Multiplexing behavioral signals in sensory representations

Julia M. Mayer[1,2] and Wiktor F. Młynarski[1,3]

[1]Faculty of Biology, LMU Munich, Germany, [2]Graduate School of Systemic Neurosciences, Munich, Germany

[3]Bernstein Center for Computational Neuroscience Munich

for correspondence: mlynarski@bio.lmu.de

Activity of sensory neurons is influenced not only by external stimuli but also by the animal's behavioral state. It is well documented that behavior influences the general properties of neural activity, such as response gain. However, it is not known whether it could affect the sensory tuning of individual neurons in a more refined way and what the functional benefit of such nuanced modulation might be. Here, we investigate this in the mouse visual cortex, where sensory response gain varies with locomotion speed. First, using numerical simulations, we demonstrate that gain modulation can multiplex behavioral information in sensory populations, without compromising the accuracy of sensory coding. To implement such multiplexing, behavioral signals should modify the sensory tuning in individual neurons. Second, we analyze neural activity in the mouse visual cortex made available by the Allen Brain Observatory. Our analysis indicates that, in agreement with the theory, locomotion-induced gain modulation can modify the tuning of sensory neurons to direction of visual motion. In that way, the visual cortex could instantiate an accurate, joint representation of sensory and movement-related signals and support computations that simultaneously require both types of information.

INTRODUCTION

With the advent of neural recordings in behaving animals it became clear that sensory neurons are continuously modulated by multiple non-sensory signals [1]. For example, neurons in the primary visual cortex (V1) of rodents can be modulated by spontaneous behavior and body movements [2, 3], location in environment [4], head and eye position [5], internal states such as arousal [6] and attention [7].

Modulation of sensory neurons by animal movement raises general questions about sensory coding during active behavior. It may intuitively appear that if neurons in visual areas are simultaneously affected by many behavioral variables, the precision of sensory representation must deteriorate. It could in principle be more beneficial to represent sensory signals and behavioral states in separate neural populations and, if necessary, combine them only in higher brain areas, not in the early stages of sensory processing.

A behavioral state that exerts a particularly strong impact on sensory activity is locomotion speed [8]. Neurons in the visual system can dramatically modulate their gain depending on the movement of the animal [9–12]. Although this effect is pronounced and well documented, its functional benefits remain unknown. A key obstacle in understanding the computational role of running speed–dependent gain modulation is the limited knowledge of its potential nuances, specifically, whether running speed can alter sensory tuning so that neurons prefer different stimuli at different speeds.

Here, we directly address both questions. We postulate a putative computational objective for neurons in the visual cortex. To satisfy this objective, behavior should modulate sensory tuning in different behavioral states. We analyze neural activity in the mouse visual cortex and provide evidence that neurons indeed modify their tuning depending on the animal's speed.

Our guiding hypothesis is that sensory neurons modulated by behavioral signals represent a joint state of the organism and its environment. Such representations could support behavioral planning and action selection, particularly in smaller nervous systems with a limited number of neurons. We used numerical simulations to ask how gain modulation can influence the simultaneous decoding of stimuli and the modulating behavioral signals by downstream areas. We considered different cascade models of sensory neurons that combine simpler sensory features. Each of those features is represented by an independent subunit that can be modulated by the behavioral state. We found that, depending on their architecture, these neuron models trade off the accuracy of encoding of sensory and behavioral variables. Surprisingly, however, behavioral modulation can have a positive effect on stimulus encoding accuracy as long as the features combined by the neuron are modulated differentially by behavior. This neuron architecture retains information about the modulating variable, thus instantiating an accurate joint representation of stimuli and the behavioral state.

To make this hypothesis specific, we focus on modulation of tuning to the visual motion direction in the visual cortex by running speed of a mouse. If the visual cortex implements a joint representation of these signals, visual neurons should change their sensory tuning depending on the running speed of the mouse. To the best of our knowledge, such nuanced modulation of sensory representations by behavioral state has not been reported previously.

We tested this prediction by analyzing two-photon calcium imaging data from the Allen Brain Observatory [13]. We computed two-dimensional tuning curves that characterize how neurons that encode direction of visual motion are modulated by running speed. In agreement with theoretical predictions, we found that the tuning of sensory neurons in the mouse visual cortex is modulated by the behavioral state. This modulation takes a specific form - running speed adjusts weights of direction-selective subunits pooled by each neuron. Using a sparse decomposition of orientation-velocity tuning curves, we extracted the shape of subunit gain as a function of the animal running speed. We found a diversity of gain modulation patterns used by cortical neurons to dynamically change the weights of sensory inputs at different locomotion speeds. These results suggest that the mouse visual cortex instantiates an accurate multiplexed representation of the behavioral state and the sensory environment.

RESULTS

Multiplexing sensory and behavioral information via gain modulation

To examine how sensory neurons can combine stimuli and behavioral signals, we consider three different neuron models of increasing complexity (Fig. 1 a-c). The first model (Fig. 1a, top) consists of a single subunit selective for a preferred direction of visual motion, whose output is transformed with a nonlinearity to generate stochastic spiking activity. The gain of the nonlinearity is modulated by a behavioral variable (for example the running speed), and this modulation can take different forms, e.g. monotonically increasing or non-monotonic (Fig. 1a, bottom, gray lines). The average activity of the neuron as a function of the stimulus direction and the running speed is characterized by two-dimensional tuning curves (two examples are plotted in Fig. 1a, bottom).

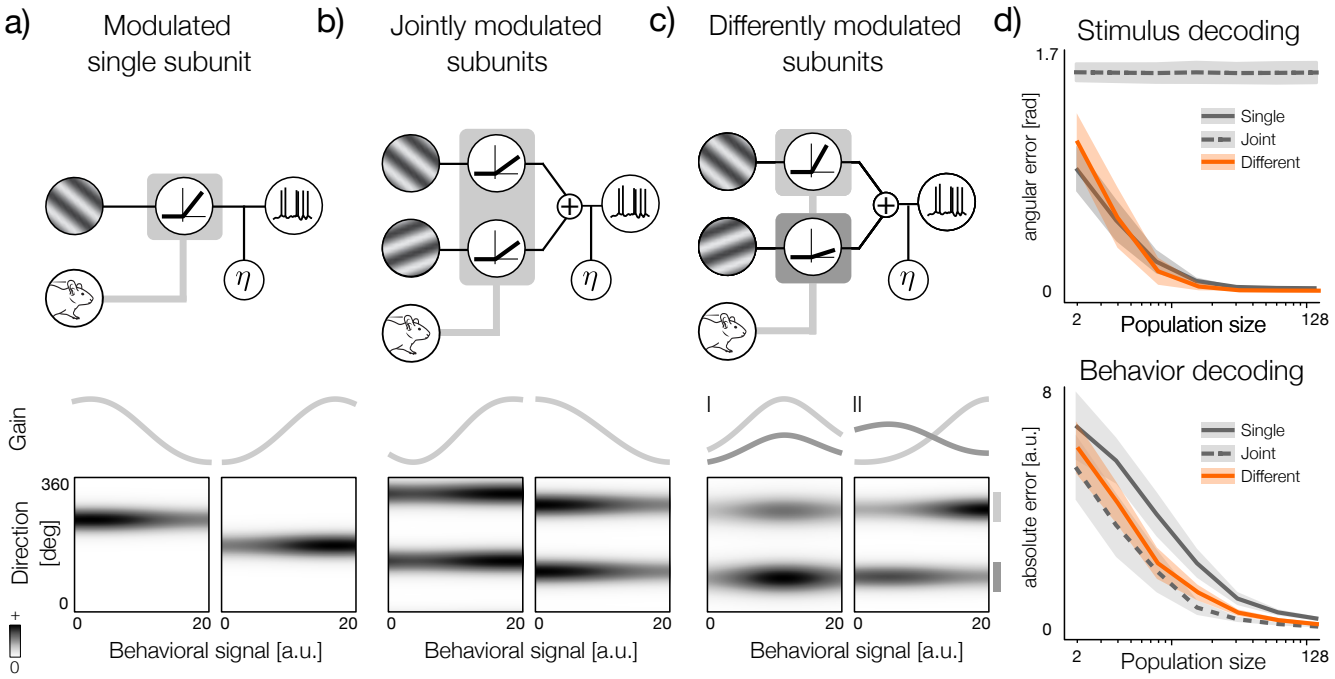

Figure 1: Multiplexing sensory and behavioral signals with gain modulation. a) Top: the simplest neuron model with a single direction-selective subunit and nonlinearity modulated by the running speed. By adjusting the nonlinearity, the behavioral variable modulates the gain of the neural response. The nonlinearity of the output is distorted by noise η Bottom: two example sensory-behavioral tuning curves generated by single-subunit neurons with different preferred direction of visual motion and behavior-controlled gain modulation. The shape of the behavioral modulator is plotted above tuning curves. **b)** The same as in a) but for a neuron model that pools multiple (here two) subunits that are jointly modulated by the behavioral signal. **c)** The same as in a) but for a neuron model that that pools multiple (here two) subunits that are independently modulated by the behavioral signal. Modulators of each subunit are plotted in the bottom panel as lines in different shades of gray. Modulators have different amplitudes (I) or very different shapes as a function of the behavioral variable (II). **d)** Top: simulated stimulus decoding from populations of different neuron models. Lines denote errors averaged over multiple random populations and shaded bars denote the standard deviation of the error. Single subunits (gray solid line and shaded area) and differently modulated multi-subunit (orange line and shaded area) models have very similar performance. Jointly modulated multiple subunit model (gray dashed line and gray are) has the constant average error of approximately $\pi/2$ regardless of population size. Bottom: simulated decoding of the behavioral variable that modulates neuronal activity. Colors and markers the same as above. Single subunit model shows lowest population performance while multiple subunit models show comparable performance.

The second neuron model combines two (or more) different subunits that prefer different directions of visual motion (Fig.1b, top). If the directions preferred by individual subunits are separated by 180 degrees, the neuron is orientation-selective. The behavioral state jointly modulates the nonlinearities of both subunits, and their gain changes simultaneously as reflected in the example tuning curves (Fig. 1b, bottom).

In the two models described above, the sensory tuning of a neuron does not change with the behavioral signal. In the third model, the behavioral state modulates the gain of each subunit in a different way (Fig.1c, top). In addition to differences in their trajectory, modulators can also have different amplitudes (Fig.1c, bottom, panels denoted with Roman numerals). The resulting tuning curves reveal that neurons belonging to this class can change their preferred stimulus direction, depending on the relative gain of each subunit (Fig.1c, bottom).

To understand how neurons of each type encode stimulus direction and the behavioral signal, we simulated population activity and performed a decoding analysis. We assumed that the downstream area randomly pools neurons of a given type and simulated decoding from multiple random populations (see Methods). This simple approach has a number of advantages. First, it allows us to study the encoding performance of a given neuron model without explicitly optimizing the population. We thus focus on distributions of solutions, instead of just a single optimum. Second, it allows us to study how well behavioral and sensory signals are encoded by different neural architectures without reference to the specific computational task that those neurons may perform. The only assumption is that these tasks require an accurate representation of the stimuli and the behavioral state. Third, random grouping requires less fine-tuning and the nervous system could easily implement such a pooling scheme [14, 15].

As expected, we found that direction-selective, single-subunit neurons are best at encoding sensory stimuli - the average error decreases fastest as a function of the population size (Fig. 1d, top panel, solid gray line). In contrast, neurons whose two subunits are jointly modulated yield the highest average error (Fig. 1d, top panel, gray dashed line) because stimulus representations are ambiguous up to 180 degrees in populations of orientation-tuned neurons. Relative performance reverses when decoding the modulating behavioral variable (Fig. 1d, bottom panel). Here, neurons with jointly modulated subunits represent the behavioral variable with the highest accuracy at a given population size. Single-subunit neurons typically require twice as large a population size to achieve the same level of decoding error. This is because, in order for behavior signals to be encoded at all, they must coincide with the stimuli represented by a neuron. Because orientation-tuned neurons are tuned to more stimuli than direction-tuned ones, they are better at encoding behavior.

Decoding analysis demonstrates that populations of direction- and orientation-specific neurons trade off the accuracy encoding of sensory inputs and behavioral signals. Interestingly, the third neuron model, in which subunits are modulated differently, navigates this trade-off and achieves good performance for both variables (Fig. 1d, top and bottom panels, orange line). Behavioral modulation differentiates responses to inputs encoded by individual subunits, thus resolving the ambiguity present in orientation-tuned cells. The resulting decoding performance is very close to that of direction-selective neurons. At the same time, because these neurons combine more than one stimulus direction, they encode behavior with higher accuracy than populations of single subunit cells of the same size. Therefore behavioral gain modulation resolves the ambiguity between individual subunits and thus tunes the sensory code while simultaneously allowing for accurate representation of the modulating behavioral input.

Estimation of joint tuning to the stimulus direction and running speed in the mouse visual cortex

The simulation analysis described above forms a specific experimental prediction: if the visual cortex uses gain modulation to instantiate a joint representation of the direction of visual motion and running speed, we should be able to identify neurons whose sensory tuning is modulated by the behavioral state. To test this hypothesis, we explored two-photon calcium imaging recordings of the mouse visual cortex made available by the Allen Brain Institute [13].

The first step of our analysis was to determine the simultaneous tuning of neurons in the mouse visual cortex to a stimulus direction and running speed, important sensory and behavioral variables. The Allen Brain Institute recordings were performed in head-fixed behaving mice whose running speed was acquired simultaneously with neural activity (Fig. 2a). Mice were exposed to drifting grating stimuli at different temporal frequencies and movement directions (see Methods). To understand how sensory tuning interacts with behavior, we needed to estimate the tuning at different running speeds. This presented a challenge because individual animals had different behavioral preferences and generated very different distributions of running speed. Since our aim was to pool and compare the data recorded in multiple sessions and animals, we could not rely on quantile binning as suggested before [9, 10]. Instead, we used 8 bins linearly spaced between 0 and 20 cm/s.

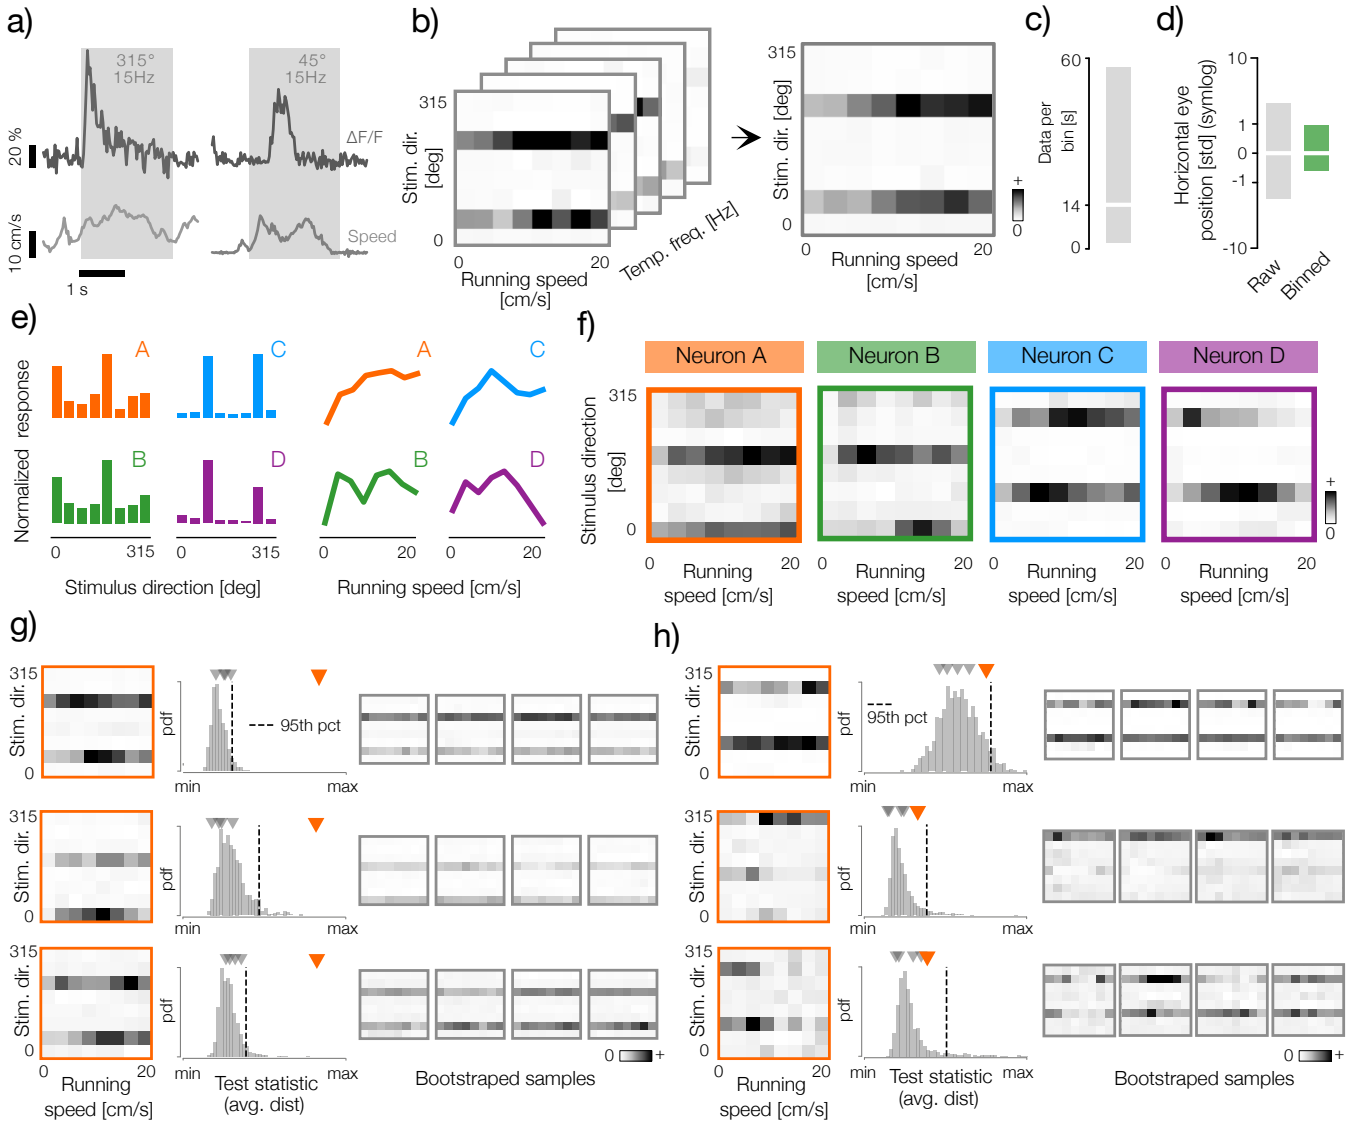

Figure 2: Analysis of tuning to stimulus direction and running speed in the mouse visual cortex. a) Example traces of the calcium imaging data (dF/F) and of mouse running speed obtained from the Allen Brain Observatory (top, dark gray lines and bottom, light gray lines respectively). Gray rectangles mark duration of the stimulus whose orientation and temporal frequency are marked with gray text. **b)** Left panel: example three-dimensional (running speed - stimulus direction - stimulus temporal frequency) tuning curve of a neuron computed using thresholded dF/F signal (see Methods). Right panel: a two-dimensional (running speed - stimulus direction) tuning curve obtained by averaging the three-dimensional tuning curve on the left across the temporal frequency dimension. **c)** Distribution of seconds of data per bin in all analyzed datasets. The horizontal white line denotes the average (14 s) and the gray rectangle denotes the interval spanned by 5th and 95th percentiles of the distribution. **d)** Distributions of the horizontal eye positions pooled across raw recordings (left, gray bar) and averaged within bins of the orientation-speed tuning curves (right, green bar). White lines denote the mean (0 deg) and rectangles denote the interval spanned by 5th and 95th percentiles. Data for each animal was standardized independently via z-scoring. The scale of the plot is symlog (i.e. $\operatorname{sign} x \log_{10}(|x|)$). The eye position data was not available for all analyzed datasets. **e)** Left panel: orientation tuning curves of four example neurons. Right panel: running-speed tuning curves of the same neurons as in the left panel (marked with corresponding letter and color). **f)** Joint orientation-speed tuning curves of four example neurons marked with corresponding letters and colors in e. **g)** Example neurons that pass the bootstrap test of velocity modulation. Left column: orientation-speed tuning curves of four example neurons (orange frames). Middle column: distributions of the test statistic of the modulation test (see Methods). Vertical dashed line denotes the significance threshold equal to the 95th percentile of the statistic distribution. Orange and gray triangles denote respectively the values of the test statistic of the original tuning curve and five example bootstrapped samples visualized in the right column. Right column: five example bootstrapped orientation-speed tuning curves. **h)** Example neurons that do not pass the bootstrap test of velocity modulation. Panel is analogous to g).

Using this binning scheme, we computed three-dimensional tensors ("tuning curves") for each neuron by pooling and averaging the normalized dF/F signal for each bin specified by the stimulus orientation, it's temporal frequency, and the running speed of the animal (Fig. 2b, left panel). Each three-dimensional tensor was then averaged across the temporal frequency axis, generating a two-dimensional direction-speed matrix, which we refer to as a direction-speed tuning curve (Fig.2b, right panel). To ensure that the tuning curves were computed using a comparable set of

stimuli, we selected only the data sets with minimal differences in stimulus frequencies presented in each direction and running speed (see Methods). In this way, we aggregated a reasonable amount of unbiased data for each direction-speed bin (average 14 seconds, Fig.2c). An additional benefit of this pooling scheme was the reduction of the eye position variability within each bin (Fig. 2d, gray and green bars, respectively).

Since the focus of our study is on adjustments of sensory code by behavioral variables, we limited our analysis only to neurons that are significantly tuned to stimulus direction, as measured by the sharpness of their tuning to stimulus direction (see Methods). Tuning of four example neurons to the stimulus direction and the running speed is shown in Fig.2e (left and right panels, respectively). Pairs of neurons denoted by A, B as well as C, D have similar sensory preference and respond with comparable strength to stimuli with the same direction. Behavioral tuning does not appear to be related to the sensory preferences of these cells. Here, different pairs of cells (A, C and B, D) respond similarly to the running speed of the animal. If these example neurons encoded behavioral and sensory signals independently of each other, their tuning could be expressed in full using these two one-dimensional tuning curves. However, inspection of the joint tuning to running speed and stimulus direction reveals that this is not the case (Fig.2f). In each example neuron, sensory tuning visibly varies across the considered range of running speed. Although neurons do not entirely shift their orientation-tuning curves, they visibly modulate the gain of their responses to preferred stimuli in different behavioral states. The gain modulation patterns are distinct in each cell, suggesting that they represent different joint features of sensory inputs and behavior.
The example neurons displayed in Fig.2e-f clearly reveal different patterns of joint direction-speed tuning. However, their tuning curves were computed using different amounts of data, which itself are variable and noisy. To ensure that the computed tuning curves reflect preferences of individual neurons, and not biases in the data distribution or noise, we performed a bootstrap analysis (Fig.2g-h). For each neuron, we randomly shuffled the data across running speed values and computed estimates of the direction-speed tuning curves (Fig.2g-h, right columns). This bootstrapping scheme preserved the sensory tuning of each cell (both to the stimulus direction and temporal frequency) and the amount of data per bin, modifying only the impact of the running speed on the response gain (see Methods for details). The majority of neurons were significantly modulated by behavior, exceeding the chance level established via bootstrapping(Fig.2g). Neurons that did not pass the significance test (e.g. examples in Fig.2h) were excluded from further analysis.

Sparse decomposition of tuning curves reveals subtle effects of running speed on sensory tuning

The shape of the orientation-speed tuning curves in the mouse visual cortex suggests a specific computational architecture of individual neurons, discussed in the first part of the results section. The data indicate that neurons combine the input of direction-selective subunits, whose gain is modulated by running speed (Fig.3a). To separate the impact of movement on these individual subunits, we fitted each tuning curve with the sparse-decomposition model (Fig.3b). The components of the model correspond directly to elements of the neuronal architecture in Fig.3a. The model represents the average neural activity summarized in a tuning curve as a sum of basis functions that correspond to sensory subunits tuned to a specific direction of visual motion (Fig.3b, second and forth panels from the left). The row of each subunit basis function is point-wise multiplied by a corresponding velocity modulator, i.e. the curve that describes how the gain of that subunit changes with the speed of the animal (Fig.3b, third and fifth panels from the left). During fitting, we placed a separate sparse penalty on individual subunits and on modulators, which increased the robustness and quality of the fits (see Methods). We excluded from further analysis cells that were not significantly well fitted by the sparse decomposition model, which we assessed by bootstrap testing (see Methods). The tuning curves of the retained neurons were well approximated by the model fits, with the average fit quality equal to 8 dB SNR (Fig. 3c).

The parameter of the fitted model identified the selectivity for direction and orientation in the mouse visual cortex, consistent with previous findings [16, 17]. The largest number of neurons had 2 significant subunits (N=753, see Methods for the definition of significant subunits). If two subunits differ by 180 degrees in their sensory preference, they instantiate orientation tuning, which was predominantly the case (e.g. Fig.3g). The second largest group of neurons (N=671) was best fitted by a single significant subunit (e.g. Fig.3e), which corresponds to direction selectivity. In other neurons, additional significant subunits typically modeled broadened sensory tuning. We therefore concentrate our analysis on direction and orientation selective cells with either 1 or 2 significant subunits, respectively.

The sparse decomposition model successfully extracted speed modulators from tuning curves of direction-selective neurons (Fig.3f). Estimated modulators take different forms e.g. increasing (Fig.3f, neuron C) or non-monotonic and preferring a specific speed (Fig.3f, neurons A, B, D). These findings are in agreement with previous studies that reported a diversity of speed tuning in the visual system of a mouse [9–11]

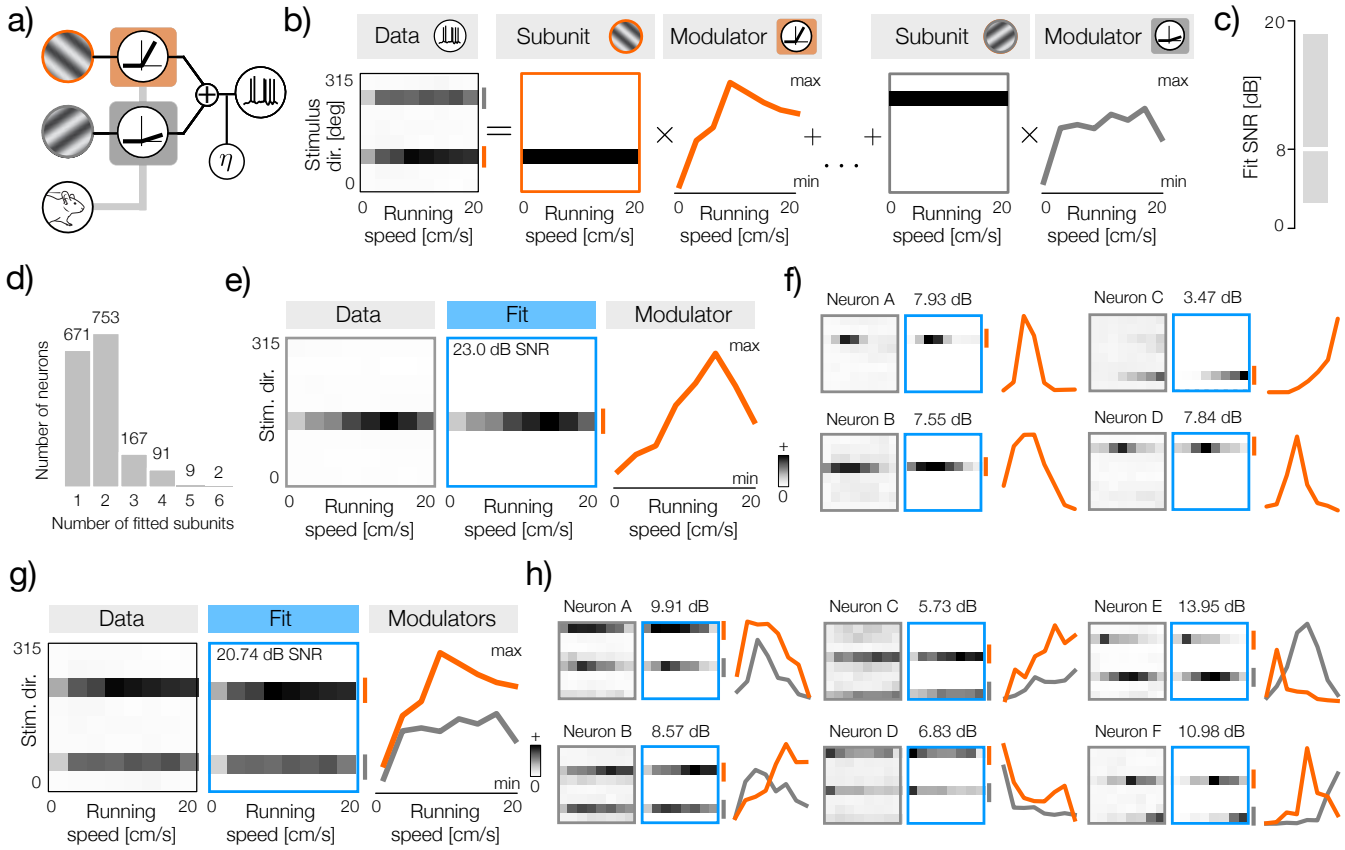

Figure 3: Sparse decomposition of direction-speed tuning curves. **a)** Conceptual schematic of a neuron where two sensory subunits (orange and gray circles) are modulated independently by the running speed of the animal. **b)** Schematic of the statistical model that decomposes stimulus-speed tuning curves (left column) into a sum of subunit basis functions (second and fourth panels from the left) multiplied by speed-controlled gain modulators (third panel from the left and right panel). The subunit basis functions determine sensory tuning of a neuron, while modulators determine modulation of each subunit by the running speed of the animal. **c)** Distribution of the SNR values of model fits to stimulus-speed tuning curves. The horizontal white line denotes the average (8 dB SNR) and the gray rectangle denotes the interval spanned by 5th and 95th percentiles of the distribution. **d)** Histogram of the number of subunits fitter to each neuron. Note that not a single neuron had 7 or 8 significant subunits. **e)** Left panel: example stimulus-speed tuning curve of a direction-selective neuron with a single sensory subunit. Middle panel: reconstruction of the tuning panel from the model fit. The single fitted sensory subunit is marked with a short vertical orange line on the right side. The SNR of the fit is equal to 23 dB. Right panel: speed modulator of the direction-selective subunit. **f)** Four example direction-selective neurons (gray frames) together with their respective fits (blue frames) marked with fitted subunits (short, vertical orange lines on the right side of each blue frame) and speed-modulators (orange lines). **g)** Same as e but for an example direction-selective neuron with two sensory subunits. The subunits and their modulators are now marked with gray and orange lines. **h)** Same as f but for six example orientation-selective neurons.

In orientation-tuned neurons, we obtain two modulators, one for each direction-selective subunit. (Fig.3h). Interestingly, these modulators can differ significantly between subunits. The differences between the modulators take various forms. Individual subunit modulators can increase or decrease at different rates, instantiating differential sensory tuning at various running speeds (Fig.3, neurons C, D). Individual modulators are also often non-monotonic and can either peak at similar running speeds (Fig.3h, neuron A) or, in a dramatic demonstration of the shift in sensory tuning, amplify responses at different speeds of locomotion (Fig.3h, neurons B, E, F).

Speed induced gain modulation impacts sensory tuning

As indicated by the theoretical analysis, behavioral and sensory signals are most accurately represented by orientation-selective neurons that combine two direction-specific subunits. Moreover, only in this neuron architecture can sensory tuning can be modified by behavior via gain modulation of subunits. To better understand these phenomena, we focus our analysis on orientation-selective neurons with two significant subunits.

In a large majority of two-subunit neurons (83%), individual subunits encode stimulus directions separated by 180 degrees, implementing the "real" orientation selectivity (Fig.4a, right panel). The joint distribution of stimulus directions preferred by individual subunits does not reveal any systematic bias (Fig.4a, left panel).

Orientation-selective cells combine individual subunits asymmetrically, with non-equal weights. We quantified this effect by averaging modulators of each subunit and taking the ratio of the stronger and weaker modulators. Although the most likely ratio is equal to 1, the average ratio of the modulator strengths of the subunit is 3.36 and spans a wide range as indicated by the distribution of the ratio logarithms (Fig.4c). The asymmetry of subunit modulators is one of the ways to remove the ambiguity between subunits separated by 180 degrees and to improve the accuracy of the sensory code by multiplexing behavioral signals.

The similarity of subunit modulators in the same neuron can be easily quantified with correlation. The correlations of the modulator pairs are broadly distributed covering a wide range from highly positive to highly negative values (Fig.4c, orange bars). Although the distribution is broad, it is clearly skewed towards more positive correlations, indicating that many modulator pairs are similar. To ensure that this pattern is not random, we computed correlations between randomly permuted modulators (Fig.4c, gray bars), which is significantly different from data (KS-test, p-value < 0.001).

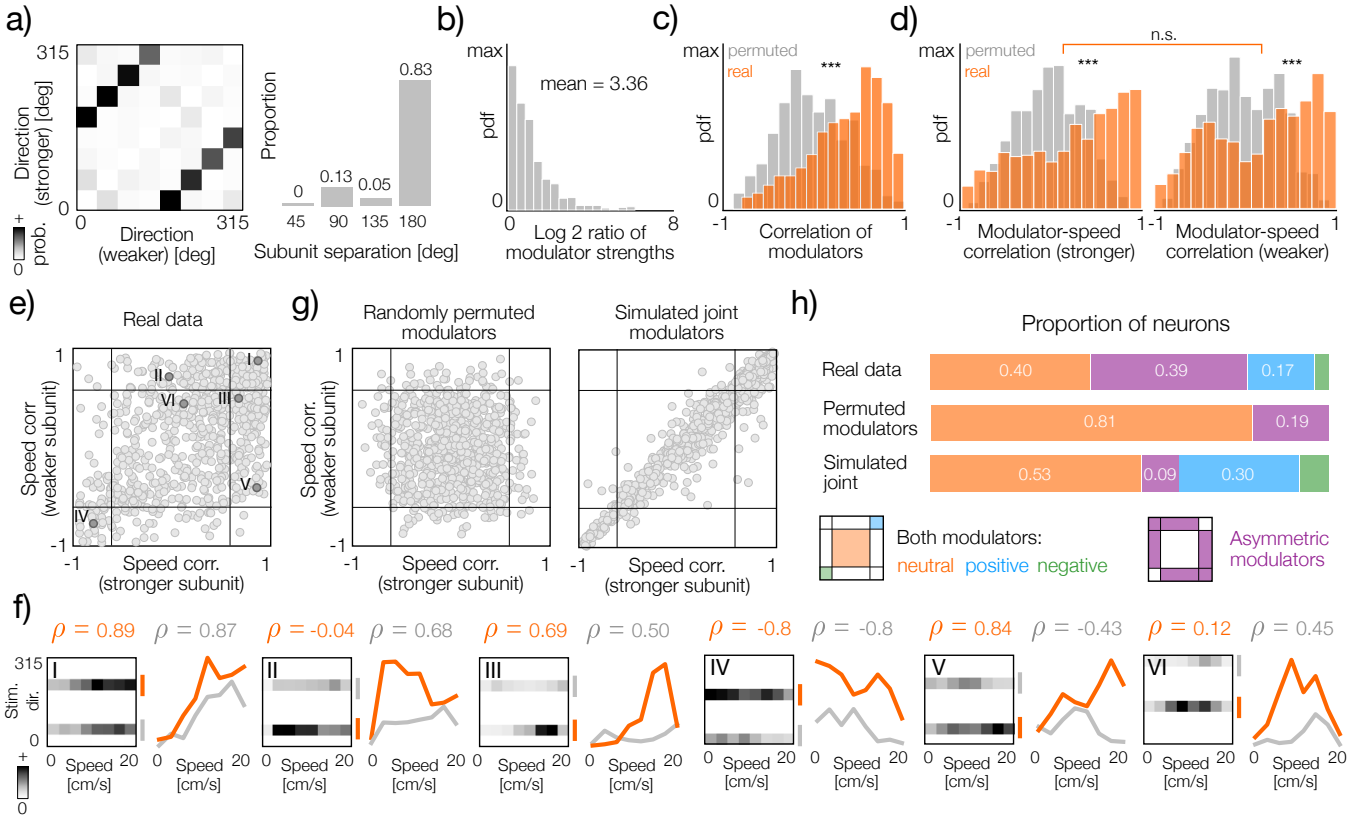

Figure 4: Behavior differentially modulates sensory subunits in orientation-selective cells. a) Left panel: joint distribution of stimulus directions preferred by subunits in two-subunit neurons. Right panel: histogram of angular separation (distance) between stimuli preferred by individual subunits in two-subunit neurons. **b)** Histogram of log-ratios of modulator strength (stronger to weaker). The mean equal to 3.36 was computed before logarithmic transform. **c)** Histograms of correlations between modulators in two-subunit neurons. Orange and grey correspond to the original data and randomly permuted modulators respectively. **d)** Histograms of correlations between modulators and running speed in stronger and weaker subunits (left and right panels respectively). Orange and grey correspond to the original data and randomly permuted modulators respectively. Correlation distributions are significantly different from permuted data in both panels (KS-test, p-value ¡ 0.001). Correlation distributions for stronger and weaker modulators are not significantly different (KS-test, p-value=0.24). **e)** Scatter plot of modulator correlations with running speed. Each gray dot corresponds to a single neuron. Thin black lines within the plot denote 5th (left and bottom line) and 95th (right and upper line) percentiles of modulator-speed correlations computed with randomly permuted modulators. Darker points with Roman numerals correspond to identically labeled tuning curves in panel f. **f)** Example direction-speed tuning curves (grayscale heatmaps) and corresponding fitted modulators of stronger and weaker subunits (orange and gray lines respectively). Correlations of the stronger and weaker modulator with the running speed are denoted as ρ in orange and gray respectively. Roman numerals correspond to points marked in panel e. **g)** The same as e but for randomly permuted modulators (left) and simulated joint modulators (right). **h)** Proportions of cells in different regions of the correlation plane for real data (top row), randomly permuted modulators (second row from the top) and simulated joint modulators (third row from the top). Color legend of areas on the correlation plane (as in e-f) is depicted in the bottom row.

To understand how individual modulators are related to the behavioral state, we computed their correlations with the running speed of the animal (Fig.4d, orange bars). In stronger and weaker subunits (Fig.4d orange bars, left and right panel, respectively), the correlations with speed are statistically indistinguishable (KS test, p-value = 0.24). In both cases, the values span a wide range of correlations, with the mode being close to a perfect correlation

(i.e. $\rho = 1$), which corresponds to a monotonic increase of the gain with running speed. The similarity of these distributions supports the notion that both modulators reflect significant biological processes and not noise. The correlation distributions for individual modulators are also significantly different from the shuffled data obtained by random permutations (Fig.4d, gray bars), suggesting that the observed pattern is not obtained due to noise.

Distributions of modulator correlations with running speed characterize tuning changes of individual subunits, but not of entire neurons. Changes in neuronal tuning induced by the running speed are reflected in the joint distribution of the speed-modulator correlations (Fig.4e). Each neuron is characterized by a point on the plane, whose horizontal and vertical axes correspond to the weaker and stronger subunit, respectively. To provide a reference and quantify joint-modulation patterns, we mark 5th and 95th percentiles of speed correlations with randomly permuted modulators (Fig.4e left and bottom, right and top black lines, respectively). The scatter plot indicates that speed correlations of both modulators are largely independent - one does not strictly determine the other. Although the highest concentration of neurons is in the top-right corner, where both modulators are strongly correlated with speed, this is a reflection of marginal distributions, not a joint coordination of modulators. If one of the modulators is strongly correlated with speed, there is a broad range of possible correlations for the second modulator. We visualize example neurons corresponding to different speed correlation patterns in Fig.4f.

To ensure that the joint distribution of modulator speed correlations is not random and to interpret it more easily, we generated two simulated joint distributions of modulator speed correlations (Fig.4g). First, we randomly permuted each modulator to completely destroy their speed-correlation or relationship with other modulator in the same neuron (Fig.4g, left panel). Second, we simulated neurons with shared modulation. To do so, for each neuron, we replaced individual modulators of individual subunits with the average of both modulators scaled to preserve the strength of each modulator and distorted with additive noise (Fig.4g, right panel; see Methods). In that way, we simulated a scenario in which subunits of each neuron are jointly modulated by the same process, and any differences between modulator fits are due to random factors in the data. Visibly, both simulated distributions differ from the data.

To quantify these differences, we computed a fraction of neurons in different areas of the correlation plane for the real data and the simulated datasets (Fig.4h). We divided the correlation plane into four areas. The first three areas correspond to modulator pairs that are simultaneously strongly positively correlated with speed, strongly negatively correlated with speed, or neutral i.e. their correlation magnitude is not significant (Fig.4h, blue, green, and orange areas, respectively). We note that if both modulators are ”neutral” with respect to running speed, they can be quite different from each other and modify the sensory tuning of a neuron at different running speeds. The fourth area corresponds to pairs of asymmetric modulators, where one modulator is strongly related to running speed and the other not (Fig.4h, bottom row, purple area). The distribution of neurons in different regions of the correlation plane is clearly different between the real data and both simulated datasets (Fig.4h, top three rows). In particular, the number of asymmetric modulators where each subunit is differently affected by the running speed is twice as high as in the permuted dataset and almost four times as high as in the simulated-joint modulator dataset. These results strongly suggest that behavior can differentially modulate individual subunits, and in this way adjusts sensory representation in different behavioral states.

Clustering reveals diversity of speed-modulation patterns in sensory neurons

Our next aim was to characterize the gain modulation patterns in more detail. As before, we focus primarily on two-subunit orientation-selective neurons. To analyze joint patterns of speed modulation, we concatenated the modulators of both subunits. We then performed a clustering analysis of such concatenated modulators. Because the speed correlation analysis suggests a continuum of tuning patterns (see Fig.4c-e), our aim was not to identify a specific number of discrete tuning ”types”, but rather to summarize their diversity and reveal general trends in the population. We therefore chose to analyze 12 clusters; however, analyses with different cluster numbers led us to similar conclusions, supporting the existence of a continuum of modulator properties.

Cluster centroids reveal a diversity of speed modulation patterns in the analyzed population (Fig.5a). In agreement with the previous analysis, the average modulator pairs in each cluster consist of one stronger and one weaker average subunit (Fig.5a, darker and lighter shades of the same color, respectively). Strong differences between centroids suggest that neurons in each cluster combine sensory features with different relative weights at different speeds of movement. In some clusters, the difference between the average modulators is very pronounced (e.g., clusters 1, 8 and 9), while in others modulator averages follow a similar trend (e.g. clusters 6, 11 and 12).

Modulator averages are representative of the general trend within each cluster; however, they may not characterize well the modulation patterns in individual neurons. In particular, due to the asymmetry of modulator strengths, the clustering outcome depends more on the stronger subunits. To explore the fine structure of individual clusters, we visualized all modulator pairs, sorted by the center of mass of the weaker modulator (Fig.5b). Within each group, the strongest modulators are relatively similar, as indicated in the weak spread of their centers of mass (Fig.5b, left

part of each panel, colored dots). However, weaker modulators are typically diverse even within the same cluster (Fig.5b, right part of each panel, colored dots). Their centers of mass are often broadly distributed spanning the entire range of running speed values (e.g. clusters 2, 3, 4, 8 and 10).

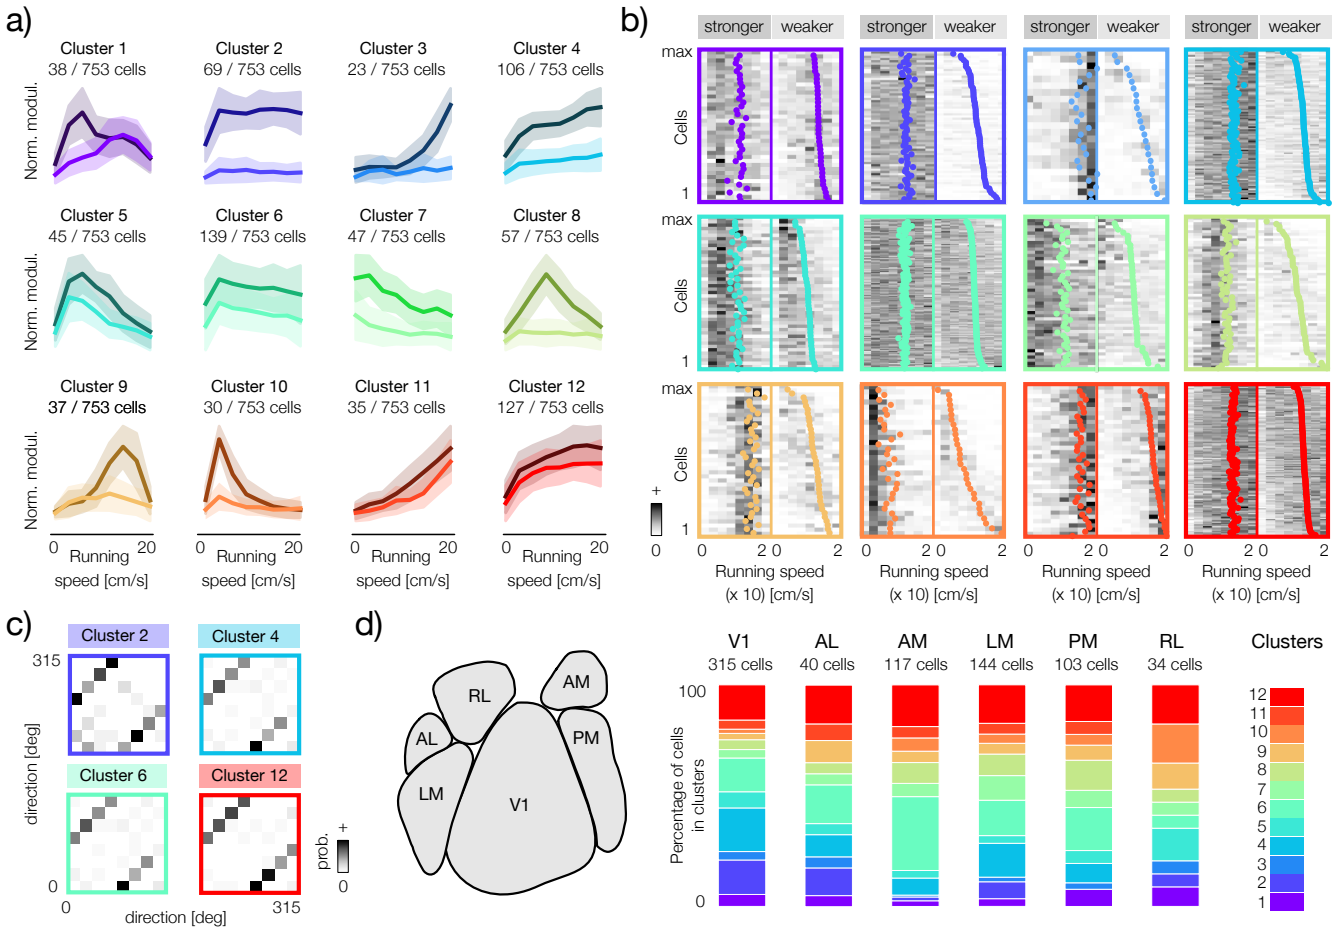

Figure 5: Clustering two-subunit neurons shows diverse effects of locomotion on sensory tuning a) Twelve clusters of speed modulators in neurons with two subunits. Bold lines correspond to cluster means (centroids) and shaded areas to cluster standard deviations. Weaker modulators are plotted with a lighter color variant. **b)** Individual modulators in each of the twelve clusters. Each panel corresponds to an identically positioned panel in a, where colors of centroids in a match the color of frames in b. Each panel consists of two parts corresponding to the stronger and weaker modulator of each cell, separated by a thin, vertical line. Colorful dots correspond to the centers of mass of each modulator. Within each cluster neurons are sorted by the center of mass of the weaker modulator. **c)** Joint distributions of direction selectivity of fitted subunits in four largest clusters. **d)** Left panel: a schematic of the anatomical division of the mouse visual cortex according to the Allen Brain Observatory (replotted from [13]). Right panel: distributions of cluster memberships in the six areas of the visual cortex. Each vertical bar corresponds to one area, the total bar height corresponds to 100 % of neurons in that area (the absolute number of neurons is denoted above each bar). Colors (right panel, color bar) correspond to clusters visualized in panels a and b.

The joint distributions of the stimulus directions preferred by the subunits were similar across the clusters, with individual subunits separated by 180 degrees with the highest probability (Fig. 5 c: histograms corresponding to four largest clusters). We did not observe systematic preferences for stimulus directions within clusters, suggesting that different patterns of speed modulation act on all neurons, regardless of their sensory tuning.

To explore the possibility that certain types of speed modulation are specific to different areas of the brain, we analyzed the probabilities of cluster membership for neurons in the six areas of the visual cortex (Fig.5d). Most of the clusters were represented in all visual areas. Upon closer inspection, certain trends become apparent. In one example, clusters 2 and 4, where modulation saturated rapidly after the onset of locomotion, were more prevalent in the primary visual cortex (V1) than in higher areas. In another example, cluster 6, characterized by a weak decrease in gain at higher speeds, was most strongly represented in the area AM. In general, these results are consistent with the notion that different modulation patterns are broadly distributed throughout the visual cortex. We also note that because of the large differences in the number of neurons analyzed within each area, it is difficult to reach systematic conclusions.

To further characterize the impact of speed on sensory representations, we analyzed the patterns of speed

modulation in single-subunit, direction-selective neurons (Fig.6. We grouped neurons into nine clusters that, as previously, we consider to be representative of broader trends, rather than corresponding discrete and isolated types. Clustering revealed various forms of modulation (Fig. 6a). The gain in the largest cluster of neurons (cluster 1) saturated rapidly after the onset of movement. In other clusters, gain increases monotonically with speed (e.g. clusters 5 and 8) or shows tuning for specific running speeds (clusters 2, 3, 4, 6, 7, and 9). The analysis of sensory tuning within the clusters did not indicate any systematic relationship between the speed modulation pattern and the preferences of the direction of stimulus motion (Fig.6b). To determine whether the prevalence of these groups of neurons is a characteristic of specific anatomical regions, we examined their distribution in the six visual cortical areas (Fig. 6c). The nine modulator clusters were present in all areas, and their relative proportions varied. Although the inter-area differences are subtle, certain trends are clearly visible. Interestingly, the primary visual cortex had the largest fraction of cells whose gain saturates even at minimal running speeds (cluster 1), which resembles a similar trend revealed by the clustering distribution of two-subunit cells (Fig.5d).

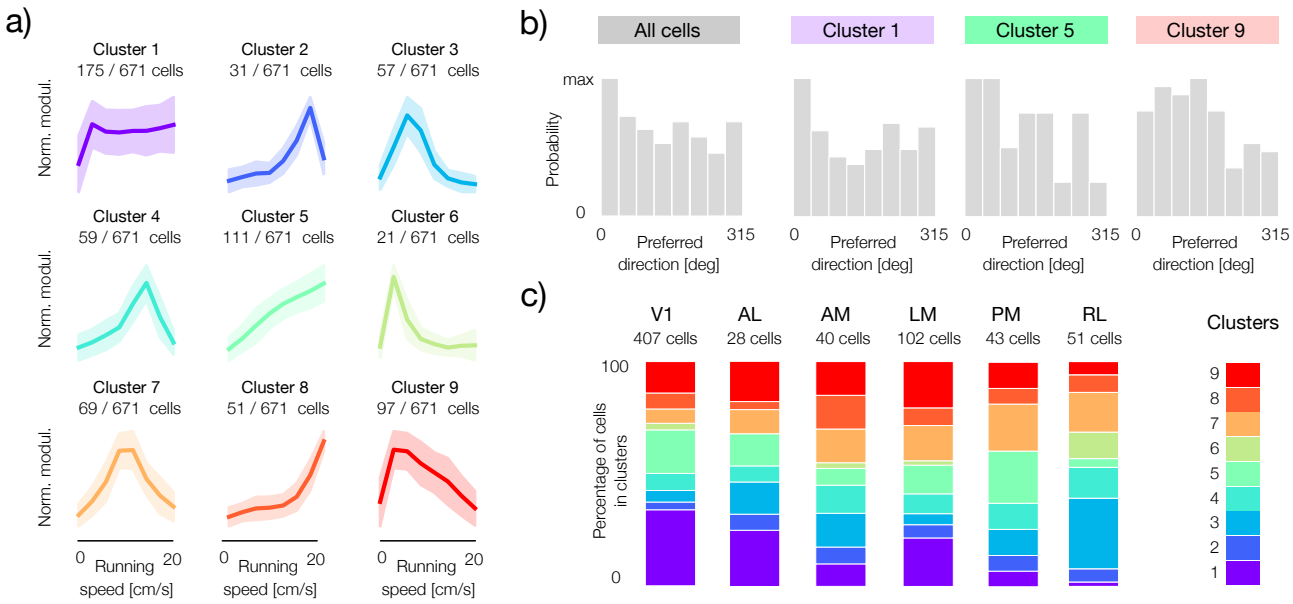

Figure 6: Clustering single-subunit neurons shows variability of running speed tuning. a) Nine clusters of speed modulators in neurons fitted with one subunit. Bold lines correspond to cluster means and shaded areas to cluster standard deviations. **b)** Distributions of preferred directions for single-subunit neurons. Distribution across all cells (left panel) and three example clusters (colored), corresponding to panel a). The clusters generally show no clear tuning bias for specific stimulus directions. **c)** Proportional distributions of modulator clusters within six areas of the visual cortex. The colored segments in each vertical bar show the percentage of neurons belonging to each cluster (color bar corresponding to panel a), for each visual area. The total number of neurons in each area is given above the bars.

Speed-direction tuning curves support accurate decoding of stimuli and behavioral signals

In the final part of the analysis, we directly tested whether the tuning curves of neurons in the visual cortex support a multiplexed representation of sensory and behavioral signals, as suggested by population models discussed in the first section. To do so, we performed a decoding analysis in which we simulated neural activity using tuning curves estimated from the experimental data (see Methods). A decoding with simulated neural activity preserves real patterns of neural tuning while simultaneously alleviating issues related to different numbers of trials and conditions in different datasets.

As in the first section, we consider three neuron architectures (Fig.7a-c). Tuning curves of direction and orientation selective neurons with independently modulated subunits were computed directly from the data (Fig. 7a and c, respectively). For each orientation selective neuron, we simulated a corresponding neuron with identical sensory and speed tuning, but with a velocity modulation shared by both subunits (Fig. 7b). We considered only the two-subunit neurons where the subunits were separated by 180 degrees instantiating the orientation-selectivity.

Using such real and data-matched tuning curves, we performed decoding from random subpopulations, as described in the first section of the results. We focused on three areas with the highest number of neurons retained. In agreement with theoretical analyses (Fig.1), populations of real orientation tuning curves enable the most accurate

simultaneous decoding of stimulus directions (Fig.7d-f, left panels) and behavioral states (Fig.7d-f, right panels). In all considered areas, the precision of stimulus decoding in orientation- and direction-selective cells is practically indistinguishable (Fig.7d, left panel). At the same time, the decoding performance of the running speed from populations of orientation-selective neurons is improved compared to that of direction-selective neurons. These results indicate that behavioral modulation improves the accuracy of sensory tuning while multiplexing information about the movement of the animal.

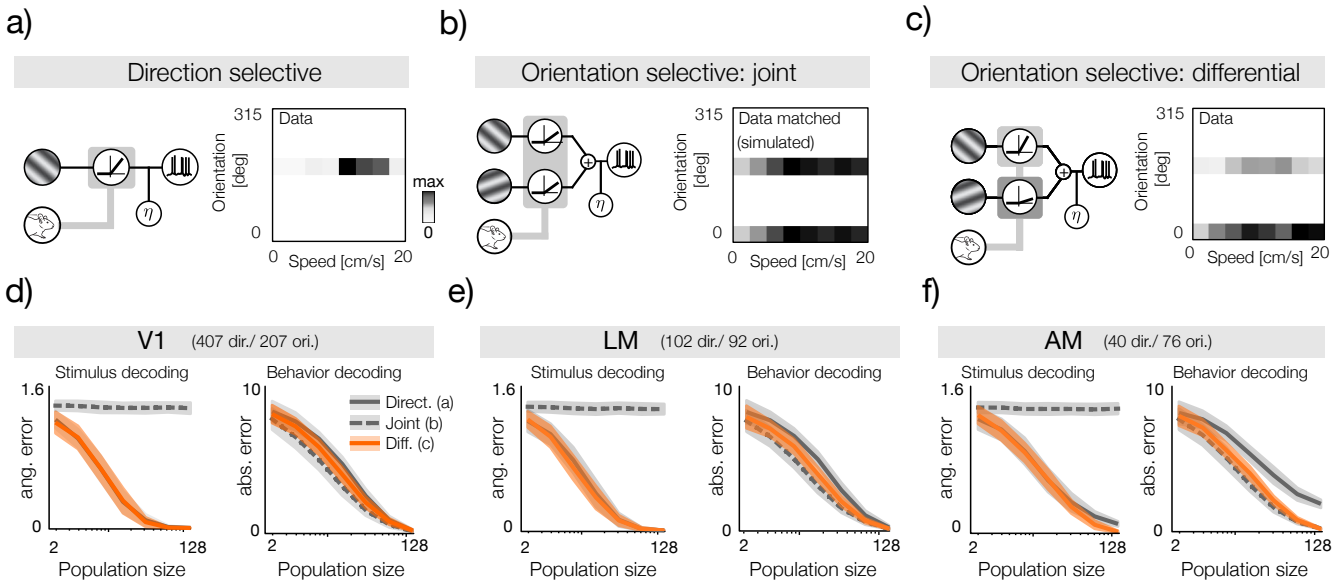

Figure 7: Simulated decoding is consistent with multiplexing of behavioral signals in sensory representations. a) A schematic of a direction-selective speed-modulated neuron (left) with an example, corresponding tuning curve fit. **b)** A schematic of an orientation-selective neuron with shared speed modulation of subunits (left) and an example corresponding tuning curve. To ensure identical modulation of both subunits, the tuning curve has been matched to c to ensure identical speed tuning. **c)** A schematic of an orientation-selective neuron with shared speed modulation of subunits (left) and an example corresponding tuning curve fit. **d)** Decoding of the stimulus direction (left panel) and running speed (behavior, right panel) using activity simulated from tuning curves of the primary visual cortex (V1; see Methods for details). In each panel the decoding error is plotted as a function of the population size. Lines correspond to average errors, shaded areas to standard deviations. Populations of direction selective neurons, orientation selective neurons with shared modulation and orientation selective neurons with differential modulation are denoted by gray solid, gray dashed and orange lines respectively. Note that for stimulus decoding (left panel) the orange and solid gray lines and areas overlap almost entirely. **e)** Same as d, but for the area LM. **f)** Same as d, but for the area AM.

DISCUSSION

In this work, we provide evidence that movement can modify the sensory tuning of visual neurons in subtle ways. This modification takes a specific form: the gain of sensory inputs pooled by subunits of each neuron is modulated differently by running speed. We propose two putative functions for this nuanced modulation. First, it resolves the ambiguity of orientation-selective neurons, enabling a significantly more accurate sensory code. Second, it allows the decoding of the modulating behavioral variable with high accuracy. In such a way, the visual cortex maintains information about the sensory environment and the behavior of the organism, supporting computations which require this information. Importantly, these benefits are accessible even with randomly pooled populations. This finding is in agreement with studies that demonstrate the utility of heterogeneity, variability, and sometimes even randomness of tuning in sensory representations and neural circuits [14, 15, 18, 19].

In our theoretical analysis, we ask which neural architectures support the decoding of the stimulus and behavioral modulator simultaneously. These decoding objectives are simultaneously very simple and quite general allowing us to remain agnostic about the specific role of computations that simultaneously require behavioral signals and sensory stimuli. Several computations of this kind have already been discussed in the field. A recent study has suggested an elegant example by demonstrating that joint tuning to running speed and spatial frequency may support the depth estimation via motion parallax [20]. Other well-known examples include the signaling of sensorimotor mismatches [21], or the estimation of the direction of movement from visual and proprioceptive signals [9]. We believe that many types of not-yet-described computations in the visual cortex may require simultaneous, joint processing of behavioral and sensory signals using mechanisms described here.

An open set of questions concerns the biological implementation of the gain modulation by running speed at the single-subunit level. Multiple studies demonstrated that the selectivity of neurons in the visual cortex cannot be easily captured by a single linear receptive field, and that even individual neurons combine multiple stimulus features [22–25]. These features, which we refer to as subunits, span a subspace in the stimulus domain. We consider it possible that individual subunits correspond to neural projections from the earlier stages of visual processing, e.g. the lateral geniculate nucleus (LGN). Because LGN neurons are modulated by running speed [11], a cortical neuron that combines inputs from two differentially modulated LGN cells could instantiate the computations described here.

The current results complement our recent study, in which we demonstrated that gain modulation by locomotion is required to maintain an accurate and efficient sensory representation [26]. The theory presented in [26] predicts that gain modulation optimized to maintain coding accuracy takes a specific form: the strength of sensory responses should increase rapidly at the start of locomotion and then saturate. Interestingly, qualitatively similar speed modulation patterns are identified by clustering analyses in this study (Fig.6a, cluster 1 and Fig.5a, cluster 2 and 4). Neurons whose subunits are modulated in a way most consistent with this theory are mostly abundant in the primary visual cortex. These results enable us to speculate that gain saturation is a signature of a "sensor-like" neuron whose role is to encode and transmit sensory information accurately across behavioral states. Other patterns of gain modulation may indicate different types of computation. Exploring these ideas further remains, of course, a subject of future work.

Gain modulation is a prominent mechanism that contributes to neuronal variability in behaving animals [27]. Single neurons can be subject to multiple internal states that affect the gain of their sensory responses [27–29]. Different normative theories propose various putative functions of gain modulation, e.g. task-dependent adaptation of sensory representations [30] or improving downstream decoders [31]. Our work suggests an additional form and role for the gain modulation by running speed. The existence of multiple modulatory variables is also a caveat. In principle, it is possible that some of the speed-induced modulation reported here is caused by other factors, unrelated to speed (e.g. attentional state or head movements). We foresee the possibility of resolving these issues through the application of dedicated statistical models that can separate different sources of variation (e.g. [27, 32, 33]). Although we cannot separate those factors with complete certainty, our results indicate that the sensory tuning of individual neurons is dynamically modulated in behaving animals regardless of the origin of the modulatory signal.

ACKNOWLEDGEMENTS

We thank Laura Busse, Lukas Meyerolbersleben and Ann Hermundstad for helpful discussions.

METHODS

Simulated populations of gain-modulated neurons

We simulated neurons tuned to an angular variable $\phi \in [-\pi, \pi]$ (e.g. stimulus direction) using von-Mises-like tuning curves of the form: $r(\phi) = \exp(\kappa * \cos(\phi_i - \phi))$, where ϕ_i is the preferred tuning direction and $\kappa = 4$ controls width of the tuning function. Each sensory tuning was normalized so that the maximal response was equal to 1. We simulated the speed tuning by combining 8 radial basis functions of the following form $\exp(\frac{(s-s_i)^2}{\sigma^2})$ where s_i denotes the position of the peak and $\sigma^2 = 10$ is the width of the basis function. The peak positions were linearly spaced on the $[0, 20]$ interval. Each speed modulator $m(s)$ was expressed as a linear combination of basis functions: $m(s) = \sum_{i=1}^{8} w_i \exp(\frac{(s-s_i)^2}{\sigma^2})$ where the weights w_i were randomly drawn from a Gaussian distribution with standard deviation equal to 0.5 To generate the response r to a particular stimulus direction ϕ and speed s the speed modulator was applied multiplicatively after exponentiation. The functional form of the response took the form: $r(s, \phi) = r(\phi) \exp(m(s)) + \eta$, where η is the additive Gaussian noise with standard deviation $\sigma = 0.2$. For two-subunit model neurons with shared modulation, we generated the same modulator for both subunits, and the neural response was summed over both subunits. For two-subunit neurons with independent modulation, separate modulators were randomly generated for each subunit. To simulate asymmetry of the modulator strength, a random number uniformly distributed on the $[0, 0.2]$ interval was subtracted from the weights w_i of the weaker subunit.

We simulated 500 random populations for the following population sizes: $2, 4, 8, 16, 32, 64, 128$ neurons. In each of the simulations we generated 500 random stimuli ϕ uniformly distributed on a circle and speed values s uniformly distributed on the $[0, 20]$ interval.

We performed the maximum-a-posteriori (MAP) decoding. First, we computed the joint posterior over the stimulus ϕ_t and speed s_t given the noisy response of the population $\vec{r}_t$, i.e. $p(\phi_t, s_t|\vec{r}_t)$. We then marginalized the joint posterior to generate two marginal posteriors $p(\phi_t|\vec{r}_t)$ and $p(s_t|\vec{r}_t)$. The maxima of these posteriors were taken to be the decoded estimates $\hat{\phi}_t$ and $\hat{s}_t$, respectively. All computations were performed numerically, using discretized values of parameters. Because the prior distributions $p(\phi_t), p(s_t)$ were both uniform, this procedure was equivalent to maximum likelihood (ML) decoding. To quantify the speed decoding error we used the squared distance averaged over samples, i.e. $\langle (s_t - \hat{s}_t)^2 \rangle_t$. To quantify the stimulus direction decoding error, we used the average angular distance $\langle \min(|\phi_t - \hat{\phi}_t|, 2\pi - |\phi_t - \hat{\phi}_t|) \rangle_t$.

Simulated decoding analyses

We performed the MAP decoding identically as with the simulated tuning curves (see above). We used speed-orientation tuning curves from three areas with the highest number of retained neurons. We used only the two-subunit cells where subunits were separated by 180 degrees (numbers of cells per area are stated in Fig.7d-f). For each area, we generated populations of the following sizes: $2, 4, 8, 16, 32, 64, 128$. For each size value, we generated 200 random populations, each sampled from tuning curves of a corresponding area with replacement. During simulation, each of the populations was used to decode the random stimulus and speed values 1000. Stimulus values corresponded to the eight directions used in experiments, while speed values corresponded to the speed-binning used to derive tuning curves. Each tuning curve was normalized so that its maximum was equal to 1. Neural responses were generated by adding a random Gaussian number with zero mean and standard deviation $\sigma = 0.2$ to the value of the tuning curve corresponding to the speed and orientation pair. To generate neurons that are data-matched but share modulation across the subunits, we replaced modulators of each of the subunits with their average.

Data preprocessing

We used data made available by the Allen Brain Observatory [34]. The datasets comprised of two-photon calcium imaging of neural populations recorded in behaving mice. The animals were exposed to drifting grating stimuli at 0.04 cpd spatial frequency, at 8 different orientations linearly spaced 45 degrees apart between 0 and 360 degrees and 5 temporal frequencies (1, 2, 4, 8, 15 Hz).

To reduce the variability of the data and focus the analysis on variation above the baseline, we thresholded the ΔF/F signal by setting all negative values to zero. Following thresholding, we normalized the activity of each cell by dividing it by its standard deviation.

To ensure the robustness of our conclusions, we applied several data filtering steps - all described in detail below. We included only datasets that met the following requirements: they had sufficient data in each bin and were representative across at least three different temporal frequencies, the joint tuning curves were well-fitted by our subunit model, the cells exhibited sensory tuning, and their modulation along the running speed axis was significant per the bootstrap test. Following all selection criteria, we obtained a dataset of 1693 neurons from 100 recording sessions.

We aimed to minimize the distortions in estimated tuning due to differences in the amount of data per bin and disparities between the range of temporal frequencies presented for each bin. To achieve this, we excluded any datasets with less than 1 second of data in any of the 64 bins of direction-speed tuning curves. Additionally, we removed datasets where more than 16 of the 64 bins were derived from averages that included fewer than 3 different temporal frequencies (out of 5 total).

Computation of tuning curves

We divided the running speed into eight linearly spaced bins that ranged from 0 and 20 cm/s. This allowed for consistent comparisons across different datasets with varied velocity distributions. Subsequently, we binned the transformed ΔF/F values of each cell according to the eight orientations, five temporal frequencies and eight running speed bins. The activity within each bin was then averaged to obtain the average response $r_{o,v,f}$ where o, v, f denote the stimulus orientation, running speed and the stimulus temporal frequency, respectively. We refer to these average activity tensors as 3D tuning curves. Averaging along the frequency dimension we obtained $r_{o,v}$, which we refer to as direction-speed tuning curve.

Tuning specificity

To assess the sensory tuning specificity of each cell, we computed their stimulus tuning curves $R^{t,s}$ by averaging each direction-speed tuning curve R^t along the speed dimension. We introduce the normalized entropy of each tuning curve $H(R^{t,s}) = \frac{1}{3}\sum_{k=1}^{8} \bar{R}_k^{t,s} \log_2(\bar{R}_k^{t,s})$, where $\bar{R}^{t,s}$ is the sensory tuning curve normalized, such that it's entries sum up to unity. The normalized entropy H is equal to 0 if the tuning curve is uniform and to 1 if the neuron responds to only one stimulus. We quantified the tuning specificity of each neuron as $S = 1 - H(R^{t,s})$. Neurons were classified as stimulus selective and included in further analysis, if their sharpness values exceeded the 95th percentile of the null distribution derived from the set of randomly shuffled tuning curves (see the description of goodness of fit assessment).

Bootstrap test of speed modulation

We performed a bootstrap test to test for the significance of speed modulation. For each neuron, we generated 1,000 3D tuning curve tensors by resampling the data that contributed to each bin (defined by stimulus orientation, stimulus temporal frequency, and running speed). For each of the bins and in each of the 1000 random tuning curves, we resampled identical number of data points, corresponding to the responses of the neuron to the same stimulus direction and temporal frequency, but to random speed values. This procedure ensured that the resulting tuning curves were velocity-invariant while preserving the direction and frequency tuning properties of the neuron, as well as the number of data samples in each bin. Each of the randomly resampled 3D tuning curves was averaged across the temporal frequency dimension, as with the original data.

We used the following test statistic to assess the modulation strength: $d_{mod}(R^t) = \sum_{l=1}^{8} d_{Euc}(R^{t,s}, R_k^t)$, where $R^{t,s}$ is the sensory tuning curve (see above), R_k^t are sensory tuning curves at each speed bin k and d_{Euc} is the Euclidean distance. Modulation of sensory tuning by running speed was considered significant if the value of the test statistic evaluated for the real tuning curve exceeded the 95th percentile of the null distribution of the test statistic generated by bootstrapping for the considered cell.

The sparse decomposition model

We designed a statistical model to parameterize the patterns of joint tuning to stimulus direction and running speed by decomposing the tuning curves of individual neurons. The model assumes that each neuron pools the responses of direction-selective subunits, where each subunit is multiplicatively modulated by running speed.

A direction-speed tuning curve of the $t-$th neuron R^t is a two-dimensional, K matrix, where $L = 8$ is the number of speed bins and $K = 8$ is the number of stimulus direction bins. Each element $R_{k,l}^t$ represents the mean response of the neuron t to a grating of a specific direction k and the speed interval l.

Basis functions ϕ^i represent sensory tuning to one of the K preferred directions and are effectively zero-matrices with entries in one row set to 1. The total number of basis functions is therefore equal to the number of possible stimulus directions i.e. 8. The model parameters $\Theta_t = \{\vec{s}_t, M_t\}$ consist of a matrix of modulators M ($\mathbb{R}^{L\times 8}$), which describe speed-bin-specific modulation, and weights $\vec{s}$ ($\mathbb{R}^I$), which represent the weight of each orientation-specific subunit. The model reconstructs speed-direction tuning curves as the sum of weighted and modulated basis functions:

$$\hat{R}_{k,l}^t = \sum_{i=1}^{8} \phi_{k,l}^i \cdot {m_{i,l}^t}^2 \cdot s_i^{t2} \tag{1}$$

, where square is taken to ensure positivity of modulators and basis function coefficients.

We estimated the model parameters for each neuron Θ_t by minimizing the cost function that combined the mean squared error (MSE) of the reconstruction between the reconstructed and actual data and regularization terms that penalized the absolute value of the subunit coefficients and modulators:

$$\text{Cost}(\Theta_t) = \frac{1}{TKL}\sum_{t=1}^{T}\sum_{k=1}^{K}\sum_{l=1}^{L}\left(X_{t,k,l} - \hat{X}_{t,k,l}\right)^2 + \lambda_s \frac{1}{TK}\sum_{t=1}^{T}\sum_{k=1}^{K} s_{t,k}^2 + \lambda_m \frac{1}{TL}\sum_{t=1}^{T}\sum_{l=1}^{L} M_{t,l}^2 \tag{2}$$

where λ_s and λ_m are the sparsity parameters for subunit coefficients and modulators, respectively. For simplicity we set them to be identical i.e. $\lambda = \lambda_s = \lambda_m$. We optimized the model for each neuron using 10 sparsity values, evenly spaced within the range [0.01, 0.1].

For each neuron, the model was fitted using 300 gradient steps, starting from five independent random initializations of parameters. We selected the fits that generate the highest signal-to-noise ratio (SNR) value.

Following the model fitting we multiplied each modulator m by the corresponding subunit coefficient s (we refer to those as scaled modulators). We found that separating these quantities during fitting resulted in better matches to the data structure; however, following the fit the separation was not useful. For each neuron, we kept only the subunits whose average value of the corresponding scaled modulators was equal to or higher than 5% of the maximal one. We refer to those subunits as "significant subunits".

Goodnes of fit assesment

For each analyzed dataset, we generated a null set of tuning curves. We did so by randomly sampling the direction-speed tuning curves of 10,000 neurons with replacement. The positions of the 64 bins of each joint tuning curve were then randomly permuted. We fitted the model to each random tuning curve using the procedure described above. For each considered value of the sparsity constraint λ, we retained the distribution of reconstruction SNRs obtained in that way. Neurons passed the goodness-of-fit test if their reconstruction quality exceeded the 95th percentile of the null distribution. For each neuron, we selected a single reconstruction corresponding to the lowest λ value at which its fit first exceeded the significance threshold.

Analyses of two-subunit neurons

For each neuron with two significant subunits, we computed the $\log_2$ of the ratio of the average stronger modulator to the average weaker modulator (Fig. 4b).

To construct null distributions of modulator correlations and modulator-speed correlations, we randomly permuted individual modulators in each neuron across the speed dimension.

To simulate joint modulation of both subunits, we computed the average of both modulators for each neuron. In the next step, we replaced individual modulators in each neuron with this average normalized such that the norm of each subunit was preserved. We added Gaussian noise with standard deviation $\sigma = 0.01$.

Clustering analyses

To identify patterns of velocity modulation, we performed K-means clustering on the shapes of the subunit modulators from all neurons that passed the filtering criteria. For each neuron with two significant subunits, we concatenated the scaled modulators (with the stronger subunit leading), resulting in a 16 single feature vector. The feature vectors were divided by their L_2 norm. The number of clusters was chosen heuristically and was equal to 9 of one-subunit neurons and 12 for two-subunit neurons.

References

1. Parker, P. R., Brown, M. A., Smear, M. C. & Niell, C. M. Movement-related signals in sensory areas: roles in natural behavior. *Trends in neurosciences* **43,** 581–595 (2020).
2. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience* **22,** 1677–1686 (2019).
3. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364,** eaav7893 (2019).
4. Saleem, A. B., Diamanti, E. M., Fournier, J., Harris, K. D. & Carandini, M. Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature* **562,** 124–127 (2018).
5. Parker, P. R., Abe, E. T., Leonard, E. S., Martins, D. M. & Niell, C. M. Joint coding of visual input and eye/head position in V1 of freely moving mice. *Neuron* **110,** 3897–3906 (2022).
6. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86,** 740–754 (2015).
7. Kanamori, T. & Mrsic-Flogel, T. D. Independent response modulation of visual cortical neurons by attentional and behavioral states. *Neuron* **110,** 3907–3918 (2022).
8. Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65,** 472–479 (2010).
9. Saleem, A. B., Ayaz, A., Jeffery, K. J., Harris, K. D. & Carandini, M. Integration of visual motion and locomotion in mouse visual cortex. *Nature neuroscience* **16,** 1864–1869 (2013).
10. Christensen, A. J. & Pillow, J. W. Reduced neural activity but improved coding in rodent higher-order visual cortex during locomotion. *Nature communications* **13,** 1676 (2022).

11. Erisken, S. *et al.* Effects of locomotion extend throughout the mouse early visual system. *Current Biology* **24,** 2899–2907 (2014).
12. Schröder, S. *et al.* Arousal modulates retinal output. *Neuron* **107,** 487–495 (2020).
13. De Vries, S. E., Siegle, J. H. & Koch, C. Sharing neurophysiology data from the Allen Brain Observatory. *Elife* **12,** e85550 (2023).
14. Maoz, O., Tkačik, G., Esteki, M. S., Kiani, R. & Schneidman, E. Learning probabilistic neural representations with randomly connected circuits. *Proceedings of the National Academy of Sciences* **117,** 25066–25073 (2020).
15. Barak, O., Rigotti, M. & Fusi, S. The sparseness of mixed selectivity neurons controls the generalization–discrimination trade-off. *Journal of Neuroscience* **33,** 3844–3856 (2013).
16. Fahey, P. G. *et al.* A global map of orientation tuning in mouse visual cortex. *BioRXiv,* 745323 (2019).
17. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience* **28,** 7520–7536 (2008).
18. Berry Ii, M. J., Lebois, F., Ziskind, A. & da Silveira, R. A. Functional diversity in the retina improves the population code. *Neural computation* **31,** 270–311 (2019).
19. Goris, R. L., Simoncelli, E. P. & Movshon, J. A. Origin and function of tuning diversity in macaque visual cortex. *Neuron* **88,** 819–831 (2015).
20. He, Y., Colas Nieto, A., Blot, A. & Znamenskiy, P. A depth map of visual space in the primary visual cortex. *bioRxiv,* 2024–09 (2024).
21. Keller, G. B., Bonhoeffer, T. & Hübener, M. Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron* **74,** 809–815 (2012).
22. Vintch, B., Movshon, J. A. & Simoncelli, E. P. A convolutional subunit model for neuronal responses in macaque V1. *Journal of Neuroscience* **35,** 14829–14841 (2015).
23. Sharpee, T., Rust, N. C. & Bialek, W. Analyzing neural responses to natural signals: maximally informative dimensions. *Neural computation* **16,** 223–250 (2004).
24. Rust, N. C., Schwartz, O., Movshon, J. A. & Simoncelli, E. P. Spatiotemporal elements of macaque v1 receptive fields. *Neuron* **46,** 945–956 (2005).
25. Schwartz, O., Pillow, J. W., Rust, N. C. & Simoncelli, E. P. Spike-triggered neural characterization. *Journal of vision* **6,** 13–13 (2006).
26. Gant, J. M. & Mlynarski, W. F. Anticipatory modulation as a unifying principle of sensory coding during locomotion. *bioRxiv,* 2025–06 (2025).
27. Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nature neuroscience* **17,** 858–865 (2014).
28. Kanamori, T. & Mrsic-Flogel, T. D. Independent response modulation of visual cortical neurons by attentional and behavioral states. *Neuron* **110,** 3907–3918 (2022).
29. Spacek, M. A. *et al.* Robust effects of corticothalamic feedback and behavioral state on movie responses in mouse dLGN. *Elife* **11,** e70469 (2022).
30. Młynarski, W. & Tkačik, G. Efficient coding theory of dynamic attentional modulation. *PLoS Biology* **20,** e3001889 (2022).
31. Haimerl, C., Ruff, D. A., Cohen, M. R., Savin, C. & Simoncelli, E. P. Targeted V1 comodulation supports task-adaptive sensory decisions. *Nature communications* **14,** 7879 (2023).
32. Liu, D. & Lengyel, M. A universal probabilistic spike count model reveals ongoing modulation of neural variability. *Advances in Neural Information Processing Systems* **34,** 13392–13405 (2021).
33. Zhu, R. J. & Wei, X.-X. Unsupervised approach to decomposing neural tuning variability. *Nature Communications* **14,** 2298 (2023).
34. De Vries, S. E. J. *et al.* A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience* **23,** 138–151. ISSN: 1546-1726. `http://dx.doi.org/10.1038/s41593-019-0550-9` (Dec. 2019).