

Conservation of preparatory neural events regardless of how movement is initiated

Antonio H. Lara¹, Gamaleldin F. Elsayed^{1,2}, John P. Cunningham^{2,5,6} and Mark M. Churchland^{1,3,4,5} *

¹ Department of Neuroscience, Columbia University Medical Center, New York, New York, USA.
York, USA.

² Center for Theoretical Neuroscience, Columbia University, New York, New York, USA

³ David Mahoney Center for Brain and Behavior Research, Columbia University Medical Center, New York, New York, USA.

⁴ Kavli Institute for Brain Science, Columbia University Medical Center, New York, New York, USA.

⁵ Grossman Center for the Statistics of Mind, Columbia University Medical Center, New York, New York, USA.

⁶ Department of Statistics, Columbia University, New York, New York, USA.

*Correspondence: mmc3502@columbia.edu

Abstract

Voluntary movement is believed to be preceded by a preparatory stage. Evidence arises from experiments where a delay separates instruction and execution cues. While this sequence emulates some real-world situations (*e.g.*, swatting a fly upon landing) movements are commonly made at a moment of one's choosing (reaching for a coffee cup) or are made reactively (intercepting a falling cup). To ascertain whether neural events are conserved across such contexts, we examined motor cortex population-level responses in monkeys when reaches were initiated either after an imposed delay, at a self-chosen time, or reactively with very low latency. We found that the same preparatory and movement-related events were conserved. However, preparation was temporally flexible and could be remarkably brief. Our findings support the existing hypothesis that preparation is an obligatory stage that achieves a consistent state prior to movement. Yet our results reveal that preparation can unfold more rapidly than previously supposed.

1 Introduction

2 Multiple lines of evidence argue that voluntary movement is preceded by a preparatory stage¹⁻⁹. Most
3 fundamentally, the voluntary reaction time (RT, the time between a sensory stimulus and the onset of an
4 evoked movement) is typically longer than expected given afferent and efferent delays, suggesting a time-
5 consuming preparatory process separating sensation from action. RTs typically become shorter when a delay
6 period separates an instruction from a go cue, presumably because preparation has time to complete before
7 the go cue^{2,4,8,10}. Neurons in many brain areas — including primary motor cortex (M1) and dorsal premotor
8 cortex (PMd) — respond selectively during the delay^{1,4,11-15}. Delay-period activity is predictive of RT variability<sup>4,8-
9 10,16</sup>, and its electrical disruption erases the RT-savings provided by the delay¹⁷. These observations are
10 consistent with the presence of a neural process that must occur before voluntary movement can be initiated,
11 yet does not itself cause movement and can therefore be completed in advance when circumstances allow.

12 Although appealing, aspects of the above interpretation remain incomplete, uncertain, or controversial. At the
13 neural level, it remains uncertain whether delay-period activity reflects an essential preparatory stage. It has
14 been argued that it does, and that movement-period neural dynamics are seeded by a preparatory state
15 reflected in delay-period activity^{10,18-21}. This hypothesis predicts that there should be a recognizable
16 progression of preparatory and movement-related activity even in the absence of a delay. Determining
17 whether this is true has been challenging. An early study revealed that at least some aspects of delay-period
18 activity are recapitulated without a delay⁶. Yet in a recent study, Ames et al.²² found that the neural state
19 during the delay-period was bypassed in the absence of a delay. This result has been variously interpreted as
20 evidence that delay-period activity is beneficial but not essential²², that delay-period activity is not preparatory
21 but suppressive²³, or that the range of acceptable preparatory states is broad⁸. Yet Ames et al. also found that
22 the initial ~50 ms of target-driven response was similar with and without a delay. As in⁶, this finding suggested
23 that early, putatively preparatory aspects of the response may be conserved.

24 Ambiguity at the neural level is underscored by behavioral results suggesting that preparation may be
25 unnecessary²³. Preparation has generally been considered to be a time-consuming process that makes a
26 sizeable contribution to the RT^{2,4,5,9,10}. Yet under certain circumstance, humans and monkeys display very short
27 RTs²⁴⁻²⁶ despite having no time to prepare in advance. Such RTs appear incompatible with a preparatory stage
28 that consumes considerable time.

29 We wished to explore the hypothesis that delay-period activity reflects a preparatory process that also occurs
30 when there is no delay, but displays enough temporal flexibility to allow very short RTs. In principle, this
31 hypothesis is readily tested via a straightforward strategy: by exploiting the delay period to identify putatively
32 preparatory neural activity, then inquiring whether and how similar activity is present in other contexts.

33 Unfortunately, this strategy is difficult or impossible to accomplish at the level of single neurons. Most neurons
34 with delay-period activity also exhibit movement-related activity. Indeed, a neuron that primarily displays
35 delay-period activity for some movements can primarily display movement-related activity for other
36 movements. Thus, in the absence of a delay, any activity that occurs between the go cue and movement onset
37 could be preparatory, movement-related, or some combination of the two. One might hope to make a
38 distinction base on timing: pinpointing a time before which activity is putatively preparatory and after which it
39 is putatively movement-related. However, the nature of such timing is unknown, and is itself one of the
40 subjects of inquiry. To identify putatively preparatory and movement-related patterns, one must therefore
41 move beyond single-neuron responses, and instead focus upon population-level properties.

42 To this end, we leveraged the recent finding²⁷ that delay-period activity and movement-related activity occupy
43 orthogonal subspaces. This allowed us to use responses during a delay to define a putatively-preparatory
44 subspace, and responses during movement to define a putatively-movement subspace. The temporal
45 separation provided by the delay period makes it possible to define these subspaces based on well-separated
46 epochs, avoiding the interval where the transition from preparation to movement is uncertain. We further
47 exploited the empirical fact that these spaces are nearly orthogonal, which allows activity in one subspace to
48 be measured independently of activity in the other subspace. Having identified the relevant subspaces, we
49 then observed activity in those subspaces when movements were initiated without an experimenter-imposed
50 delay. Specifically, monkeys performed reaches in three interleaved contexts. The ‘cue-initiated’ context
51 employed a standard delay period (which was used was to define the spaces), the ‘self-initiated’ context
52 allowed the monkey freedom to decide when to move, and the ‘quasi-automatic’ context required monkeys to
53 reactively intercept a moving target with no time to prepare in advance.

54 We found that both preparatory-subspace and movement-subspace patterns of neural activity were conserved
55 across contexts. For example, quasi-automatic reaches, although executed with very short RTs, exhibited the
56 same sequence of preparatory and movement-related events as self-initiated reaches. However, the time-
57 course of preparatory-subspace events was remarkably flexible across contexts: preparatory-subspace activity
58 developed slowly in anticipation of self-initiated reaches, but lead movement-subspace activity by only a few
59 tens of milliseconds for quasi-automatic reaches. Our results support the hypothesis that a conserved
60 preparatory process is present regardless of whether there exists an experimenter-imposed delay. However,
61 the temporal flexibility of that process is considerable; it can arise well before movement onset, but can also
62 consume surprisingly little time. In agreement with^{24,26}, such brevity argues against the idea that movement
63 preparation necessarily involves time-consuming cognitive or high-level planning processes. It is more likely
64 that preparatory activity plays a straightforward and mechanistic role: initializing the circuits that are about to
65 produce descending movement commands.

66 **Results**

67 We trained two monkeys (Ba and Ax) to execute reaches in three contexts, which differed regarding how
68 movement initiation is cued, and how much time is allowed for preparation. In all contexts, reaches began at a
69 central touch-point and were made to targets arranged radially in the vertical-horizontal plane. The cue-
70 initiated context (**Fig. 1a**) emulated the standard instructed-delay paradigm: a variable delay period (0–1000
71 ms) separated target onset from an explicit go cue. We analyzed trials with delays >400 ms; shorter delays
72 were included to encourage timely and consistent preparation. In the self-initiated context (**Fig. 1b**) monkeys
73 were free to reach upon target presentation, but waiting longer yielded proportionally larger rewards up to a
74 limit at 1200 ms. Growing reward size was mirrored by growing target size. Thus, available reward was always
75 cued directly; monkeys did not need to estimate elapsed time. Reward (and target size) ceased growing at the
76 moment of movement onset. In the quasi-automatic context (**Fig. 1c**) the target moved rapidly along a radial
77 path towards the screen's edge. The target moved as soon as it appeared, and monkeys had to intercept it
78 mid-flight. Trials for the three contexts were randomly interleaved. The color of the central touch-point and
79 target specified the context: red, blue and yellow indicated the cue-initiated, self-initiated and quasi-automatic
80 contexts. Monkeys successfully completed the majority of trials for all three contexts: 93% and 95% of cue-
81 initiated trials (monkey Ba and Ax respectively), 94% and 97% of self-initiated trials, and 93% and 93% of quasi-
82 automatic trials.

83 *Reaction times*

84 RTs were measured based on the moment when hand velocity crossed 1% of its peak (the hand was typically
85 held very steady before movement, allowing this low threshold). In the cue-initiated context, RTs for monkey
86 Ba were 269 ± 50 ms (mean \pm s.d.) and RTs for monkey Ax were 251 ± 37 ms (**Fig. 2a,b, red traces**). These RTs
87 are on the brisk side of the range reported in prior studies, consistent with the goal of the cue-initiated
88 context: to encourage monkeys to prepare during the delay and reach promptly after the go cue. Despite these
89 brisk RTs, monkeys were almost always successful in waiting for the go cue; reaches during the delay occurred
90 on <1% of trials for both monkeys. Similarly, RTs <150 ms were exceedingly rare, consistent with a go cue
91 whose timing could not be anticipated.

92 In the self-initiated context, reaches could be made immediately after target onset, but monkeys typically
93 waited at least 600 ms as this garnered a larger reward. Yet monkeys rarely waited until the time of maximum
94 reward (at 1200 ms). We define RT in the self-initiated context as the interval between target and movement
95 onset. Thus, RT has a unified definition across contexts: the time of movement initiation relative to the time

96 when movement was first permitted. Self-initiated RTs were $951 \text{ ms} \pm 132 \text{ ms}$ (mean \pm s.d.) and $1012 \text{ ms} \pm 106$
97 ms (**Fig. 2a,b, blue**). The considerable RT variability in the self-initiated context is unlikely to reflect uncertainty
98 regarding reward size, which was directly conveyed by target size. RT variability presumably reflects the natural
99 tension between a desire for large reward and a desire for immediate reward, with different factors
100 dominating on different trials.

101 The quasi-automatic context evoked particularly short RTs: $221 \pm 32 \text{ ms}$ and $208 \text{ ms} \pm 27 \text{ ms}$ (**Fig. 2a,b, yellow**).
102 Notably, RTs were shorter in the quasi-automatic context than in the cue-initiated context (on average by 48
103 and 43 ms) even though the quasi-automatic context provided no time to prepare in advance. Despite such
104 time pressure, monkeys almost never ‘jumped the gun’. For example, reaches were very rare during the 100
105 ms interval immediately after target onset (<1% of trials for both monkeys). This was unsurprising: target
106 location was unknown ahead of time, making it impossible to successfully exploit a strategy of anticipatory
107 reaches. Furthermore, any small movement before target onset resulted in an aborted trial, further
108 discouraging anticipatory movements or adjustments. The observed RTs are thus legitimate responses to
109 target onset. Given this, it is notable that RTs could be in the 170-200 ms range. Muscle activity (documented
110 further below) leads this moment by $\sim 80 \text{ ms}$. Thus, muscle activity could begin surprisingly quickly (as little as
111 90 ms) following target onset.

112 During training, monkeys showed short RTs immediately upon encountering moving targets in the quasi-
113 automatic context. This observation is consistent with the suggestion that moving targets evoke intercepting
114 movements with an almost innate short latency²⁶. For this reason, we refer to such reaches as ‘quasi-
115 automatic’. However, we stress that this term should not be taken to imply that quasi-automatic reaches are
116 necessarily a different class of movement at the level of motor cortex. Whether or not behavior across the
117 three contexts is sub-served by a similar set of preparatory and movement-related neural events is a
118 fundamental question of this study.

119 *Reach kinematics*

120 To aid comparison of neural activity across contexts, we wished to ensure that any observed differences were
121 not trivially due to differences in the physical reaches themselves. Care was thus taken to ensure that reach
122 trajectories were similar across contexts (**Fig. 2c,d**). For example, for the quasi-automatic context, we adjusted
123 the initial target position and its velocity so that reach extent was similar to that in the other two contexts. The
124 resulting match in kinematics was very good, with only a few slight differences: for many reach directions,
125 quasi-automatic reaches had a slightly greater extent, and a slightly higher corresponding peak velocity (on

126 average 7% and 9% higher relative to cue-initiated movements). It may seem surprising that quasi-automatic
127 reaches did not display dramatically higher velocities compared to the other two contexts. However, we find
128 that monkeys typically reach rapidly even when not required to, presumably out of a desire to obtain reward
129 quickly once the decision to move has been made. Movement duration was indeed similar (~200 ms) across
130 contexts (**Fig. 2e,f**).

131 We were interested in whether, in the quasi-automatic context, monkeys might begin reaches before they are
132 fully specified⁵. For example, monkeys might initially lift their hand off the screen in a generic fashion and only
133 later adjust their reach towards the target. This could lead to an RT that measures a non-specific response, and
134 would thus not be comparable to RTs measured in the other contexts. To avoid this concern, we measured RT
135 based on the magnitude of the velocity vector in the plane of the task, ignoring depth (distance off the screen).
136 Furthermore, inspection of behavior revealed that monkeys did not adopt the above strategy. Trajectories in
137 depth were very similar across contexts. For example, **Fig. 2g,h** shows example average trajectories in the
138 horizontal / depth plane, with an expanded scale for depth.

139 Thus, our results confirm that target-directed reaches can be made without a delay period or a long RT^{24,26}. We
140 wondered whether some price in accuracy might be paid for such short-latency movements (as might be
141 expected if preparation were hurried). There was indeed a small increase in reach-direction variability in the
142 quasi-automatic context, relative to the cue-initiated context. Assessed via the circular standard deviation²⁸ of
143 the initial reach direction, variability was 17% and 11% higher (monkey Ba and Ax; $p < 0.001$ for both).
144 Consistent with the possibility of hurried preparation, quasi-automatic trials with shorter-than-median RTs
145 showed slightly greater variability versus trials with longer-than-median RTs (4.9% and 3.0%; $p < 0.05$ for
146 monkey Ba, NS for monkey Ax). In absolute terms these effects were small – the standard deviation of reach
147 direction was never higher than 4.3 degrees. Still, the quasi-automatic context had slightly more reaches that
148 missed the target: an increase from 1.5% (cue-initiated) to 3.3% (quasi-automatic) for monkey Ba, and from
149 0.9% to 2.4% for monkey Ax. In summary, very low-latency reaches tended to be slightly less accurate, but
150 were otherwise very similar to reaches in the other two contexts.

151 *Muscle activity*

152 As with kinematics, muscle activity (EMG) should ideally be similar across contexts to aid comparisons of neural
153 activity. Individual-trial EMG measurements are ‘noisy’ (due to the spiking variability of motoneurons) but the
154 trial-averaged EMG pattern can be accurately estimated by filtering the rectified EMG and averaging across
155 many trials. This process is very similar to that used for the neural data, and employed the same 20 ms

156 standard-deviation Gaussian filter. Trial-averaged EMG patterns were very similar across contexts
157 (**Supplementary Fig. 1**). During movement, the median correlation between the self-initiated and cue-initiated
158 contexts was 0.98 and 0.98 (monkey Ba and Ax respectively; activity compared in a 300 ms window starting 50
159 ms before movement onset; median taken across muscles). The median correlation between quasi-automatic
160 and cue-initiated contexts was 0.97 and 0.96. We also assessed the strength of each muscle's selectivity as the
161 standard deviation, across directions, of temporally-averaged EMG. Such selectivity was similar across
162 contexts, but was slightly higher for the quasi-automatic context: by 7.2% and 8.8% (monkey Ba and Ax)
163 relative to the cue-initiated context. This is consistent with the slightly higher extent/velocity for many
164 directions in the quasi-automatic context. Muscle-activity selectivity was very slightly lower in magnitude for
165 the self-initiated context (by 1.9% and 3.0%) relative to the cue-initiated context, consistent with slightly lower
166 velocities for some reach directions (visible upon close inspection of **Fig. 2e,f**).

167 Many muscles displayed activity while holding the central touch point, before target or movement onset. Such
168 activity was necessary to support the outstretched arm. This affords an opportunity to ask whether putatively
169 preparatory neural activity relates to overt changes in muscle tone. As in many prior studies, we found little
170 change in EMG during intervals when preparation would be expected to occur. We examined EMG during the
171 delay period of the cue-initiated context, and during the interval from target onset until 250 ms before
172 movement onset during the self-initiated context. EMG changed little or not at all following target onset, and
173 showed little or no selectivity for reach direction (**Supplementary Fig. 1**). Specifically, for the cue-initiated
174 context, the across-condition variance of EMG (*i.e.*, selectivity for reach direction) during the delay was 0.76%
175 and 2.8% of the variance during movement. For the self-initiated context, the values were 0.51% and 1.2%.
176 Thus, any preparatory process that might be present in this time-range has minimal direct impact on muscle
177 activity.

178 *Single-neuron responses during the cue-initiated context*

179 The spikes of well-isolated neurons were recorded from the arm region of motor cortex, including M1 and the
180 immediately adjacent region of PMd (129 and 172 neurons for monkey Ba and Ax). Spike-trains were filtered
181 and trial-averaged to yield an estimate of firing rate as a function of time (**Figure 3** shows examples responses
182 for three neurons). Filtering employed a narrow Gaussian kernel (20 ms SD) to ensure that multi-phasic and
183 temporally structured response features²⁹ were not lost. Filtering was performed after concatenating two
184 epochs that were time-locked to target and movement onset, respectively. Filtering the concatenated spike-
185 trains yielded a continuous estimate of rate with no discontinuity. However, it should be kept in mind that the
186 temporal interval between target and movement onset was variable across trials. To ensure the continuous

187 estimate of rate was representative, we chose analysis epochs based on typical behavioral performance. For
188 example, in the quasi-automatic context, epochs were selected such that target onset and movement onset
189 were separated by an interval equal to the mean RT.

190 Many neurons showed activity that varied with reach direction during the delay period of the cue-initiated
191 context (*e.g.*, **Figure 3**, *red traces*). For this context, we defined a 450 ms ‘delay epoch’, beginning 50 ms after
192 target onset. Variation of delay-epoch firing rate with target direction was significant (ANOVA, $p < 0.05$) for the
193 majority of neurons (74/129 and 88/172 for monkey Ba and Ax). We also defined a 300 ms ‘movement epoch’,
194 starting 50 ms before movement onset (just after EMG began to change) and ending just after the hand landed
195 on the target. Variation of movement-epoch firing rate with target direction was also significant for the
196 majority of neurons (116/129 and 144/172).

197 As in prior studies, delay-period activity suggests a preparatory process. A natural question – relevant to
198 interpretation of activity in the other two contexts – is when putatively preparation-related activity transitions
199 to movement-related activity. Presumably this must happen following the go cue but before movement. Yet
200 the moment of this hypothesized transition is difficult or impossible to determine via inspection of single-
201 neuron responses. For example, the neuron illustrated in **Figure 3a** shows multiple response phases following
202 the go cue, including a peak ~75 ms before movement onset (higher for rightwards reaches; *lighter red traces*)
203 and a subsequent peak during movement (higher for leftwards reaches; *darker red traces*). Should one
204 consider the first peak to be a final strengthening of a preparation-related response, or the beginning of a
205 movement-related response? Is such a distinction even meaningful? These questions are challenging because
206 there is no easily identifiable moment when putatively preparatory activity ends and movement-related
207 activity begins. Similar ambiguity was present even for neurons with simpler response patterns. The neuron
208 illustrated in **Figure 3c** (*red traces*) exhibits activity just before movement that is (approximately) a magnified
209 version of delay-period activity. Again, it is unclear whether such activity reflects the culmination of
210 preparation-related activity, or a movement-related burst.

211 These uncertainties highlight a known limitation of the instructed-delay paradigm: although activity between
212 target onset and the go cue is suggestive of a preparatory process, it is more challenging to interpret neural
213 events between the go cue and movement onset⁶. One might wish to define events before a specified time as
214 putatively preparatory, and events after that time as movement-related. Yet when that time should be (or
215 even whether the transition happens at a discrete time) is not clear from inspection of single-neuron
216 responses. This creates a problem: if a delay period is necessary to identify activity as putatively preparatory,
217 how can we test whether putatively preparatory events occur in the absence of a delay? This challenge will
218 become particularly relevant when examining responses during the quasi-automatic context.

219 *Single-neuron responses during the self-initiated versus cue-initiated context*

220 In the self-initiated context, movement initiation is never demanded by an unpredictable go cue. Rather, the
221 monkey chooses when to reach and can potentially anticipate that choice in advance. Given this, how does
222 activity in the self-initiated context relate to that in the cue-initiated context? There exist at least three
223 possibilities. First, if delay-period activity is primarily suppressive, then a similar pattern of activity should be
224 present in the self-initiated context. That pattern should be strongest shortly after target onset (when reaches
225 should be most strongly suppressed because they yield little reward) and should wane as the time of
226 movement approaches. Second, if delay-period activity is primarily preparatory, then during the self-initiated
227 context, a similar pattern should grow with time as movement nears. Any strengthening could be gradual –
228 starting hundreds of milliseconds before movement onset – or rapid – occurring just before movement-related
229 activity. Third, activity in the self-initiated context could look quite unlike activity during the cue-initiated
230 context. Self-initiated movements likely involve a larger role of anterior areas, including the supplementary
231 motor area³⁰. If so, the role of motor cortex might be reduced or altered. This could impact pre-movement
232 activity, movement-related activity, or both.

233 Single-neuron responses followed the second prediction: the patterns of pre-movement activity in the self-
234 initiated context grew with time and came to resemble the patterns of delay-period activity in the cue-initiated
235 context. For example, in **Figure 3a,c**, the ordering of traces ~250 ms before movement onset is similar for the
236 cue-initiated (*red traces*) and self-initiated (*blue traces*) contexts. The pattern of pre-movement activity in the
237 self-initiated context became stronger rather than weaker with time with time. Across all neurons, the median
238 correlation between self-initiated and cue-initiated activity patterns was low during the first 150 ms after
239 target onset (median $r = 0.39$ and $r = 0.16$ for the two monkeys), reflecting the fact that early activity was
240 typically weak in the self-initiated context. This correlation became stronger as movement onset approached
241 (median $r = 0.86$ and $r = 0.74$, using a 150 ms window ending 50 ms before movement onset). During
242 movement, the correlation was quite high (median $r = 0.92$ and $r = 0.87$). Such strengthening of activity agrees
243 with related results in rodent³¹, and is inconsistent with a suppressive process that relaxes to eventually allow
244 movement. These observations will be further quantified by population-level analyses below.

245 *Single-neuron responses during the quasi-automatic context*

246 Before considering putatively preparatory events, we note that movement-related responses in the quasi-
247 automatic context closely resembled those in the other two contexts. This can be appreciated by visually
248 comparing, across rows in **Figure 3**, activity from just before movement onset until the end of the trial. These

249 example neurons were representative: the median correlation between movement-epoch activity patterns
250 during the quasi-automatic and cue-initiated contexts was 0.85 and 0.85 (monkey Ba and Ax). The primary
251 difference in movement-epoch response patterns was a tendency for some features to be slightly magnified in
252 the quasi-automatic context (*e.g.*, the central peak in **Fig. 3b**). This observation is consistent with the slight
253 increase in reach speed, and with the slight increase in the magnitude of muscle activity. The similarity in
254 movement-related activity across contexts was not a given. Because there may be subcortical contributions to
255 very short-latency movements^{25,26}, movement-related cortical activity could have been different or reduced
256 during the quasi-automatic context relative to the other two contexts.

257 The similarity of movement-related responses makes it sensible to ask whether those responses are preceded
258 by similar patterns of preparatory activity. Is the very first portion of the response in the quasi-automatic
259 context potentially preparatory? Or does that initial response simply constitute the beginning of movement-
260 related activity? We saw no way of addressing this question using individual-neuron analyses. Consider the
261 neurons illustrated **Figure 3a,c**. For both, the very first pattern of activity to emerge in the quasi-automatic
262 context resembled that seen shortly before movement onset in the cue-initiated context. Interpretation thus
263 hinges on whether activity at that time in the cue-initiated context reflects the culmination of preparation or
264 the beginning of movement-related activity. As discussed above, this is difficult or impossible to infer from
265 individual-neuron responses. We therefore turn to analyses that leverage population-level properties.

266 *Segregating preparatory and movement responses at the population level*

267 We employed a recently developed analytical method to segregate the population response into putatively
268 preparatory and movement-related response patterns. This method leverages the observation that the
269 correlation structure between neurons changes dramatically between delay-period and movement-period
270 epochs²⁷. Specifically, the 'neural dimensions' that best capture delay-epoch activity do not capture
271 movement-epoch activity, and vice versa. That observation was unexpected; it occurs because neurons with
272 related response properties during preparation become unrelated during movement, something predicted by
273 no existing model²⁷. Yet the finding has considerable utility from the standpoint of the present study. Although
274 putatively preparatory and movement-related processes are not separable at the single-neuron level, they are
275 potentially separable at the population level. Our strategy is to use the cue-initiated context to identify a set of
276 neural dimensions that captures putatively preparatory activity and an orthogonal set of dimensions that
277 captures movement-related activity. These dimensions can then be used to examine the population response
278 in the other contexts. If delay-period activity indeed reflects preparation, then dimensions that capture delay-
279 period activity in the cue-initiated context may similarly capture preparatory processing in the other two

280 contexts. Alternatively, if delay-period activity reflects some non-preparatory process (*e.g.*, suppression) or a
281 process specific to the presence of an experimenter-imposed delay, then the dimensions that captured delay-
282 period activity will either not capture activity during the other two contexts, or will capture activity with
283 structure that is very different from that observed during the delay.

284 Based on neural responses during the cue-initiated context only, we isolated twelve preparatory dimensions,
285 collectively the ‘preparatory subspace’. The preparatory subspace captured 80% of firing rate variance (*i.e.*,
286 firing-rate structure across all neurons) during the delay epoch, but only 3% of variance during the movement
287 epoch. We isolated twelve movement dimensions (collectively the ‘movement subspace’) which together
288 captured 85% of variance during the movement epoch, but only 3% of variance during the delay epoch. The
289 above percentages are for monkey Ba and were similar for monkey Ax (72% versus 4% and 83% versus 3%).
290 The ability to achieve this near-perfect segregation is not a general feature of neural data. It is a consequence
291 of the dramatic change in covariance between the delay and movement epochs.

292 We projected the population response, for the cue-initiated context, onto the preparatory and movement
293 dimensions. This revealed putatively preparatory and movement-related activity patterns (**Figure 4, middle**).
294 Each projection is a weighted sum of single-neuron responses (weights are the elements of the vector defining
295 the dimension). Yet unlike a generic linear readout, weights were optimized to capture response structure.
296 Projections are thus not only readouts but also building blocks of single-neuron responses, much as for
297 principal component analysis. Indeed, the dimensions we found span a space similar to the top principal
298 components. Using these building blocks, it becomes possible to estimate putatively preparatory and
299 movement-related contributions to each neuron’s response. For example, the response of neuron 88 (**Fig. 4,**
300 *rightmost column*) is accurately approximated as the sum of a preparatory-subspace pattern (a weighted sum
301 of preparatory projections, *orange*) and a movement-subspace pattern (a weighted sum of movement
302 projections, *purple*). The reconstruction is not perfect – the pattern in **Figure 4 (right)** differs slightly from the
303 true response in **Figure 3a** – but is quite good ($R^2 = 0.93$). High reconstruction accuracy reflects the high
304 proportion of firing-rate variance captured. Do the dimensions found using the cue-initiated context – in
305 particular the preparatory dimensions – similarly capture variance during the other two contexts?

306 *Reconstruction of neural responses across contexts*

307 The dimensions found using the cue-initiated context continued to capture a high percentage of response
308 structure in the self-initiated and quasi-automatic contexts. In the self-initiated context, the total variance
309 captured by both preparatory and movement subspaces was 89% and 84% (monkey Ba and Ax) of that in the

310 cue-initiated context. In the quasi-automatic context, the variance captured was 87% and 84% of that in the
311 cue-initiated context. The high percentage of captured variance is reflected in accurate reconstruction of
312 neural responses in the self-initiated and quasi-automated contexts. For example, the response of neuron 88
313 (**Fig. 3a**) was accurately reconstructed not only in the cue-initiated context (**Fig. 4, right**) but also in the self-
314 initiated (**Fig. 5a**) and quasi-automatic (**Fig. 5b**) contexts.

315 These successful reconstructions involved contributions from both preparatory and movement subspaces. For
316 example, the reconstruction for neuron 88 included a robust preparatory-subspace pattern during the self-
317 initiated context (**Fig. 5a, orange**) and a short-lived but strong preparatory-subspace pattern during the quasi-
318 automatic context (**Fig. 5b, orange**). Within this preparatory-subspace patterns, the ordering of conditions was
319 similar across contexts: for neuron 88 the pattern was most positive for rightwards reaches (*light traces*) and
320 most negative for leftwards reaches (*dark traces*). To facilitate quantitative comparison, for each neuron we
321 measured the preparatory pattern 100 ms before movement onset, yielding a vector with one value per
322 direction. This vector captures the directionality of the preparatory pattern. To assess whether directionality
323 was similar across contexts, for each neuron we regressed the preparatory patterns for the self-initiated and
324 quasi-automatic contexts against that observed for the cue-initiated context. If the two patterns are the same,
325 then regressing one versus the other will yield a slope of one. In contrast, an average slope of zero would
326 indicate no consistent relationship between the preparatory patterns across contexts.

327 When comparing self-initiated and cue-initiated contexts, slopes were strongly positive (**Fig. 5c, d black bars in**
328 *left subpanel*). When comparing quasi-automatic and cue-initiated contexts, slopes were again strongly
329 positive (**Fig. 5c,d dark bars in right subpanel**). For monkey Ba, the slope was slightly greater than unity,
330 consistent with preparatory-subspace activity being slightly stronger in the quasi-automatic context. A
331 potential concern is that this strong similarity might not be specific to the preparatory pattern. For example, if
332 neurons have similar directionality at all times, then similarity would be high even when comparing a
333 preparatory pattern in one context and a movement pattern in another context. This was not the case: there
334 was no consistent relationship between the preparatory-subspace contribution in the cue-initiated context and
335 the movement-subspace contribution (assessed 150 ms after movement onset) in the other two contexts (**Fig.**
336 **5c,d, gray bars**). In summary, the dimensions that captured delay-period activity also made strong
337 contributions to firing rates during all three contexts, and had a similar pattern across all three contexts. The
338 nature of that pattern will be further investigated below.

339 *Temporal evolution of responses in preparatory and movement dimensions*

340 If preparatory-subspace activity is truly preparatory, then it should exhibit a time-course consistent with that
341 role. Is this true across all three contexts? For each time, we measured the across-condition variance (the
342 strength of selectivity) of preparatory-subspace activity. This variance reflects the size of the envelope
343 describing the *orange* patterns in **Figure 4**. We refer to this measure as the ‘preparatory-subspace occupancy’.
344 An important question is whether putatively preparatory events consistently precede movement-related
345 events. We therefore similarly computed the movement-subspace occupancy. Movement-subspace occupancy
346 (**Fig. 6, purple traces**) had a similar time-course across contexts: it was negligible until ~110 ms before
347 movement onset and reached a peak just after movement onset (the peak occurred between 30 and 80 ms
348 after movement onset across both monkeys and all contexts).

349 Preparatory-subspace occupancy (**Fig. 6, orange traces**) followed a very different time course for each context.
350 In the cue-initiated context, there was an initial rapid rise that was sustained (at a lower level) throughout the
351 delay period. Preparatory-subspace activity then declined rapidly just before movement onset, reaching
352 baseline levels around the time the reach began. It is worth stressing that, by construction, preparatory
353 subspace occupancy is high during the delay period of the cue-initiated context. However, no further structure
354 is imposed; occupancy could have declined following the go cue, could have stayed the same, or could have
355 become stronger.

356 In the self-initiated context, the rise in preparatory subspace occupancy following target onset was weaker
357 (monkey Ba) or much weaker (monkey Ax) than for the cue-initiated context. Preparatory-subspace occupancy
358 remained weak from 200–400 ms after target onset. Occupancy then grew as movement approached, and
359 reached a peak before movement onset (at 120 ms and 160 ms for monkey Ba and Ax respectively). Occupancy
360 at that time was then similar to occupancy in the cue-initiated context at the same time (slightly greater for
361 monkey Ba and slightly smaller for monkey Ax). These observations concur with the hypothesis that, in the self-
362 initiated context, monkeys do not consistently prepare their reach immediately following target onset, but
363 instead wait until nearer the time they choose to move. Whether the ramp of increasing occupancy reflects
364 ramping on individual trials cannot be inferred from the present data. It is equally plausible that preparation
365 has a sudden onset that is variable relative to movement onset, resulting in a ramp in the averaged data.

366 In the quasi-automatic context, preparatory-subspace occupancy rose rapidly following target onset and was
367 short-lived: occupancy peaked 70 ms and 80 ms (monkey Ba and Ax respectively) before movement onset and
368 then declined. The magnitude of this peak in preparatory-subspace occupancy was similar to, but slightly
369 higher than, the peak observed in the other two contexts just before movement onset. These observations are

370 consistent with the hypothesis that a preparatory stage is present even for low-latency intercepting reaches.
371 However, preparation appears to be very rapid: preparatory subspace occupancy precedes movement-
372 subspace occupancy by only a few tens of milliseconds: 33 ms for monkey Ba and 42 ms for monkey Ax.
373 (Latency was measured as the time occupancy crossed a 10% threshold, and a shorter filter was used to
374 minimize the influence of filtering on latency, see *methods*).

375 Comparing between monkeys, there was one obvious difference in the time-course of preparatory subspace
376 occupancy. For monkey Ax, the initial, post-target peak during the cue-initiated context was larger than at any
377 other time, for any context. (Occupancy is plotted in normalized form, and thus the large initial peak for the
378 cue-initiated context necessarily means that all other peaks are plotted with values below unity.) In contrast,
379 for monkey Ba, the post-target and pre-movement peaks in the cue-initiated context were closer in magnitude.
380 When comparing the peak just before movement across contexts, the two monkeys were more similar. For
381 both monkeys, the pre-movement peak in preparatory subspace occupancy was a similar size across contexts,
382 and was slightly larger for the quasi-initiated context. We now ask whether the events within that subspace are
383 conserved across contexts.

384 *Preparatory and movement events in state space*

385 **Figure 7** plots ‘snapshots’ of projections onto a two-dimensional preparatory subspace (**a**) and a two-
386 dimensional movement subspace (**b**). (**Supplementary Fig. 2** shows similar snapshots for monkey Ax). Within
387 each snapshot, each trace plots the evolution of the neural state for one reach direction over a 150 ms period,
388 beginning at the indicated time. For this task, preparatory-subspace activity was quite low-dimensional: the
389 first two dimensions captured much more variance than subsequent dimensions. *E.g.*, for monkey Ba, the third
390 preparatory dimension captured only 14% as much variance as the first. Thus, the preparatory subspace
391 projections in **Figure 7a** give a reasonably complete view of the preparatory state, and subsequent
392 quantification is based on those dimensions. Movement subspace activity was considerably higher
393 dimensional: there were many dimensions with structure that was clearly not noise. The projections in **Figure**
394 **7b** thus yield only a partial view. Subsequent quantification therefore employed all twelve dimensions.

395 There was a remarkable consistency, across contexts, in the patterns of the neural trajectories. The most
396 notable differences across contexts regarded not the patterns *per se*, but the time-course of preparatory-
397 subspace events. In the cue-initiated context, target onset prompted preparatory-subspace activity to become
398 strongly selective for reach direction (*red traces* in **Fig 7a** separate upon target onset). The resulting pattern
399 was sustained throughout the delay period, then collapsed near the time of movement onset. In the self-

400 initiated context, target onset prompted a weaker separation of preparatory-subspace neural states (*blue*
401 *traces* in **Fig 7a** separate less than *red traces* upon target onset). As time neared the onset of the self-initiated
402 reach, the preparatory-subspace pattern became more robust until it was approximately as strong as that in
403 the cue-initiated context. For the quasi-automatic context, the preparatory-subspace pattern was very short-
404 lived: it grew rapidly following target onset then immediately collapsed prior to movement onset. However,
405 while present, the preparatory-subspace pattern during the quasi-automatic context closely resembled that in
406 the other two contexts (compare across contexts in the third-to-last column of **Fig 7a**). For example, the
407 dependence of the neural state on reach direction was similar across contexts (*lighter / darker* traces indicate
408 rightwards / leftwards movements).

409 Movement-subspace patterns were very similar across contexts, in both their pattern and their timing. Target
410 onset produced essentially no separation of movement-subspace states for the cue-initiated or self-initiated
411 contexts. This is consistent with finding that target onset produced little or no change in EMG activity.

412 Movement-subspace states started to differentiate between reach directions ~110 ms before movement
413 onset. This occurred at a similar time and in a similar way across contexts. During movement, the neural state
414 evolved according to rotational dynamics, as previously reported¹⁹ and in a manner predicted by neural
415 network models³². As for such models, rotational dynamics were present in a subset of dimensions; the
416 dimensions shown here were chosen specifically to capture such dynamics for the cue-initiated context, and
417 naturally captures similar dynamics for the other two contexts.

418 Comparing between subspaces reinforces and extends the results described in **Figure 6**. Across all contexts,
419 preparatory-subspace activity always emerged before movement-subspace activity began. Preparatory-
420 subspace activity and movement-subspace activity then showed considerable overlap: the former declined as
421 the latter emerged. Just before and during that period of overlap, the pattern of preparatory subspace activity
422 was similar across all three contexts. We explore this finding quantitatively below.

423 *Quantification of similarity across contexts*

424 A central question of this study is whether similar movement-subspace events are preceded by similar
425 preparatory subspace events. To quantify movement-subspace similarity, we measured the correlation (per
426 dimension) between time-evolving patterns measured during a 150 ms window starting at movement onset.
427 These correlations were high for all comparisons. Comparing the self-initiated and cue-initiated contexts, for
428 the two dimensions shown in **Figure 7**, the correlation was 0.96 and 0.97 for monkey Ba and 0.97 and 0.95 for
429 monkey Ax. Across all twelve movement dimensions, correlations ranged from 0.96 to 0.99 for monkey Ba and

430 from 0.94 to 0.99 for monkey Ax. Correlations were similarly high when comparing the quasi-automatic and
431 cue-initiated contexts. For the two dimensions shown in **Figure 7**, the correlation was 0.95 and 0.93 for
432 monkey Ba and 0.91 and 0.97 for monkey Ax. Across all movement dimensions, correlations ranged from 0.88
433 to 0.99 for monkey Ba and from 0.90 to 0.98 for monkey Ax. These results agree with the similarity of
434 movement-epoch activity across contexts observed at the individual-neuron level. How similar are the
435 preceding patterns of preparatory subspace activity?

436 To address this question, we focused on the preparatory subspace state at a specific moment: 70 ms before
437 movement onset. At this moment, movement-related activity is just starting to emerge. We have previously
438 hypothesized that the preparatory state seeds movement-related dynamics^{10,18,19,21,27,32}. Under this hypothesis,
439 the preparatory state when movement-related activity begins is critical. This hypothesis thus predicts that the
440 preparatory subspace state at that time should be very similar across contexts, given that subsequent patterns
441 of movement-related activity are similar. We refer to that potentially critical preparatory state as the ‘final’
442 preparatory state; after that moment, movement-subspace activity becomes strong and preparatory-subspace
443 activity declines to near-baseline levels.

444 The final preparatory state was similar across contexts. For each reach direction, the neural states across
445 contexts formed a cluster (these are grouped via covariance ellipses in **Fig. 8a,b**). Clusters were quite tight for
446 monkey Ba and somewhat less so for monkey Ax. To quantify the tightness of clustering – *i.e.*, the similarity of
447 states across contexts – we computed the correlation, between contexts, of the set of preparatory-subspace
448 states (one state per reach direction, with two dimensions describing each state). The correlation between cue-
449 initiated and self-initiated contexts was 0.99 (95% c.i.: 0.98–0.99, $p < 0.001$) and 0.95 (95% c.i.: 0.86–0.98, $p <$
450 0.001) (monkey Ba and Ax). The correlation between cue-initiated and quasi-automatic contexts was 0.99 (95%
451 c.i.: 0.97–1.0, $p < 0.001$) and 0.91 (95% c.i.: 0.76–0.97, $p < 0.001$). The correlation between self-initiated and
452 quasi-automatic contexts was 0.99 (95% c.i.: 0.97–1.0, $p < 0.001$) and 0.92 (95% c.i.: 0.79–0.97, $p < 0.001$).

453 To assess the time-course of similarity, at each time we computed the covariance, for a pair of contexts,
454 between the neural states in the preparatory space. Covariance reflects both similarity and strength, and is
455 thus expected to peak at a time when preparatory patterns are both similar and robust. When comparing the
456 self-initiated and cue-initiated contexts (**Fig. 8c,d; blue**) covariance rose as movement approached, peaking 120
457 ms and 130 ms (monkey Ba and Ax) before movement onset. This is consistent with what can be observed in
458 earlier figures: in the preparatory subspace, the pattern of states in the self-initiated context generally
459 resembles that in the cue-initiated context, but is weaker until the time of movement onset nears.

460 When comparing the quasi-automatic and cue-initiated contexts (**Fig. 8c,d; blue**) covariance rose rapidly,
461 peaking 80 ms and 90 ms (monkey Ba and Ax) before movement onset. These peaks occur just after activity in

462 the movement-subspace first begins to change, which occurred 90 ms (monkey Ba) and 94 ms (monkey Ax)
463 prior to movement onset in the quasi-automatic context. The narrowness of the peak underscores that the
464 similarity in preparatory subspaces states was short-lived; it was high for only a few tens of milliseconds,
465 around the time that movement-subspace activity was beginning to develop. Thus, while the pattern of
466 preparatory subspace activity in the quasi-automatic context comes to closely match that in the cue-initiated
467 context, this similarity occurs late (just as movement-subspace activity is developing) and is not sustained. This
468 is consistent with a preparatory process that is observed across all contexts, but that unfolds very rapidly in the
469 quasi-automatic context.

470 *Relative timing of movement-related events*

471 Our subspace-based analysis method isolates a movement subspace that is, by construction, occupied during
472 movement for the cue-initiated context. However, our method imposes no additional constraints on the timing
473 of movement-subspace events: they could begin well before movement onset, at the time of movement onset,
474 or after movement onset. We were particularly interested in the relationship between movement-subspace
475 occupancy and the onset of muscle activity. Does movement-subspace occupancy occur with timing
476 appropriate given a role in producing descending commands that cause muscle activity? For both movement-
477 subspace occupancy and EMG, we assessed latency by measuring the moment when activity surpassed 10% of
478 its peak. To minimize the impact of filtering on latency, these analyses employed a 10 ms Gaussian filter (rather
479 than 20 ms for all other analyses) for both neural and EMG data.

480 Across monkeys and contexts, the movement subspace always became occupied just before the onset of
481 changes in EMG, with an average latency of 21 ms. For comparison, the conduction delay from cortex to
482 muscles, assessed via spike-triggered averages, can be as little as 6 ms from the time of a spike to the peak of
483 the EMG response³³. This delay would be slightly reduced (to ~4 ms for the lowest-latency neurons) when
484 considering the beginning rather than the peak EMG response. Thus, activity in the movement subspace rises
485 early enough to potentially account for the onset of muscle activity. This was consistently true across contexts,
486 although with slight variability. The latency between the onset of movement-subspace activity and muscle
487 activity was, for monkey Ba and Ax, 27 and 20 ms (cue-initiated), 33 and 22 ms (self-initiated) and 19 and 6 ms
488 (quasi-automatic). These exact latencies should be interpreted with some caution: latencies are notoriously
489 difficult to assess because high thresholds overestimate latency while low thresholds are sensitive to noise.
490 Still, our best estimates indicate that, if cortico-motoneurons draw from movement-subspace activity, the
491 onset of such activity occurs early enough to plausibly account for the onset of muscle activity.

492 **Discussion**

493 *Is delay-period activity a reflection of motor preparation?*

494 Early studies generally viewed delay-period activity as preparatory, but noted that directional selectivity often
495 reverses between delay and movement epochs, suggesting a suppressive role³⁴. Subsequent experiments
496 revealed that delay-period and movement-related activity patterns typically differ^{6,18,27,35}, ruling out the
497 hypothesis that preparation involves a subthreshold version of movement-related activity. A different
498 preparatory role for delay-period activity was suggested: serving as the initial state of a neural dynamical
499 system whose evolution produces movement^{10,18,21}. In support, one can directly observe that the phase and
500 amplitude of movement-related dynamics flow from the state achieved during the delay¹⁹. Under this
501 hypothesis, preparatory activity is a necessary precursor to movement-related activity. Yet a recent study
502 yielded mixed evidence regarding the presence of a consistent preparatory state with and without a delay²².
503 That mixed evidence highlighted the longstanding uncertainty regarding whether delay-period activity
504 represents a true preparatory process, a facilitatory but non-obligatory process, or a suppressive process
505 specific to an artificial imposed delay²³.

506 Our results reveal that the neural process present during a delay-period is not specific to that situation, but is
507 consistently observed in other contexts. This putatively preparatory process has the following properties. First,
508 activity occupies a neural subspace orthogonal to that occupied during movement. Second, such activity
509 consistently occurs before activity in the movement-related subspace. Third, regardless of the presence of an
510 imposed delay period, the neural state in the preparatory subspace achieves a similar movement-specific state
511 before movement onset. That similarity is maximal at the critical moment when movement-subspace activity is
512 just beginning. These results essentially rule out the hypothesis that delay-period activity is primarily
513 suppressive. The suppressive hypothesis cannot explain the presence of preparatory-subspace activity in the
514 quasi-automatic context, or the rising profile of preparatory-subspace occupancy in the self-initiated context.

515 That said, our results do not prove that preparatory-subspace activity is preparatory – they only show that it
516 follows the major predictions of that hypothesis. Proving that hypothesis would require specifically perturbing
517 activity in that subspace and observing the impact on behavior – something not currently feasible. That said, it
518 is known that a non-specific disruption of premotor cortex activity, at the end of the delay period, impacts RT
519 in a manner consistent with disruption of a preparatory process¹⁷. Given that evidence and the present
520 observations, we tentatively interpret preparatory subspace activity as preparatory and ask what conclusions
521 might follow.

522 *Does preparation necessarily consume time?*

523 Early behavioral investigations leveraged and supported the assumption that it takes considerable time to
524 ‘plan’ or ‘specify’ the desired movement^{2,5}. Influenced by this framework, subsequent physiology and modeling
525 studies proposed that preparation involves the time-evolving strengthening and shaping of neural activity
526 directly specifying movement parameters^{9,36,37}. We have argued that movement is specified more implicitly, by
527 achieving a preparatory state that seeds movement dynamics^{18,19,21,27,32}, but found evidence that it takes time
528 (100-200 ms) to consistently prepare¹⁰. Thus, a time-consuming preparatory process has often been considered
529 to be a major determinant of RT. This traditional framework has enjoyed explanatory power, and has
530 motivated successful comparisons of trial-to-trial RT variability with trial-to-trial variability of putatively
531 preparatory activity^{4,8-10,16}.

532 Yet there have been compelling recent arguments against the necessity of a time-consuming preparatory
533 process^{23,26}. The present study supports those arguments. RTs in the quasi-automatic context were on average
534 221 and 208 ms, and were frequently 170-200 ms on individual trials. These short RTs occur despite the
535 inability to prepare in advance, and cannot be explained by anticipation: monkeys had no fore-knowledge of
536 target direction and did not attempt to ‘jump the gun’. Given a delay of at least 50 ms for visual information to
537 reach motor cortex, and an afferent delay of at least 75 ms (including the sizeable lag between muscle activity
538 and movement onset), there cannot exist an obligatory preparatory process that necessarily takes 100-200 ms
539 to complete. That conclusion is further supported by the neural data. In the quasi-automatic context,
540 preparatory subspace activity lead movement-subspace activity by only ~40 ms, and the preparatory-subspace
541 state came to match that in the cue-initiated context in ~70 ms. These findings rule out the idea of a slow,
542 cognitive planning process that must complete before movement. These findings support our prior proposal
543 that preparatory activity is necessary to seed movement-generating dynamics. However, the development of
544 such activity can occur much faster than previously supposed.

545 Nevertheless, the influential idea that motor preparation tends to consumes time may have some merits. It
546 may be that preparation often, or even typically, spans time. In the self-initiated context, putatively
547 preparatory activity begins hundreds of milliseconds before movement-related activity. This raises a central
548 question: if preparation can be fast, why is it ever extended? Why do monkeys not simply wait to prepare until
549 just before movement onset? We can only speculate, but the ability to rapidly and consistently achieve the
550 correct preparatory state may not be something that can be counted on in all real-world situations, especially
551 for less familiar or more challenging movements. The motor system may thus have developed the conservative
552 strategy of preparing in advance when possible, allowing time for errors to be corrected before movement
553 generation begins^{10,17}. We did indeed find that accuracy was slightly reduced in the quasi-automatic context, as

554 would be expected if movement is sometimes triggered before preparation has fully converged on the
555 appropriate preparatory state.

556 *Putatively preparatory and movement-related processes overlap*

557 Preparatory subspace activity overlapped with movement subspace activity by slightly more than 100 ms. This
558 overlap is consistent with (and indeed required by) the hypothesis that preparatory-subspace activity seeds
559 movement-subspace dynamics. Aspects of the overlap explain a seeming discrepancy between our results and
560 the recent finding of Ames et al.²² that the neural state for no-delay trials does not pass through the state
561 achieved during the delay of long-delay trials. This might seem at odds with our finding that a consistent
562 preparatory-subspace state is achieved across all contexts. In fact, our results are fully compatible.
563 Preparatory-subspace activity in the quasi-automatic context achieves its maximal match with that in the cue-
564 initiated context slightly after movement-subspace activity emerged. Thus, at the moment the match is
565 achieved, the full neural state contains both preparation-related and movement-related contributions. The
566 neural state at this moment will therefore not match that during the delay-period of the cue-initiated context,
567 when there is no contribution from movement-related dimensions.

568 With this conflict resolved, the present results support and extend two key conclusions of Ames et al. First, the
569 initial response to target onset can be similar with and without an imposed delay (compare the initial
570 development of preparatory subspace activity between cue-initiated and quasi-automatic contexts in **Fig. 7**).
571 However, this early response is unlikely to be an inevitable visual response: it is considerable weaker in the
572 self-initiated context. Second, when under time pressure, the neural state does not momentarily pause at a
573 stable state prior to the onset of movement-related activity (also see⁸). Indeed, in the quasi-automatic context,
574 events are so compressed that preparatory-subspace activity is still developing as movement-subspace activity
575 is beginning.

576 *Preparing versus deciding*

577 Although our data argue against the conception of preparation as an intrinsically slow, cognitive process, they
578 are quite consistent with the idea that slow cognitive processes influence preparatory activity. It is well
579 established that preparatory activity in a variety of brain regions can reflect decisions regarding when or where
580 to move^{31,38-42}. Such decisions can sometimes unfold slowly or vacillate with time. In motor cortex, preparatory

581 subspace activity may therefore sometimes evolve slowly simply because the overall movement goal is being
582 decided slowly.

583 In the present study, the rising strength of pre-movement activity in the self-initiated context is somewhat
584 reminiscent of the rise of choice-related activity in decision tasks. However, in the present case the target was
585 always fully specified; target choice did not become more certain with time. Thus, strengthening pre-
586 movement activity is unlikely to be related to target choice *per se*, and is more likely to reflect preparation to
587 execute a choice that was clear from the outset (also see^{31,43}). This suggests that having a clear movement goal
588 does not necessarily mean that low-level preparatory processes are fully engaged. Whether or not preparatory
589 activity develops may depend on whether it is reasonably likely that movement will be initiated soon.
590 Consistent with this interpretation, studies that use a fixed, predictable delay typically find that delay-period
591 activity ramps up with time (e.g.,⁹) while studies that use an unpredictable delay tend to find delay-period
592 activity that reaches a rough plateau after a burst following target onset (e.g.,¹⁰).

593 Thus, the processes of deciding what to do, preparing to do it, and actually initiating, may occur with variable
594 timing relative to one another. This hypothesis is potentially relevant to the finding that there exist neural
595 events that are predictive of movement initiation, yet precede movement by more than the typical reaction
596 time, and also precede self-report of the decision to initiate movement^{43,44}. This is consistent with our finding
597 that preparatory subspace activity, in the self-initiated context, develops hundreds of milliseconds before
598 movement onset, and potentially before a definitive choice to execute movement has been made.

599 *Cortical involvement despite fast RTs*

600 Reaching can involve very rapid, nearly involuntary corrections that are likely to have a subcortical
601 contribution²⁵. It has thus been suggested that entire movements may sometimes be produced subcortically,
602 perhaps with minimal cortical involvement²⁶. In particular, a loud, startling sound can release a pre-planned
603 movement (the 'StartReact' phenomenon) with EMG-based RTs of 70-100 ms. While this short latency is due in
604 part to the use of a highly salient auditory stimulus (which suffers less sensory delay than a visual stimulus), it
605 also depends on the subcortically generated startle reflex. Given that the triggering impetus likely arises
606 subcortically, it has been suggested that movement generation itself may not depend on cortical involvement.
607 Yet a recent study argued against a reduction of cortical involvement in StartReact⁴⁵. The authors instead
608 interpreted StartReact as a subcortical triggering of movement-generating dynamics that span cortical and
609 subcortical circuits in the same way as conventionally triggered movements. Our findings indicate that cortical
610 events are not necessarily slow, and are thus compatible with this view. EMG activity in the quasi-automatic

611 context could begin changing as early as 90 ms following target onset, with a mean of 150 ms and 130 ms for
612 the two monkeys. This is not as fast as during StartReact, but is only slightly slower when one considers the
613 additional sensory delay incurred by a visual stimulus (also note that in StartReact, movements are planned
614 ahead of time while in the quasi-automatic context they are not). The fastest RTs in the quasi-automatic
615 context are thus likely to be near the physiological limit. Yet we saw no evidence of reduced cortical
616 involvement. Indeed, the patterns of cortical movement-related activity in the quasi-automatic context were
617 very similar to those in the other two contexts. As in⁴⁵, we do not suggest an absence of sub-cortical
618 involvement, merely a conservation of cortical involvement.

619 The conservation of neural events across contexts in motor cortex should not be taken to imply that events
620 other cortical or subcortical areas will be similarly conserved. Monkeys were clearly aware of the differences
621 between contexts, and behaved appropriately. It must therefore be the case that some brain areas perform
622 different computations in different contexts as necessary to initiate movement at the appropriate time. We
623 have indeed observed that neural activity in the supplementary motor area differs across contexts not only
624 during movement, but even before target onset (unpublished observations) and one suspects that this will be
625 true of a variety of cortical and subcortical areas. However, motor cortex appears to be playing a more
626 mechanical role: similar movements are driven by similar patterns of movement-related activity, following
627 similar patterns of preparatory activity, across a broad range of timing constraints.

628 **Methods**

629 *Subjects and task*

630 Subjects were two adult male macaque monkeys (*Macaca mulatta*) aged 10 and 14 years and weighing 11 – 13
631 Kg at the time of the experiments. Daily fluid intake was regulated to maintain motivation to perform the task.
632 All procedures were in accord with the US National Institutes of Health guidelines and were approved by the
633 Columbia University Institutional Animal Care and Use Committee.

634 Subjects sat in a primate chair facing an LCD display and performed reaches with their right arm while their left
635 arm was comfortably restrained. The timing of stimulus presentation was controlled using a photodetector
636 (Thorlabs) to track individual frames on the display, such that the exact moment (within ~1 ms) of all events
637 was known. This allows accurate assessment of reaction times and neural response latencies relative to visual
638 events. Hand position was monitored using an infrared optical system (Polaris; Northern Digital) to track (~0.3
639 mm precision) a reflective bead temporarily affixed to the third and fourth digits. Each trial began when the
640 monkey touched and held a central touch-point. Touch-point color indicated context (**Fig. 1**). After the touch-
641 point was held for 450 – 550 ms (randomized) a colored 10 mm diameter disk (the target) appeared in one of
642 eight possible locations radially arranged around the touch point. Target distance was 130 mm for cue and self-
643 initiated contexts and 40 mm for the quasi-automatic context (**Fig. 1**). Trials for different contexts / directions
644 were interleaved using a block-randomized design.

645 In the cue-initiated context, after a variable delay period (0–1000ms) the target suddenly grew in size to 30 mm
646 and the central touch point simultaneously disappeared. These events served as the go-cue, instructing the
647 monkey to make the movement. Reaches were successful if they were initiated within 500 ms of the go cue,
648 had a duration < 500 ms, and landed within an 18 mm radius window centered on the target. Juice was
649 delivered if the monkey held the target, with minimal hand motion, for 200 ms (this criterion was also shared
650 across all three contexts).

651 In the self-initiated context, the target slowly and steadily grew in size, starting upon its appearance and ending
652 when the reach began. Growth continued to a maximum size of 30 mm, which was achieved 1200 ms after
653 target appearance (most reaches occurred before this time). The reward for a correct reach grew exponentially
654 starting at 1 drop and achieved a maximum of 8 drops after 1200 ms. Monkeys were free to move as soon as
655 the target appeared. However, monkeys essentially always waited longer in order to obtain larger rewards. In
656 rare instances where no movement was detected 1500 ms after target onset, the trial was aborted and flagged
657 as an error. Requirements for reach duration and accuracy were as for the cue-initiated context.

658 In the quasi-automatic context, the target moved radially away from the central touch-point at 25 cm / s.
659 Motion began immediately upon target appearance. Target motion ended if a reach succeeded in bringing the
660 hand to the target mid-flight. If the target was not intercepted (*e.g.*, if reach initiation was too slow) then the
661 target continued moving until off the screen. Target speed and initial location (40 mm from the touch-point)
662 were titrated, during training, such that the target was typically intercepted ~130 mm from the touch-point
663 (the same location as the targets for the other two contexts). For successful interception, reaches had to land
664 within an elliptical acceptance window (16 mm by 20 mm radius, with the long axis aligned with target
665 motion). If the target was successfully intercepted, it grew in size to 30 mm and reward was delivered after the
666 hold period.

667 Movement onset (and thus the RT) was measured based on hand speed: the magnitude of velocity in the plane
668 of the movements (not including depth). We considered the time of movement onset to be the first moment
669 when speed exceeded a threshold, set to 1% of average peak speed for that target location. To ensure this
670 measure was robust on a given trial, we also measured 'backwards' from the moment of peak speed, and
671 found the last moment when peak speed was below a threshold. On rare occasions, these 'forwards' and
672 'backwards' measurements disagreed (*e.g.*, due to an unusual velocity profile) in which case the time of
673 movement onset was considered undefined and the RT was not analyzed.

674 *Neural and muscle recordings*

675 After subjects became proficient in the task, we performed sterile surgery to implant a head restraint. At the
676 same time, we implanted a recording chamber centered over the arm area of primary motor cortex (M1) and
677 the dorsal premotor cortex (PMd) of the left hemisphere. Chamber positioning was guided by structural
678 magnetic resonance images taken shortly before implantation. We used intracortical microstimulation to
679 confirm that our recordings were from the forelimb region of motor cortex (biphasic pulses, cathodal leading,
680 250 μ S pulse width delivered at 333 Hz for a total duration of 50 ms). Microstimulation typically evoked
681 contractions of the shoulder and upper-arm muscles, at currents from 5 μ A – 60 μ A depending on the location
682 and cortical layer. We recorded single-neuron responses using traditional tungsten electrodes (FHC) or one or
683 more silicon linear-array electrodes (V-probes; Plexon) lowered into cortex using a motorized microdrive. For
684 tungsten-electrode recordings, spikes were sorted online using a window discriminator (Blackrock
685 Microsystems). For linear-array recordings spikes were sorted offline (Plexon Offline Sorter). We recorded all
686 well isolated task-responsive neurons; no attempt was made to screen for neuronal selectivity for reach
687 direction or any other response property. Spikes were smoothed with a Gaussian kernel with standard

688 deviation of 20 ms and averaged across trials to produce peri-stimulus time histograms. For measurements of
689 latency, we used a 10 ms Gaussian to minimize the impact of filtering on latency.

690 We recorded electromyogram (EMG) activity using intramuscular electrodes from the following
691 muscles: lower and upper aspects of the trapezius, medial, lateral and anterior aspects of the deltoid, medial
692 and outer aspects of the biceps, brachialis, pectoralis and latismus dorsi. The triceps were minimally active and
693 were not recorded. EMG signals were bandpass filtered (10 – 500 Hz), digitized at 1kHz, rectified, smoothed
694 with a Gaussian kernel with standard deviation of 20 ms, and averaged across trials to produce peri-stimulus
695 time histograms.

696 *Data pre-processing prior to population analyses*

697 As in our previous work, we employed two pre-processing steps¹⁹. First, the responses of each neuron were
698 soft-normalized so that neurons with high firing rates had approximately unity firing-rate range (normalization
699 factor = firing rate range+5). This step ensures that subsequent dimensionality reduction (see below) captures
700 the response of all neurons, rather than a handful of high firing-rate neurons. Second, the responses for each
701 neuron were mean-centered at each time as follows: we calculated the mean activity across all conditions of
702 each neuron at each time point, and subtracted this mean activity from each condition's response. This step
703 ensures that dimensionality reduction focuses on dimensions where responses are selective across conditions,
704 rather than dimensions where activity varies in a similar fashion across all conditions⁴⁶.

705 *Identifying preparatory and movement dimensions*

706 We recently developed a method that leverages the finding that neural responses in the delay-epoch are
707 nearly orthogonal to responses in the movement epoch²⁷. The method identifies one set of preparatory
708 dimensions and an orthogonal set of movement dimensions. Briefly, we define two matrices based on data
709 from the cue-initiated context only: $P \in \mathbb{R}^{N \times CT}$ which holds preparatory epoch responses and $M \in$
710 $\mathbb{R}^{N \times CT}$ which holds movement-epoch neural responses. N is the number of neurons recorded, C is the
711 number of reach directions and T is the number of time points. The method seeks to find a set of preparatory
712 dimensions, \hat{Q}_{prep} , that maximally capture the percentage of variance of P and an orthogonal set of
713 movement dimensions, \hat{Q}_{move} , that maximally capture the percentage of variance of M . We compute the
714 preparatory and movement-epoch covariance matrices $C_{prep} = cov(P)$ and $C_{move} = cov(M)$ and optimize
715 the following objective function:

$$716 \quad [\hat{Q}_{prep}, \hat{Q}_{move}] = \underset{[Q_{prep}, Q_{move}]}{\operatorname{argmax}} \frac{1}{2} \left(\frac{\operatorname{Tr}(Q_{prep}^T C_{prep} Q_{prep})}{\sum_{i=1}^{d_{prep}} \sigma_{prep}(i)} + \frac{\operatorname{Tr}(Q_{move}^T C_{move} Q_{move})}{\sum_{i=1}^{d_{move}} \sigma_{move}(i)} \right)$$

$$717 \quad \text{subject to } Q_{prep}^T Q_{move} = \mathbf{0}, \quad Q_{prep}^T Q_{prep} = I, \quad Q_{move}^T Q_{move} = I$$

718 where $\sigma_{prep}(i)$ is the i^{th} singular value of C_{prep} , and $\sigma_{move}(i)$ is the i^{th} singular value of C_{move} . Q_{prep} and
 719 Q_{move} are the bases for the preparatory and movement subspaces respectively and $\operatorname{Tr}(\cdot)$ is the matrix trace
 720 operator. The term $\operatorname{Tr}(Q_{prep}^T C_{prep} Q_{prep})$ reflects the preparatory-epoch data variance captured by the
 721 preparatory subspace, and $\operatorname{Tr}(Q_{move}^T C_{move} Q_{move})$ reflects the movement-epoch data variance captured by
 722 the movement subspace. We chose the dimensionality of Q_{prep} to be 12 (i.e., $Q_{prep} \in \mathbb{R}^{N \times 12}$), which
 723 captured $\sim 80\%$ of preparatory-epoch variance (the remaining variance had very little structure and appeared
 724 to be primarily sampling noise). Similarly, we chose the dimensionality of Q_{move} to be 12, which captured \sim
 725 85% of movement-epoch variance. Results were robust with respect to the choice of dimensionality.
 726 The optimization objective is normalized (by the singular values) to be insensitive to the relative dimensionality
 727 and amount of response variance in the two subspaces. This normalization is particularly important in our case
 728 since movement activity is stronger and typically has higher dimensionality than the preparatory activity. For
 729 visualization, we need to choose two dimensions spanned by Q_{prep} and Q_{move} to define the plotted projections
 730 (e.g., in **Fig. 7**). For Q_{prep} we chose the basis so that the top two dimensions captured the most variance (with
 731 all others ranked accordingly). For Q_{move} the basis was chosen using the jPCA method¹⁹ to capture movement-
 732 related oscillatory activity patterns.

733 *Projections and reconstructions*

734 For a given time t and for condition θ , the projection of the population response onto the k^{th} preparatory
 735 dimension is simply a weighted sum of all single-neuron responses: $x_k^{prep}(t, \theta) = \sum_{n=1}^N Q_{n,k}^{prep} r_n(t, \theta)$ where
 736 $Q_{n,k}^{prep}$ is the element in the n^{th} row and k^{th} column of Q_{prep} (see previous section) and $r_n(t, \theta)$ is the
 737 response of the n^{th} neuron. This is illustrated in Figure 4, where the orange weights w are taken from Q_{prep} .
 738 The projection onto each movement dimension is defined analogously. The response of a given neuron can
 739 then be reconstructed as: $\hat{r}_n(t, \theta) = \sum_{k=1}^{12} Q_{n,k}^{prep} x_k^{prep}(t, \theta) + \sum_{k=1}^{12} Q_{n,k}^{move} x_k^{move}(t, \theta)$. The two sums on the
 740 right-hand side are the preparatory and movement contributions respectively, which we can term $\hat{r}_n^{prep}(t, \theta)$
 741 and $\hat{r}_n^{move}(t, \theta)$. These are the preparatory and movement patterns for neuron n .

742 *Subspace occupancy*

743 Our measurement of subspace occupancy is equivalent to the variance explained metric often computed in the
744 context of principal component analysis. For the preparatory subspace, occupancy was computed as:

745 $occupancy^{prep}(t) = \sum_{k=1}^{12} var_{\theta}(x_k^{prep}(t, \theta))$ where var_{θ} indicates taking the variance across conditions (i.e.,

746 directions). Because dimensions are orthonormal this is equivalent to computing the variance for each

747 neuron's preparatory pattern and then summing: $\sum_{k=1}^{12} var_{\theta}(x_k^{prep}(t, \theta)) = \sum_{n=1}^N var_{\theta}(\hat{r}_n^{prep}(t, \theta))$.

748 Movement subspace occupancy was defined analogously. To estimate the sampling error of the subspace

749 occupancy we used a bootstrap procedure. We created 1000 surrogate populations by redrawing with

750 replacement from the original population. We computed the subspace occupancy for each surrogate

751 population, and for each time computed the standard deviation across these 1000 measures.

752 *Latency of physiological events*

753 To measure the latency of preparatory- and movement-subspace occupancy, we filtered the spike trains of all

754 neurons using a Gaussian kernel with 10 ms standard deviation. We recomputed the preparatory and

755 movement dimensions using these data and calculated the subspace occupancy as before. We measured the

756 latency as the first moment in time in which occupancy reached 10% of peak occupancy.

757 Similarly, to calculate the latency of the EMG with respect to movement onset, we filtered EMG activity of all

758 muscles using a Gaussian kernel with a 10 ms standard deviation. We then performed PCA on the EMG activity

759 for each context separately and projected the corresponding EMG responses onto the first PC. We measured

760 the latency as the first moment in which activity in the first PC reached 10% of peak activity.

Acknowledgments

We thank Yanina Pavlova and Sean Perkins for technical support. This work was supported by the Sloan Foundation, the Simons Foundation (SCGB#325171 and SCGB#325233), the Grossman Center for the Statistics of Mind, the McKnight Foundation, an NIH Directors New Innovator Award (DP2 NS083037), NIH/NSF CRCNS R01NS100066, the Kavli Foundation, a Klingenstein-Simons Fellowship, the Searle Scholars Program, an NIH Postdoctoral Fellowship (F32 NS092350) and the Gatsby Charitable Trust. We thank Andrew Zimnik for helpful discussions and performing some analyses of behavior.

Author Contributions

A.H.L. and M.M.C. designed the experiments. A.H.L. recorded and analyzed the data. G.F.E. and J.P.C. developed and implemented the method for segregating preparatory and movement subspaces. All authors contributed to analysis choices and interpretation. A.H.L. and M.M.C. wrote the manuscript with help from G.F.E. and J.P.C.

Competing Financial Interests

The authors declare no competing financial interests.

References

- 1 Tanji, J. & Evarts, E. V. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J Neurophysiol* **39**, 1062-1068 (1976).
- 2 Rosenbaum, D. A. Human movement initiation: specification of arm, direction, and extent. *J Exp Psychol Gen* **109**, 444-474 (1980).
- 3 Wise, S. P. The primate premotor cortex: past, present, and preparatory. *Annu Rev Neurosci* **8**, 1-19 (1985).
- 4 Riehle, A. & Requin, J. The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behav Brain Res* **53**, 35-49 (1993).
- 5 Ghez, C. *et al.* Discrete and continuous planning of hand movements and isometric force trajectories. *Exp Brain Res* **115**, 217-233 (1997).
- 6 Crammond, D. J. & Kalaska, J. F. Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *J Neurophysiol* **84**, 986-1005 (2000).
- 7 Rickert, J., Riehle, A., Aertsen, A., Rotter, S. & Nawrot, M. P. Dynamic encoding of movement direction in motor cortical neurons. *J Neurosci* **29**, 13870-13882, doi:29/44/13870 [pii]10.1523/JNEUROSCI.5441-08.2009 (2009).
- 8 Michaels, J. A., Dann, B., Intveld, R. W. & Scherberger, H. Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J Neurosci* **35**, 11415-11432, doi:10.1523/JNEUROSCI.1714-15.2015 (2015).
- 9 Bastian, A., Schoner, G. & Riehle, A. Preshaping and continuous evolution of motor cortical representations during movement preparation. *Eur J Neurosci* **18**, 2047-2058 (2003).
- 10 Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G. & Shenoy, K. V. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci* **26**, 3697-3712 (2006).
- 11 Weinrich, M., Wise, S. P. & Mauritz, K. H. A neurophysiological study of the premotor cortex in the rhesus monkey. *Brain* **107 (Pt 2)**, 385-414 (1984).
- 12 Godschalk, M., Lemon, R. N., Kuypers, H. G. & van der Steen, J. The involvement of monkey premotor cortex neurones in preparation of visually cued arm movements. *Behav Brain Res* **18**, 143-157 (1985).
- 13 Wise, S. P. & Kurata, K. Set-related activity in the premotor cortex of rhesus monkeys: effect of triggering cues and relatively long delay intervals. *Somatosens Mot Res* **6**, 455-476 (1989).
- 14 Kurata, K. Distribution of neurons with set- and movement-related activity before hand and foot movements in the premotor cortex of rhesus monkeys. *Exp Brain Res* **77**, 245-256 (1989).
- 15 Crammond, D. J. & Kalaska, J. F. Neuronal activity in primate parietal cortex area 5 varies with intended movement direction during an instructed-delay period. *Exp Brain Res* **76**, 458-462 (1989).
- 16 Afshar, A. *et al.* Single-trial neural correlates of arm movement preparation. *Neuron* **71**, 555-564, doi:10.1016/j.neuron.2011.05.047 (2011).
- 17 Churchland, M. M. & Shenoy, K. V. Delay of movement caused by disruption of cortical preparatory activity. *J Neurophysiol* **97**, 348-359 (2007).
- 18 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* **68**, 387-400, doi:doi:10.1016/j.neuron.2010.09.015 (2010).
- 19 Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51-56, doi:10.1038/nature11129 (2012).

- 20 Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annual Review of Neuroscience* **36**, 337-359, doi:10.1146/annurev-neuro-062111-150509 (2013).
- 21 Churchland, M. M. & Cunningham, J. P. A Dynamical Basis Set for Generating Reaches. *Cold Spring Harb Symp Quant Biol* **79**, 67-80, doi:10.1101/sqb.2014.79.024703 (2014).
- 22 Ames, K. C., Ryu, S. I. & Shenoy, K. V. Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron* **81**, 438-451, doi:10.1016/j.neuron.2013.11.003 (2014).
- 23 Wong, A. L., Haith, A. M. & Krakauer, J. W. Motor Planning. *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, doi:10.1177/1073858414541484 (2014).
- 24 Haith, A. M., Pakpoor, J. & Krakauer, J. W. Independence of Movement Preparation and Movement Initiation. *J Neurosci* **36**, 3007-3015, doi:10.1523/JNEUROSCI.3245-15.2016 (2016).
- 25 Day, B. L. Subcortical visuomotor control of human limb movement. *Adv Exp Med Biol* **826**, 55-68, doi:10.1007/978-1-4939-1338-1_5 (2014).
- 26 Perfiliev, S., Isa, T., Johnels, B., Steg, G. & Wessberg, J. Reflexive limb selection and control of reach direction to moving targets in cats, monkeys, and humans. *Journal of Neurophysiology* **104**, 2423-2432, doi:10.1152/jn.01133.2009 (2010).
- 27 Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* **7**, 13239, doi:10.1038/ncomms13239 (2016).
- 28 Berens, P. CircStat: A Matlab Toolbox for Circular Statistics. *Journal of Statistical Software* **31** (2009).
- 29 Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol* **97**, 4235-4257 (2007).
- 30 Passingham, R. E. *The frontal lobes and voluntary action*. (Oxford University Press, 1993).
- 31 Murakami, M., Vicente, M. I., Costa, G. M. & Mainen, Z. F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci*, doi:10.1038/nn.3826 (2014).
- 32 Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat Neurosci* **18**, 1025-1033, doi:10.1038/nn.4042 (2015).
- 33 Fetz, E. E. & Cheney, P. D. Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *J Neurophysiol* **44**, 751-772 (1980).
- 34 Wise, S. P., Weinrich, M. & Mauritz, K. H. Movement-related activity in the premotor cortex of rhesus macaques. *Prog Brain Res* **64**, 117-131 (1986).
- 35 Kaufman, M. T. *et al.* Roles of Monkey Premotor Neuron Classes in Movement Preparation and Execution. *Journal of Neurophysiology* **104**, 799-810, doi:10.1152/jn.00231.2009 (2010).
- 36 Bastian, A., Riehle, A., Erlhagen, W. & Schoner, G. Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport* **9**, 315-319 (1998).
- 37 Erlhagen, W. & Schoner, G. Dynamic field theory of movement preparation. *Psychol Rev* **109**, 545-572 (2002).
- 38 Thura, D., Beauregard-Racine, J., Fradet, C. W. & Cisek, P. Decision making by urgency gating: theory and experimental support. *J Neurophysiol* **108**, 2912-2930, doi:10.1152/jn.01071.2011 (2012).
- 39 Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801-814, doi:S0896-6273(05)00059-0 [pii]10.1016/j.neuron.2005.01.027 (2005).

- 40 Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* **22**, 9475-9489 (2002).
- 41 Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *eLife* **4**, e04677, doi:10.7554/eLife.04677 (2015).
- 42 Guo, Z. V. *et al.* Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179-194, doi:10.1016/j.neuron.2013.10.020 (2014).
- 43 Schurger, A., Sitt, J. D. & Dehaene, S. An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci U S A* **109**, E2904-2913, doi:10.1073/pnas.1210467109 (2012).
- 44 Libet, B., Gleason, C. A., Wright, E. W. & Pearl, D. K. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain* **106 (Pt 3)**, 623-642 (1983).
- 45 Dean, L. R. & Baker, S. N. Fractionation of muscle activity in rapid responses to startling cues. *J Neurophysiol* **117**, 1713-1719, doi:10.1152/jn.01009.2015 (2017).
- 46 Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* **3**, doi:10.1523/ENEURO.0085-16.2016 (2016).

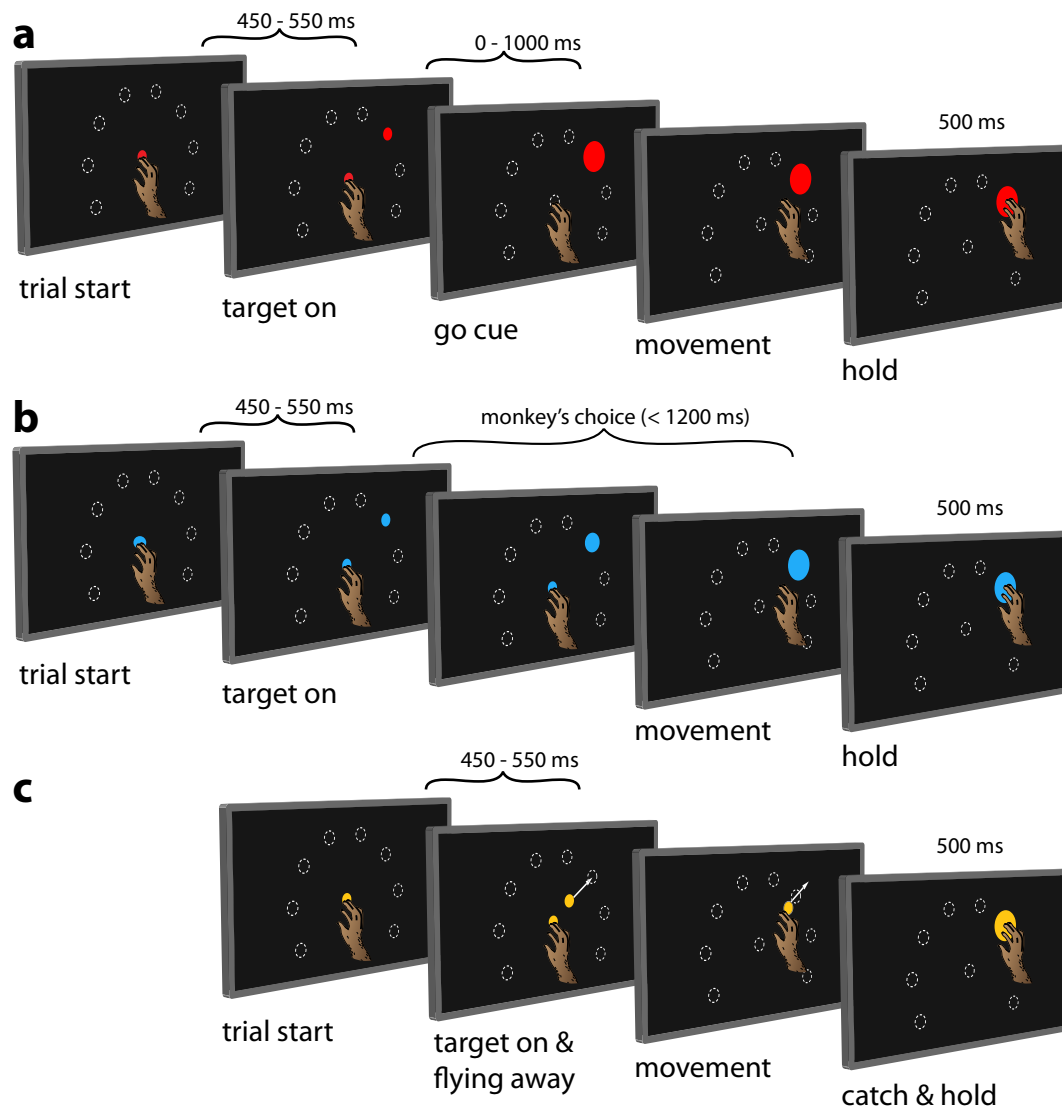


Figure 1. Behavioral task. Monkeys performed the same set of reaches under three initiation contexts. **a)** Cue-initiated context. Trials started when the monkeys touched a red central point on the screen. After a brief delay (450 - 550 ms) a red target appeared in one of eight possible locations (white dashed circles, not visible to the monkey) 130 mm from the touch point. After a variable delay period (0 - 1000ms) the target suddenly increased in size providing the go cue to initiate the reach. **b)** Self-initiated context. Trials began as above, but the central point was blue. Subsequently, a small blue target appeared and gradually grew in size. Monkeys were free to initiate the reach as soon as the target appeared on the screen. However longer waiting times were rewarded with larger amounts of juice. **c)** Quasi-automatic context. The central point was yellow. Yellow targets appeared in one of eight possible locations. The initial appearance of the target was 40 mm from the touch point. Immediately after appearing, the target moved radially outward. Monkeys had to initiate the reach quickly in order to intercept the target before it reached the edge of the screen and disappeared. Targets were intercepted at a location near the location of the targets in the other two tasks (dashed circles).

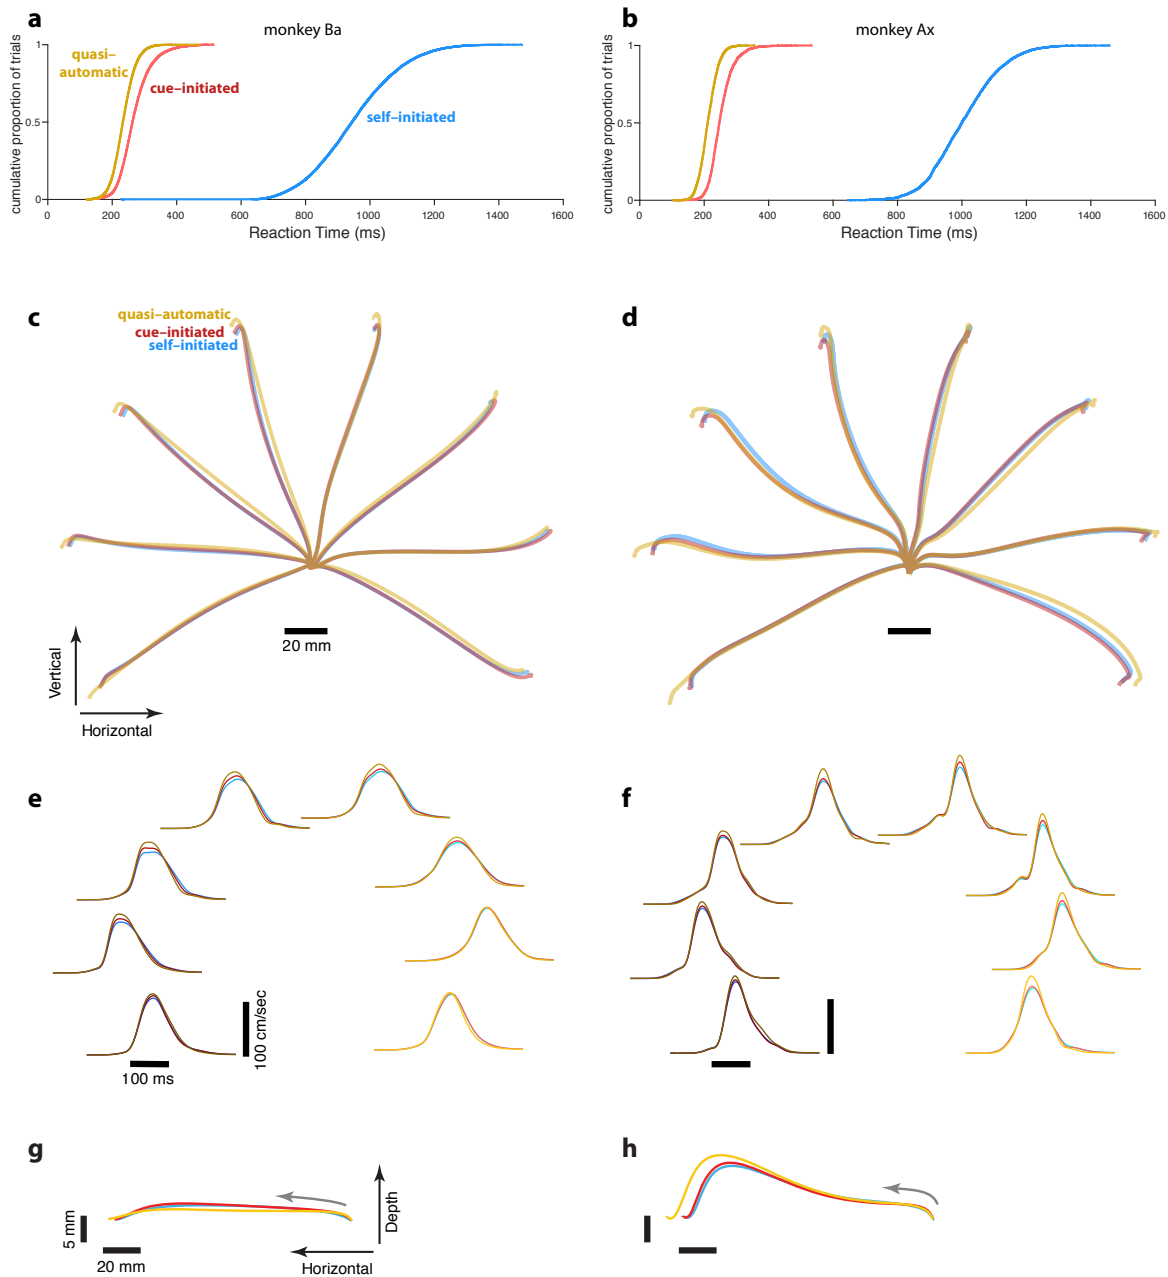


Figure 2. Reaction time distributions and reach kinematics. **a,b)** cumulative reaction time distributions for the three contexts. Trials are pooled across all recordings. **c,d)** Average reach trajectories for the eight targets in the three contexts for monkeys Ba and Ax respectively. **e,f)** Average speed profiles in the three contexts, using the same color coding as above. Additionally, the shade of each line indicates reach direction. This same shade-coding (*light traces* for rightwards reaches and *dark traces* for leftwards reaches) is preserved in subsequent figures. **g,h)** Average reach trajectories for one example reach direction (leftward) with depth shown on an expanded scale to allow closer examination of trajectories in that dimension. Gray arrows indicate direction in which the hand traveled.




Figure 3. Responses of three example motor cortex neurons. Each column shows the responses of a single neuron for the three initiation contexts. Each trace plots the trial-averaged firing rate for one reach direction (same color scheme as in Fig. 2). Gray shaded regions indicate the delay and movement epochs, used to define the preparatory and movement dimensions in subsequent analyses. All traces contain data that was aligned to target onset for the left-hand side of the trace, and to movement onset for the right-hand side of the trace. Individual trials had these two intervals spliced together before filtering and averaging. The sizes of these intervals and the moment of splicing were chosen to imitate the typical timing between target onset and movement onset (keeping in mind that this was variable across trials). For the cue-initiated context, the left-hand side contains data from -200 – 450 ms relative to target onset (only trials with delays >400 ms were analyzed). The right-hand side contains data from -350 ms before movement to 400 ms post movement. The indicated time of the go cue is based on the mean reaction time. For the self-initiated context, spliced averages were computed using the same timing as above, to aid visual comparison. For the quasi-automatic context, the first 150 ms of the response is aligned to the target onset and the subsequent response is aligned to movement onset. Splicing was performed so that the interval from target onset to movement onset matched the mean reaction time. All vertical calibration bars indicate 20 spikes/s.




Figure 4. Preparatory and movement contributions to neuronal responses. The population of neural responses (*left column*) can be linearly ‘decoded’ into preparatory and movement-related projections (*middle column*). $w_{n,k}$ is the decode weight from the n^{th} neuron to the k^{th} projection, and the collection of $w_{\cdot,k}$ such weights is the k^{th} neural dimension. Empirically, it is possible to find orthogonal dimensions that segregate preparatory and movement-related response patterns: one set of projections shows structure only during preparation (*orange traces in middle column*), while the other shows structure only around the time of the movement (*purple traces in middle column*). Because dimensions are chosen to maximally capture data variance (*i.e.*, the structure of firing rates), individual-neuron responses can be reconstructed from the projections. Reconstruction employs the same weights that defined the projections (*e.g.*, if neuron 88 contributed to the first preparatory projection with weight $w_{88,1}$, then that first preparatory projection contributes to the reconstruction of neuron 88 with the same weight). The weighted sum of the preparatory projections yields that neuron’s ‘preparatory pattern’ (*orange traces in right column*) and the weighted sum of movement projections yields that neuron’s ‘movement pattern’ (*purple traces in right column*). The full reconstruction is the sum of these two patterns (which describe the tuned aspects of the neuron’s response) plus a time-varying mean that captures any untuned trends in the overall mean firing rate with time (not shown). The success of the approximation can be appreciated by comparison with the true response in Fig. 3a.

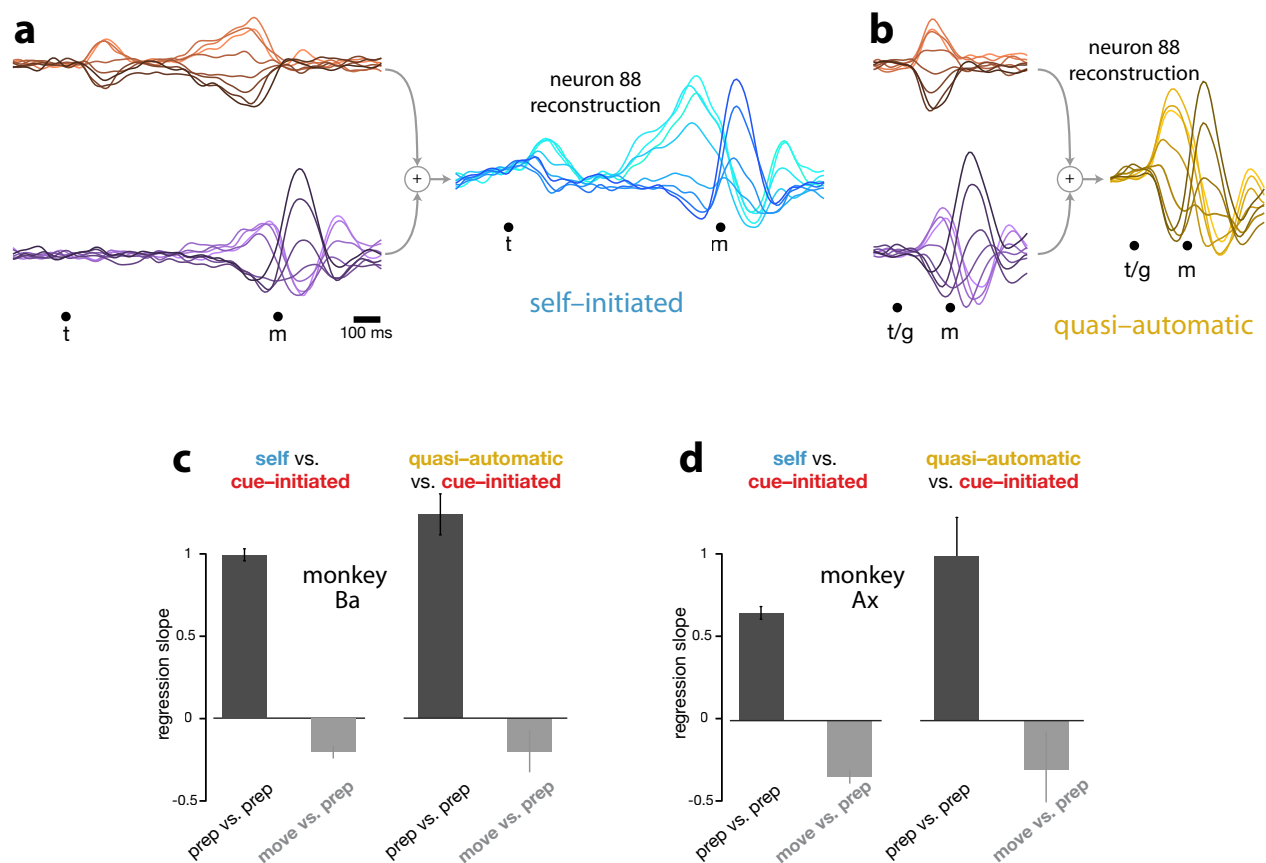


Figure 5. Reconstruction of single-neuron responses in the self-initiated and quasi-automatic contexts, using preparatory and movement dimensions found for the cue-initiated context. a). Example reconstruction for the self-initiated context. The population response was projected onto the preparatory and movement dimensions (found using the cue-initiated context). We then used those projections to reconstruct the preparatory (*orange*) and movement (*purple*) patterns that contributed to each neuron’s response. These patterns are shown for neuron 88. The full reconstruction (*blue*) is the sum of the preparatory and movement-related patterns, plus the across-condition mean response (which has no tuning, but captures the overall mean rate at each time). This reconstructed response can be compared with the true response in Fig. 3a. **b).** As in **a**, but for the quasi-automatic context. The reconstructed response (*yellow*) can be compared with the true response in Fig. 3a. **c).** Summary of the degree to which a given neuron’s preparatory-subspace patterns were similar across contexts. Data are for monkey Ba. For each neuron, we took the value of the preparatory subspace pattern for each direction 100 ms before movement onset. These values form a vector. To compute the similarity of the vector for the self-initiated context with that for the cue-initiated context, we regressed the former versus the latter and took the slope. The same was done for the quasi-automatic context versus the cue-initiated context. Dark bars show the average slope across neurons \pm SEM. As a comparison, we repeated this analysis but regressed the movement pattern during the self-initiated and quasi-automatic contexts versus the preparatory pattern during the cue-initiated context (*light gray bars*). The movement pattern was assessed 150 ms after movement onset. **d).** As in **c**, but for monkey Ax.

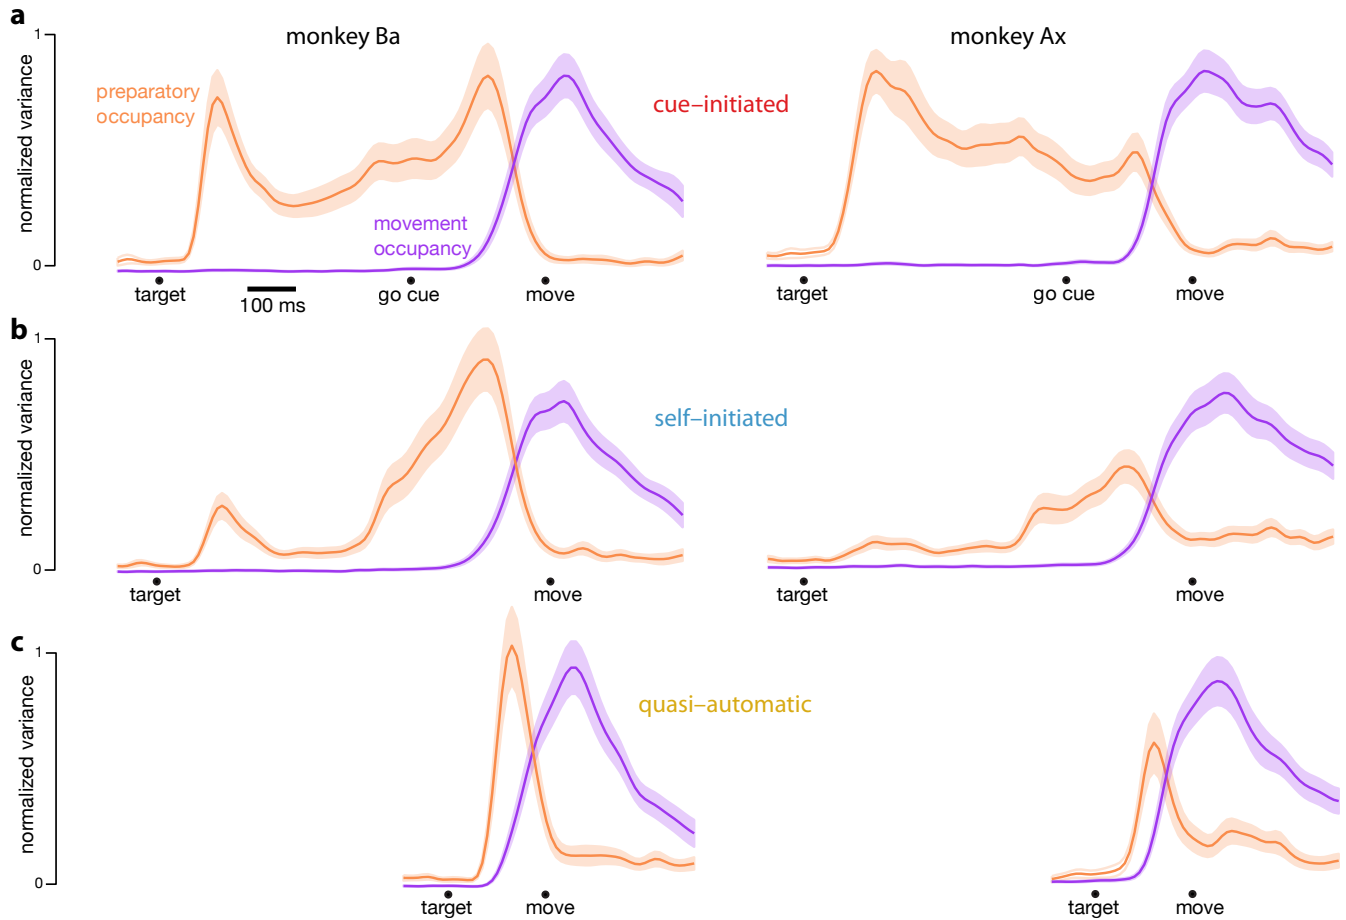


Figure 6. Preparatory and movement-subspace occupancy. a) Preparatory and movement subspace occupancy for the cue-initiated context. The two columns show results for monkeys Ba and Ax. Occupancy was calculated as the sum of the across-condition variance in the preparatory and movement dimensions respectively. Preparatory-subspace occupancy, across all three contexts, was normalized by the highest value attained in the cue-initiated context. Movement-subspace occupancy was similarly normalized. The shaded region denotes the standard deviation of the sampling error (equivalent to the standard error) computed via bootstrap (*methods*). b,c) Occupancy during the self-initiated and quasi-automatic contexts respectively.




Figure 7. Snapshots of neural population state in the preparatory and movement subspaces for monkey Ba. Responses for cue-initiated (*red*), self-initiated (*blue*) and quasi-automatic (*yellow*) contexts projected onto the top two preparatory dimensions (**a**) and top two movement dimensions (**b**). Those dimensions were found using the cue-initiated context data only. Trace colors correspond to target direction (same color scheme as in **Figure 3**). Each snapshot shows the neural state in that subspace, for all eight directions, across a 150 ms window. Snapshot labeled 'baseline' begins 150 ms before target onset. Snapshot labeled 'target' plots data starting at target onset. For the cue-initiated and self-initiated contexts, the subsequent three snapshots show activity in 150 ms increments, still aligned to target onset. Snapshots labeled 'move -120' start 120 ms before movement onset, with data aligned to movement onset. Subsequent panels begin at the indicated time.

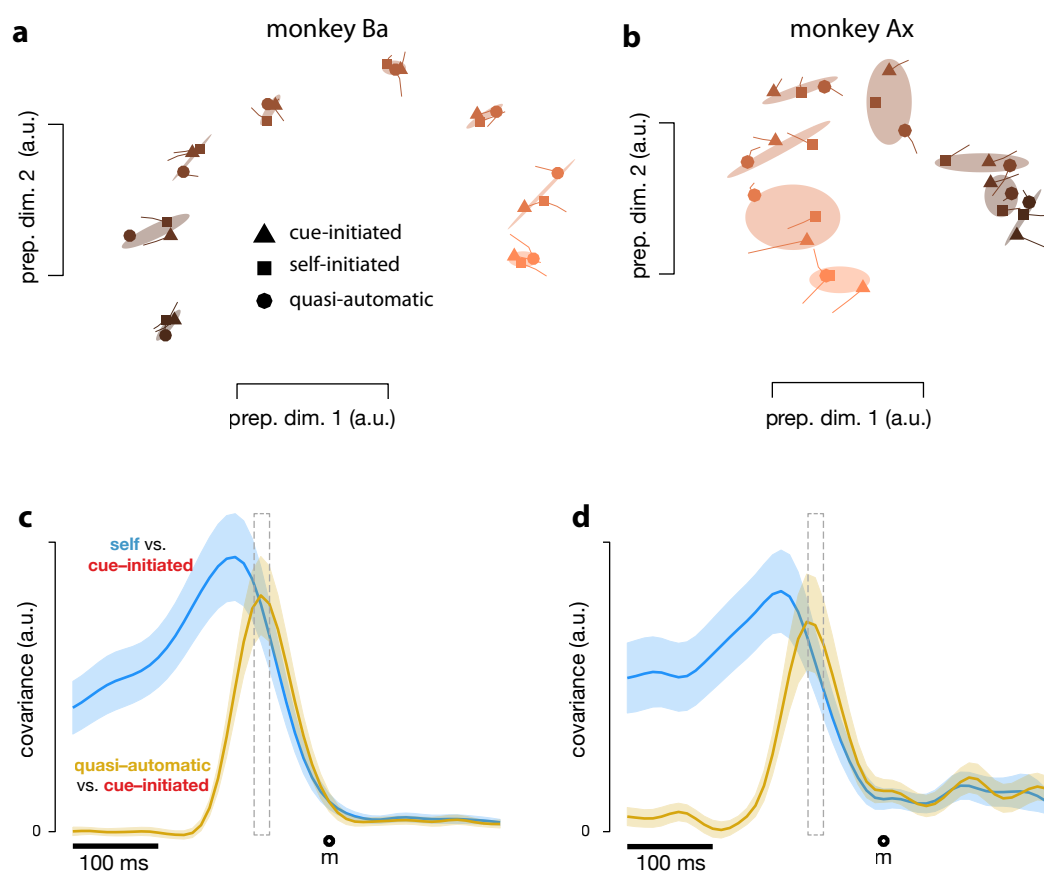
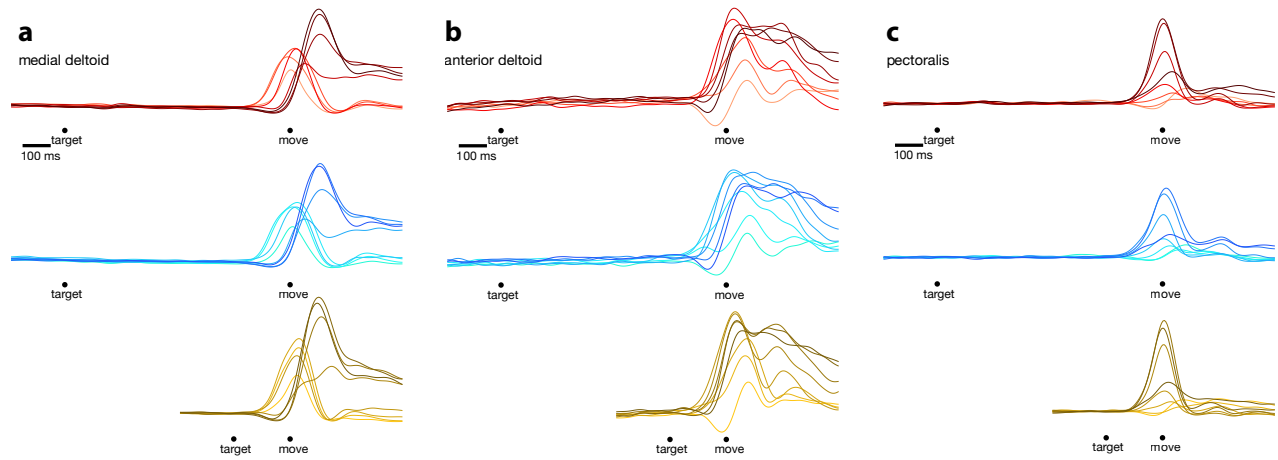



Figure 8. Preparatory subspace activity just before movement onset. **a)** Data for monkey Ba. Each marker denotes the neural state in a two-dimensional preparatory subspace: the two dimensions that captured the most variance, as in Fig. 7a. Markers indicate the state 70 ms before movement onset. Tails plot 20 ms of activity leading up to that time. The three shapes show states for the three contexts. Shaded regions plot a covariance ellipse for each triplet of states. A different symbol shade is used for each target direction (*light* for right, *dark* for left). **b)** As in **a** but for monkey Ax. **c)** Quantification of the time-course of the similarity in the pattern of preparatory states between contexts. Data is for monkey Ba. *Blue trace* plots the covariance between the preparatory pattern in the self-initiated context and that in the cue-initiated context. *Yellow trace* plots the covariance between the preparatory pattern in the quasi-automatic context and that in the cue-initiated context. The covariance is high when patterns are both strong and similar. Note that the units of the vertical scale are arbitrary (for reference, the correlation peaks close to one). Gray dashed window of time indicates the 20 ms time range (from 90 to 70 ms before movement onset) shown in **a** and **b**. The shaded regions denote the standard deviation of the sampling error (equivalent to the standard error) computed via bootstrap (*methods*).



Supplementary Figure 1. Responses from three muscles of the upper arm. Each column shows responses of a single muscle for the three initiation contexts. Each trace denotes the trial-averaged firing rate for one reach direction. Individual EMG records were filtered and rectified before smoothing with a Gaussian and averaging. Alignment of individual trials to task events is as described in Fig. 3 of the main text.



Supplementary Figure 2. Snapshots of neural population state in the preparatory and movement subspaces for monkey Ax. All plotting conventions are as in Figure 7 of the main text.