imagine or attempt strategies in the population during trials using different sides, showing only units responsive to one or both strategies.

(C and D) Distribution of the degree of specificity to the left or right side in the population for different strategies.

(E and F) Distribution of the degree of specificity to the hand or shoulder in the population during trials using different sides.

(G and H) Distribution of the degree of specificity to attempted/imagined movements compared to speaking.

several questions about how intent is coded in human AIP and to better understand how the motor variables are coded with respect to each other.

Effector Specificity in PPC

Classically, the regions around the IPS have been viewed as organizing around the control of different effectors such as the eye, hand, and arm. In a recent challenge to the centrality of an effector-based organization, Medendorp and colleagues have found that effector specificity in the BOLD response of fMRI is much more pronounced between the hand and eye than the hand and other body parts, arguing that effectors as such are not differentiated in the planning regions of PPC (Heed et al., 2011a). In line with these results, we found significant numbers of neurons tuned to movements of the hand and shoulder in a small patch of AIP. However, unlike the response at the level of voxels, the neural response to each effector was functionally segregated. Thus, while our results challenge the idea of strict anatomical segregation of effector representations across cortical areas, we do find local functional segregation of effectors within a cortical field. The current findings suggest that effector specificity at the global anatomical scale could be thought of in terms of relative emphasis rather than strict specialization in humans.

In non-human primates (NHPs), a global organization for eye and arm movements is supported by greater planning activity of single neurons for reaches in the parietal reach region (PRR) and saccades in the lateral intraparietal region (Cui and Andersen, 2007; Hwang et al., 2012; Quian Quiroga et al., 2006; Snyder et al., 1997, 1998). Reversible inactivation of PRR produces reach-specific deficits and LIP a bias toward saccade deficits (Christopoulos et al., 2015; Kubanek and Snyder, 2015; Yttri et al., 2014). A grasp-specific deficit has been reported for AIP (Gallese et al., 1994). These results indicate that, at a global level,

there is functional specificity by effector in non-human primates, and fMRI studies in humans suggest a similar global specialization. However, these areas also communicate with one another. For instance, inactivation produces a reach deficit in PRR when reaches are made alone but both reach and saccade deficits when combined hand-eye movements are made (Hwang et al., 2014). Thus, the degree of effector overlap in AIP in human may reflect the coordination of movement and communication between effector-specific areas.

An advantage of our human study is that the participant can perform a large number of tasks by verbal instruction. In NHP studies, the animals must be trained for long periods and thus the number of tasks and task variables are generally limited per study. Interestingly, area LIP has been studied by a number of groups using a number of different tasks. As a result, LIP has been found to modulate activity for tasks examining movement planning, attention, categorization, and decision making, resulting in a variety of proposals for its function (Andersen and Cui, 2009). It may be that the large number of variables to which human AIP is selective may be a reflection of the versatility of using different tasks and that both human and NHP PPC areas are modulated by a very large number of variables. Indeed, several NHP studies in AIP have reported overlapping populations of cells tuned to grasp type and reach target consistent with mixed selectivity between effectors as presented here (Asher et al., 2007; Fattori et al., 2009; Lehmann and Scherberger, 2013, 2015).

Differences between effector segregation in human and NHP studies of PPC may be a result of possible lack of homologies between human AIP of the current study and AIP of NHPs. In fact, we do not know the extent or number of grasp-related areas defined by single-neuron recordings in human IPS and whether there are grasp regions in humans that do not exist in NHPs. Finally, the lack of strict anatomical segregation of effectors
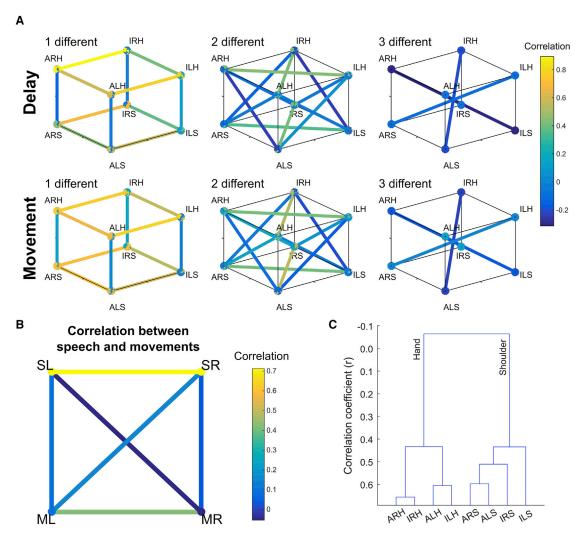
Figure 5. Functional Relationships between Movement Conditions

(A) Similarity between population-level neural responses for each movement condition. Pairwise comparisons are separated by the number of motor dimensions that differ in the comparison (left to right) and task phase (movement or delay). Similarity measured as the pairwise correlation between movement conditions (ALH, attempt left hand; ILH, imagine left hand; ARH, attempt right hand; IRH, imagine right hand; ALS, attempt left shoulder; ILS, imagine left shoulder; ARS, attempt right shoulder; IRS, imagine right shoulder). See also Figure S5 for results of individual sessions.

(B) Correlations between four movement types: left and right movements (averaged across both strategies), and speech (SL, speak left; SR, speak right; ML, movement left; MR, movement right).

(C) Dendrogram summarizing the structure apparent in (A), namely strong segregation by effector.
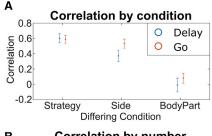
may point toward a global topographic organization governed around more behaviorally meaningful aspects of behavior such as manipulation, reaching, climbing, and defense (Graziano and Aflalo, 2007; Jastorff et al., 2010). Whichever possibilities outlined above account for the large number of variable encodings in human AIP, an exciting aspect of our results is that they open the possibility of decoding movements of many body parts from one small patch of cortex.
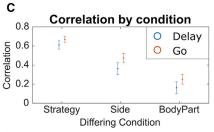
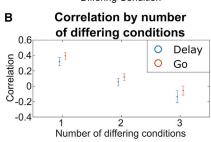## Asymmetric Coding of Motor Variables and Functional Segregation of Body Parts

Recently there has been increased interest in not only the types of variables that are coded in a cortical region, but also how

these variables are coded with respect to each other in an effort to understand the underlying logic of the computations performed within a cortical field (Fusi et al., 2016; Raposo et al., 2014). For instance, several papers have shown that higher cortical areas like PPC and prefrontal cortex may employ a computational strategy by which response variables are randomly mixed (Raposo et al., 2014; Rigotti et al., 2013). While such a coding scheme can give rise to complex and difficult to interpret representations at the level of single neurons, the population code is information rich and enables simple linear classifiers to decode any variable of interest. In these papers, it was shown that response variables were randomly distributed across neurons, as illustrated in Figure 3D. Our data provide insights
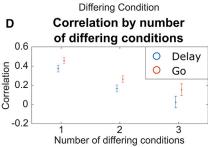
Movements above and below the level of injury

A

Correlation by condition

B

Correlation by number of differing conditions

Movements below the level of injury

C

Correlation by condition

D

Correlation by number of differing conditions

Figure 6. Segregation by Body Part

(A) Average correlation between movement conditions differing by exactly one task variable and grouped by the differing conditions (e.g., for strategy, the average correlation of all movement condition pairs differing only by strategy). Intervals represent the 95% confidence intervals.

(B) For movements above and below the level of injury, average correlation between movement conditions in the Delay and Go phases grouped by the number of differing traits (average of each cube, Figure 5A). Intervals represent the 95% confidence intervals in the correlations.

(C and D) Same as (A) and (B) but with shoulder shrug movements replaced with shoulder abduction movements (a movement below the level of injury).

into understanding population coding by demonstrating that in human AIP certain response features can be seemingly randomly distributed across the population, while others are not. In particular, we find that coding for body part is uncorrelated in the sense that, across the population, knowing that a neuron is tuned to shoulder movements provides little to no information about whether the neuron is tuned to hand movements (or speech; Figure 4). This is true even if you know other attributes of the movement, such as whether the movement was imagined or attempted or performed with the right or left side of the body. In contrast, when comparing within the same body part, knowing a neuron is tuned to movements of the right side makes it highly likely that the neuron will be tuned to the left side as well. The same is true for imagined and attempted movements. Thus, while some variables seem to be randomly distributed across the population (e.g., body part) the relationship between other variables (e.g., body side, mental strategy) is organized in relationship to a third variable (body part). This effectively leads to functional segregation of body part at a population level. Such functional segregation between body parts is very similar in principle to the relationship between planning and execution-related activity that has recently been described in frontal motor areas (Churchland et al., 2010; Kaufman et al., 2014) where planning activity fails to excite subspaces that are hypothesized to produce muscle output.

But why are some variables functionally overlapping while others are functionally segregated? One possible answer is computational savings. Overlapping activity at the level of the population may be rooted in shared computational resources. For example, many computations related to planning and executing grasps including object affordance processing as well as basic kinematic processing would be similar for the right and left hand. Motor imagery has also been hypothesized to engage internal models used for sensory estimation during overt execution (see below) and thus imagery and execution should
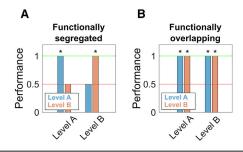
rely on largely overlapping computations. Thus, despite the potential computational benefits to random mixing of variables (Fusi et al., 2016), the computational savings of overlapping resources for certain classes of computations may outweigh losses in the total information the population encodes.

Another possibility is that the highly overlapping representations provide part of the neural substrate through which transfer of learning occurs. Motor skills learned with one hand frequently result in improvements in performance with the other hand (Amemiya et al., 2010). Likewise, use of motor imagery is found to improve performance during motor execution (Dickstein and Deutsch, 2007). One possibility is that overlapping networks would be able to facilitate this sort of transfer of learning. For example, repeatedly imagining a movement with the right hand would recruit a similar network as executing a movement with the right hand, making any neural adaptation from learning the movement more likely to transfer between the strategies.

Despite the greater functional overlap between body side and strategy, it is important to note that all the tested movement conditions are still differentiable from each other (Figure 8). Interestingly, this greater overlap for body side may explain why patients with motor deficits often "mirror" movements in a contralateral limb. In cerebral palsy, for example, patients making a grasp with their left hand often mirror the movements with their right hand (Kuhtz-Buschbeck et al., 2000).

A point of note is that the movements selected in this study (hand squeezes and shoulder shrugs) are not necessarily the best exemplars of movements of the respective body parts. Different combinations of hand or shoulder movements may have slightly more or less overlap. Bilateral symmetric hand squeezes and shoulder shrugs occur more naturally than "squeeze-and-shrug" actions which may not be part of the natural movement repertoire. The statistical frequency with which different body parts are moved together could also affect the degree of functional overlap between the body parts. A better understanding of how different exemplars of movements across different effectors relate will be important in understanding the functional organization of motor actions in AIP.
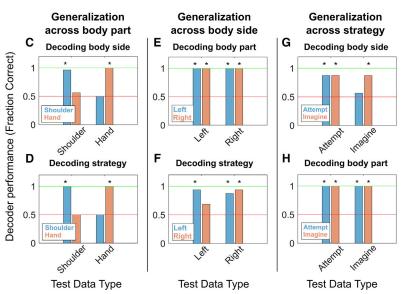
Figure 7. Representations of Variables Generalize across Side and Strategy but Not Body Part

(A and B) Schematic illustrating expected classifier behavior if variables are functionally segregated (A) versus overlapping (B). (A) Functional segregation within a variable (e.g., body part) implies that a classifier trained to differentiate the levels of one dimension (e.g., right from left) will not generalize across the levels of the dimension of interest (e.g., from shoulder to hand) resulting in chance performance. (B) In contrast, functional overlap implies generalization resulting in above-chance performance when comparing classifier performance across levels.

(C) Performance of decoders trained on data split by body part for classifying the body side. Blue/orange bars represent the performance of the decoder trained on shoulder/hand movement data. Horizontal axis labels represent which body part's data each decoder was tested on. Performance was measured as the fraction of trials accurately classified by the decoder, with in-sample performance determined by cross-validation. Asterisks represent performance significantly different from chance, as determined by a rank shuffle test. The red line represents chance performance level (0.5) while the green line represents perfect performance (1.0).

(D) Similar to (C) but decoding strategy instead of body side.

(E and F) Similar to (C) but with data split by body side and decoding for body part (E) and strategy (F), respectively.

(G and H) Similar to (C) but with data split by strategy and decoding for body side (G) and body part (H), respectively.

Attempted and Imagined Movements in Human AIP after Long-Term Injury

In this study, we looked at neural coding of imagined and attempted actions above and below the level of injury in a paralyzed individual. By current theory, imagined movements may represent the simulation of an internal model of the arm, a model that also forms the basis for sensory estimation during overt forms of behavior (Gail and Andersen, 2006; Jeannerod, 1995; Mulliken et al., 2008). In broad strokes, this theory predicts that neural representations of imagined and overt movements should have a high degree of similarity given the shared neural substrate, but also be different given the absence of movement during imagery (Jeannerod, 1995; Munzert et al., 2009). Our results support this view insofar as we demonstrate the high degree of functional overlap between imagined and attempted movements. However, we also show that neural differences between imagined and executed movements persist and are highly similar even after long-term injury and disuse (see Figures 4 and 5). This observation that neural coding differences persist even despite injury preventing overt movement during attempted actions is inconsistent with the proposal that the primary difference between imagined and executed movements is the actual movement itself (Jeannerod, 1995). Further, the patterns of similarities and differences in how the population codes mental strategy and body side—for instance, the preference for attempted over imagined movements for the right but not left side of the

body (Figure 4A versus 4B)—demonstrate that higher-order population structure is conserved following injury. This suggests that preservation of motor intention signals enables successful BMI
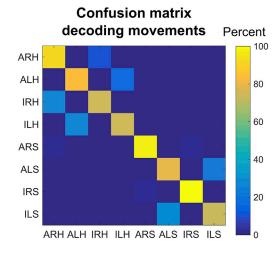

Figure 8. All Movement Variables Decodable from the Population
Confusion matrix for cross-validated classification of the eight movement conditions (ALH, attempt left hand; ILH, imagine left hand; ARH, attempt right hand; IRH, imagine right hand; ALS, attempt left shoulder; ILS, imagine left shoulder; ARS, attempt right shoulder; IRS, imagine right shoulder).

functionality many years post-injury (Aflalo et al., 2015). A better understanding of how different cortical subregions maintain representations of motor intent post-injury may help inform choice of implant sites as a function of time post injury.

These results demonstrate for the first time that networks activated during attempted actions are highly overlapping with networks activated during imagined actions at the level of populations of individual neurons and that the correspondence between actions is body part specific (hand and shoulder). However, there is a symmetry in our results such that networks activated during right-hand actions are highly overlapping with networks activated for left-hand actions, and the correspondence between right and left actions are strategy specific (e.g., right-side actions look more like left-side actions using the same strategy). In other words, the relationship between imagined and attempted actions is similar in basic form to the relationship between left- and right-sided actions, although the degree of overlap is greater for strategy.

The current experiment was performed in the larger context of a brain-machine interface clinical (BMI) trial. We have previously shown that a paralyzed patient can use motor imagery to control a robotic limb (Aflalo et al., 2015). Is the use of motor imagery the best method for the user of a BMI to modulate their own neural activity? Alternatively, it is possible that attempted movements somehow better engage or otherwise enable the user to control an external device. Here we show that the distinction between imagined and attempted actions is preserved, even in limbs for which no movement is possible. Future work is needed to determine whether these differences translate into performance differences during closed-loop neural control.

Orofacial Coding in Human AIP

We included speech conditions in which N.S. spoke "left" and "right" as a third fundamentally different movement. A smaller proportion of neurons were tuned more weakly to speech acts, demonstrating that not all actions are coded in an equivalent manner in AIP (Figures 2 and 4). This task was not designed to understand the functional significance of "speech" tuned units, but one possibility is that these neurons code for orofacial movements and may form the building blocks for more complex coordinated movements of behavioral relevance such as coordinated movement of the hand to the mouth for feeding or tearing open a bag of chips with your mouth. It is also possible that such orofacial tuning coordinates "grasping" actions performed with the mouth (Jastorff et al., 2010).

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral setup
  - Physiological recordings
  - Task procedure

- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Unit selection
  - Linear analysis 1
  - AUC analysis
  - Linear analysis 2
  - Test of population bias in representing motor variables
  - Degree of specificity
  - Correlation between representations
  - Decoder analysis
  - Condition classification
  - Unit quality classification
- DATA AND SOFTWARE AVAILABILITY
- ADDITIONAL RESOURCES

REFERENCES

Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., Shanfield, K., Hayes-Jackson, S., Aisen, M., Heck, C., et al. (2015). Neurophysiology. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. Science 348, 906–910.

Amemiya, K., Ishizu, T., Ayabe, T., and Kojima, S. (2010). Effects of motor imagery on intermanual transfer: a near-infrared spectroscopy and behavioural study. Brain Res. 1343, 93–103.

Andersen, R.A., and Buneo, C.A. (2002). Intentional maps in posterior parietal cortex. Annu. Rev. Neurosci. 25, 189–220.

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron 63, 568–583.

Andersen, R.A., Essick, G.K., and Siegel, R.M. (1987). Neurons of area 7 activated by both visual stimuli and oculomotor behavior. Exp. Brain Res. 67, 316–322.

Asher, I., Stark, E., Abeles, M., and Prut, Y. (2007). Comparison of direction and object selectivity of local field potentials and single units in macaque posterior parietal cortex during prehension. J. Neurophysiol. 97, 3684–3695.

Astafiev, S.V., Shulman, G.L., Stanley, C.M., Snyder, A.Z., Van Essen, D.C., and Corbetta, M. (2003). Functional organization of human intraparietal and frontal cortex for attending, looking, and pointing. J. Neurosci. 23, 4689–4699.

Balint, R. (1909). Seelenlahmung des "Schauens," optische Ataxie, raumliche Storung der Aufmerksamkeit. Monatsschr. Psychiatr. Neurol. 25, 51–81.

Beurze, S.M., de Lange, F.P., Toni, I., and Medendorp, W.P. (2009). Spatial and effector processing in the human parietofrontal network for reaches and saccades. J. Neurophysiol. 101, 3053–3062.

Brainard, D.H. (1997). The psychophysics toolbox. Spat. Vis. 10, 433–436.

Christopoulos, V.N., Bonaiuto, J., Kagan, I., and Andersen, R.A. (2015). Inactivation of parietal reach region affects reaching but not saccade choices in internally guided decisions. J. Neurosci. 35, 11719–11728.

Churchland, M.M., and Cunningham, J.P. (2014). A dynamical basis set for generating reaches. Cold Spring Harb. Symp. Quant. Biol. 79, 67–80.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., and Shenoy, K.V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron 68, 387–400.

Connolly, J.D., Andersen, R.A., and Goodale, M.A. (2003). FMRI evidence for a 'parietal reach region' in the human brain. Exp. Brain Res. 153, 140–145.

Cui, H., and Andersen, R.A. (2007). Posterior parietal cortex encodes autonomously selected motor plans. Neuron 56, 552–559.

Culham, J.C., Danckert, S.L., DeSouza, J.F.X., Gati, J.S., Menon, R.S., and Goodale, M.A. (2003). Visually guided grasping produces fMRI activation in dorsal but not ventral stream brain areas. Exp. Brain Res. 153, 180–189.

Dickstein, R., and Deutsch, J.E. (2007). Motor imagery in physical therapist practice. Phys. Ther. 87, 942–953.

Fattori, P., Breveglieri, R., Marzocchi, N., Filippini, D., Bosco, A., and Galletti, C. (2009). Hand orientation during reach-to-grasp movements modulates neuronal activity in the medial posterior parietal area V6A. J. Neurosci. 29, 1928–1936.

Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. Curr. Opin. Neurobiol. 37, 66–74.

Gail, A., and Andersen, R.A. (2006). Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J. Neurosci. 26, 9376–9384.

Gallese, V., Murata, A., Kaseda, M., Niki, N., and Sakata, H. (1994). Deficit of hand preshaping after muscimol injection in monkey parietal cortex. Neuroreport 5, 1525–1529.

Gallivan, J.P., McLean, D.A., Smith, F.W., and Culham, J.C. (2011). Decoding effector-dependent and effector-independent movement intentions from human parieto-frontal brain activity. J. Neurosci. 31, 17149–17168.

Gallivan, J.P., McLean, D.A., Flanagan, J.R., and Culham, J.C. (2013). Where one hand meets the other: limb-specific and action-dependent movement plans decoded from preparatory signals in single human frontoparietal brain areas. J. Neurosci. 33, 1991–2008.

Gerardin, E., Sirigu, A., Lehéricy, S., Poline, J.B., Gaymard, B., Marsault, C., Agid, Y., and Le Bihan, D. (2000). Partially overlapping neural networks for real and imagined hand movements. Cereb. Cortex 10, 1093–1104.

Graziano, M.S.A., and Aflalo, T.N. (2007). Rethinking cortical organization: moving away from discrete areas arranged in hierarchies. Neuroscientist 13, 138–147.

Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J. Neurophysiol. 84, 401–414.

Harris, K.D., Quiroga, R.Q., Freeman, J., and Smith, S.L. (2016). Improving data quality in neuronal population recordings. Nat. Neurosci. 19, 1165–1174.

Heed, T., Beurze, S.M., Toni, I., Röder, B., and Medendorp, W.P. (2011a). Functional rather than effector-specific organization of human posterior parietal cortex. J. Neurosci. 31, 3066–3076.

Hinkley, L.B.N., Krubitzer, L.A., Padberg, J., and Disbrow, E.A. (2009). Visual-manual exploration and posterior parietal cortex in humans. J. Neurophysiol. 102, 3433–3446.

Holmes, G. (1918). Disturbances of visual orientation. Br. J. Ophthalmol. 2, 449–468.

Hwang, E.J., Hauschild, M., Wilke, M., and Andersen, R.A. (2012). Inactivation of the parietal reach region causes optic ataxia, impairing reaches but not saccades. Neuron 76, 1021–1029.

Hwang, E.J., Hauschild, M., Wilke, M., and Andersen, R.A. (2014). Spatial and temporal eye-hand coordination relies on the parietal reach region. J. Neurosci. 34, 12884–12892.

Jastorff, J., Begliomini, C., Fabbri-Destro, M., Rizzolatti, G., and Orban, G.A. (2010). Coding observed motor acts: different organizational principles in the parietal and premotor cortex of humans. J. Neurophysiol. 104, 128–140.

Jeannerod, M. (1995). Mental imagery in the motor context. Neuropsychologia 33, 1419–1432.

Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. 17, 440–448.

Klaes, C., Kellis, S., Aflalo, T., Lee, B., Pejsa, K., Shanfield, K., Hayes-Jackson, S., Aisen, M., Heck, C., Liu, C., and Andersen, R.A. (2015). Hand shape representations in the human posterior parietal cortex. J. Neurosci. 35, 15466–15476.

Kubanek, J., and Snyder, L.H. (2015). Reward-based decision signals in parietal cortex are partially embodied. J. Neurosci. 35, 4869–4881.

Kuhtz-Buschbeck, J.P., Sundholm, L.K., Eliasson, A.-C., and Forssberg, H. (2000). Quantitative assessment of mirror movements in children and adolescents with hemiplegic cerebral palsy. Dev. Med. Child Neurol. 42, 728–736.

Lehmann, S.J., and Scherberger, H. (2013). Reach and gaze representations in macaque parietal and premotor grasp areas. J. Neurosci. 33, 7038–7049.

Lehmann, S.J., and Scherberger, H. (2015). Spatial Representations in Local Field Potential Activity of Primate Anterior Intraparietal Cortex (AIP). PLoS ONE 10, e0142679.

Levy, I., Schluppeck, D., Heeger, D.J., and Glimcher, P.W. (2007). Specificity of human cortical areas for reaches and saccades. J. Neurosci. 27, 4687–4696.

McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., and Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. Neuron 83, 202–215.

Mountcastle, V.B., Lynch, J.C., Georgopoulos, A., Sakata, H., and Acuna, C. (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. J. Neurophysiol. 38, 871–908.

Mulliken, G.H., Musallam, S., and Andersen, R.A. (2008). Forward estimation of movement state in posterior parietal cortex. Proc. Natl. Acad. Sci. USA 105, 8170–8177.

Munzert, J., Lorey, B., and Zentgraf, K. (2009). Cognitive motor processes: the role of motor imagery in the study of motor representations. Brain Res. Brain Res. Rev. 60, 306–326.

Murata, A., Gallese, V., Luppino, G., Kaseda, M., and Sakata, H. (2000). Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. J. Neurophysiol. 83, 2580–2601.

Pouzat, C., Mazor, O., and Laurent, G. (2002). Using noise signature to optimize spike-sorting and to assess neuronal classification quality. J. Neurosci. Methods 122, 43–57.

Prado, J., Clavagnier, S., Otzenberger, H., Scheiber, C., Kennedy, H., and Perenin, M.T. (2005). Two cortical systems for reaching in central and peripheral vision. Neuron 48, 849–858.

Quian Quiroga, R., Snyder, L.H., Batista, A.P., Cui, H., and Andersen, R.A. (2006). Movement intention is better predicted than attention in the posterior parietal cortex. J. Neurosci. 26, 3615–3620.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. 17, 1784–1792.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167–170.

Snyder, L.H., Grieve, K.L., Brotchie, P., and Andersen, R.A. (1998). Separate body- and world-referenced representations of visual space in parietal cortex. Nature *394*, 887–891.

Tibshirani, R., Ealther, G., and Hastie, T. (2001). Estimating the number of clusters in a dataset via the Gap statistic. J.R. Stat. Soc. *63*, 411–423.

Ungerleider, L.G., and Mishkin, M. (1982). Two cortical visual systems. In Analysis of Visual Behavior, D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield, eds. (MIT Press), pp. 549–585.

Yttri, E.A., Wang, C., Liu, Y., and Snyder, L.H. (2014). The parietal reach region is limb specific and not involved in eye-hand coordination. J. Neurophysiol. *111*, 520–532.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
MATLAB R2016a	MathWorks	http://www.mathworks.com
Other		
Neuroport System	Blackrock Microsystems	http://blackrockmicro.com/
EMG	B&L Engineering	http://www.bleng.com/, S.N. A3735

CONTACT FOR RESOURCE SHARING

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Carey Y. Zhang (cyzhang@caltech.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Subject N.S. is a 59-year-old female tetraplegic 7 years post-injury and has a C3-C4 spinal lesion (motor complete), having lost control and sensation in her hands but retaining movements and sensations in her upper trapezius. In this paper we refer to contraction of the upper trapezius as "shoulder movements" as short-hand for the resulting shoulder shrugging movement. The studies were approved by the California Institute of Technology, University of California, Los Angeles, and Casa Colina Centers for Rehabilitation Internal Review Boards. Informed consent was obtained from the participant N.S. after the nature of the study and possible risks were explained. Study sessions occurred at Casa Colina Centers for Rehabilitation.

METHOD DETAILS

Behavioral setup
All tasks were performed with N.S. seated in her motorized wheel chair. Tasks were displayed on a 27-inch LCD monitor in a lit room. The monitor was positioned so that the screen occupied approximately 40 degrees of visual angle. Stimulus presentation was controlled using the Psychophysics Toolbox (Brainard, 1997) for MATLAB. No eye fixation was required or enforced.

Physiological recordings
Subject N.S. was implanted with two 96-channel Neuroport arrays (Blackrock Microsystems model numbers 4382 and 4383) in putative homologs of area AIP and Brodmann's Area 5d. Array placement was determined based on preoperative fMRI (Aflalo et al., 2015) and the array was placed at Talairach coordinate [-36 lateral, 48 posterior, 53 superior]. Neural activity was amplified, digitized, and recorded with the Neuroport neural signal processor (NSP). The Neuroport System, comprising the arrays and NSP, has received FDA clearance for < 30 days acute recordings. We received FDA IDE clearance (IDE #G120096, G120287) to extend the duration of the implant for the purposes of a brain-machine interface clinical study using signals from posterior parietal cortex.

   During recording, thresholds for action potential detection were set at −4.5 times the root-mean-square after high pass filtering (250 Hz cut-off) the full-bandwidth signal sampled at 30 kHz in the Central software suite (Blackrock Microsystems). Each waveform was composed of 48 samples (1.6 ms) with 10 samples prior to triggering 38 samples after. Single and multiunit activity was sorted by k-medoids clustering using the gap criteria to estimate the total number of clusters (Tibshirani et al., 2001). Clustering was performed on the first n principal components, where n was selected to account for 95% of waveform variance. Post hoc review of sorted unit statistics showed that channels were sorted with between 2-4 principal components (see Figure S2A). Results of offline sorting were reviewed and adjusted if deemed necessary following standard practice (Harris et al., 2016). Only neurons recorded from the array implanted in putative AIP were analyzed. Pooling across all versions of the task, on average 93 sorted units were recorded from N.S. per session. Furthermore, to avoid bias, all spike sorting was performed prior to any analysis and blind to a unit or channel's response during the task. We used several metrics to quantify sort quality (see Figures S2B–S2F) including 1) the percentage of interspike intervals (ISIs) shorter than 3ms, 2) the signal-to-noise ratio (SNR) of the mean waveform, 3) the between spike projection distance (Pouzat et al., 2002), 4) the modified coefficient of variation of the ISI (CV2), and 5) the cluster isolation distance (Harris et al., 2000) of each sorted cluster.

We recorded electromyogram (EMG) activity over the right trapezius muscle using B&L Engineering EMG electrodes. Raw analog EMG activity was fed into the NSP, aligned with neural signals, and sampled at 2 kHz. Signals were band-pass filtered (5^{th} order Butterworth filter with cut-off frequencies of 10 and 250Hz), full-wave rectified, and smoothed (box-car, 50ms window).

Task procedure

Several versions of a delayed movement task were constructed to determine the extent of tuning to control strategy within the neural populations recorded from AIP. In the primary task (Figure 1A), N.S. was cued for 2.5 s to what strategy (imagine or attempt), side (left or right), and body part (hand or shoulder) to use, e.g., attempting to squeeze the right hand. In total there were eight possible actions which were pseudorandomly interleaved such that each condition was performed once before repetition. After a delay of 1.5 s, N.S. was cued to perform the cued action. Between each trial there was a 3 s inter-trial interval (ITI). Hand movements were hand squeezes while shoulder movements were shoulder shrugs (contraction of the trapezius). We ran 64 trials (8 trials per condition) on each session. This task was run over the course of 4 non-consecutive days. In total 357 units were recorded across the four recording sessions. Unless otherwise indicated, all figures were generated from data collected from this version of the task.

In a separate set of sessions, we repeated the experiment with the modification that shoulder shrugging movements were replaced with shoulder abduction in the frontal plane. Attempted shoulder abduction resulted in no overt movement and thus allowed us to compare body part representations exclusively below the level of injury. Six sessions run over the course of 6 non-consecutive days were recorded resulting in 629 recorded units. Each session contained 64 trials (8 trials per condition).

Movements of the shoulder and hand are frequently made together during natural behavior. We modified the delayed movement task by adding "speak left" and "speak right" as two actions unrelated to any hand or shoulder movements. To avoid overly long data collection sessions (as determined by patient feedback), we minimized the number of conditions by splitting sessions into either hand or shoulder movements exclusively resulting in 6 conditions pseudorandomly interleaved (Imagine Left, Imagine Right, Attempt Left, Attempt Right, Speak Left, Speak Right). Three sessions were recorded for the hand and the shoulder separately, with each session containing 72 trials (12 trials for each condition). In total 299 units were recorded for sessions using the hand while 228 units were recorded for sessions using the shoulder.

QUANTIFICATION AND STATISTICAL ANALYSIS

All analyses were performed using MATLAB 2016a.

Unit selection

Analyses were performed on all units regardless of sort quality for statistical power and ease of presentation. To ensure that such pooling did not bias the conclusions of this paper, we performed the analyses separately on well isolated versus potentially multi-unit activity and found the results to be similar (see Figure S1 for more details and also Unit quality classification below for how high-quality single units were identified). Units were pooled across days assuming independent populations across recording days. Analysis of separate days was also performed to demonstrate stability of results across sessions (Figure S5). Units with mean firing rates less than 1.5 Hz were excluded from the analysis so that low firing rate effects would be minimized.

Linear analysis 1

We used a linear regression analyses to quantify tuning to each condition (e.g., for experiment 1, 8 total conditions from the possible combinations of the 2 strategies, 2 body parts, and 2 body sides). We created a design matrix consisting of indicator variables for each condition (e.g., the indicator for right attempted shoulder movements would consist of a vector where data points associated with right attempted shoulder movements would be assigned a 1, while all other conditions and the baseline samples would be assigned a 0). We then estimated firing rate as a linear combination of these indicator variables:

$$FR = \sum_c \beta_c X_c + \beta_0$$

where FR is the firing rate, X_c is the vector indicator variable for condition c, β_c is the estimated scalar weighting coefficient for condition c, and β_0 is a constant offset term. In such a model, the estimated beta coefficients represent the expected firing rate changes from baseline for each condition. Tuning to each condition (Figure 2A) was based on the p value of the t statistic for each associated beta coefficient. This definition of whether a unit is significantly tuned to a condition is used as an inclusion criterion for some analyses, with the significance level (e.g., $p < 0.05$, uncorrected or Bonferroni corrected) depending on the specific analysis. The significance level of the differences between the number of units tuned to each condition was calculated using a two-sided Wilcoxon rank sum test on the distribution of p values for each pair of conditions (Figure S3).

AUC analysis

We performed a ROC analysis to quantify tuning strength for each condition. For each unit, strength of tuning was summarized as the area under the curve (AUC) when comparing each condition's Go or Delay neural response to baseline. The AUC values can range from 0 to 1, with 1 indicating that every go/delay measurement is *greater than* every baseline measurement (excitatory

response) and 0 indicating that every go/delay measurement is *less than* every baseline measurement (inhibitory response). To summarize the population (Figure 2B), we combined excitatory and inhibitory responses by reversing condition labels for AUC responses below 0.5. The separated responses are reported in Figure S4. Only units with significant AUC (p < 0.05, permutation test) are included in the population average and thus we answer the question ''what is the average strength of tuning for each condition for units with significant tuning to the condition.'' Hence this measure is descriptive and not a statistical assessment of significant tuning in the population. Pairwise differences between the AUC for each condition were calculated using a two-sided Wilcoxon rank sum test on the significant AUC values (Figure S3).

Linear analysis 2

Above we performed a linear regression analyses where firing rate was modeled as a function of each condition response. Here we perform an additional linear analysis where firing rate is modeled as the linear combination of indicator variables for each motor variable (strategy, body side, and body part) and their interaction. All temporal epochs were identical; however, the design matrix was updated to reflect the new model:

$$FR = \beta_1 Strategy + \beta_2 BodySide + \beta_3 BodyPart + \beta_4 Strategy * BodySide + \beta_5 Strategy * BodyPart + \beta_6 BodySide * BodyPart + \beta_0$$

Each unit was classified as being tuned to a term if the p value of the corresponding beta coefficient was significant (i.e., p < 0.05, uncorrected). To examine the effect of the different motor variables on firing rate patterns across the population we performed a MANOVA test on the linear beta coefficients of the model. All units were used in the test (regardless of whether they showed tuning to a variable or not).

Test of population bias in representing motor variables

To examine whether there were systematic biases for the three different motor variables (strategy, body side, body part) across the population we performed a MANOVA test. The baseline firing rate of each neuron (taken during the intertrial interval) was subtracted from the firing rate of the neuron during the Go phase, and this baseline-subtracted firing rate was used in the test. All units were used in the test (regardless of whether they showed tuning to a variable or not).

In comparing between the 8 movement conditions and the speaking conditions, we performed a t test on the baseline subtracted firing rates, with the 8 movement conditions pooled into one group and the 2 speaking conditions pooling into a second group. Once again, all units were used regardless of their tuning.

Degree of specificity

We used a degree of specificity analysis to begin to understand how the different movement attributes were encoded relative to each other in individual neurons. This allowed us to test for population level tendencies for e.g., exclusive activation for one effector versus the other or whether tuning tended to be overlapping. We computed the degree of specificity for each motor dimension (e.g., left versus right, imagine versus attempt, hand versus shoulder) as the normalized difference in beta values computed for each condition $((|\beta_1| - |\beta_2|)/(|\beta_1| + |\beta_2|))$. Beta value extraction is described above (see Linear Model 1). The degree of specificity ranges from -1 to 1. A value of 1 indicates exclusive modulation for comparison variable 1, a value of -1 indicates exclusive modulation for comparison variable 2, while a value of 0 indicates comparable activation to both variables. Degree of specificity was only computed if at least one beta coefficient in the comparison was significant (p < 0.05, uncorrected). We used a Wilcoxon signed rank test to determine whether the median of the distribution of specificity values was significantly shifted from 0.

Correlation between representations

We used a population correlation analysis to measure the similarity between the neural representations of each condition. For each neuron, the beta coefficient for each condition (see Linear Model 1) was normalized by the corresponding 95% confidence interval to ensure a common scale proportional to signal-to-noise. The normalized beta values for each unit and for each condition were used to create a vector summarizing the population response. We used the correlation between these vectors to measure the similarity in neural space between conditions. Only units with a significant beta coefficient for at least one condition (p < 0.05, Bonferroni corrected) were included in the analyses.

We used hierarchical clustering (agglomerative hierarchical cluster tree; using the built-in MATLAB 2016a linkage and dendrogram functions with unweighted average correlation as the measure) to summarize the structure in the patterns of correlation between the different conditions (McKenzie et al., 2014).

Decoder analysis

Functional segregation of body parts should result in minimal shared representations of other motor variables across different body parts. Conversely, the lack of functional segregation between different body sides (or strategy) should lead to comparatively greater shared representation of the other motor variables across different body sides (or strategies). To test this, we trained a linear classifier (linear discriminant with equal diagonal covariance matrices) to differentiate between the two levels of one motor variable, restricting the training data to one level of a second motor variable and then applying the classifier to the other level. Cross-validated performance restricted to the first level was also computed as this provided an upper bound on classification accuracy given the signal

to noise of our data. For instance, we trained a classifier to differentiate left from right movements on shoulder movement trials and tested it on hand movement trials, also computing cross-validated performance within shoulder trials (Figure 7C). For features, we used firing rates from the first 2 s of the "Go" phase for units with significant tuning to any of the eight movement conditions. Only units significantly tuned to at least one of the 8 movement conditions (p < 0.05, Bonferroni corrected) were included in the analysis. Classifier performance was determined to be above chance if its performance was greater than 95% of decoders trained on randomly shuffled data (1000 shuffles).

Condition classification

We performed classification analyses using linear discriminant analysis with equal diagonal covariance matrices for each condition. Classification accuracy was estimated using stratified leave-one out cross-validation. Classification features were constructed using the first 2 s of the "Go" phase. Only units with significant tuning to any of the eight movement conditions were used (p < 0.05, Bonferroni corrected). Significant units were estimated from the training data and applied to the test data for each fold of the cross-validation routine to avoid peaking effects. Units were pooled across days for analyses. For each fold (1000 repetitions), one feature per condition was randomly sampled as the test data with all other samples used for training.

Unit quality classification

For the analysis in Figure S1, we needed to separate spike sorted units into high-quality single units and multi-units. This classification was done based on the cluster isolation distance (Harris et al., 2000) with a threshold of $10^{1.6}$ dividing high-quality single units and multi-units. This threshold was chosen based on visual inspection of the distribution of all cluster isolation distance values (see Figure S2).

DATA AND SOFTWARE AVAILABILITY

Data and MATLAB analysis scripts available upon request from Carey Y. Zhang (cyzhang@caltech.edu).

ADDITIONAL RESOURCES

This study was conducted as part of NIH clinical trial NCT01958086.