SEE CORRECTION FOR THIS ARTICLE
PNAS RESEARCH ARTICLE | NEUROSCIENCE

# An abstract categorical decision code in dorsal premotor cortex

Gabriel Diaz-deLeon[a,b], Manuel Alvarez[b], Lucas Bayones[b], Antonio Zainos[b], Jerónimo Zizumbo[b], Sergio Parra[b], Sebastián Pujalte[b], Ranulfo Romo[b,c,1,2], Román Rossi-Pool[b,c,1], and Victor de Lafuente[a,1]

Contributed by Ranulfo Romo; received August 24, 2022; accepted November 4, 2022; reviewed by Omri Barak and Xiao-Jing Wang

**The dorsal premotor cortex (DPC) has classically been associated with a role in preparing and executing the physical motor variables during cognitive tasks. While recent work has provided nuanced insights into this role, here we propose that DPC also participates more actively in decision-making. We recorded neuronal activity in DPC while two trained monkeys performed a vibrotactile categorization task, utilizing two partially overlapping ranges of stimulus values that varied on two physical attributes: vibrotactile frequency and amplitude. We observed a broad heterogeneity across DPC neurons, the majority of which maintained the same response patterns across attributes and ranges, coding in the same periods, mixing temporal and categorical dynamics. The predominant categorical signal was maintained throughout the delay, movement periods and notably during the intertrial period. Putting the entire population's data through two dimensionality reduction techniques, we found strong temporal and categorical representations without remnants of the stimuli's physical parameters. Furthermore, projecting the activity of one population over the population axes of the other yielded identical categorical and temporal responses. Finally, we sought to identify functional subpopulations based on the combined activity of all stimuli, neurons, and time points; however, we found that single-unit responses mixed temporal and categorical dynamics and couldn't be clustered. All these point to DPC playing a more decision-related role than previously anticipated.**

dorsal premotor cortex | categorical decision | inter-trial coding | continuum responses | temporal signals

## Significance

Dorsal premotor cortex's role in somatosensory processing has been generally limited to movement, but our current results add to growing evidence favoring a more abstract function. We recorded DPC's activity in two monkeys trained in a vibrotactile categorization task of two distinct physical attributes, and found an abstract decision signal throughout the population, underpinned by purely temporal signals. Neurons maintained significant responses regardless of attribute or range; population signals converged to identical categorical representations, as shown when cross projected between contexts. Importantly, reemergence of the decision signal during the intertrial period also suggests a feedback role. These results support that DPC plays a larger role during decision-making and outcome feedback, regardless of the stimulus attributes that triggered the decision report.

Daily, humans and animals learn to recognize commonalities between distinct subjects and produce unifying categories and concepts in a natural fashion, instead of storing every individual experience. This process, defined as categorization, permits the approximation of quick decisions necessary in daily life, e.g., labeling the worth of an item as expensive or slowing down when we realize we are driving too fast in a school zone. Once learned, these categories can be generalized, assimilating unrelated groups, which results in the expansion of the previous category's definition. Thus, the understanding of novel topics and experiences, and matching them with the most appropriate decisions, depends on prior knowledge of the category in question. Furthermore, this categorization process also depends on the contextual situation; under different conditions, the same category could have different meanings. For example, acceptable driving speeds depend on context, given that they differ between a city street and a highway. The interval values that delimit each category are modulated in a context-dependent manner.

Complex phenomena such as categorization and decision-making rely on a variety of underlying mechanisms. Although it is known that specific processes depend on the decision to be made (1–3), it is believed that some common task-resolving mechanisms could emerge as well (4, 5). Neural coding from frontal lobe areas demonstrates a high capacity for abstraction and generalization (4, 6–9), which may even be generalized across different sensory modalities (10). In contrast, sensory cortices respond faithfully or code specific attributes of the stimuli (11). A previous study in individual neurons in the primary somatosensory cortex (S1), in which trained monkeys had to categorize three different sensory features of a vibrotactile stimulus, showed that some neurons code a particular feature of the stimuli, while others code all features indiscriminately (12). Thus, S1 neurons represent different aspects of the stimuli, but not its category.

Employing the same vibrotactile categorization task (VCT), here we studied the neuronal activity recorded from the dorsal premotor cortex (DPC). This cortical region was historically associated with the planning of physical movement and response (13–20). Several studies suggest that the activity in DPC can also be related to an abstracted representation of the final decision, making this area a prime candidate to study the categorical representation that emerges within the cortex. In particular, during a temporal pattern discrimination task, neurons of this area codify stimulus patterns in abstracted categories, which were signaled in several different epochs of the task, such as the working memory,

Author contributions: R.R., R.R.-P., and V.d.L. designed research; M.A., A.Z., R.R., and V.d.L. performed research; G.D.-d.L., L.B., J.Z., S. Parra, S. Pujalte, and R.R.-P. analyzed data; and G.D.-d.L., L.B., J.Z., S. Parra, S. Pujalte, R.R., R.R.-P., and V.d.L. wrote the paper.

Reviewers: O.B., Technion-Israel Institute of Technology; and X.-J.W., New York University.

The authors declare no competing interest.



[1]To whom correspondence may be addressed. Email: ranulfo.romo@gmail.com, romanr@ifc.unam.mx, or lafuente@unam.mx.

[2]Present address: El Colegio Nacional, 06020 Mexico City, Mexico.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2214562119/-/DCSupplemental.

Published December 5, 2022.

comparison, and decision-making periods (4). Moreover, a variety of heterogeneous activity with a broad diversity and mixed selectivity dynamics is elicited (21). To give sense to this diversity, the neural dynamics were investigated at the population level, employing the collective activity of the DPC network during the decision process (22). This population-level approach was also used to analyze DPC responses during a visual task, showing that preparatory and movement dynamics were coded orthogonally in the population (20, 23). Importantly, at different times, the same neurons modify their responses to generate these separate coding subspaces (20). These recent findings provide relevant evidence to revindicate the role of DPC during abstract categorization, working memory, and decision-making. Therefore, the participation of DPC in more cognitive processes should be further explored.

In this regard, and since it is known that categories serve as a mechanism for uniting various dissimilar objects, we posed two initial questions to face this exploration: is it possible that separate members of the same category are represented similarly in the neuronal activity? Does the brain unite separate categorical representations, or maintain them independently in isolated subpopulations? To address these questions, we trained two monkeys to perform the VCT, which consisted in evaluating the magnitude of a given stimulus as "high" or "low" based on the physical attribute that was varied (frequency or amplitude). For each attribute, two different superimposed ranges were trained. Importantly, stimulus magnitudes shared by sets of the same physical attribute should be categorized as "low" in one while "high" in the other.

Characterizing the breadth of responses that appear among the population's single neurons, we observed an overwhelming categorical representation that was maintained throughout the delay period, lasting until the intertrial period. Further, all observable sensory dynamics appeared to be correlates of the categorical responses, consistently across the two physical attributes and the two superimposed sets. Then, we transit to considering how a common population of single units behave for differing stimulus attributes. This would allow us to see whether neurons were committed to a single attribute, or whether they would present with a more abstract response pattern that unifies the distinct physical attributes. We found that most neurons presented similar response patterns for both attributes, often coding for both attributes at the same moment.

Although single-unit studies still yield interesting findings in this area, little is known about the aggregate activity that emerges across the entire population when the VCT is performed. To address this issue, we used dimensionality reduction techniques to better understand the heterogeneity of dynamics occurring within the population recorded from DPC and find the most representative underlying dynamics (20, 22, 24–26). Once again, we observed a consistently maintained categorical representation, and the dynamics that emerge from different sets are entirely relatable using populations of neurons recorded between different sets. Moreover, two orthogonal dynamics were observed, one during the delay and the other one during the movement/intertrial periods. Despite this, we found that these orthogonal population responses are represented by the same neural substrate; instead of clusters of neurons associated with each orthogonal dynamic, we find a continual transition between these signals. These results demonstrate that DPC wholly leverages its heterogeneous activity to form a highly efficient abstract decision categorical code, which serves to categorize different stimulus attributes regardless of their physical attributes.

## Results

**Vibrotactile Categorization Task.** Two monkeys (*Macaca mulatta*) were trained to categorize a vibrotactile stimulus as high or low for two different physical attributes: frequency and amplitude, each in two different overlapping ranges: short frequency range (SFR= 10 to 30 Hz in steps of 2 Hz) and short amplitude range (SAR = 20 to 80 μm in steps of 6 μm); long frequency range (LFR= 14 to 78 Hz in steps of 8 Hz) and long amplitude range (LAR = 42 to 136 μm in steps of 12 μm), resulting in a total of four different predetermined sets (Fig. 1 *A*–*C*) (12). Sets were presented in a random order. The monkeys reported their choice by reaching and pressing one of the two buttons with their left free hand. The experimental setup was designed so that the angular difference between the buttons (11°) was expected to be statistically indistinguishable in the activity of motor or premotor neurons (*SI Appendix*). Monkeys' performance was consistent for all sets: frequency short (SFR, n = 61 sessions, 90.3 ± 0.16%) and long (LFR, n = 16 sessions, 82.9 ± 1.2%, Fig. 1*B*); amplitude short (SAR, n = 48 sessions, 85.4

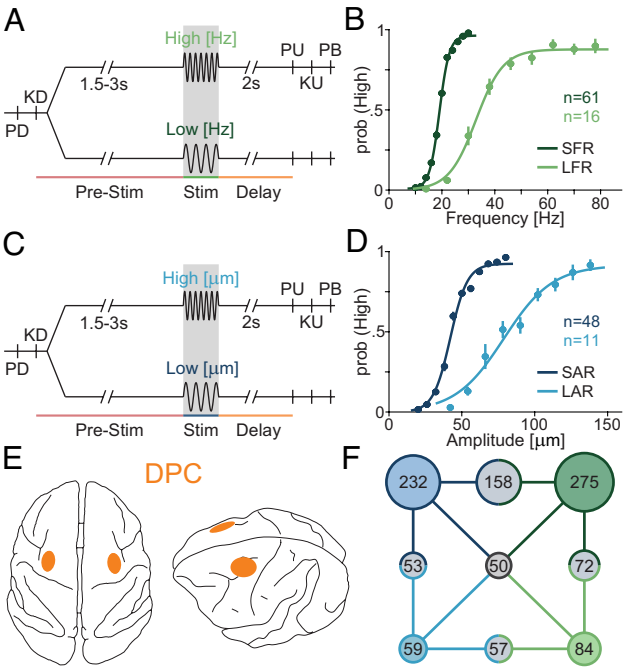


**Fig. 1.** Task Schematic, Psychophysical Performance, Data Sets, and Recording Sites. (*A* and *C*) Task schematic for categorization of frequencies (*A*) and amplitudes (*C*). Trials start when the mechanical probe lowers, indenting the glabrous skin of one fingertip of the monkeys' restrained right hand (probe down event, "PD"). The animals respond by placing their left hand on an immovable key (key down event, "KD"). After KD, a variable period (1.5 to 3 s) is presented, followed by the presentation of a single stimulus lasting 0.5 s. After stimulation, a fixed delay period of 2 s is presented, followed by the probe up event ("PU"). PU serves as the "go" cue for monkeys to remove their hand from the key (key up event, "KU") and report their decision using one of the two push buttons placed in front of him (push button event, "PB"). Correct answers were rewarded with a few drops of juice. (*B* and *D*) Psychometric performance for all categorization sets: (*B*) Psychometric curves for the two frequency range sets, short (dark green) and long (light green). The short frequency set (SFR) had 12 stimuli that varied between 10 and 30 Hz, in steps of 2 Hz, with the middle stimulus (20 Hz) repeated once for each category (61 sessions). The LFR set had 10 stimuli that varied between 14 and 78 Hz, in steps of 8 Hz, with the middle stimulus (46 Hz) repeated once for each category (16 sessions). (*D*) Psychometric curves for two amplitude range sets, short (dark blue; SAR) and long (light blue; LAR). The SAR set had 12 stimuli that varied between 20 and 80 μm, in steps of 6 μm, with the middle stimulus (50 μm) repeated once for each category (48 sessions). The LAR set had 10 stimuli that varied between 42 and 138 μm, in steps of 12 μm, with the middle stimulus (90 μm) repeated once for each category (11 sessions). (*E*) Top (*Left*) and lateral (*Right*) views of a monkey brain, with recorded DPC area highlighted (yellow). (*F*) Diagram of populations of neurons recorded per set. Colors are consistent with those in panels *B* and *D*. Corner circles with solid colors represent sets, intersecting circles with two mixed colors represent populations recorded in both sets, and the black circle in the middle is the population of neurons recorded in all the four sets.

± 0.35%) and long (LAR, n = 11 sessions, 81.1 ± 1.41%, Fig. 1*D*). This is represented by the rightward shift of the lighter colored psychometric curve as well as the saturation at values of 0 and 1. All data (355 neurons) were recorded from DPC (Fig. 1*E*), where 275 neurons were recorded during the SFR set, 232 neurons during the SAR set, 84 during the LFR set, and 59 during the LAR set (Fig. 1*F*). Additionally, a control version of the SFR set was presented to the monkeys. In this set, at the beginning of each trial, a bright light located within each button turns on, indicating the correct (rewarded) answer for that trial. The rest of the trial develops identically to those of the active SFR set. A total of 175 units were recorded during this control task.

**Single Neuron Responses during the Vibrotactile Categorization Task.** Previous studies in DPC have shown that its neural activity is heterogeneous, spanning from purely categorical to purely temporal (20, 22, 27). To study this diversity, we characterized the coding and response properties of each neuron over the course of the VCT. Exemplary neurons are shown in Fig. 2 and *SI Appendix*, Fig. S1. We can observe a neuron with a purely categorical persistent response to the SFR set in Fig. 2 *A*, *Top*; it was silent for high stimulation frequencies but maintained the response during low ones. It is important to notice, in the raster plot and rate profile (Fig. 2 *A*, *Bottom*), that the categorical response emerged at the end of the stimulation period and was maintained throughout the delay. Remarkably, this neuron had a nearly identical response to amplitude categorization (SAR set, Fig. 2*B*). This phenomenon was consistently observed across neurons; in general, most had the same responses regardless of the physical attribute tested.

To quantify the categorical nature of neuron responses, we computed the area under the receiver operating characteristic curve (AUROC; *SI Appendix*) at each time (28), between the firing rate distributions corresponding to high and low stimulus categories. For some neurons (Fig. 2 *C* and *D*), categorical representation during the SFR and SAR sets significantly emerges at the end of the stimulation period and persists throughout the delay until after the PU event. Its activity represents the high stimulus categories silently, characterized by an AUROC < 0.5.

Given that a categorical code implies a transformation from the sensory representation, we wondered whether pure sensory information is still present in DPC neurons. To elucidate this question, we computed the category ($I_{CAT}$) and stimulus ($I_{STIM}$) mutual information for each physical attribute, for each single neuron (*SI Appendix*). These two calculations yielded identical results (Fig. 2 *E* and *F*) for their values as well as their significant time windows ($P < 0.05$), suggesting that both of them are measuring the same underlying dynamics. This foreshadows the finding that this DPC population contains little to no information of the stimuli's physical parameters. Note that $I_{STIM}$ could go up to more than double than $I_{CAT}$ (which is bounded by 1 bit, *SI Appendix*); if stimuli information were available beyond what is provided by the categorical representation, $I_{STIM}$ should surpass $I_{CAT}$. Moreover, this is also found for the long stimulus sets (SFR and LAR), where we can again observe that this neuron preferentially responds to low classes for both physical attributes (*SI Appendix*, Fig. S1 *A* and *B*). These results indicate that this exemplary neuron maintains a significant categorical representation, regardless of the physical attribute or stimulation range studied.

Other DPC neurons recorded in the SFR condition also exhibited a broad repertoire of dynamics that vary over the task period (*SI Appendix*, Fig. S2). For example, panel *A* shows a neuron that responds after the go-cue, with slightly categorical dynamics in

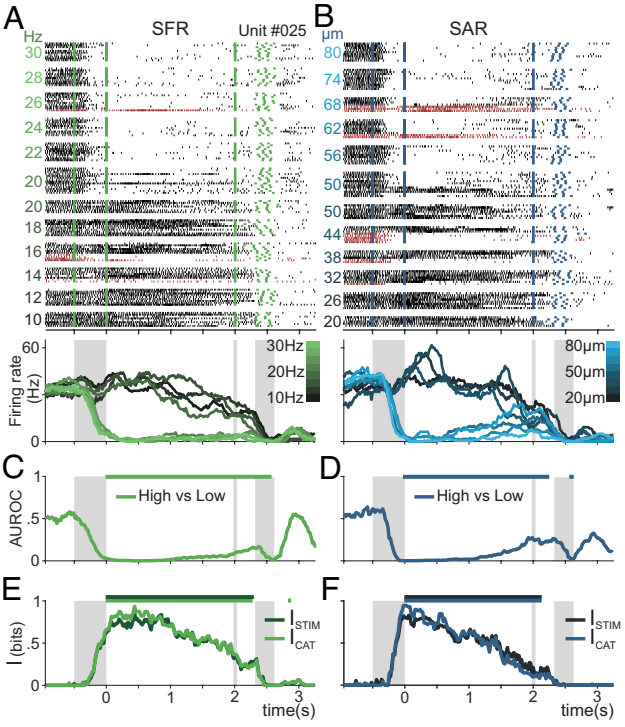


**Fig. 2.** A Single Neuron's Activity during Frequency and Amplitude Categorization. (*A* and *B*) Raster plots of a single neuron recorded in the SFR (*A*, *C*, and *E*) and SAR (*B*, *D*, and *F*) sets. Black and red tick marks represent spikes in trials with correct and incorrect answers, respectively. Green (*A*) and blue (*B*) tick marks indicate psychophysical events. Trials are grouped by class (stimulus intensity), with intensity values marked on the left side. From *Left* to *Right*, the first green and blue tick marks that occur at −0.5 s indicate stimulus onset, and the second marks at t = 0 s indicate stimulus offset. After a fixed 2 s delay period, the third green and blue tick marks indicate PU, while the fourth represents KU and the fifth represents PB. Figures below raster plots represent firing rate averages for all the hit trials associated with each of the eight visualized classes. Some classes were excluded for clarity in the visualization. Large gray rectangles represent the stimulation period, while thin rectangles mark PU, and the third and last rectangle marks the average movement period, beginning with KU and ending with PB. Firing rate curves range from dark gray for low category stimuli to bright green (*A*, frequency) or bright blue (*B*, amplitude) for high category stimuli. (*C* and *D*) Time-dependent differential AUROC (high versus low) values were taken between the firing rate distributions of correct trials. AUROC < 0.5 or AUROC > 0.5 indicates an increased firing rate response for classes of category low or high, respectively. Significant windows, based on permutation tests and correcting for multiple comparisons, are marked above the curve ($P < 0.01$). (*E* and *F*) Time-dependent category ($I_{CAT}$, light green and blue) and stimulus mutual information ($I_{STIM}$, dark green and blue) curves ($P < 0.05$).

the movement period that become even more pronounced in the intertrial period. However, in *B*, there is a neuron whose response is slightly categorical during the stimulation period but whose clearest categorical response occurs only after movement, during the intertrial period. The neuron in *C* has pure temporal dynamics up until the go-cue, becoming categorical during the movement period, and then presents mixed coding in the intertrial period. The neuron in panel *D* mirrors the one shown in Fig. 2 and *SI Appendix*, Fig. S1, with a persistent categorical code throughout the delay period but with a preference for the low rather than high category. Interestingly, its categorical coding is quenched after the go-cue and movement period. The neuron in *SI Appendix*, Fig. S2*E* demonstrates a slight categorical response during the delay and after the movement period, with clear temporal modulation. Finally, panel *F* shows a neuron that represents purely temporal dynamics, with almost no variation between the classes throughout the entire task period. It is important to note that following the movement period, dynamics continue to emerge in several

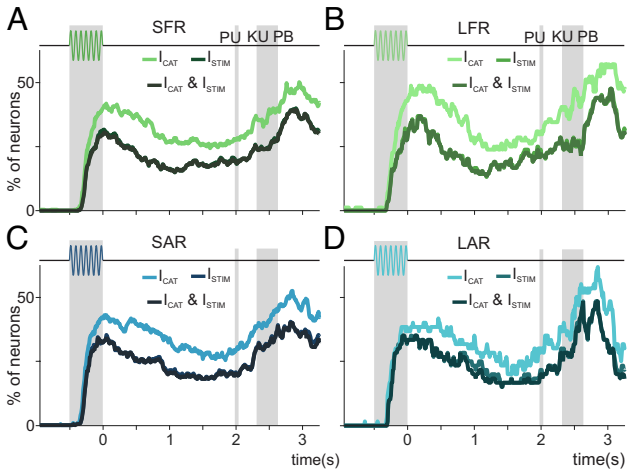


**Fig. 3.** Single Neuron Coding during Frequency and Amplitude Categorization. (*A* and *B*) Neuron population coding for the stimulus frequency (*A* and *B*) and stimulus amplitude sets (*C* and *D*). Population proportions recorded for frequency short range (SFR; *A*, n = 275) and long range (LFR; *B*, n = 84) and for amplitude short range (SAR; *C*, n = 232) and long range (LAR; *D*, n = 54) with significant category information ($I_{CAT}$), stimulus information ($I_{STIM}$), and both information types at once ($I_{CAT}$ and $I_{STIM}$).

neurons (*SI Appendix*, Fig. S2 *B–E*, see *SI Appendix*, Fig. S5*B* [neurons: #085, #154, #169, #150, #143, #127]) in a manner that is independent of those present during the motor report. This signal is maintained even later than the reward delivery (~40 ms after push button), during the intertrial period. In summary, as opposed to ref. 12, DPC neurons display a diverse range of categorical and temporal dynamics in their responses.

**Single Neuron Coding Dynamics.** Since we now have a picture of single neuron responses in DPC, we proceed to analyze how the entire population behaves as a function of time. For each metric ($I_{CAT}$ and $I_{STIM}$) and stimulus set, Fig. 3 shows the proportion of all neurons yielding significant results across the entire population. Notably, in Fig. 3*A* for SFR, the response pattern is similar across the mutual information metrics: the categorical representation emerges at the end of the stimulation period, it is maintained persistently throughout the delay with a slightly smaller proportion, and then the representation picks up during the movement period. Intriguingly, the greatest proportion of neurons with significant coding is found after movement, in the intertrial period. The peak during the stimulus period corresponds to the emergence of a categorical decision, where the subject decides its report, and the persistent dynamic during the delay corresponds to the subject's maintenance of the decision in working memory. However, a high proportion of this categorical coding persists into the intertrial period. This was previously found in DPC during a different task (4, 22), and is probably related to learning or reward processes, such as integrating the evaluation of choice outcomes into future decisions (29, 30). Note that this same response pattern was observed for the population recorded during the LFR condition (Fig. 3*B*).

We now wonder if the categorical ($I_{CAT}$) and the stimulus information ($I_{STIM}$) can account for separate underlying dynamics. In all panels of Fig. 3, the proportion of neurons with significant $I_{CAT}$ is greater than that of $I_{STIM}$, suggesting that categorical coding is more abundant. However, the percentage of neurons with both significant $I_{CAT}$ and $I_{STIM}$ ($I_{CAT}$ and $I_{STIM}$) almost perfectly overlaps with the curve for only $I_{STIM}$. As previously discussed for individual neurons (Fig. 2*A*), this result suggests that categorical information is the only one present in the DPC network. These findings remained true for all sets, regardless of attribute or range, exhibiting a general feature of this cortical area. Given the prior association between DPC and a particular role dedicated to the motor report, we compared results with those from the light control task described above. Since movements were the same in the control as in our other tasks if decision coding dynamics could actually be interpreted as movement or its preparation, they should also appear in the control. This was not the case, as can be appreciated in *SI Appendix*, Fig. S3. We used AUROC, variance tests and information to compare the significant activity that could be found in the population recorded for our experimental SFR set (*SI Appendix*, Fig. S3*A*) to its control version (*SI Appendix*, Fig. S3*B*). For example, if we take a look over the results for the measures of $I_{CAT}$ in SFR in Fig. 3*A* (and AUROC in *SI Appendix*, Fig. S3*A*) and compare it with those obtained in *SI Appendix*, Fig. S3*B*, we can observe that there is an evident loss of categorical information across the entire population throughout the control task. This suggests that the sustained activity observed during experimental sets might be better interpreted as a persistent decision signal of which category each stimulus belongs to, rather than as the preparation of a particular movement. However, it remains unclear whether a single neuron performs common operations between different sets, or if the network assigns specific roles to individual neurons based on the set being tested.

To test this, we analyzed the neurons recorded for more than one set in order to understand how the DPC network codes the information during the categorization of different stimulus attributes. When comparing SFR and SAR, we calculated the information tuning as a function of time for both stimulus sets, and then compared each time bin to determine whether the single neurons were coding information related to a single or both stimulus attributes. Fig. 4*A* shows the coding as a function of time for each neuron recorded in both sets (SFR/SAR, n = 158). We observed that most neurons code both attributes at some point throughout the task. To further elaborate, we computed the population proportions with significant coding as a function of time for individual or both attributes (Fig. 4*C*). Note, that during the task, the percentage of neurons that codes both features ($I_{both,CAT}$) is much higher than those with a single-feature coding ($I_{freq,CAT}$ or $I_{amp,CAT}$). These results strengthen our hypothesis of the network's convergent mechanism for stimulus categorization. Further, a similar abstract

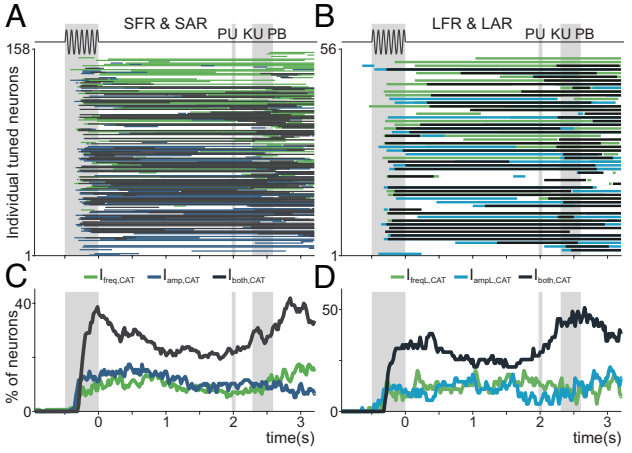


**Fig. 4.** Abstract Information Coding of Stimulus Attributes. (*A* and *B*) Information tuning for all neurons recorded in the frequency and amplitude sets. (*A*) Neurons recorded (n = 158) during the SFR and SAR. (*B*) Neurons recorded for both long stimulus attribute ranges (LFR and LAR; n = 54). Lines indicate individual neurons with significant information during only one set (blue and green) or both sets (gray). (*C* and *D*) Percentage of neurons, recorded during short-range (*C*) or long-range sets (*D*), with only frequency (green), amplitude (blue), or dual information coding (dark gray).

population coding was observed for the neurons recorded during both long ranges of stimulus sets (Fig. 4 *B* and *D*, LFR/LAR, n = 56). These results indicate that the network reused the coding dynamics of most of the individual units, regardless of whether the stimulus varied based on amplitude or frequency. However, we wondered whether this occurs when considering the emerging activity from the entire population of neurons regardless of each stimulus feature. We extended this analysis to consider if we could decode, which attribute was being varied, from the activity of the neurons recorded in two attributes of similar ranges. In *SI Appendix*, Fig. S4 *A* and *B*, it can be seen that the attribute cannot be consistently decoded, despite multiple points of comparison ($Cat_{LOW}$, $Cat_{HIGH}$, all trials). This means that the activity of the population of neurons is almost entirely preserved between the sets of differing attributes, serving as further evidence of the conserved roles for individual single units. However, we wondered whether this might instead be possible when considering the emerging activity from the entire population of neurons.

**Generic Population Dynamics across Stimulus Attributes.** To understand the signals that were obfuscated in the heterogeneity of single neuron responses, dimensionality reduction methods were recently applied to compute the relevant population dynamics (25, 31–33). Here, we employed demixed-principal component analysis (dPCA) to find the population axes that maximize the amount of explained variance associated with a particular task parameter (22, 24). To compute these representative axes (dPCs), we calculated the marginalized covariance matrices which summarize the population coactivity with respect to each task parameter. These decoding axes yield dimensions that summarize the population-level dynamics, allowing us to understand how the network is representing the different parameters associated with the task. Moreover, this approach allowed us to separate the population coding dynamics from the pure temporal signals that explained most of the variance.

We begin by considering the four stimulus sets separately (SFR, SAR, LFR, and LAR) and computing the dPCs with the marginalized activity with respect to the stimulus identity. In Fig. 5 panels *A* and *B* (*Top* and *Bottom*), we plotted the first two stimulus-dPCs calculated for SFR and SAR, where each one has eight classes projected. In a similar manner to the single neuron analysis, the stimulus-axes for SFR and SAR showed a categorical representation that emerges halfway through the stimulation period and maintain itself throughout the delay period until after the GO cue. It is only during the movement period that the dPCs began to diverge: the first dPC maintains its categorical representation

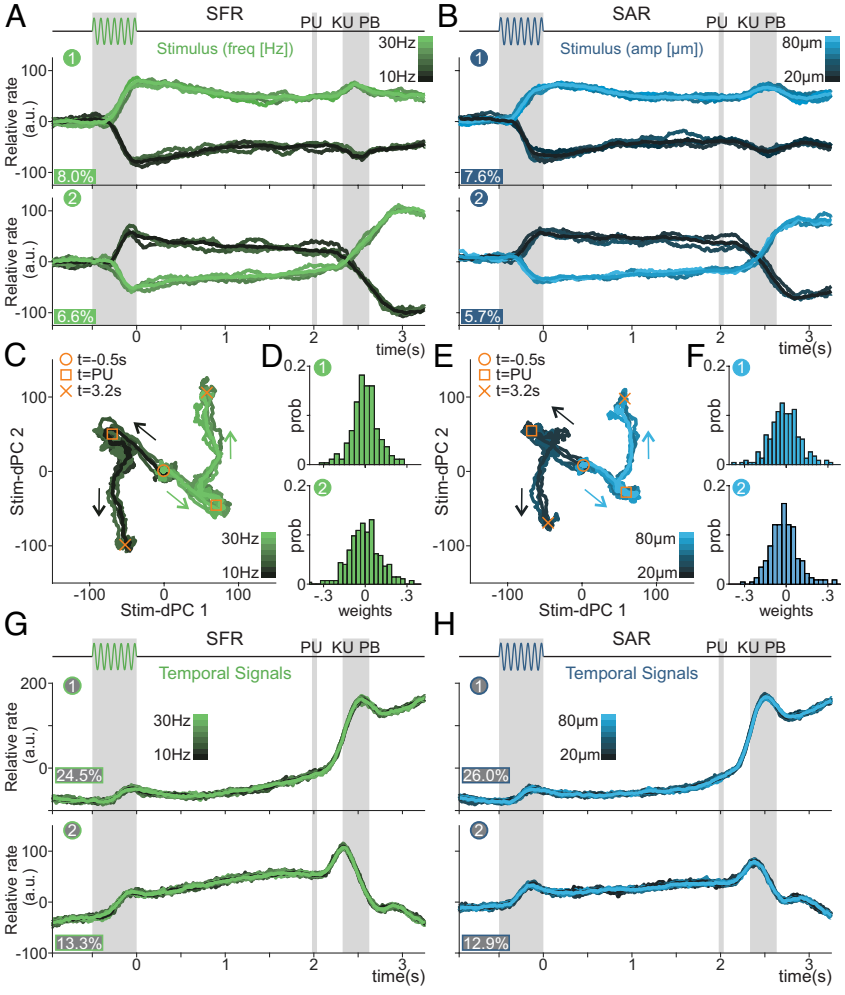


**Fig. 5.** Stimulus and Temporal Population Dynamics. (*A* and *B*) Projections of class averages over the first and second decoding stimulus-axes of the short-range population for frequency (*A*, SFR, n = 275) and amplitude (*B*, SAR, n = 232). Classes ranging from dark (lowest stimulus) to light (highest stimulus). The ordinal number of each dPC is shown in a circle; explained variances are shown as percentages. (*C* and *E*) Phase diagram of the same classes projected across the first two decoding axes for frequency (*C*) and amplitude (*E*). The yellow circle marks t = −0.5 s (SO), the yellow square marks the end of the delay (PU, t = 2 s), and the yellow x marks t = 3.2 s (intertrial period). (*D* and *F*) Distributions of the neuronal weights for the first two decoding axes of short range for frequency (SFR, *D*) and amplitude (SAR, *F*). (*G* and *H*) The first two temporal population signals for the short-range population for frequency (SFR, *G*) and amplitude (SAR, *H*).

until well after the GO cue and movement period (Fig. 5 *A* and *B*, *Top*), while the second dPC exhibits its strongest categorical representation during the intertrial period (Fig. 5 *A* and *B*, *Bottom*). Remarkably, this same response pattern was observed for the LFR and LAR sets (*SI Appendix*, Fig. S5 *A* and *B*, *Top* and *Bottom*). Since each dPC was identified using separate data sets, the congruence across these results strongly suggests that a common, generic coding mechanism was being employed.

If we consider both stimulus-dPCs together, it is possible to understand how the network manages to differentiate the traces related to each category during the different task periods. In panels *C* and *E* of Fig. 5, we plot two-dimensional phase diagrams obtained from the first two dPCs associated with frequency and amplitude. All traces start together at the center of the phase diagram, but soon afterward, the traces for the stimulus related to categories low and high begin to diverge, consolidating into what could be separate decision paths. Importantly, after the PU event (square symbol), the population evolution changed drastically. This suggests that the population dynamics between these two task periods are almost orthogonal (20, 23, 34). Intriguingly, this orthogonal population signal endures well into the intertrial period. Moreover, we also found that this same behavior remains consistent for the long-range sets as well (*SI Appendix*, Fig. S5 *C* and *E* for LFR and LAR, respectively).

Considering the contributions of each neuron for these two dPC axes, we find that the weights are normally distributed around 0, indicating that the network balances positive and negative coding equally (Fig. 5 *D* and *F* and *SI Appendix*, Fig. S5 *D* and *F*). We then computed the pure temporal population responses, employing the condition-independent marginalized matrices to calculate the first two temp-dPCs (24). For all sets, these two signals remained unaltered and explained more than 42% of the whole variance (Fig. 5 *G* and *H* and *SI Appendix*, Fig. S5 *G* and *H*). This high percentage suggests that they constitute an essential underlying mechanism for this categorization process (27). When considered together, these results suggest that the same coding and temporal population dynamics are exploited by DPC neurons during the categorization of different stimulus attributes and ranges. Taking this line of inquiry a step further, and to tie back to the last analysis of the previous section, we performed the dPCA calculations on the shared sets of neurons recorded for two different attributes within similar ranges, combining the activity of both populations as an extra parameter for dPCA to take into consideration. We found that the attribute information was not decodable for any of the ranges (*SI Appendix*, Fig. S4 *C* and *D*), while the variance explained by the attribute axes is low. Intriguingly, despite adding an additional parameter, the stimulus components are preserved (*SI Appendix*, Fig. S4 *E* and *F*).

Finally, we decided to go further and test the universality of these abstracted population codes. For this, we produced the decoding axes for a common population of neurons recorded in both the SFR and SAR sets and then projected the SAR population activity over the SFR decoding axes (Fig. 6*A*), and vice versa (Fig. 6*B*). The axes (dPC) that were optimized to decode stimulus

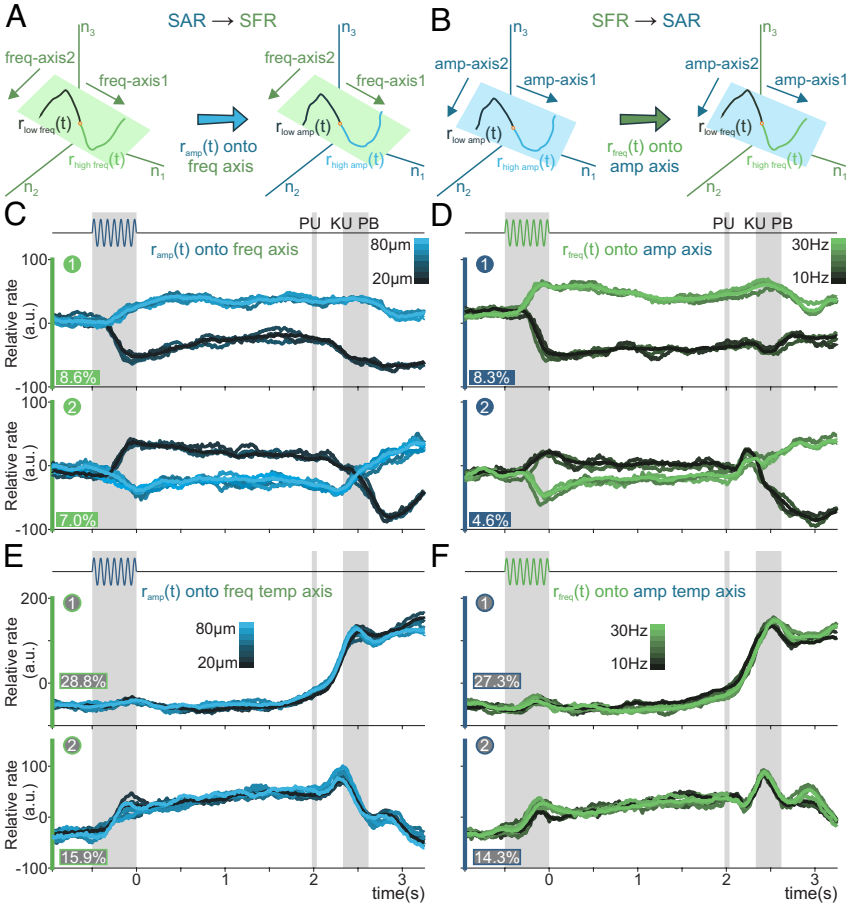


**Fig. 6.** Abstract Temporal and Categorical Coding Population Dynamics. (*A* and *B*) Schematic depicting the projection of the SAR activity over the SFR decoding axes (*A*), and vice versa (*B*) (n = 158). (*C* and *E*) Projection of classes from the SAR population over the first two SFR freq-axes (*C*), and the first two SFR temporal-axes (*E*). Axes computed with the SFR dynamic. The numbers in green circles, as well as the color of the y-axis, mark the source of the decoding axes. Explained variance of the SAR dynamic for each SFR freq-dPC is included below. (*D* and *F*) Projection of classes from the SFR population over the first two blue circles, as well as the color of the y-axis, mark the source of the decoding axes. Explained variance of the SFR dynamic for each SAR amp-dPC is included below.

identity during the SFR set were employed to project the population activity recorded during the SAR set. This procedure is only possible because the neurons were recorded during both sets: SFR and SAR. We also performed the same procedure for the common population between the LAR and LFR sets (*SI Appendix*, Fig. S6 *A* and *B*). The projection of SAR-population activity over the first and second SFR-decoding axes gives rise to the same categorical representations as shown in Fig. 5: the categorical decision coding lasts through the delay and grows in the intertrial period (Fig. 6*C*). This remains consistent when projecting the LAR-population activity over the LFR-decoding axes (*SI Appendix*, Fig. S6*C*). Moreover, a nearly identical categorical representation appears when we project the SFR-population activity over the SAR-decoding axes (Fig. 6*D*), as well as when we project the LFR-population activity over the LAR-decoding axes (*SI Appendix*, Fig. S6*D*). Furthermore, the temporal signals for the SAR-population activity projected over SFR-temporal axes (Fig. 6*E*) mirror the SFR-population activity projected over the SAR-temporal axes (Fig. 6*F*). Analogous results were also found for long-range sets (*SI Appendix*, Fig. S6 *E* and *F*). Moreover, these signals are also remarkably similar to the pure decoding axes that were obtained for the entire populations in Fig. 5 and *SI Appendix*, Fig. S5. These results strengthen the evidence that the DPC network uses a universal population dynamic during this task. The same coding and temporal population responses are evoked regardless of the stimulus' physical attribute or range.

**Mixed Single Neuron Responses Generate the Population Dynamics.** To further characterize the DPC responses, we asked whether the different population dynamics are generated from separate groups of neurons. In other words, is it possible to divide the neural responses into clusters associated with the population responses observed in Figs. 5 and 6? To tackle this question, we paired the neuronal weight distributions to create scatterplots for three different pairs of decoding axes: stim-dPC-1 and stim-dPC-2 (*SI Appendix*, Fig. S7 *A*, *Left*), stim-dPC-1 and temp-dPC-1 (*SI Appendix*, Fig. S7 *A*, *Middle*), and temp-dPC-1 and temp-dPC-2 (*SI Appendix*, Fig. S7 *A*, *Right*). In all the three cases, the 18 example neurons (*SI Appendix*, Fig. S7 *B*, *Right* and *Bottom*) are distributed in a seemingly random manner, although there are some intuitive differences. For example, neuron units #182 and #007 are relatively closer in the plane created with the stim-dPCs (*Left*), than the one created with the temp-dPCs (*Right*). Intuitively, the response pattern to the varying stimuli is similar for these two example neurons, although they may have two completely different temporal signals supporting their categorical response. This initial analysis supports what other studies have found in other tasks (22, 24, 35): that the population of neurons recorded in DPC demonstrates a mixing of temporal and coding dynamics to varying degrees.

To corroborate these results with a different approach, we used a nonlinear dimensionality reduction technique known as Uniform Manifold Approximation and Projection (UMAP) (36, 37). This method was recently applied to identify grid cells in the entorhinal cortex and reveal the toroidal geometry underlying their responses 38. Using this technique, we projected the concatenated firing rate from all stimuli related to the SFR set, into two-dimensional space, providing us with a visual representation of the data that could potentially detect functional clusters of neuronal activity (Fig. 7*A*). Under the assumption that such clusters could exist, we first estimated the probability density functions associated to each UMAP dimension (*SI Appendix*, Fig. S7*B*, above and to the right of the UMAP plane). These nearly unimodal distributions suggest that there is only a single, centralized cluster in the UMAP plane.

To further address this result, we calculated the peak clustering index 39 (*SI Appendix*) of each point and found two outliers (*Inset* in *SI Appendix*, Fig. S5*B*). When graphing these density peaks (square symbols in *SI Appendix*, Fig. S7), it is possible to note that they do not separate different groups, suggesting that they belong to the same cluster. Moreover, to visualize this single, central cluster, we underlaid a contour density plot onto the neurons plotted across the UMAP plane (Fig. 7*A*) and observed that our cluster is indeed a single, continuous entity. We sought to characterize how groups of neurons vary across the UMAP plane, so we calculated a double-Gaussian ($\sigma = 0.4$) weighted average of the activity of all neighboring neurons at certain points (X-marks) throughout the

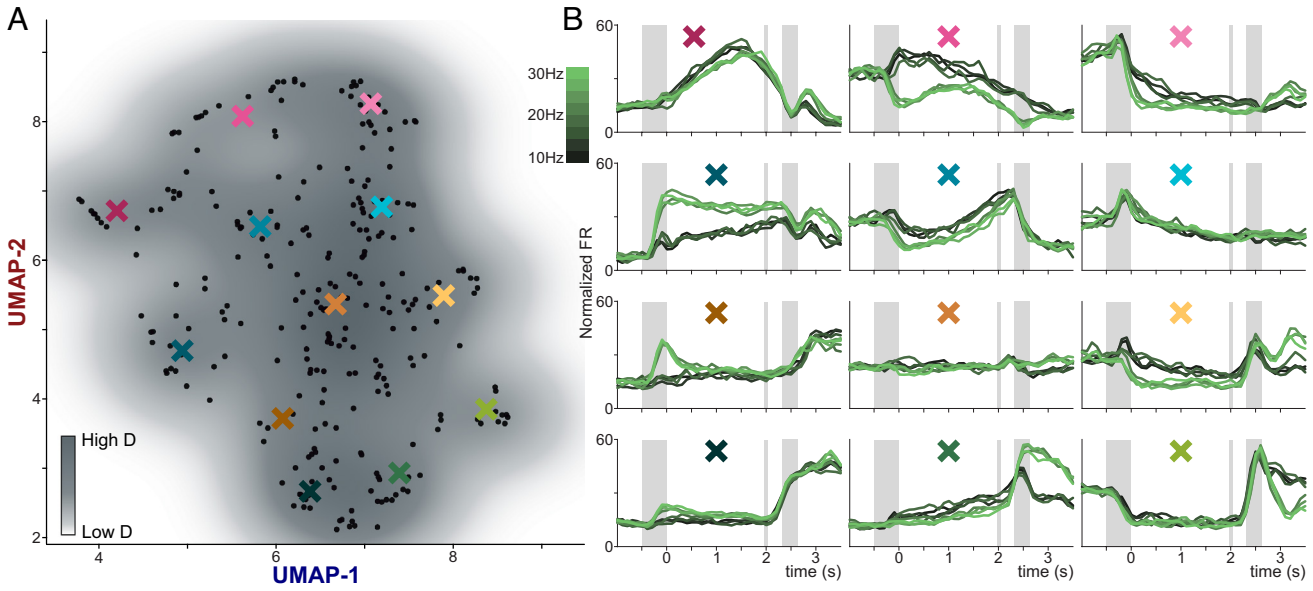


**Fig. 7.** Mixed Dynamics Interlink Categorical and Temporal Neural Responses. (*A*) UMAP plane with a density contour plot. Each black point represents a single neuron. The colored X marks are the centers of the double-Gaussian ($\sigma = 0.4$) weighted averages presented in the panels in (*B*). (*B*) Each subpanel presents the weighted average of neuronal activity and is presented in back-transformed arbitrary units that roughly correspond to the firing rate values (Hz).

UMAP plane (Fig. 7*B*). Observing the results, we can see that categorical dynamics occur to varying degrees across many of the averages, along with very clear portions of temporal signals. Although these clearly marked and separated periods of categorical versus temporal dynamics are only emerging due to the averaging of several neurons, it provides a unique opportunity to observe how integrally mixed these two regimes of dynamics are. Notably, the intertrial period demonstrates both categorical and temporal dynamics that have been preserved through the averaging procedure, suggesting the relevance of this DPC signal during this task. From this result, we can conclude that the temporal and categorical dynamics are inherently interrelated, making it impossible to isolate separable functional groups. Ergo, even if the population signals were obtained with orthogonal readouts, they were constructed from the same neuronal substrate of heterogeneous and mixed signals. This suggests that single-unit responses exist within a continuum between extreme, pure responses instead of belonging to separate groups for each response type.

## Discussion

In this work, we sought to understand whether the DPC contributes to the elaboration of a generalized, abstract categorical decision code. Two trained monkeys performed the VCT, where vibrotactile stimuli had to be categorized as either "high" or "low." Two features of the experimental design dictate the core questions of our research: 1) stimuli could vary in either of the two different physical attributes: frequency and amplitude; and 2) values for these attributes varied across two superimposed ranges, so that around half of the stimuli switched categories according to context. These two features were developed to test for abstractions in the neural code, and they were motivated by the following type of questions: does DPC generalize its response across both the two ranges and the two physical attributes? This would be the most abstract code possible for this task. Analyzing the neuronal activity of DPC at the single and population level, we found that regardless of attribute or range, the abstract decision is coded by the network employing a unique neuronal population—a common neural substrate—from which an abstract code emerges to represent both categories. From a computational perspective, this phenomenon is highly flexible and efficient, given that the same underlying population of neurons arrives at an appropriate decision independently of context. Although the dynamics observed in DPC are not enough to resolve the task by themselves, since they only represent the emerging final decision, the results shown above add to evidence that DPC is not only limited to the planning of physical movement. Importantly, this decision is also coded during the intertrial period, suggesting that this signal is relevant for future trials. Our results suggest that it is also involved in the construction and reevaluation of a complex and abstract categorical decision code.

The notions of a singular abstract code and of a common neural substrate for it are supported by the evidence from d-PCA (24) and UMAP (36, 37). After studying the information carried by single units, these two population approaches were used to examine how this information is distributed across DPC neurons. First, we applied the dimensionality reduction of d-PCA to find the most prominent signals in our population. The foremost finding delivered by d-PCA is that, for all the four combinations of range and physical attribute studied, the emergent population signals are the same. This is true for the categorical coding signals, as well as for their underpinning temporal dynamics. This might most clearly be observed when we projected the neural responses to one attribute over the components extracted while the other was presented; doing this swap made no difference to the signals. The implications of this observation are twofold: when we consider that these population signals are obtained by combining all neuronal responses according to a set of weights, on the one hand, this makes evident that the population signals in each context converged to the same abstract categorical code; on the other hand, the weight distributions served as a first indication that there were no separate subpopulations. To further test this, UMAP was brought to bear on the full dynamics of the neural population. Density-based clustering on the results of this nonlinear algorithm also showed no presence of subpopulations. Moreover, as it was stated in a previous DPC work (27) and further corroborated here, with UMAP and d-PCA approaches, temporal signals constitute the basic infrastructure over which this abstract code unfolds. Thus, it appears that the singular abstract code emerged from the dynamics of a neural substrate common to all contexts.

Nevertheless, where does DPC fit in the process of constructing this decision? And what signals are available to it? Moreover, are there any upstream areas that are intermediate in their abstraction, generalizing for one feature but not the other? And how would areas produce this abstraction? Do they consolidate subpopulations of neurons that each generalize a certain feature? Or do abstract signals emerge from a common neural substrate? To discuss this, we should consider first where in the somatosensory hierarchy DPC is thought to belong and what the areas upstream to it do. In a variety of purported hierarchies of the somatosensory processing network in macaques, DPC sits between the sensory areas and the motor output 11,40, together with other premotor association cortices (1). Regarding the sensory areas, the secondary somatosensory cortex (S2) would be immediately upstream to DPC 28,41; while S1, further away, serves as the entrance into the cortex. Previous work using a VCT has demonstrated how faithful the representation of the physical attributes (amplitude and frequency) is in S1 (12). These findings have been confirmed in a wide variety of tasks (4, 11, 28), as well as several animal models 42,43. On the contrary, recent results from our group found that S2 could be a suitable candidate for the site of transformation from sensory representations to more abstract signals 41, since both types of responses were observed to have similar proportions in this area. From those findings, S2 appears to be an essential driver of sensory abstraction and the main source of inputs available for DPC to construct its decision categorical code. With this in mind, we do not mean to suggest that DPC is the only area which plays a key role in the formation of this code. Decision-related activity is known to be widely distributed across frontal lobe areas (1), and most are probably also recipients of the signals generated by S2 and equivalent areas from other senses. Nevertheless, previous findings make DPC an interesting target for the study of our research questions (4, 11).

Classically, the role of DPC has been limited to preparing physical motor reports 14,18,19,44, mental rehearsal of known physical movements 13,45, and serving as an aide to the primary motor cortex (M1) 46–48. However, in this work, we put forth three arguments about why its results should not be interpreted as motor planning or preparation: the design of the experimental setup, the contrast between neural responses in the active versus control tasks, and the orthogonality of the categorical signals before and after the decision report. First, as it is stated in Methods (*SI Appendix*), our experiment was carefully designed to minimize variations in activity related to push button movement. The results presented above of strongly differentiated categorical signals would be inconsistent with the motor planning interpretation only, according to previous studies on the coding of movement direction in cortical populations. Second, the light control for the SFR set shows an almost complete

black out in the neural responses versus the active set. Here, besides not being significant categorical coding in the delay, there are also no responses during the movement period, which indicates that there is no ability to differentiate the subject's movement based on the categorical response. Third, there was an additional interesting observation found through our population analyses, reminiscent of previous works (20, 23). Although there was decision-related information before and after the movement (during the intertrial period), dPCA showed that these two decision signals are arranged in an orthogonal fashion. This might indicate that they have different functions. While the activity during working memory maintains the decision information, the activity after movement could be directly related to an evaluation of the decision to ascertain whether it was rewarded (30). This would mean that DPC is involved in a reemergence of the decision, with no relationship to movement whatsoever; if the decision-related activity in DPC occurs independently from movement, then its role must be far more nuanced than previously proposed. This could be seen as part of a feedback loop that can continually reinforce the somatosensory processing network's performance based on recent experience.

However, if previous evidence has shown that abstract categorical decision signals emerge as early as S2, what would be the function of DPC's decision signal? The answer may lie in the temporal dynamics of the signal. If the frontal lobe does not directly produce these abstract signals, it might be more concerned with a) providing the dynamical infrastructure to sustain abstract information persistently in working memory, b) associating it with relevant contextual information, c) broadcasting the decision for motor execution, and d) distributing feedback about the outcome afterward. In our results, points a) and d) might be the most evident; one during the delay between stimulus presentation and decision report, and the other during the intertrial period. In any case, we would like to highlight a) as a proximate computational answer to the question of DPC's function: it would seem that this area is not computing the category—that might happen in S2—but rather developing it into an abstract, persistent signal.

In brief, the results presented here show that a common code is employed to maintain the categorization of differing physical attributes each presented in two partially overlapping ranges. Here, we would like to mention that this mechanism for categorical coding is similar to those that have been described for other, different tasks 49. In addition, we have demonstrated through dimensionality reduction techniques that this singular categorical code emerges from the entire DPC's neuronal population and not from isolated functional groups. We think that this could be favored thanks to the broad range of heterogeneous responses—a mixture between categorical and temporal dynamics—that single units tend to present (21), which is also known to improve neural network performance 50.

The evidence presented here supports the attractor model of working memory: the work has been written from that perspective. Our phase portraits of population activity and our examples of single-unit's persistent activity are in line with such a model. However, alternative models have been proposed, like the activity-silent model, about which there is active discussion regarding their theoretical implications and relation to experimental data. For an outstanding review on attractor models of mnemonic persistent activity, see ref. 51. Given that the decision code in our data emerges explicitly (not from hidden variables) from the heterogeneous activity of units (21) and their population trajectories, the activity-silent model does not seem to apply to our experimental data. Specifically, our observation of four distinct attractors in phase space, corresponding in pairs to two orthogonal decision dimensions, appears even harder to reconcile. Additionally, the firing-rate distributions of the delay and the foreperiod are equal, as reported for other tasks (26): this also clashes with an argument for the activity-silent model based on energetic efficiency 51. Nonetheless, further and more theoretical analysis should be performed. For example, the activity-silent model might be particularly interesting to consider when studying intertrial dynamics like the ones glimpsed in our work, since its theorized biological substrate could enable history dependence across trials 51.

To conclude, we hope our current work highlights the ability of frontal lobe neurons to adapt strategies to a variety of task conditions. This adaptability could be further tested by asking whether the categorical representation of stimuli duration would also be the same as for the attributes studied here. Furthermore, we have observed highly abstracted signals in DPC's activity, but the origin of these signals remains to be examined under the structure of the VCT. Note, we think that most of the ideas discussed throughout also apply to other premotor (associative) areas involved in perceptual decision-making, not only DPC in particular. Finally, we believe that for future avenues of research, it would be greatly informative to study other kinds of data, such as local field potentials (LFPs), to gain further understanding of how the network utilizes all possible strategies to coordinate categorical decision-related responses.

## Materials and Methods

Two monkeys were trained to label the intensity of different physical attributes (frequency or amplitude) of a vibrotactile stimulus as "high" or "low" (Fig. 1 and *SI Appendix*). Neuronal recordings were obtained in the DPC, either contralateral (left hemisphere) or ipsilateral (right hemisphere), while the monkeys performed the categorical tasks. The animals were handled in accordance with the standards of the NIH and Society for Neuroscience. All protocols were approved by the Institutional Animal Care and Use Committee of the Instituto de Fisiología Celular, Universidad Nacional Autónoma de México.

**Data, Materials, and Software Availability.** Data files are publicly available at Zenodo (https://zenodo.org/record/7293152) (52).

**ACKNOWLEDGMENTS.** We thank Hector Diaz for his technical assistance. This work was supported by grants: PAPIIT-IN205022 (to R.R.-P.) and PAPIIT-IG200521 (to V.d.L.) from the Dirección de Asuntos del Personal Académico de la Universidad Nacional Autónoma de México; CONACYT-319347 (to R.R.-P.), CONACYT-319212 (to V.d.L.), and CB2014-20140892 (to R.R.) from Consejo Nacional de Ciencia y Tecnología; IBRO Early Career Award 2022 (to R.R.-P.) from International Brain Research Association. G.D.-d.L. (fellowship CONACYT-964544), S. Parra and J.Z. are doctoral students from Programa de Doctorado en Ciencias Biomédicas, UNAM. L.B. is a postdoctoral fellow (Postdoctoral fellowship CONACYT-838783).

Author affiliations: [a]Instituto de Neurobiología, Universidad Nacional Autónoma de México, Querétaro 76230, Mexico; [b]Instituto de Fisiología Celular, Departamento de Neurociencia Cognitiva, Universidad Nacional Autónoma de México, 04510, Mexico City, Mexico; and [c]Centro de Ciencias de la Complejidad, Universidad Nacional Autónoma de México, Mexico City, 04510, Mexico

1. A. Hernández *et al.*, Decoding a perceptual decision process across cortex. *Neuron* **66**, 300–314 (2010).
2. V. de Lafuente, R. Romo, Neuronal correlates of subjective sensory experience. *Nat. Neurosci.* **8**, 1698–1703 (2005).
3. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
4. R. Rossi-Pool *et al.*, Emergence of an abstract categorical code enabling the discrimination of temporally structured tactile stimuli. *Proc. Natl. Acad. Sci. U.S.A.* **113**, E7966–E7975 (2016).

5. E. D. Remington, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019.e5 (2018).
6. J. A. Cromer, J. E. Roy, E. K. Miller, Representation of multiple, independent categories in the primate prefrontal cortex. *Neuron* **66**, 796–807 (2010).
7. A. Sarma, N. Y. Masse, X.-J. Wang, D. J. Freedman, Task-specific versus generalized mnemonic representations in parietal and prefrontal cortices. *Nat. Neurosci.* **19**, 143–149 (2015).
8. W. Chaisangmongkon, S. K. Swaminathan, D. J. Freedman, X.-J. Wang, Computing by robust transience: How the fronto-parietal network performs sequential, category-based decisions. *Neuron* **93**, 1504–1517.e4 (2017).
9. S. J. Goodwin, R. K. Blackman, S. Sakellaridi, M. V. Chafee, Executive control over cognition: Stronger and earlier rule-based modulation of spatial category signals in prefrontal cortex relative to parietal cortex. *J. Neurosci.* **32**, 3499–3515 (2012).
10. J. Vergara, N. Rivera, R. Rossi-Pool, R. Romo, A neural parametric code for storing information of more than one sensory modality in working memory. *Neuron* **89**, 54–62 (2016).
11. R. Romo, R. Rossi-Pool, Turning touch into perception. *Neuron* **105**, 16–33 (2020).
12. M. Alvarez, A. Zainos, R. Romo, Decoding stimulus features in primate somatosensory cortex during perceptual categorization. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 4773–4778 (2015).
13. P. Cisek, J. F. Kalaska, Neural of correlates of mental rehearsal in dorsal premotor cortex. *Nature* **431**, 993–996 (2004).
14. P. Cisek, J. F. Kalaska, Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814 (2005).
15. J. I. Glaser, M. G. Perich, P. Ramkumar, L. E. Miller, K. P. Körding, Population coding of conditional probability distributions in dorsal premotor cortex. *Nat. Commun.* **9**, 1–14 (2018).
16. M. T. Kaufman *et al.*, Roles of monkey premotor neuron classes in movement preparation and execution. *J. Neurophysiol.* **104**, 799–810 (2010).
17. M. M. Churchland *et al.*, Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
18. B. M. Dekleva, K. P. Kording, L. E. Miller, Single reach plans in dorsal premotor cortex during a two-target task. *Nat. Commun.* **9**, 3556 (2018).
19. M. Wang *et al.*, Macaque dorsal premotor cortex exhibits decision-related activity only when specific stimulus–response associations are known. *Nat. Commun.* **10**, 1793 (2019).
20. G. F. Elsayed, A. H. Lara, M. T. Kaufman, M. M. Churchland, J. P. Cunningham, Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239 (2016).
21. M. Rigotti *et al.*, The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
22. R. Rossi-Pool *et al.*, Decoding a decision process in the neuronal population of dorsal premotor cortex. *Neuron* **96**, 1432–1446.e7 (2017).
23. M. T. Kaufman, M. M. Churchland, S. I. Ryu, K. V. Shenoy, Cortical activity in the null space: Permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
24. D. Kobak *et al.*, Demixed principal component analysis of neural population data. *Elife* **5**, e10989 (2016).
25. F. Carnevale, V. de Lafuente, R. Romo, O. Barak, N. Parga, Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron* **86**, 1067–1077 (2015).
26. J. D. Murray *et al.*, Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 394–399 (2017).
27. R. Rossi-Pool *et al.*, Temporal signals underlying a cognitive process in the dorsal premotor cortex. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 7523–7532 (2019).
28. V. de Lafuente, R. Romo, Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 14266–14271 (2006).
29. J. L. Pardo-Vazquez, V. Leboran, C. Acuña, A role for the ventral premotor cortex beyond performance monitoring. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 18815–18819 (2009).
30. A. Akrami, C. D. Kopec, M. E. Diamond, C. D. Brody, Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).
31. G. F. Elsayed, J. P. Cunningham, Structure in neural population recordings: An expected byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318 (2017).
32. J. Wang, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
33. R. B. Ebitz, B. Y. Hayden, The population doctrine in cognitive neuroscience. *Neuron* **109**, 3055–3068 (2021).
34. E. M. Trautmann *et al.*, Accurate estimation of neural population dynamics without spike sorting. *Neuron* **103**, 292–308 (2019).
35. D. Kobak, J. L. Pardo-Vazquez, M. Valente, C. K. Machens, A. Renart, State-dependent geometry of population activity in rat auditory cortex. *Elife* **8**, e44526 (2019).
36. L. McInnes, J. Healy, N. Saul, L. Großberger, UMAP: Uniform manifold approximation and projection. *J. Open Source Softw.* **3**, 861 (2018).
37. S. Parra *et al.*, Hierarchical unimodal processing within the primary somatosensory cortex during a bimodal detection task. (2022). *PNAS*, doi.org/10.1073/pnas.2213847119
38. R. J. Gardner *et al.*, Toroidal topology of population activity in grid cells. *Nature* **602**, 123–128 (2022).
39. A. Rodriguez, A. Laio, Clustering by fast search and find of density peaks. *Science* **344**, 1492–1496 (2014).
40. R. Romo, E. Salinas, Touch and go: Decision-making mechanisms in somatosensation. *Annu. Rev. Neurosci.* **24**, 107–137 (2001).
41. R. Rossi-Pool, A. Zainos, M. Alvarez, G. Diaz-deLeon, R. Romo, A continuum of invariant sensory and behavioral-context perceptual coding in secondary somatosensory cortex. *Nat. Commun.* **12**, 2000 (2021).
42. A. Fassihi, A. Akrami, V. Esmaeili, M. E. Diamond, Tactile perception and working memory in rats and humans. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 2331–2336 (2014).
43. J. H. Siegle *et al.*, Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021), 10.1038/s41586-020-03171-x.
44. B. Pesaran, M. J. Nelson, R. A. Andersen, Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron* **51**, 125–134 (2006).
45. V. Papadourakis, V. Raos, Neurons in the macaque dorsal premotor cortex respond to execution and observation of actions. *Cereb. Cortex* **29**, 4223–4237 (2019).
46. D. Thura, P. Cisek, Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* **81**, 1401–1416 (2014).
47. É. Coallier, T. Michelet, J. F. Kalaska, 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 (2015).
48. S. Bestmann *et al.*, Dorsal premotor cortex exerts state-dependent causal influences on activity in contralateral primary motor and dorsal premotor cortex. *Cereb. Cortex* **18**, 1281–1291 (2008).
49. K. Mohan, O. Zhu, D. J. Freedman, Interaction between neuronal encoding and population dynamics during categorization task switching in parietal cortex. *Neuron* **109**, 700–712.e4 (2021).
50. N. Perez-Nieves, V. C. H. Leung, P. L. Dragotti, D. F. M. Goodman, Neural heterogeneity promotes robust learning. *Nat. Commun.* **12**, 5791 (2021).
51. X.-J. Wang, 50 years of mnemonic persistent activity: Quo vadis? *Trends Neurosci.* **44**, 888–902 (2021).
52. M. Alvarez *et al.*, An abstract categorical decision code in dorsal premotor cortex (2022), https://zenodo.org/record/7293152.