bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Parallel planning through an optimal neural subspace in motor cortex

Nicolas Meirhaeghe[1], Alexa Riehle[1,2], Thomas Brochier[1]

[1]Institut de Neurosciences de la Timone (INT), UMR 7289, CNRS, Aix-Marseille Université, Marseille 13005, France
[2]Institute of Neuroscience and Medicine (INM-6), Jülich Research Centre, Jülich 52428, Germany

Corresponding author
Nicolas Meirhaeghe, Ph.D.
Institut de Neurosciences de la Timone
UMR 7289, CNRS, Aix-Marseille Université,
Marseille 13005, France
Email: nicolas.meirhaeghe@univ-amu.fr

Acknowledgements
NM was funded by a postdoctoral fellowship from the European Molecular Biology Organization (EMBO ALTF 329-2021). TB and AR were supported by the RIKEN-CNRS Collaborative Research Agreement, the Helmholtz portfolio theme “Supercomputing and modeling for the human brain” (SMHB), and the ANR-GRASP (France). The authors wish to thank Manuel Zaepffel and Margaux Duret for participating in the data collection.

Author contributions
TB and AR conceived the project. AR processed the behavioral and electrophysiological data. NM performed all the analyses and wrote the manuscript. All authors contributed to the interpretation of the data. TB supervised the project.

Competing interest
The authors declare no competing interest.

Data and code availability
The data and code needed to generate the figures will be made available upon publication.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Summary

Preparatory activity in motor cortex may facilitate movement execution by representing movement parameters (*representational theory*) or by initializing movement-related dynamics to an optimal state (*dynamical system theory*). Here, we confronted these two theories using neural data from non-human primates. We analyzed the structure of preparatory activity in a task where animals needed to plan two movements simultaneously. Contrary to what the representational theory predicted, we did not find evidence for the concurrent representation of the two movements. Instead, our data revealed that parallel planning was achieved by adjusting preparatory activity to an intermediate state that served as an optimal initial condition to generate both movements. This optimization quantitatively explained fluctuations in the animals' behavior, and directly supported the dynamical system theory. Together, these results uncover a simple mechanism for planning multiple movements in parallel, and shed light on an enduring debate about the nature of preparatory activity in motor cortex.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Introduction

To act efficiently, we often plan our movements ahead. Motor planning has historically been studied in experiments where subjects are given a preparatory period before executing a movement specified in advance (Day et al., 1989; Rosenbaum, 1980; Wise, 1985). Behaviorally, a well-established result is that motor planning enables movements to be initiated more quickly (Crammond and Kalaska, 2000; Ghez et al., 1997; Riehle and Requin, 1989). At the neural level, the benefit of motor planning on reaction time has been linked to preparatory activity in premotor and primary motor cortex (Afshar et al., 2011; Bastian et al., 2003; Churchland and Shenoy, 2007a; Michaels et al., 2015; Riehle and Requin, 1989, 1993; Tanji and Evarts, 1976; Weinrich et al., 1984). Indeed, earlier work has shown that preparatory activity in these brain regions is strongly modulated by key parameters (e.g., direction, speed, extent) of the upcoming movement (Churchland et al., 2006a; Even-Chen et al., 2019; Godschalk et al., 1985; Kurata, 1993; Messier and Kalaska, 2000; Riehle et al., 1994). These findings have led to the *representational theory* of movement preparation which considers motor plans as parametric representations of movement features in motor cortex (Requin et al., 1991).

More recent work has argued in favor of an alternative model of movement preparation grounded in the theory of dynamical systems (Erlhagen and Schöner, 2002; Scott, 2012; Sussillo et al., 2015; Todorov and Jordan, 2002; Versteeg and Miller, 2022). According to the so-called *initial condition (IC) hypothesis*, preparatory neural activity does not represent movement parameters, but instead reflects a controlled dynamical process whose role is to optimize the initial state from which movement-related neural dynamics unfolds (Churchland et al., 2006b, 2010; Hennequin et al., 2014; Kao et al., 2021). Multiple lines of evidence support this alternate view of motor planning. First, preparatory activity was shown to be more dynamic than previously thought (Bastian et al., 2003; Churchland et al., 2006b; Hatsopoulos et al., 2007; Rickert et al., 2009). Second, preparatory and movement-related activity can significantly differ in terms of tuning properties (Churchland et al., 2010), arguing against a purely parametric view. Third, causally perturbing the preparatory state shortly before movement initiation selectively delays movement (Churchland and Shenoy, 2007a). Finally, trial-to-trial fluctuations in the preparatory state are strongly predictive of reaction times (Afshar et al., 2011; Michaels et al., 2015; Pandarinath et al., 2018; Riehle and Requin, 1993).

One appeal of the dynamical system theory is that it provides a conceptually straightforward interpretation of preparatory and motor-related signals that are notoriously difficult to parse out at the level of individual neurons (Batista et al., 2007; Churchland and Shenoy, 2007b; Churchland et al., 2012; Fetz, 1992; Scott, 2008). Indeed, the notion of initial condition has contributed to elucidating a number of

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

computational questions at the neural population level, from the relationship between movement preparation and execution (Elsayed et al., 2016; Kaufman et al., 2014), to the role of preparatory activity in motor learning (Golub et al., 2018; Sadtler et al., 2014; Sun et al., 2020; Vyas et al., 2018, 2020a), and the logic behind the neural control of timed movements (Remington et al., 2018; Sohn et al., 2019; Wang et al., 2017). Nevertheless, there remains extensive debate about which of the representational or the dynamical system theory should be adopted to account for preparatory activity in motor cortex (Omrani et al., 2017).

One major challenge in arbitrating between these two theories is that they need not be mutually exclusive. In fact, one direct prediction of the IC hypothesis is that preparatory activity corresponding to a specific movement *should* co-vary with the associated kinematic parameters (Churchland et al., 2010; Shenoy et al., 2011). Although this realization has encouraged the settling of the debate, one alternative path forward is to identify novel experimental conditions in which the two models might diverge. Here we propose to focus on the case of multi-movement planning in which not one, but multiple movements need to be *simultaneously* prepared. Multi-movement planning poses an interesting challenge for the motor system, and can be used to expose diverging predictions of the two models. The representational theory indeed predicts that all potential movements should be concurrently represented in preparatory activity. By contrast, according to the dynamical system theory, if each movement is associated with its own initial condition (IC), preparing for all possible movements should be achieved by reaching an intermediate IC ideally located in-between the individual ICs.

In support of the representational view, prior studies on multi-movement planning suggest that concurrent movement representations co-exist in motor cortex (Bastian et al. 2003; Cisek and Kalaska 2005; Thura and Cisek 2014; Dekleva et al. 2016). These studies, however, typically involved movements that were associated with different spatial locations (i.e., hand reaches toward multiple possible targets) and therefore could not disambiguate motor plans from the visual representations of the movement goal (Cisek, 2012; Shen and Alexander, 1997; Wong and Haith, 2017). Moreover, the parametric nature of reaching movements used in these tasks may have indirectly contributed to biasing the results toward a representational view. In the present study, we developed a multi-movement planning task in which monkeys had to execute one of two possible grasping movements based on a non-spatial cue. Grasping movements, unlike reaching movements, are not represented on a continuum but rather correspond to discrete movement categories, making them more suitable to confront the predictions of the two models of motor planning. Overall, we found that our data were inconsistent with the representational theory, but supported instead an augmented view of the initial condition hypothesis extended to the case of multi-movement planning.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Results

Task and behavior

Two monkeys (L and N) performed an instructed delayed reach-to-grasp task (Brochier et al., 2018; Milekovic et al., 2015; Riehle et al., 2013). The animals were trained to grasp an object using 2 possible hand grips (side grip, SG, or precision grip, PG) and subsequently pull and hold the object using 2 possible force levels (low force, LF, or high force, HF). On each trial, the animal had to wait for two successive instructions separated by a 1-s delay before initiating their movement (**Figure 1A**; Methods). The grip and force instructions were displayed via a square of 4 light-emitting diodes (LEDs) as follows: the two leftmost (resp., rightmost) LEDs instructed SG (resp., PG), while the top (resp., bottom) LEDs instructed HF (resp., LF). There were two main conditions in this task. In the “grip-cued” condition, the grip instruction was provided first, followed by the force instruction. In the “grip-uncued” condition, the force instruction was provided first, followed by the grip instruction (**Figure 1B**). In both conditions, the second instruction also served as the imperative Go signal for the animal to initiate its movement; we therefore refer to the first instruction as the “Cue”, the second instruction as the “Go”, and the 1-s delay between them as the “preparatory period”. Note that the terminology used here to describe the two task conditions differs from previous studies (Brochier et al., 2018; Milekovic et al., 2015; Riehle et al., 2013) but is more relevant to the objectives of the present study.

In this task, when grip instruction is provided first (grip-cued), animals can plan the desired grip in advance during the preparatory period. In contrast, when grip instruction is provided last (grip-uncued), animals do not know the desired grip until the time of Go, and are therefore left uncertain about which of the two grips to plan. We used this task to study motor planning related to a single grip versus two simultaneous grips. We chose to focus on the grip, as opposed to the force, because it was the most relevant parameter to plan the initial phase of the movement, i.e., reaching toward and grasping the object. The force level, which was only relevant for the late phase of the movement, was used to match task contingencies (“Cue” signal revealing partial information, followed by the “Go” signal revealing full information) between the two conditions.

In sum, the task was composed of 2 conditions, with 4 trial types each (2 grips x 2 forces). The grip-cued condition was composed of PG-HF and PG-LF (hereafter collectively referred to as “PG-cued”), and SG-HF and SG-LF (“SG-cued”) trials, while the grip-uncued condition was composed of HF-PG and LF-PG (“PG-uncued”) and HF-SG and LF-SG (“SG-uncued”) trials. Note that the 4 trial types were identical across conditions in terms of final movement, but differed only in the order that the grip/force

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

information was provided. The 4 trial types were randomly interleaved within each condition, and the conditions were performed in separate blocks of trials within the same behavioral session (Methods).

At the end of training, animals were proficient in all conditions and trial types (average success rate: 92% for monkey L; 97% for monkey N; **Table S1**). These high success rates indicate that the monkeys were able to interpret and use the instruction cues to rapidly alternate between the desired actions. Critically, we confirmed that animals used the grip information (when available) to plan their movement ahead. Indeed, reaction times (RT; defined between the time of Go and movement initiation) were shorter in the grip-cued compared to the grip-uncued condition (RT = mean+sem; RT_{cued} = 155±3 ms, RT_{uncued} = 263±4 ms, $p<10^{-10}$ for monkey L; RT_{cued} = 303±5 ms, RT_{uncued} = 363±8 ms, $p<10^{-9}$ for monkey N; ANOVA on RT testing for the main effect of grip-cued versus grip-uncued; **Figure 1C**), and this effect was robustly observed across animals and behavioral sessions (paired *t*-test on across-session RT_{cued} versus RT_{uncued}, $t(55)=12.6$, $p<10^{-10}$; **Figure 1D**). These results are in line with numerous previous studies showing a beneficial effect of movement preparation on reaction time (Ames et al., 2014; Churchland and Shenoy, 2007a; Crammond and Kalaska, 2000; Ghez et al., 1997; Riehle, 2005; Riehle and Requin, 1989, 1993; Zaepffel and Brochier, 2012).

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

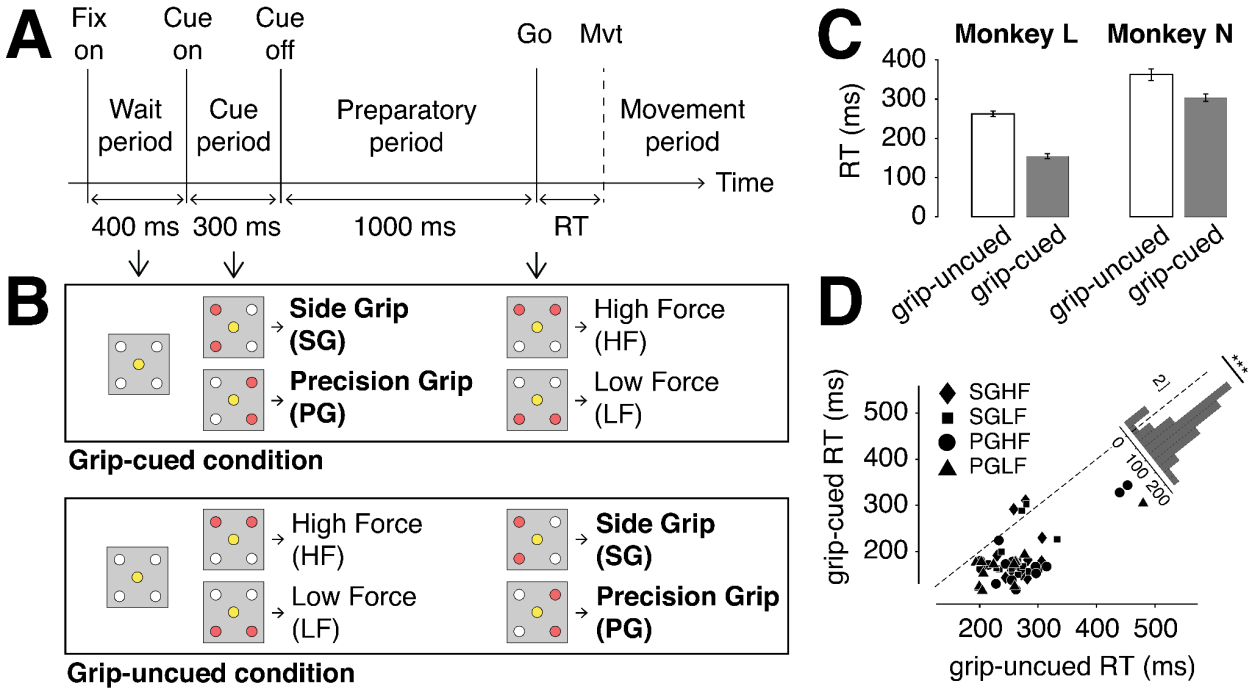

Figure 1. Task and behavior. (A) Trial structure. Each trial started with a 400-ms waiting period during which only the center LED was on. During the subsequent cue period, 2 of the 4 peripheral LEDs were illuminated for 300 ms to provide the first instruction. The cue was then turned off for 1000 ms during the preparatory period. At the end of the preparatory period, 2 other peripheral LEDs were illuminated to provide the second instruction and simultaneously signal the Go. By definition, the reaction time (RT) was defined as the time between movement initiation (vertical dashed line) and the Go signal (Methods). **(B)** Experimental conditions. In the grip-cued condition (top), the grip instruction (side grip, SG or precision grip, PG) was provided first, followed by the force instruction (high force, HF or low force, LF). In the grip-uncued condition (bottom), the force instruction was provided first, followed by the grip instruction. Instructions were provided as follows: the two leftmost (resp., rightmost) LEDs were used to instruct SG (resp., PG), and the two top (resp., bottom) LEDs were used to signal HF (resp., LF). **(C)** Reaction times for the two conditions in a typical session. In both monkeys, reaction times were shorter in the grip-cued (filled bars) compared to the grip-uncued condition (empty bars). Error bars represent 95% confidence intervals. **(D)** Average RT in the grip-cued condition plotted as a function of the average grip-uncued RT. Each point represents data from the same behavioral session. Different symbols (diamond, square, circle, triangle) indicate the trial type (resp., SGHF, SGLF, PGHF, PGLF). Inset: distribution of the difference between grip-uncued RT and grip-cued RT across sessions. This difference was significantly greater than zero (paired *t*-test, $t(55)=12.6$, $p<10^{-10}$).

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Neural encoding of grip information during single-grip planning

To investigate the neural basis of motor planning in this task, we analyzed spiking activity in the primary (M1) and dorsal premotor (PMd) cortex recorded from chronically-implanted multi-channel arrays (Brochier et al., 2018). Prior studies using this task have explored how the spatio-temporal structure of local field potentials and spiking activity in PMd/M1 relate to various aspects of animals' behavior (Denker et al., 2018; Milekovic et al., 2015; Riehle et al., 2013, 2018; Torre et al., 2016). None of these studies have raised the question of multi-grip planning. Here, we sought to directly compare preparatory activity in the grip-cued and grip-uncued conditions to confront the predictions of the two leading theories of motor planning. Accordingly, we focused our analyses on the task epoch related to the hand reaching preparation and execution.

Our epoch of interest included the 300-ms presentation of the Cue (providing grip information in grip-cued, and force information in grip-uncued), followed by the 1-s preparatory period preceding the Go signal, and the subsequent movement period. In this epoch, single neurons had widely varied activity profiles (**Figure 2A, red and blue lines**). In the grip-cued condition, some neurons encoded the grip type shortly after the Cue (**Figure 2A (iv)**), or right before the Go (**Figure 2A (ii)**), while others remained sensitive to the grip throughout the preparatory period (**Figure 2A (i, iii)**). By contrast, in the grip-uncued condition, grip encoding only emerged after the Go (i.e., when the grip information became available; **Figure S1**). These findings are in keeping with prior studies showing that PMd/M1 neurons are tuned to upcoming movement parameters (Churchland et al., 2006a; Even-Chen et al., 2019; Godschalk et al., 1985; Kurata, 1993; Messier and Kalaska, 2000; Riehle et al., 1994), i.e., here the grip type.

Since grip information is the relevant movement parameter, grip encoding may serve as a proxy for studying the formation of the two grip-specific motor plans. Accordingly, we quantified the strength and temporal evolution of grip encoding within a trial. To do so we turned to population-level analyses, by considering population activity as a collection of states (i.e., neural trajectory) evolving in a high-dimensional space where each dimension represents the activity of one neuron (Sohn et al., 2020; Vyas et al., 2020b). In this state space, we computed the distance between the trajectories associated with the two grip-cued conditions (**Figure 2B**, Methods). This distance sharply increased around 150 ms following the presentation of the Cue, and remained above its baseline throughout the preparatory period in both monkeys (**Figure 2C, 2D**). As a point of comparison, we computed the distance between the two force-specific trajectories in the grip-uncued condition. We found that force information was not as strongly encoded as grip information during the preparatory period (**Figure S1**). These results are consistent with motor cortex encoding the relevant grip information during the planning of the upcoming hand reach (Fluet et al., 2010; Milekovic et al., 2015).

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

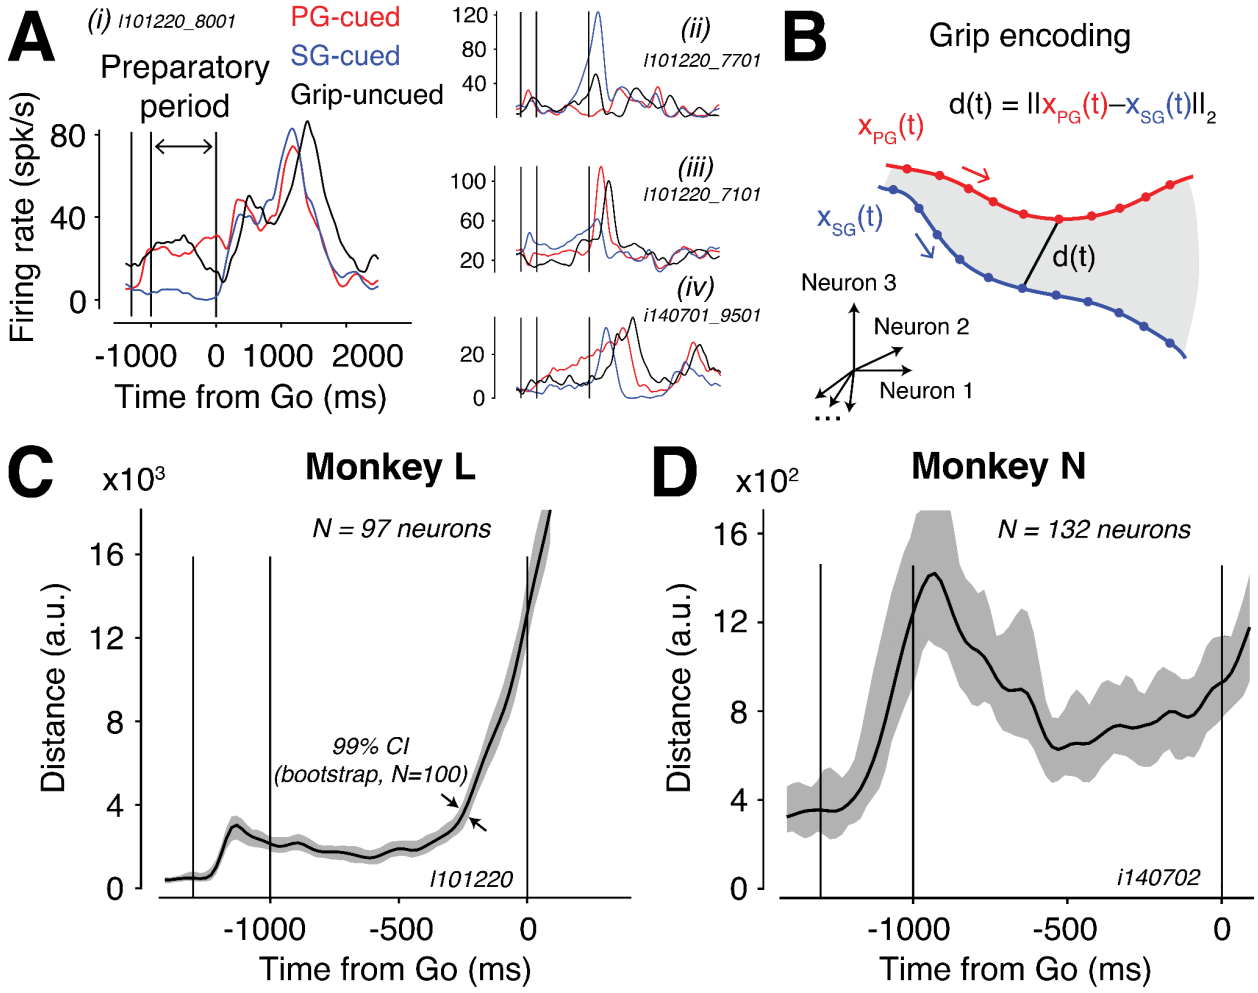

Figure 2. Neural encoding of grip information (A) Firing rate of four example neurons from the two monkeys aligned to the Go signal. Neural activity is color-coded by condition; red for PG-cued (here, PGHF), blue for SG-cued (SGHF), and black for grip-uncued (HFPG). Firing rates were obtained by binning (w_{bin} = 20 ms) and smoothing (sd_{kernel} = 40 ms) spike counts averaged across trials of the same condition. Vertical lines represent Cue on, Cue off, and Go. **(B)** Schematic of population trajectories in the state space. To quantify grip encoding, we computed the instantaneous euclidean distance between the PG-cued (red) and the SG-cued (blue) trajectory (Methods). **(C)** Distance between the trajectories as a function of time in the preparatory period (aligned to the Go signal). Shaded areas denote the 99% CI obtained from bootstrap resampling of trials (N=100). **(D)** Same as (C) for monkey N.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Two-grip planning is inconsistent with the representational model

The emergence of grip encoding in the grip-cued condition may be interpreted in two different ways. According to the representational view of motor planning (Bastian et al., 2003; Cisek, 2006; Cisek and Kalaska, 2005; Kalaska et al., 1997), grip encoding reflects the fact that preparatory neural activity gradually represents the upcoming hand movement parameter, which in turn expedites the process of initiating the grip-specific hand reach once the Go is provided. Alternatively, in the dynamical system view (Churchland et al., 2010; Vyas et al., 2020b), the separation between the two grip trajectories reflects the evolution toward two grip-specific initial conditions (ICs), which constitute the optimal states from which movement-related dynamics can be generated (Afshar et al., 2011; Ames et al., 2014; Hennequin et al., 2014; Kao et al., 2021).

Although the representational and dynamical system theories are both consistent with the emergence of grip encoding in the grip-cued condition, they make different predictions for the grip-uncued condition in which the animal has to prepare not one, but *two* potential grips simultaneously. In this case, the representational theory predicts that both grips should be concurrently represented in preparatory activity; that is, the two motor plans associated with each grip should coexist until the desired grip is revealed at the time of Go. By contrast, the dynamical system theory predicts that the grip-uncued preparatory activity should evolve toward an intermediate state optimized for both grips simultaneously and located in-between the two grip-cued initial conditions. In other words, a single initial condition optimized for the two movements should emerge.

We first sought to examine the predictions of the representational theory at the level of single neurons. Classically, the representational model has been tested in tasks where monkeys prepare movements toward several possible visual targets (Cisek and Kalaska, 2005; Dekleva et al., 2016, 2018; Rickert et al., 2009; Thura and Cisek, 2014). In these tasks, motor cortical neurons show strong directional tuning, i.e., they tend to fire maximally for a preferred target direction, and be nearly silent for opposite directions. In this case, when animals prepare movements toward multiple potential targets, the representational theory predicts the co-activation of distinct subsets of neurons tuned to each target direction (Cisek and Kalaska, 2005). In our task, we made the choice of using grip type rather than target direction as the relevant movement parameter to dissociate motor planning from the spatial goal representation. The resulting challenge is that, contrary to target direction, grip type is not as clearly separable at the single-neuron level: neurons are not selectively active for one grip, but instead respond to both grips with varying levels of activity (**Figure 2A**). We thus had to extend the predictions of the representational theory to this particular case. We reasoned that if both grips are represented in the grip-uncued condition, a subset of neurons should respond as in the PG-cued condition, while another subset of neurons

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

should respond as in the SG-cued condition ("concurrent representation" hypothesis; **Figure 3A, left**). Alternatively, if the two grips are not explicitly represented simultaneously, we should expect single neurons to exhibit activity levels in the grip-uncued condition that are unrelated to their responses during single-grip planning, i.e., in the PG-cued and SG-cued conditions ("null" hypothesis; **Figure 3A, middle**).

Visual inspection of single neurons provided initial evidence against the concurrent representation hypothesis: many of the rate profiles in the grip-uncued condition appeared to reach a value at the time of Go which differed from either of the two grip-cued conditions (**Figure 2A, black line**). To quantify this effect systematically across neurons, we computed for every grip-selective neuron (Methods) the ratio $\beta = \frac{r_{PG}^{cued} - r^{uncued}}{r_{PG}^{cued} - r_{SG}^{cued}}$. In this expression, r_{PG}^{cued} and r_{SG}^{cued} designate the firing rate in the PG-cued and SG-cued condition, respectively, and r^{uncued} designates the firing rate in the grip-uncued condition associated with either of the two force levels, which were tested separately. By definition, β was close to 0 if r^{uncued} was close to r_{PG}^{cued}, and β was close to 1 if r^{uncued} was close to r_{SG}^{cued}. If both grips are concurrently represented in the grip-uncued condition, we expected the distribution of β across neurons to be bimodal, with modes at 0 and 1, respectively reflecting the neural representation of PG and SG.

In line with our initial observations, we found that the distribution of β computed 200 ms before Go did not display the bimodal shape predicted by the concurrent representation hypothesis, but instead peaked at an intermediate value between 0 and 1 (**Figure 3A, top right**). As a validation of our analysis, we also computed β using neural activity *after* Go, i.e., when grip information was known. In that case, we expected and observed two peaks at 0 and 1, respectively associated with the PG-uncued and SG-uncued condition (**Figure 3A, bottom row**). These results indicate that, prior to Go, the population did not represent the two grips simultaneously, but instead presented a mixture of responses that reflected neither of the two grips.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

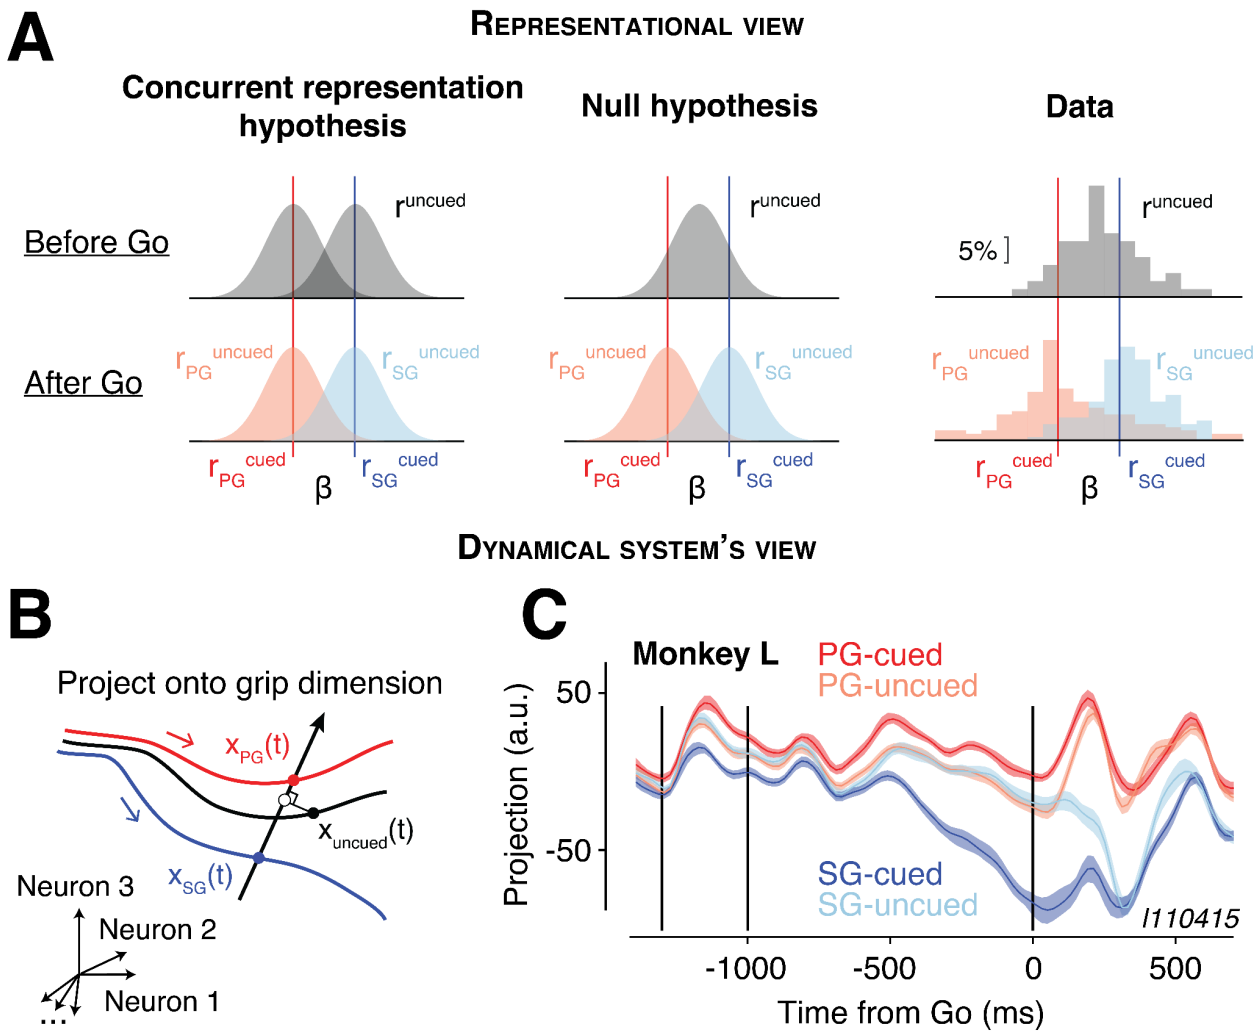

Figure 3. Neural dynamics in the grip-uncued condition. (A) Test of the representational theory. Top: for each grip-selective neuron, we computed the ratio $\beta = \frac{r_{PG}^{cued} - r^{uncued}}{r_{PG}^{cued} - r_{SG}^{cued}}$, where r_{PG}^{cued}, r_{SG}^{cued} and r^{uncued} designate the firing rate of the neuron 200 ms before Go respectively in the PG-cued, SG-cued, and grip-uncued condition. A grip selective neuron was defined as a neuron whose preparatory activity significantly differed between PG-cued and SG-cued (Methods). According to the "concurrent representation hypothesis" (left), if both grips are represented in the population, the distribution of β (black distribution) should be bimodal, with a peak at 0 (associated with the representation of PG, red vertical line) and a peak at 1 (associated with the representation of SG, blue vertical line). Alternatively, if firing rates in the grip-uncued condition are unrelated to either grip-cued condition ("null hypothesis", middle), the distribution of β should not present any particular bias toward either grip representation, and should have a single peak between 0 and 1. The empirical distribution of β ("data", right) rejected the concurrent representational hypothesis. Bottom: as a control, we verified that, 400 ms after Go, the distribution of β associated with the PG-uncued (red distribution) and SG-uncued (blue distribution) condition peaked respectively at 0 and 1, as predicted by both hypotheses. **(B)–(C)**

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Test of the dynamical system theory. **(B)** Schematic of the projection onto the grip dimension (black vector) in the state space. At any time point, the grip dimension was defined as the unit vector pointing from the SG-cued state ($x_{SG}(t)$, filled blue circle) to the PG-cued state ($x_{PG}(t)$, filled red circle). The schematic shows the orthogonal projection (open black circle) of the grip-uncued state ($x_{uncued}(t)$, filled black circle) onto the grip dimension. **(C)** Temporal evolution of the projections onto the grip dimension within the preparatory and movement epochs. The dark red and blue represent the projections of the PG-cued and SG-cued trajectories, respectively. The projection of the grip-uncued trajectories is shown separately for the PG-uncued (pale red) and the SG-uncued (pale blue) condition. Data from one session is shown; see Figure S2 for more sessions.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

The optimal subspace hypothesis: an augmented view of the initial condition hypothesis

Next, we tested the predictions of the alternative "initial condition" (IC) hypothesis based on the dynamical system view of motor planning. In the framework of dynamical systems (Vyas et al., 2020b), preparatory activity should not be interpreted as a parametric representation of the upcoming movement. Rather, it should be construed as a controlled dynamical process converging to an optimal state from which movement-related dynamics can efficiently unfold (Afshar et al., 2011). Although the IC hypothesis was originally introduced for single planned movements, one can extend it to the case of multi-movement planning. Considering that each movement is associated with its own initial condition (single-movement IC), we can formulate two predictions: the preparatory state associated with multiple alternatives should 1) lie within the subspace containing all the single-movement ICs, and 2) be located "in-between" the single-movement ICs to rapidly converge to one of them when the desired movement is finally prescribed. We refer to this augmented view of the IC hypothesis as the "optimal subspace hypothesis". Note that this term has previously been introduced in the context of single planned movements (Churchland et al., 2006b), but has a different meaning here, since it is used to extend the IC hypothesis to multi-movement planning.

To test the predictions of the optimal subspace hypothesis, we extended our previous population-level analyses. First, we defined the subspace that contained the preparatory states associated with each single grip. Since there were only two grips, this subspace was one-dimensional, and was defined by the unit vector connecting the SG-cued IC to the PG-cued IC in the state space (**Figure 3B**); we refer to this dimension as the "grip dimension". Note that because the two grip-cued trajectories did not remain parallel throughout the preparatory period (**Figure 2C,D**), the grip dimension was defined at every time point (Methods). We then projected the grip-uncued trajectory onto the grip dimension (**Figure 3B**). As a reference, we also projected the two grip-cued trajectories on the grip dimension to evaluate the proximity of the grip-uncued state relative to the two grip-cued states as a function of time in the trial.

Before Cue presentation, the preparatory states largely overlapped, which was expected since no grip information was available at that time. Following the Cue, the two grip-cued states rapidly diverged and remained separated throughout the preparatory and movement period, consistent with our previous distance analysis (**Figure 2C,D**). The grip-uncued state had a qualitatively similar temporal profile, but notably remained *in-between* the two grip-cued states until ~100 ms after Go (**Figure 3C**). The trajectory then separated from this intermediate state into two grip-uncued trajectories (PG-uncued and SG-uncued) which rapidly merged with the corresponding

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

grip-cued trajectories (PG-cued and SG-cued, respectively) about 200-250 ms after Go. The same observations were made in the two animals, and confirmed using both force levels of the grip-uncued condition (**Figure S2**). This analysis demonstrates that, before Go, the grip-uncued activity evolved toward an intermediate preparatory state along the dimension that contained the two grip-cued preparatory states, and rapidly converged to one of them once the grip information was provided.

Optimization of initial conditions explains inter-session variability

So far, our results are consistent with the optimal subspace hypothesis: preparatory activity in the grip-uncued condition reaches an intermediate state between the two grip-specific initial conditions. This intermediate state might result from an optimization process facilitating the execution of the two potential grips. To firmly establish this result, however, we ought to demonstrate a tighter relationship between the preparatory state and the animal's behavior. In particular, the exact position of the grip-uncued initial condition along the grip dimension should be predictive of animals' tendency to favor the preparation of one grip versus the other. For instance, if the grip-uncued IC is slightly closer to one of the two grip-cued ICs, say PG-cued (as can be seen in **Figure 3C**), then this should confer a slight reaction time benefit to PG relative to SG. In the following, we therefore sought to leverage the inter-session variability in the neural and behavioral data to test this stronger prediction of the optimal subspace hypothesis.

For each session, we computed a contrast measure (λ_{neural}) to quantify the proximity of the grip-uncued IC relative to the two grip-cued ICs along the grip dimension (Methods). By definition, λ_{neural} was bounded by -1 and 1, and was negative if the grip-uncued IC was closer to the PG-cued IC, and positive if the grip-uncued IC was closer to the SG-cued IC (**Figure 4A**). According to the optimal subspace hypothesis, λ_{neural} closer to -1 should confer a slight advantage to the planning of PG relative to SG (**Figure 4A, middle top**). That is, there should be a smaller RT difference between PG-uncued and PG-cued trials compared to SG-uncued and SG-cued trials. Conversely, λ_{neural} closer to +1 should confer a slight advantage to SG over PG (**Figure 4A, middle bottom**), with a smaller RT difference between SG-uncued and SG-cued compared to PG-uncued and PG-cued. To assess this effect in behavior, we computed a second contrast measure ($\lambda_{behavioral}$) which quantified RT differences between the grip-cued and grip-uncued conditions for each grip separately (Methods). By definition, $\lambda_{behavioral}$ was close to -1 if the RT difference between the cued and uncued condition was smaller for PG than SG, and +1 if the RT difference was smaller for SG than PG.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

When we plotted λ_{neural} against $\lambda_{behavioral}$, we found a strong positive correlation between the two contrast measures across sessions (Pearson correlation, ρ=0.78; $p<10^{-5}$; **Figure 4B**). This effect was specific to the grip dimension, since the correlation was abolished when we projected the preparatory states onto a random dimension (ρ=0.03; p=0.86). When we extended our correlation analysis using neural activity at various time points of the cue and preparatory periods, we found that the correlation between λ_{neural} and $\lambda_{behavioral}$ became positive around 100ms after the Cue presentation (**Figure 4C**). Moreover, the correlation was strongest when λ_{neural} was computed right at the time of Go , i.e., at the time most relevant for setting up the initial condition for the upcoming movement. Together, these results provide compelling evidence that the preparatory state in the grip-uncued condition was optimized to fall in-between the two grip-cued initial conditions, and that the extent to which this optimization favored one grip over the other was quantitatively reflected in the animals' behavior across sessions.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

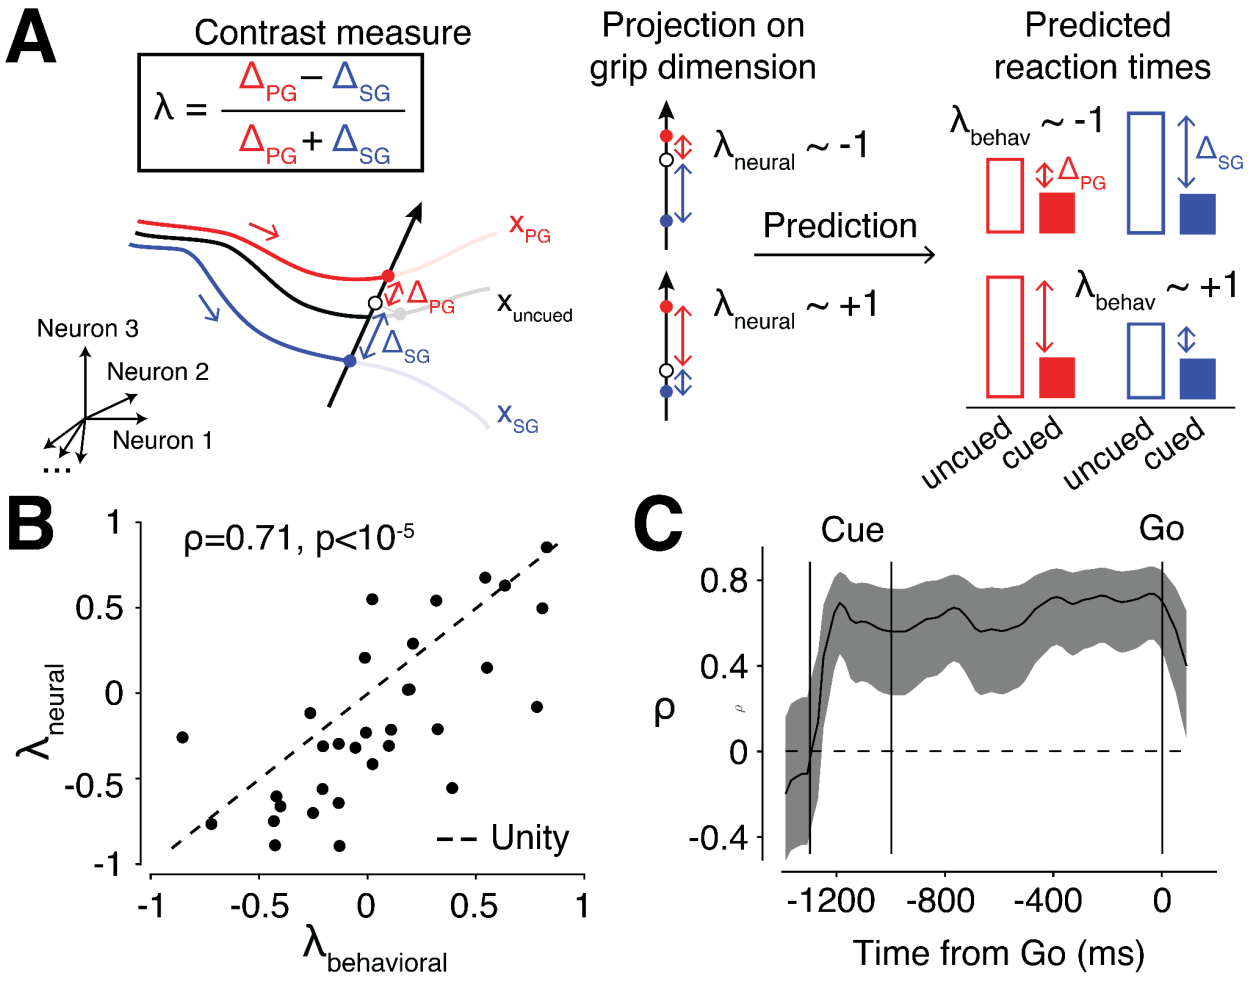

Figure 4. Predicting session-by-session variability. (A) Predictions of the optimal subspace hypothesis. Left: to evaluate the relative positions of the PG-cued, SG-cued and grip-uncued preparatory states along the grip dimension, we defined for each experimental session a contrast measure (λ_{neural}) as follows: $\lambda_{neural} = \frac{\Delta_{PG}-\Delta_{SG}}{\Delta_{PG}+\Delta_{SG}}$, where Δ_{PG} and Δ_{SG} represented the distance of the grip-uncued projection to the PG-cued and SG-cued state, respectively. Middle: λ_{neural} close to -1 (top) corresponded to the grip-uncued projection close to the PG-cued state; λ_{neural} close to +1 (bottom) corresponded to the grip-uncued projection close to the SG-cued state. Right: to assess the influence of the preparatory state on behavior, we defined another contrast measure ($\lambda_{behavioral}$) using the same expression as above, except Δ_{PG} and Δ_{SG} now represented the difference in reaction time between the grip-uncued and grip-cued condition for PG and SG, respectively. The optimal subspace hypothesis predicted that λ_{neural} and $\lambda_{behavioral}$ should correlate. **(B)** Correlation between λ_{neural} and $\lambda_{behavioral}$. Each dot represents one session. Sessions were pooled across animals. The dashed line shows the unity line. **(C)** Correlation between λ_{neural} and $\lambda_{behavioral}$ as a function of time in the preparatory period.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

The optimal subspace hypothesis holds at the single-trial level

Our previous analysis was based on trial-averaged activity, and did not take into account the variability that occurred across trials of a given session. As a final test of the optimal subspace hypothesis, we sought to verify its predictions down at the single-trial level. Let us consider a particular trial in which the projection of the grip-uncued preparatory state along the grip dimension is slightly biased toward the PG-cued IC (**Figure 5A**). Since we (arbitrarily) defined the grip dimension as pointing from the SG-cued IC to the PG-cued IC, this bias corresponds to a *large* projection onto the grip dimension. According to the optimal subspace hypothesis, this bias should provide a "head-start" to the execution of PG, and thus have a *beneficial* effect on the reaction time when the grip type revealed at the time of Go is PG, i.e., in the *PG-uncued* condition. As a result, the optimal subspace hypothesis predicts a *negative* correlation between trial-by-trial projections and RTs in the PG-uncued condition, i.e., *larger* projections onto the grip dimension lead to *shorter* reaction times (**Figure 5B, top**). The hypothesis predicts the opposite effect in the SG-uncued condition: there should be a *positive* correlation between trial-by-trial projections and RTs, i.e., *larger* projections onto the grip dimension lead to *longer* reaction times (**Figure 5B, bottom**).

To test these predictions, we estimated single-trial neural dynamics using a 200-ms sliding window, and projected individual trajectories onto the grip dimension (Methods). Because estimating single-trial activity is prone to noise, we focused on a high-yield session with a large number (n=97) of simultaneously recorded neurons, and a large number (n=141) of successful trials. Trial-by-trial preparatory activity appeared to be highly variable along the grip dimension, particularly near the time of Go (**Figure S3**). The amount of trial-by-trial variance accounted for by the grip dimension was approximately 10 times larger than expected by chance (**Figure S3**). For each individual trial, we plotted the neural projection immediately prior to Go as a function of the reaction time, separately for the PG-uncued and SG-uncued conditions (**Figure 5C**). As predicted by the optimal subspace hypothesis, we found a negative correlation in the PG-uncued condition (Pearson correlation, ρ_{PG}=-0.55; $p<10^{-5}$), and a positive correlation in the SG-uncued condition (ρ_{SG}=0.49; $p<10^{-4}$). This result held across recorded sessions, although the effect was weaker likely due to the lower number of neurons in these sessions (one-sided *t*-test on ρ across sessions, mean+sem, ρ_{PG}=-0.37$\pm$0.06; $t(31)$=-5.97, $p<10^{-6}$; ρ_{SG}=0.09$\pm$0.05, $t(31)$=1.73, $p<0.05$; **Figure 5C, inset**). Moreover, the same patterns of correlation between RT and trial-by-trial projections were found for the grip-cued conditions (**Figure S3**), reinforcing the idea that the grip dimension was key in controlling the initial condition for planning the movements. Altogether, these results confirm the tight relationship predicted by the optimal subspace hypothesis between deviations of the preparatory state along the grip dimension and animals' reaction times down at the single-trial level.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

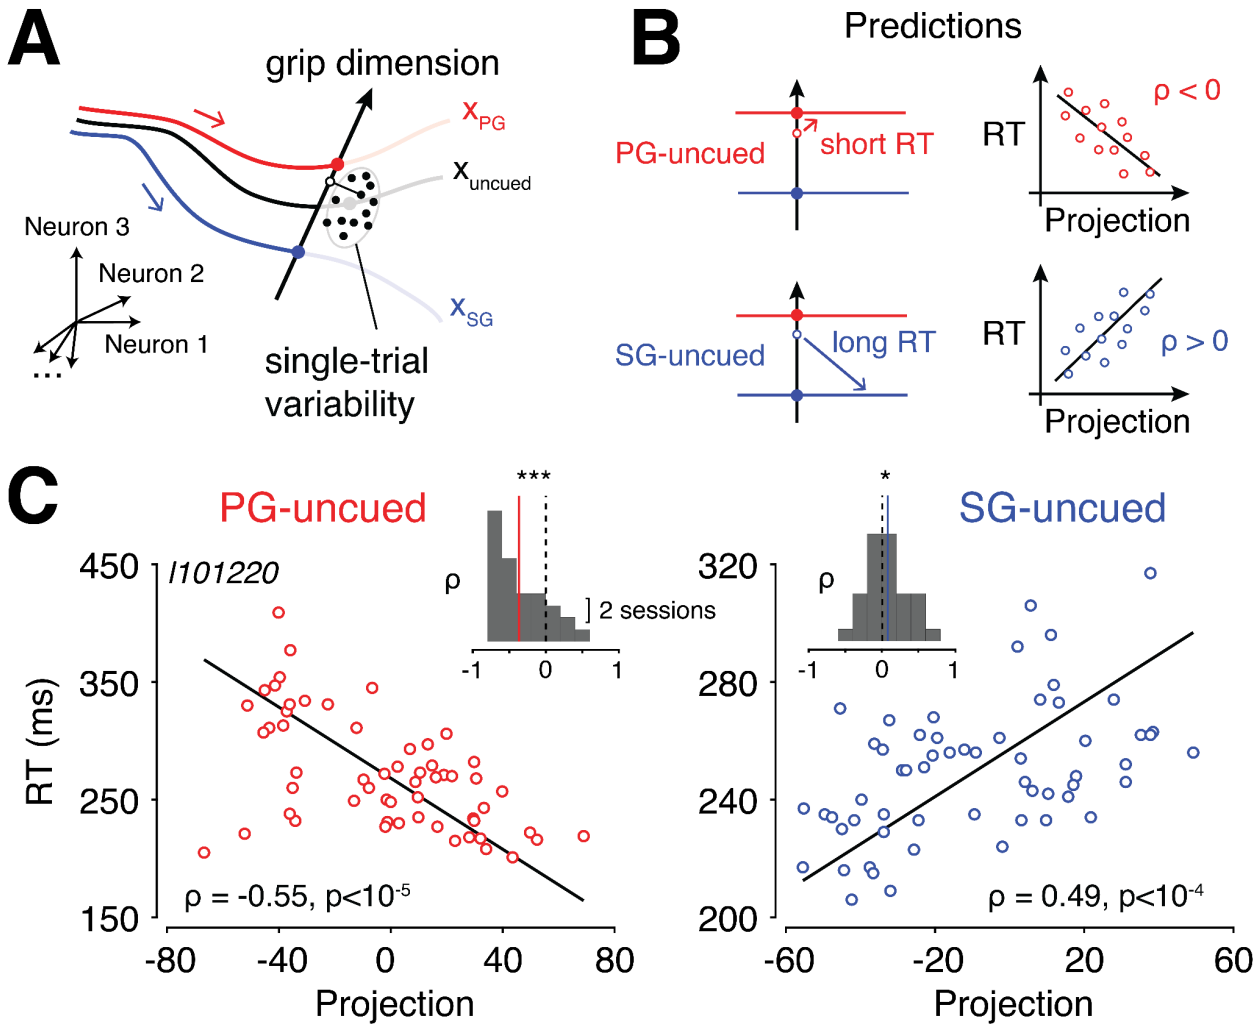

Figure 5. Predicting trial-by-trial variability. (A) Schematic of single-trial projections onto the grip dimension. Large filled circles represent trial-averaged states, while smaller circles represent individual trials. **(B)** Single-trial predictions of the optimal subspace hypothesis. One particular grip-uncued trial biased toward the PG-cued state is shown. If the grip revealed at Go is PG (i.e., PG-uncued condition, in red at the top), the reaction time (RT) for this trial is expected to be short. This predicts a negative correlation between RT and the projection (right). In contrast, if the grip revealed at Go is SG (i.e., SG-uncued condition, in blue at the bottom), the reaction time for this trial is expected to be long. This predicts a positive correlation between RT and the projection. **(C)** Correlation between RT and single-trial projections for the PG-uncued (left) and SG-uncued (right) condition in one high-yield session. Inset: Summary of correlation values across sessions (*p<0.05, ***p<0.001).

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Discussion

In this study, we analyzed preparatory neural activity in the motor cortex of monkeys planning two movements in parallel. Contrary to the predictions of existing models based on the representational view of motor planning (Cisek, 2006; Cisek and Kalaska, 2005; Crammond and Kalaska, 2000), we did not find evidence for concurrent neural representations of the different possible movements. Instead, our results supported an alternative model based on an augmented view of the "initial condition hypothesis" inspired by the theory of dynamical systems (Afshar et al., 2011; Churchland et al., 2006b, 2010; Shenoy et al., 2013). We found that the preparatory state associated with two possible movements lies within an "optimal subspace" containing the preparatory states associated with each movement planned separately. This intermediate state serves as the optimal initial condition to rapidly reconfigure motor cortical dynamics for the execution of either one of the two movements.

A number of previous studies on multi-movement planning have reported concurrent neural representations of motor plans (Bastian et al., 2003; Cisek and Kalaska, 2005; Coallier et al., 2015; Dekleva et al., 2016; Klaes et al., 2011; Thura and Cisek, 2014). These studies typically involved visually-guided movements that were associated with different spatial locations (i.e., hand reaches toward multiple potential targets). It is therefore possible that preparatory activity observed in these tasks reflected a spatial/directional tuning to the targets (Cisek, 2012; Shen and Alexander, 1997), which may have contributed to biasing the results toward a representational view (Cisek, 2006, 2007). Other studies employing a task similar to ours did find evidence for the co-activation of neurons tuned to different grip types (Baumann et al., 2009; Fluet et al., 2010). This result is not necessarily at odds with ours. Indeed, we did find that grip-selective neurons were active in the preparatory period of the grip-uncued condition. However, we showed that the activity level reached by grip-selective neurons did not reflect either one of the two levels associated with the grip-cued condition.

Our neural findings shed new light on a large body of behavioral studies showing an "averaging effect" of pre-planned movements when faced with two possible alternatives (Arai et al., 2004; Chapman et al., 2010; Chou et al., 1999; Gallivan and Chapman, 2014; Ghez et al., 1997; Hudson et al., 2007; Stewart et al., 2014). This effect has typically been attributed to two concurrent motor plans competing and blending during movement execution (Cisek, 2007; Stewart et al., 2013). Our results offer a different explanation: rather than representing the two motor plans in parallel, the motor cortex selects a single preparatory state which achieves a tradeoff between the two possible movements and optimizes task performance (Alhussein and Smith, 2021; Gallivan et al., 2015; Haith et al., 2015; Wong and Haith, 2017). In the case of reaching

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

movements, because target directions are parametrically organized in space, this intermediate state may naturally lead to a reach aimed at the average of the two potential targets. However, the fact that an intermediate state also emerges for non-parametric and categorically-distinct grasping movements suggests that this state truly reflects the optimization of an initial condition, rather than an intermediate motor plan per se.

To our knowledge, this study constitutes the first validation of the initial condition hypothesis in the context of multi-movement planning. Previous studies were mostly restricted to single planned movements (Afshar et al., 2011; Elsayed et al., 2016; Even-Chen et al., 2019; Kaufman et al., 2014; Vyas et al., 2018), or multiple movements planned in rapid sequence (Ames et al., 2014, 2019; Zimnik and Churchland, 2021). By generalizing the notion of initial condition to movement preparation under uncertainty, our results reinforce the idea that motor planning can and should be seen as a dynamical process optimized to generate appropriate neural dynamics for movement execution (Churchland and Shenoy, 2007a; Kao et al., 2021). Although the anatomical substrate underlying this optimization process is beyond the scope of our study, the thalamo-basal ganglia-cortical loop is a natural candidate (Athalye et al., 2020; Kao et al., 2021). Low-dimensional inputs to the motor cortex (Dubreuil et al., 2022; Logiaco et al., 2021; Sauerbrei et al., 2020) could for instance serve to adjust the initial state within the optimal subspace dimensions (Beiran et al., 2021; Sohn et al., 2020). Further investigations will be needed to elucidate this circuit-level question.

Our study investigated motor planning associated with only two movements. While we cannot claim that our results will generalize, we can formulate a testable prediction for cases with more alternatives. For N possible movements, we predict that preparatory activity will be located in the state space so as to minimize the distance to the N initial conditions corresponding to each movement. This optimization may be facilitated if the movements are naturally organized along a parametric continuum (e.g., reaching directions); in this case, the appropriate IC may correspond to one of the existing ICs located in-between the others. For non-parametric movements, it may be more challenging for motor cortex to find the appropriate initial condition, which would result in larger variability in the position of the intermediate IC (as can be seen from our contrast measure analysis, Figure 4B). Future studies could also test whether one movement being more likely than the others biases the location of the preparatory state toward the associated initial condition (Dekleva et al., 2018), or more generally, if each initial condition is weighted by the probability that each movement will be executed.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Methods

Experimental procedures

All procedures were approved by the local ethical committee (C2EA 71; authorization A1/10/12) and conformed to the European and French government regulations. Experiments involved two naive, awake, behaving monkeys (species: *Macaca mulatta*; ID: L and N; sex: female and male; weight: 6 and 7 kg; age: 6 years old). Monkeys were implanted with a 100-channel chronic Utah array (Blackrock Microsystems, Salt Lake City, UT, USA) in the motor cortex contralateral to the working hand (right hemisphere for both monkeys). The exact array location can be found in (Brochier et al., 2018). Data were recorded using the 128-channel Cerebus acquisition system (Blackrock Microsystems, Salt Lake City, UT, USA). Analysis of behavioral and neural data was performed using MATLAB (Mathworks, MA).

Behavioral task

Two monkeys performed an instructed delayed reach-to-grasp-and-pull task previously described in (Riehle et al., 2013). Briefly, the animals were trained to grasp an object (stainless steel parallelepiped, 40 mm x 16 mm x 10 mm, angled at 45° from the vertical and located about 20 cm away from the monkey) using 2 possible hand grips (side grip, SG, and precision grip, PG) and subsequently pull and hold the object using 2 possible force levels (low force, LF, and high force, HF). A 10 mm x 10 mm square of 4 red light-emitting diodes (LEDs) and one yellow LED at the center was used to display the two visual cues (C1 and C2) that served as task instructions. The instructions were coded as follows: the bottom two LEDs coded for LF, the top two for HF, the leftmost two for SG, and the rightmost two for PG. By definition, PG was obtained by placing the tips of the index and the thumb in a groove on the upper and lower sides of the object, respectively, and SG by placing the tip of the thumb and the lateral surface of the index on the right and left sides, respectively. An electromagnet placed inside the apparatus was used to change the object's effective weight (100g or 200g) to require a LF (magnet off) or HF (magnet on), respectively.

Trial structure

The structure of a trial was as follows: the animal started from a home position with their working hand pressing down on a pressure-sensitive switch. After a fixed 400-ms delay, the central LED was illuminated to indicate the start of a new trial. After another 400 ms, the first instruction (C1) was presented for 300 ms, followed by a 1-s preparatory period with only the central LED on. At the end of the preparatory period, the second

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

instruction (C2) was presented and also served as the imperative GO signal. At that point, the animal needed to (1) release the switch, (2) reach for the object with the appropriate grip and (3) pull and hold it with the appropriate force for 500 ms. The animal was subsequently rewarded (mixture of apple sauce and water) if both the grip type and the force level used were correct. To initiate a new trial, the monkey had to return their working hand to the home position and press the switch.

Experimental conditions

The task included 4 trial types (2 grips x 2 forces), namely SG-LF, SG-HF, PG-LF and PG-HF, which were randomly interleaved across trials. To manipulate movement preparation in this task, we varied the order in which the grip and force information were provided. Specifically, in the “grip-cued” condition, C1 provided the grip instruction, while C2 provided the force instruction. Conversely, in the “grip-uncued” condition, C1 provided the force instruction, while C2 provided the grip instruction. As a result, grip-cued and grip-uncued trials were identical in terms of movement, but differed in the order that the grip/force information necessary to plan the movement was given. Animals performed the task in short sessions (blocks of uninterrupted trials) alternating between the two conditions.

Neural recordings

Spiking data were recorded during multiple behavioral sessions (n=8 for monkey N, n=24 for monkey L). Each session had on average n=89 simultaneously recorded neurons, and n=142 successful trials. Spikes were sorted offline using Offline Spike Sorter, version 3, Plexon Inc., Dallas, TX, USA). Spike clusters which were separated significantly from each other and with less than 1% of inter-spike intervals (ISIs) of 2 ms and less were considered as single units (single-unit activity, SUA), whereas less well separated clusters and/or more than 1% of 2 ms ISIs were considered as multi-unit (multi-unit activity, MUA) recordings. In all analyses, we included both SUA and MUA.

Analysis of neural activity

To plot the response profile of individual neurons (Figures 2A), we smoothed averaged spike counts in 20-ms bins using a Gaussian kernel with a standard deviation of 40 ms.

Euclidean distance analysis

To quantify the strength of grip encoding, we computed the instantaneous euclidean distance between the neural trajectories associated with the PG-cued and SG-cued conditions (Figure 2C–D).

Mathematically, the distance was computed as follows:

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

$$d(t) = ||x^{PG}(t) - x^{SG}(t)||_2 = \sqrt{\sum_{i=1}^{N} (r_i^{PG}(t) - r_i^{SG}(t))^2}$$

where $r_i^{PG}(t)$ and $r_i^{SG}(t)$ represent the instantaneous firing rates of neuron i, and $x^{PG}(t)$ and $x^{SG}(t)$ represent the instantaneous population state in the PG-cued and SG-cued condition, respectively.

Neural tuning analysis

To quantify the relationship between firing rates in the grip-cued and grip-uncued conditions, we computed the index $\beta = \frac{r_{PG}^{cued} - r^{uncued}}{r_{PG}^{cued} - r_{SG}^{cued}}$. In this expression, r_{PG}^{cued} and r_{SG}^{cued} designate the firing rate in the PG-cued and SG-cued condition, respectively, and r^{uncued} designates the firing rate in the grip-uncued condition associated with either of the two force levels (HF or LF). This ratio was computed only for grip-selective neurons, i.e., neurons whose firing rates differed significantly (as assessed by non-overlapping 95% confidence intervals obtained via standard bootstrapping, N=100 repeats) between PG-cued and SG-cued conditions. This ensured that the denominator $r_{PG}^{cued} - r_{SG}^{cued}$ did not go to zero.

We computed the β index using activity computed in a 20-ms bin centered at two different time points of the trial: 200 ms before Go, and 400 ms after Go (Figure 3A). Before Go, we treated the two force conditions separately, and pooled them to plot the distribution of across β neurons. After Go, we treated the PG-cued and SG-cued separately.

Projection onto the grip dimension

To evaluate the relationship between the population state in the grip-uncued condition relative to the two grip-cued conditions, we defined the “grip dimension” which separated the PG-cued and SG-cued trajectories. The grip dimension was defined at each time point as the unit vector connecting the SG-cued (trial-averaged) state to the PG-cued (trial-averaged) state. We then projected the neural state associated with each condition (PG-cued, SG-cued and grip-uncued) onto the grip dimension to assess the proximity of the grip-uncued state to the other two states (Figure 3B). To obtain confidence intervals, we used a standard bootstrapping (resampling trials with replacement), with 100 repeats. Similarly, we projected session-by-session (Figure 4) and trial-by-trial (Figure 5) states onto the grip dimension to correlate with animals’ reaction times.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Supplementary material

	SG-HF	SG-LF	HF-SG	HF-PG	LF-SG	LF-PG	PG-HF	PG-LF
Monkey L	99%	95%	96%	93%	75%	92%	97%	94%
Monkey N	99%	100%	78%	100%	99%	97%	100%	100%

Table S1. Success rates for individual conditions for the two animals (n=24 sessions for monkey L, 8 for monkey N). Success rates are averaged across sessions.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

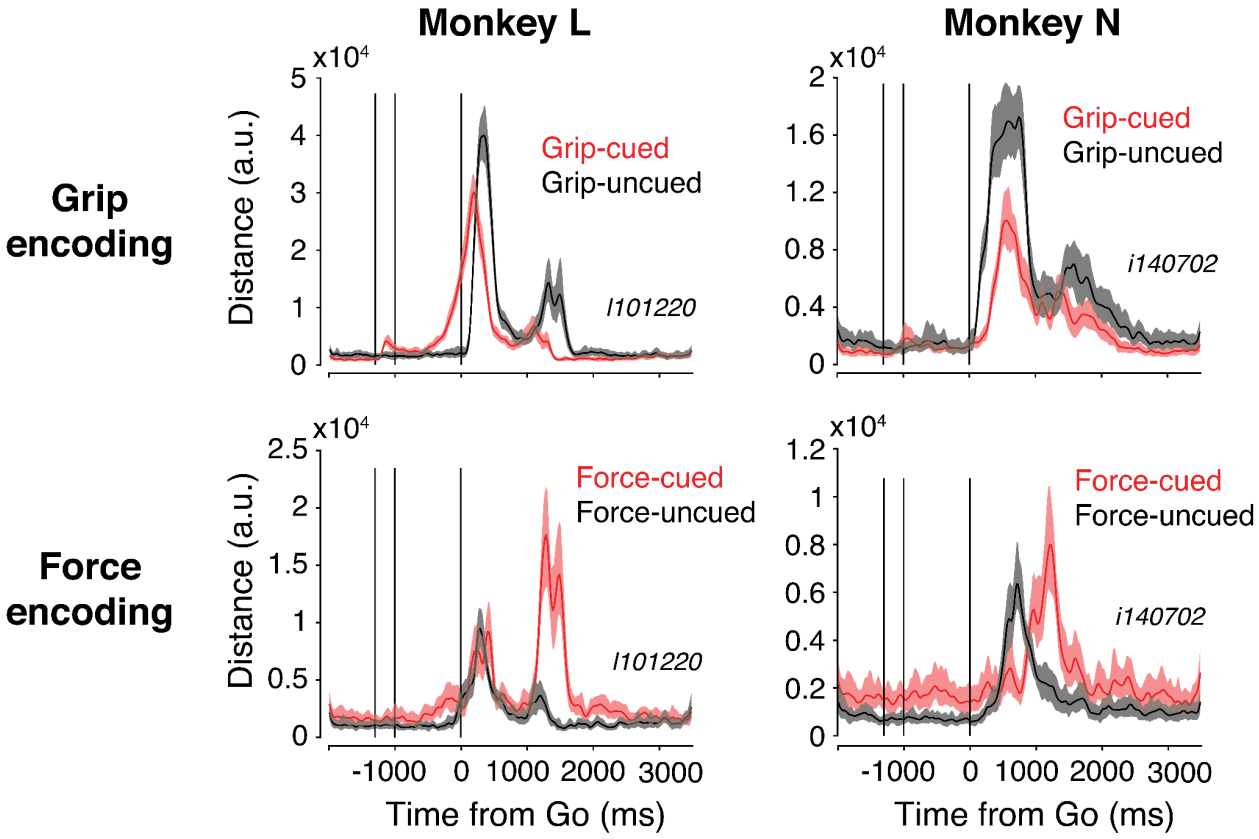

Figure S1. Related to Figure 2C. Comparison of grip and force encoding across the two monkeys (left for monkey L, right for monkey N). Top: we computed the instantaneous euclidean distance between the two grip-specific trajectories for the grip-cued (red) and grip-uncued (black) condition. Bottom: we computed the instantaneous euclidean distance between the two force-specific trajectories for the force-cued (i.e., grip-uncued, in red) and force-uncued (i.e., grip-cued, in black) condition.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

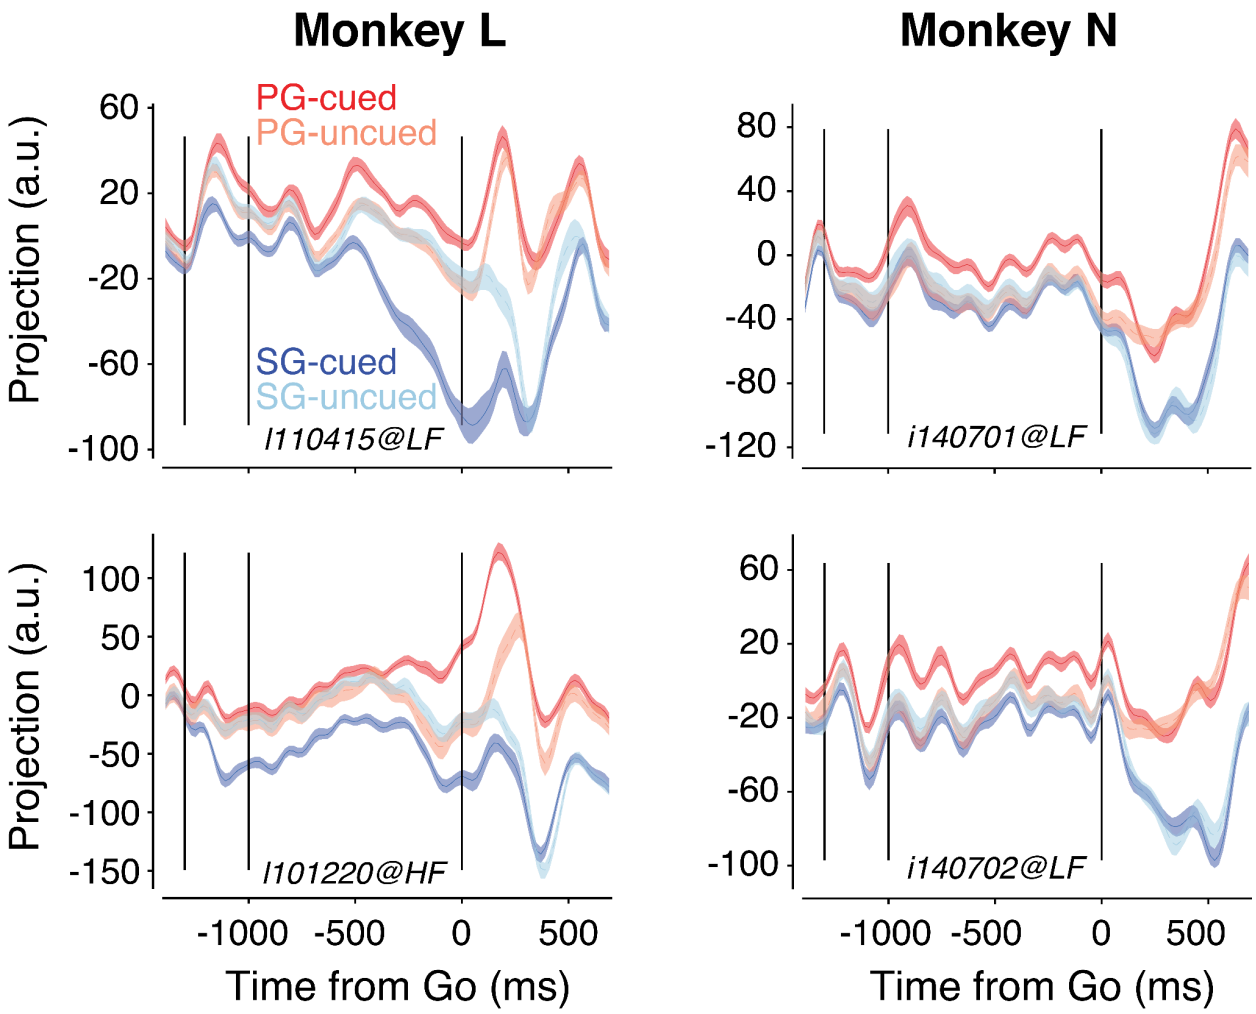

Figure S2. Related to Figure 3C. Projection of the neural trajectories associated with the PG-cued (dark red), SG-cued (dark blue), PG-uncued (light red), and SG-uncued (light blue) condition onto the grip dimension for different sessions and animals. For the PG-uncued and SG-uncued conditions, we fixed the level of force (HF for bottom left panel, LF for other panels).

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

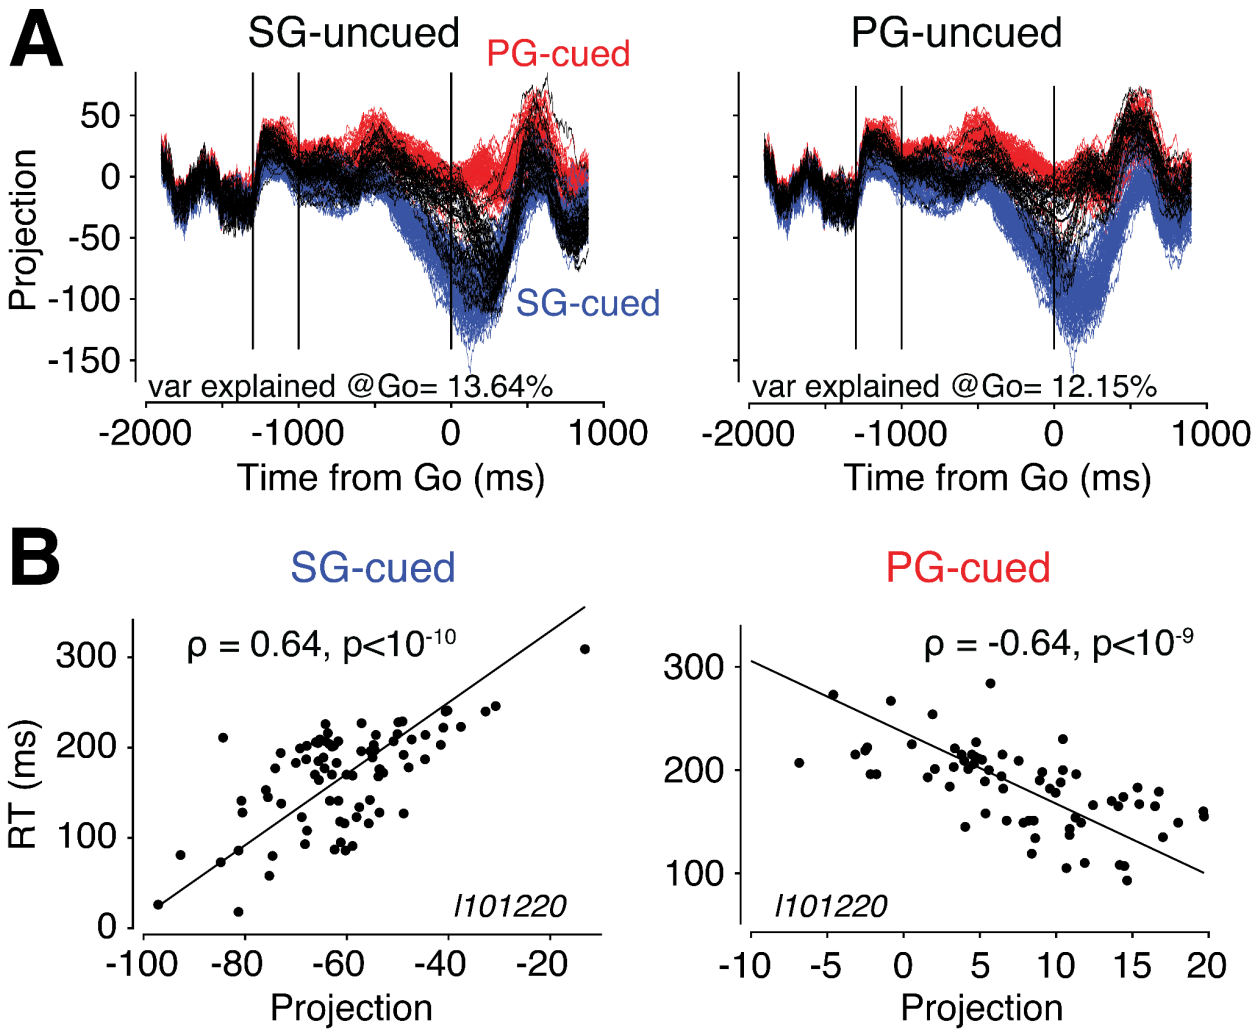

Figure S3. Related to Figure 5C. (A) Single-trial projections of SG-uncued (left) and PG-uncued onto the grip dimension. Each black line represents a single trial of the grip-uncued condition (fixed at HF); red and blue lines represent single trials of the PG-cued and SG-cued condition. Single-trial activity was computed using a 200-ms sliding window to bin the spikes. The amount of variance across trials at the time of Go (200-ms window preceding Go) along the grip dimension is indicated at the bottom of each panel. By comparison, the amount of variance along a random direction was around 1%. (B) Correlation between reaction time and single-trial projections onto the grip dimension right before Go for the SG-cued (left) and PG-cued (right) condition.

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

References

Afshar, A., Santhanam, G., Yu, B.M., Ryu, S.I., Sahani, M., and Shenoy, K.V. (2011). Single-Trial Neural Correlates of Arm Movement Preparation. Neuron *71*, 555–564. .

Alhussein, L., and Smith, M.A. (2021). Motor planning under uncertainty. Elife *10*. https://doi.org/10.7554/eLife.67019.

Ames, K.C., Ryu, S.I., and Shenoy, K.V. (2014). Neural Dynamics of Reaching following Incorrect or Absent Motor Preparation. Neuron *81*, 438–451. .

Ames, K.C., Ryu, S.I., and Shenoy, K.V. (2019). Simultaneous motor preparation and execution in a last-moment reach correction task. Nat. Commun. *10*, 2718. .

Arai, K., McPeek, R.M., and Keller, E.L. (2004). Properties of saccadic responses in monkey when multiple competing visual stimuli are present. J. Neurophysiol. *91*, 890–900. .

Athalye, V.R., Carmena, J.M., and Costa, R.M. (2020). Neural reinforcement: re-entering and refining neural dynamics leading to desirable outcomes. Curr. Opin. Neurobiol. *60*, 145–154. .

Bastian, A., Schöner, G., and Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. Eur. J. Neurosci. *18*, 2047–2058. .

Batista, A.P., Santhanam, G., Yu, B.M., Ryu, S.I., Afshar, A., and Shenoy, K.V. (2007). Reference frames for reach planning in macaque dorsal premotor cortex. J. Neurophysiol. *98*, 966–983. .

Baumann, M.A., Fluet, M.-C., and Scherberger, H. (2009). Context-specific grasp movement representation in the macaque anterior intraparietal area. J. Neurosci. *29*, 6436–6448. .

Beiran, M., Meirhaeghe, N., Sohn, H., Jazayeri, M., and Ostojic, S. (2021). Parametric control of flexible timing through low-dimensional neural manifolds.

Brochier, T., Zehl, L., Hao, Y., Duret, M., Sprenger, J., Denker, M., Grün, S., and Riehle, A. (2018). Massively parallel recordings in macaque motor cortex during an instructed delayed reach-to-grasp task. Sci Data *5*, 180055. .

Chapman, C.S., Gallivan, J.P., Wood, D.K., Milne, J.L., Culham, J.C., and Goodale, M.A. (2010). Reaching for the unknown: Multiple target encoding and real-time decision-making in a rapid reach task. Cognition *116*, 168–176. .

Chou, I.H., Sommer, M.A., and Schiller, P.H. (1999). Express averaging saccades in monkeys. Vision Res. *39*, 4200–4216. .

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Churchland, M.M., and Shenoy, K.V. (2007a). Delay of movement caused by disruption of cortical preparatory activity. J. Neurophysiol. *97*, 348–359. .

Churchland, M.M., and Shenoy, K.V. (2007b). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J. Neurophysiol. *97*, 4235–4257. .

Churchland, M.M., Santhanam, G., and Shenoy, K.V. (2006a). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J. Neurophysiol. *96*, 3130–3146. .

Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., and Shenoy, K.V. (2006b). Neural variability in premotor cortex provides a signature of motor preparation. J. Neurosci. *26*, 3697–3712. .

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

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51. .

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: a computational model. J. Neurosci. *26*, 9761–9770. .

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. Philos. Trans. R. Soc. Lond. B Biol. Sci. *362*, 1585–1599. .

Cisek, P. (2012). Making decisions through a distributed consensus. Curr. Opin. Neurobiol. *22*, 927–936. .

Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron *45*, 801–814. .

Coallier, É., Michelet, T., and Kalaska, J.F. (2015). Dorsal premotor cortex: neural correlates of reach target decisions based on a color-location matching rule and conflicting sensory evidence. J. Neurophysiol. *113*, 3543–3573. .

Crammond, D.J., and Kalaska, J.F. (2000). Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J. Neurophysiol. *84*, 986–1005. .

Day, B.L., Rothwell, J.C., Thompson, P.D., Maertens de Noordhout, A., Nakashima, K., Shannon, K., and Marsden, C.D. (1989). Delay in the execution of voluntary movement by electrical or magnetic brain stimulation in intact man. Evidence for the storage of motor programs in the brain. Brain *112 (Pt 3)*, 649–663. .

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Dekleva, B.M., Ramkumar, P., Wanda, P.A., Kording, K.P., and Miller, L.E. (2016). Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. Elife *5*. https://doi.org/10.7554/eLife.14316.

Dekleva, B.M., Kording, K.P., and Miller, L.E. (2018). Single reach plans in dorsal premotor cortex during a two-target task. Nat. Commun. *9*, 3556. .

Denker, M., Zehl, L., Kilavik, B.E., Diesmann, M., Brochier, T., Riehle, A., and Grün, S. (2018). LFP beta amplitude is linked to mesoscopic spatio-temporal phase patterns. Sci. Rep. *8*, 5200. .

Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F., and Ostojic, S. (2022). The role of population structure in computations through neural dynamics. Nat. Neurosci. 1–12. .

Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. Nat. Commun. *7*, 13239. .

Erlhagen, W., and Schöner, G. (2002). Dynamic field theory of movement preparation. Psychol. Rev. *109*, 545–572. .

Even-Chen, N., Sheffer, B., Vyas, S., Ryu, S.I., and Shenoy, K.V. (2019). Structure and variability of delay activity in premotor cortex. PLoS Comput. Biol. *15*, e1006808. .

Fetz, E.E. (1992). Are movement parameters recognizably coded in the activity of single neurons? Behav. Brain Sci. *15*, 679–690. .

Fluet, M.-C., Baumann, M.A., and Scherberger, H. (2010). Context-specific grasp movement representation in macaque ventral premotor cortex. J. Neurosci. *30*, 15175–15184. .

Gallivan, J.P., and Chapman, C.S. (2014). Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. Front. Neurosci. *8*, 215. .

Gallivan, J.P., Barton, K.S., Chapman, C.S., Wolpert, D.M., and Flanagan, J.R. (2015). Action plan co-optimization reveals the parallel encoding of competing reach movements. Nat. Commun. *6*, 7428. .

Ghez, C., Favilla, M., Ghilardi, M.F., Gordon, J., Bermejo, R., and Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. Exp. Brain Res. *115*, 217–233. .

Godschalk, M., Lemon, R.N., Kuypers, H.G., and van der Steen, J. (1985). The involvement of monkey premotor cortex neurones in preparation of visually cued arm movements. Behav. Brain Res. *18*, 143–157. .

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Golub, M.D., Sadtler, P.T., Oby, E.R., Quick, K.M., Ryu, S.I., Tyler-Kabara, E.C., Batista, A.P., Chase, S.M., and Yu, B.M. (2018). Learning by neural reassociation. Nat. Neurosci. *21*, 607–616. .

Haith, A.M., Huberdeau, D.M., and Krakauer, J.W. (2015). Hedging your bets: intermediate movements as optimal behavior in the context of an incomplete decision. PLoS Comput. Biol. *11*, e1004171. .

Hatsopoulos, N.G., Xu, Q., and Amit, Y. (2007). Encoding of movement fragments in the motor cortex. J. Neurosci. *27*, 5105–5114. .

Hennequin, G., Vogels, T.P., and Gerstner, W. (2014). Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements. Neuron *82*, 1394–1406. .

Hudson, T.E., Maloney, L.T., and Landy, M.S. (2007). Movement planning with probabilistic target information. J. Neurophysiol. *98*, 3034–3046. .

Kalaska, J.F., Scott, S.H., Cisek, P., and Sergio, L.E. (1997). Cortical control of reaching movements. Curr. Opin. Neurobiol. *7*, 849–859. .

Kao, T.-C., Sadabadi, M.S., and Hennequin, G. (2021). Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. Neuron *109*, 1567–1581.e12. .

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

Klaes, C., Westendorff, S., Chakrabarti, S., and Gail, A. (2011). Choosing goals, not rules: deciding among rule-based action plans. Neuron *70*, 536–548. .

Kurata, K. (1993). Premotor cortex of monkeys: set- and movement-related activity reflecting amplitude and direction of wrist movements. J. Neurophysiol. *69*, 187–200. .

Logiaco, L., Abbott, L.F., and Escola, S. (2021). Thalamic control of cortical dynamics in a model of flexible motor sequencing. Cell Rep. *35*, 109090. .

Messier, J., and Kalaska, J.F. (2000). Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. J. Neurophysiol. *84*, 152–165. .

Michaels, J.A., Dann, B., Intveld, R.W., and Scherberger, H. (2015). Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. J. Neurosci. *35*, 11415–11432. .

Milekovic, T., Truccolo, W., Grün, S., Riehle, A., and Brochier, T. (2015). Local field potentials in primate motor cortex encode grasp kinetic parameters. Neuroimage *114*, 338–355. .

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Omrani, M., Kaufman, M.T., Hatsopoulos, N.G., and Cheney, P.D. (2017). Perspectives on classical controversies about the motor cortex. J. Neurophysiol. *118*, 1828–1848. .

Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815. .

Remington, E.D., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. Neuron *98*, 1005–1019.e5. .

Requin, J., Brener, J., and Ring, C. (1991). Preparation for action. Handbook of Cognitive Psychophysiology: Central and Autonomic Nervous System Approaches. *745*, 357–448. .

Rickert, J., Riehle, A., Aertsen, A., Rotter, S., and Nawrot, M.P. (2009). Dynamic encoding of movement direction in motor cortical neurons. J. Neurosci. *29*, 13870–13882. .

Riehle, A. (2005). Preparation for action: one of the key functions of the motor cortex. In In Riehle, A. & Vaadia, E. (Eds), Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions, CRC-Press, Boca Raton, FL,.

Riehle, A., and Requin, J. (1989). Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J. Neurophysiol. *61*, 534–549. .

Riehle, A., and Requin, J. (1993). 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. .

Riehle, A., MacKay, W.A., and Requin, J. (1994). Are extent and force independent movement parameters? Preparation- and movement-related neuronal activity in the monkey cortex. Exp. Brain Res. *99*, 56–74. .

Riehle, A., Wirtssohn, S., Grün, S., and Brochier, T. (2013). Mapping the spatio-temporal structure of motor cortical LFP and spiking activities during reach-to-grasp movements. Front. Neural Circuits *7*, 48. .

Riehle, A., Brochier, T., Nawrot, M., and Grün, S. (2018). Behavioral Context Determines Network State and Variability Dynamics in Monkey Motor Cortex. Front. Neural Circuits *12*, 52. .

Rosenbaum, D.A. (1980). Human movement initiation: Specification of arm, direction, and extent. J. Exp. Psychol. Gen. *109*, 444–474. .

Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu,

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426. .

Sauerbrei, B.A., Guo, J.-Z., Cohen, J.D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., and Hantman, A.W. (2020). Cortical pattern generation during dexterous movement is input-driven. Nature *577*, 386–391. .

Scott, S.H. (2008). Inconvenient truths about neural processing in primary motor cortex. J. Physiol. *586*, 1217–1224. .

Scott, S.H. (2012). The computational and neural basis of voluntary motor control and planning. Trends Cogn. Sci. *16*, 541–549. .

Shen, L., and Alexander, G.E. (1997). Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. J. Neurophysiol. *77*, 1195–1212. .

Shenoy, K.V., Kaufman, M.T., Sahani, M., and Churchland, M.M. (2011). A dynamical systems view of motor preparation: implications for neural prosthetic system design. Prog. Brain Res. *192*, 33–58. .

Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. Annu. Rev. Neurosci. *36*, 337–359. .

Sohn, H., Narain, D., Meirhaeghe, N., and Jazayeri, M. (2019). Bayesian Computation through Cortical Latent Dynamics. Neuron https://doi.org/10.1016/j.neuron.2019.06.012.

Sohn, H., Meirhaeghe, N., Rajalingham, R., and Jazayeri, M. (2020). A network perspective on sensorimotor learning. Trends Neurosci. https://doi.org/10.1016/j.tins.2020.11.007.

Stewart, B.M., Baugh, L.A., Gallivan, J.P., and Flanagan, J.R. (2013). Simultaneous encoding of the direction and orientation of potential targets during reach planning: evidence of multiple competing reach plans. J. Neurophysiol. *110*, 807–816. .

Stewart, B.M., Gallivan, J.P., Baugh, L.A., and Flanagan, J.R. (2014). Motor, not visual, encoding of potential reach targets. Curr. Biol. *24*, R953–R954. .

Sun, X., O'Shea, D.J., Golub, M.D., Trautmann, E.M., Vyas, S., Ryu, S.I., and Shenoy, K.V. (2020). Skill-specific changes in cortical preparatory activity during motor learning.

Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. Nat. Neurosci. *18*, 1025–1033. .

Tanji, J., and Evarts, E.V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. J. Neurophysiol. *39*, 1062–1068. .

Thura, D., and Cisek, P. (2014). Deliberation and commitment in the premotor and

bioRxiv preprint doi: https://doi.org/10.1101/2022.07.09.499417; this version posted July 10, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

primary motor cortex during dynamic decision making. Neuron *81*, 1401–1416. .

Todorov, E., and Jordan, M.I. (2002). Optimal feedback control as a theory of motor coordination. Nat. Neurosci. *5*, 1226–1235. .

Torre, E., Quaglio, P., Denker, M., Brochier, T., Riehle, A., and Grün, S. (2016). Synchronous Spike Patterns in Macaque Motor Cortex during an Instructed-Delay Reach-to-Grasp Task. J. Neurosci. *36*, 8329–8340. .

Versteeg, C., and Miller, L. (2022). Dynamical Feedback Control: Motor cortex as an optimal feedback controller based on neural dynamics.

Vyas, S., Even-Chen, N., Stavisky, S.D., Ryu, S.I., Nuyujukian, P., and Shenoy, K.V. (2018). Neural Population Dynamics Underlying Motor Learning Transfer. Neuron *97*, 1177–1186.e3. .

Vyas, S., O'Shea, D.J., Ryu, S.I., and Shenoy, K.V. (2020a). Causal Role of Motor Preparation during Error-Driven Learning. Neuron *106*, 329–339.e4. .

Vyas, S., Golub, M.D., Sussillo, D., and Shenoy, K.V. (2020b). Computation Through Neural Population Dynamics. Annu. Rev. Neurosci. *43*, 249–275. .

Wang, J., Narain, D., Hosseini, E.A., and Jazayeri, M. (2017). Flexible timing by temporal scaling of cortical responses. Nat. Neurosci. https://doi.org/10.1038/s41593-017-0028-6.

Weinrich, M., Wise, S.P., and Mauritz, K.H. (1984). A neurophysiological study of the premotor cortex in the rhesus monkey. Brain *107 (Pt 2)*, 385–414. .

Wise, S.P. (1985). The primate premotor cortex: past, present, and preparatory. Annu. Rev. Neurosci. *8*, 1–19. .

Wong, A.L., and Haith, A.M. (2017). Motor planning flexibly optimizes performance under uncertainty about task goals. Nat. Commun. *8*, 14624. .

Zaepffel, M., and Brochier, T. (2012). Planning of visually guided reach-to-grasp movements: inference from reaction time and contingent negative variation (CNV). Psychophysiology *49*, 17–30. .

Zimnik, A.J., and Churchland, M.M. (2021). Independent generation of sequence elements by motor cortex. Nat. Neurosci. https://doi.org/10.1038/s41593-021-00798-5.