Intermediate intrinsic diversity enhances neural population coding

Shreejoy J. Tripathy[a,b], Krishnan Padmanabhan[b,c,1], Richard C. Gerkin[b,c], and Nathaniel N. Urban[b,c,d,2]

[a]Program in Neural Computation, [b]Center for the Neural Basis of Cognition, and [c]Department of Biological Sciences, Carnegie Mellon University, Pittsburgh, PA 15213; and [d]Department of Neuroscience, University of Pittsburgh, Pittsburgh, PA 15213

Edited by Eve Marder, Brandeis University, Waltham, MA, and approved April 1, 2013 (received for review December 7, 2012)

Cell-to-cell variability in molecular, genetic, and physiological features is increasingly recognized as a critical feature of complex biological systems, including the brain. Although such variability has potential advantages in robustness and reliability, how and why biological circuits assemble heterogeneous cells into functional groups is poorly understood. Here, we develop analytic approaches toward answering how neuron-level variation in intrinsic biophysical properties of olfactory bulb mitral cells influences population coding of fluctuating stimuli. We capture the intrinsic diversity of recorded populations of neurons through a statistical approach based on generalized linear models. These models are flexible enough to predict the diverse responses of individual neurons yet provide a common reference frame for comparing one neuron to the next. We then use Bayesian stimulus decoding to ask how effectively different populations of mitral cells, varying in their diversity, encode a common stimulus. We show that a key advantage provided by physiological levels of intrinsic diversity is more efficient and more robust encoding of stimuli by the population as a whole. However, we find that the populations that best encode stimulus features are not simply the most heterogeneous, but those that balance diversity with the benefits of neural similarity.

generalized linear models | intrinsic biophysics | neural variability | stimulus coding | ion channels

Biological systems including brains must function efficiently under many constraints, including constraints on the numbers of individual neurons dedicated to a given task. Brain function therefore depends on an appropriate division of labor, with specific neurons dedicated to different functions. For example, different types of retinal ganglion cells represent visual information at different timescales (1), and distinct classes of cortical interneurons play diverse roles in coordinating network activity (2). Whereas attempts to understand how distinct classes of cells encode information have proven successful (1), the importance of within-type variability remains poorly understood (3, 4) although has recently become a topic of great interest (5–8).

Although neuron-to-neuron variability is often viewed as an epiphenomenon of biological imprecision (3, 4), having neurons of the same type that respond to different stimulus features may improve stimulus encoding. This variability may be leveraged to improve functions such as stimulus encoding if heterogeneous output of neurons of a single type is collectively used for population coding. Such populations of neurons could efficiently represent complex stimuli by collectively covering the relevant stimulus space (1, 9, 10). Network interactions could further increase the efficiency of information transmission by decorrelating neural responses and reducing the redundancy between their outputs (11–13). In contrast, eliminating redundancy (also referred to as biological degeneracy, ref. 14) may make stimulus coding less robust to noise or damage (15), thus we hypothesized that an optimal coding strategy would require balancing diversity with feature similarity or overlap.

Although theorists have previously explored this issue (12, 16, 17), analysis of the function of the diversity of real populations of neurons requires overcoming methodological hurdles associated with studying cell-to-cell variability (3, 4). Cell-level differences (that are typically averaged away) must be captured and quantified. Once these differences have been quantified, one must compare the functional output of populations differing in their variability. In the context of neural coding these issues translate to answering the questions: What properties of neurons determine their response to stimuli? How are these properties distributed? And how do these distributions of properties influence the encoding of stimuli by populations? Although previous experimental approaches have identified neuron diversity using standard receptive field analyses, these typically do not describe the full complexity of neural responses to stimuli (18–20), nor do they allow the source of the response heterogeneity to be identified as either synaptic or intrinsic. In addition, simplistic readouts of population spiking output may underestimate the richness of the underlying neural code (1, 10, 21). Our approach allows the influence of intrinsic diversity to be isolated from synaptic differences and captures the full potential of these diverse populations for stimulus encoding.

Specifically, we developed measures of neuronal population diversity based on statistical generalized linear models (18, 22) that accurately reproduce the responses of recorded individual olfactory bulb mitral cells (MCs). These cells have been shown to express significant biophysical variability from neuron to neuron (5–7). We then used the framework of model-based stimulus decoding (18, 23) to compare how populations varying in their diversity optimally encode varieties of stimuli. This approach enables us to determine whether specific advantages arise from the intrinsic diversity of these neurons, and how MC populations balance the competing benefits of diversity and feature similarity.

Results

Statistical Neuron Models Capture Mitral Cell Response Diversity. We generated models of individual MCs from data collected during in vitro whole-cell recordings in which somatic current injection of broad-band-filtered noise (5) evoked action potential trains (Fig. 1 *A* and *B*; $n = 44$ neurons). Synaptic transmission was blocked pharmacologically, so that differences in the cells' spiking responses reflected only differences in their intrinsic firing properties (e.g., due to biophysical conductances and/or morphology). Each neuron's spiking response to input current was fit by a generalized linear model (GLM). GLMs extend stimulus-based reverse correlation or linear–nonlinear–Poisson (LNP) models (20, 24) by including terms that describe how a neuron's spike probability is modulated via its previous spikes (18, 22). Here each GLM had

Author contributions: S.J.T., K.P., R.C.G., and N.N.U. designed research; S.J.T. and K.P. performed research; S.J.T. and R.C.G. analyzed data; and S.J.T. and N.N.U. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]Present Address: Salk Institute for Biological Studies, La Jolla, CA 92037.

[2]To whom correspondence should be addressed. E-mail: nurban@cmu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1221214110/-/DCSupplemental.

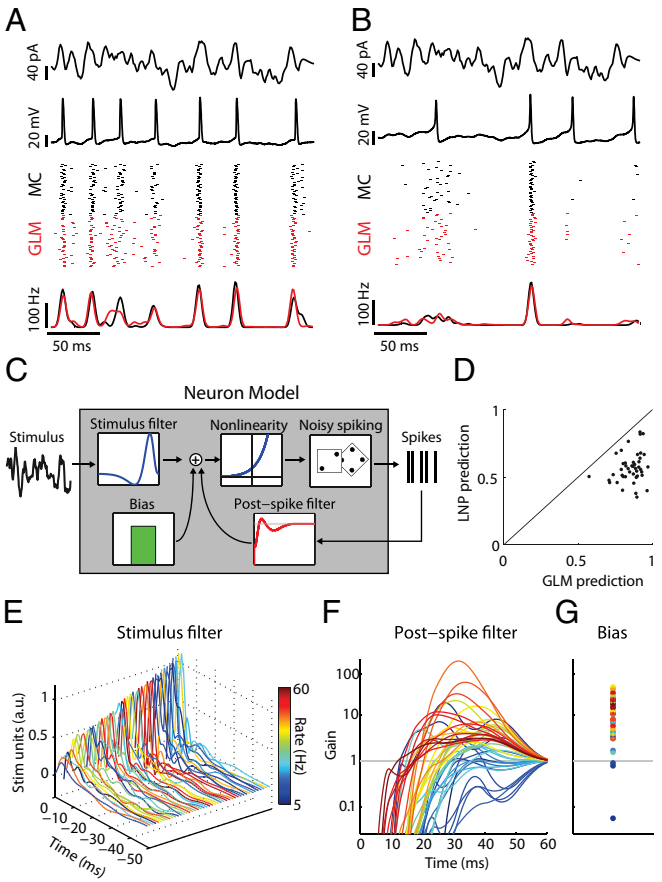

Fig. 1. Simple models capture mitral cell stimulus-evoked responses and intrinsic diversity. (*A*) MC intrinsic properties are probed using filtered broadband stimuli (first row) injected somatically to evoke changes in membrane voltage (second row). Spike rasters (third row; black) and peristimulus time histograms (PSTH) (fourth row; black) for repeated stimulus presentations (n = 40 trials) show that this MC spikes to the stimulus with temporal jitter and displays a coarse stimulus preference. Model neuron rasters (*Third Row*, red) and PSTH (fourth row, red) show that the model accurately predicts MC activity on novel stimuli. (*B*) Same as *A* but for a different neuron. (*C*) Structure of the GLM neuron model. Model parameters describe a temporal stimulus filter, a postspike filter, and a constant bias term. An exponential nonlinearity defines an instantaneous spike rate and is used to draw noisy spikes. (*D*) GLM models accurately predict 86 ± 11% (mean ± SEM) of stimulus-evoked activity across all MCs, computed as the correlation coefficient between MC and model PSTH. For all neurons, the GLM fits were better than LNP models. (*E–G*) Model parameters for all MCs. Each line corresponds to parameters for a unique neuron and is colored by mean firing rate. (*E*) Temporal stimulus filters model differential stimulus specificity of neurons. (*F*) Exponentiated postspike filters, plotted as a multiplicative gain in spike probability following a spike at t = 0 ms. Values less (greater) than 1 indicate a decreased (increased) spike probability. (*G*) Bias terms also show considerable variation. Same *y* axis as *F*.

a constant (bias) term to match baseline firing, a linear stimulus filter determining the neuron's stimulus preference, and a spike history function capturing the neuron's refractory and bursting properties (Fig. 1*C*).

This approach captures the spiking responses of neurons without explicitly modeling the many ion channels expressed by individual cells (3, 5, 6) (Fig. 1 *A* and *B*). Furthermore, GLMs modeled MC activity better than a simpler model that did not include spike history effects (LNP; Fig. 1*D* and Fig. S1), indicating that postspike refractory and bursting effects substantially contribute to action potential generation in these neurons. Because the parameters of the GLM model emergent physiological features of the recorded neurons, comparing GLM parameters across neurons illustrates the diversity among MCs (Fig. 1 *E–G*). For example, the diversity reflected in postspike (i.e., spike history) filters potentially corresponds to a recently characterized variability in the rebound depolarization current of these neurons recorded in vivo (6). Furthermore, the interaction of each MC's GLM parameters defines how it responds to stimuli and dictates the complex stimulus features that each neuron best encodes. We note that the efficacy of the GLM approach in capturing MC responses was not specific to the precise stimulus statistics delivered to the neurons here (Fig. S2).

Diversity Enables Efficient Stimulus Representation. Because the GLM approach captures the intrinsic diversity across MCs, different model MCs generate unique spike trains when presented the same dynamic stimulus (Fig. 2*B* and Fig. S3). We used this model-based approach to ask which features of these individual neurons influence the amount of information about the stimulus that each neuron captures (Fig. 2*A*). Quantifying the quality of stimulus representation using information theoretic methods (23, 24), we found that neuron information rates were strongly correlated with firing rate (r = 0.87; Fig. 2*C*), in line with previous findings (25). However, we note that we found examples of neurons that had identical firing rates and yet differed almost twofold in their information rates, suggesting the importance of additional factors other than firing rate governing the amount of transmitted information. For example, neurons with spike times that were reliable across stimulus repeats and spikes that were strongly stimulus driven (i.e., minimal contributions from bias or spike-history terms) were more informative per spike (Fig. 2*D* and Fig. S3*C*). We note that the large range and diversity of firing rates observed among the MCs here is concordant with those found in vivo (26).

We extended this information-based framework to examine how populations of recorded MCs encode a common stimulus, considering two broad possibilities. First, stimuli might be best encoded by groups of highly similar neurons, suggesting that averaging across the population of recorded neurons can compensate for unreliable spiking in any single neuron (10). Alternatively, stimuli might be best encoded by groups of heterogeneous neurons, suggesting that maximizing the representation of temporal features of the stimuli is important (12, 27). We specifically chose to study how diverse groups collectively represent an identical stimulus to mimic features of the olfactory bulb, where 25–50 sister MCs projecting to the same glomerulus (5) each receive highly correlated stimulus- and respiration-driven synaptic input (26, 28–30).

We created populations of uncoupled virtual mitral cells by randomly selecting groups of model neurons (i.e., fit from the recorded MCs). Spiking responses in these virtual populations were then simulated using the GLM models, enabling us to probe ensemble responses to many more stimuli than could be delivered during experimental recordings. The neurons in these synthetic populations varied in the diversity of their GLM parameters, allowing us to determine how neuronal diversity influences the encoding of fluctuating stimuli. To this end, we used Bayesian model–based decoding, which optimally reconstructs the input to a population (i.e., its "perceived stimulus") given its ensemble response (18, 23). This approach solves the high-dimensional problem of interpreting dynamic population responses (13, 23) without making undue simplifications or assumptions about the nature of the neural code (5, 10, 21). However, we note that we could have instead focused on alternative metrics of population output instead of stimulus representation efficacy.

We first used the analysis described above on populations consisting of pairs of simulated neurons. Homogeneous pairs, composed of two copies of the same model neuron (with identical stimulus filtering properties), encoded 73 ± 11% more

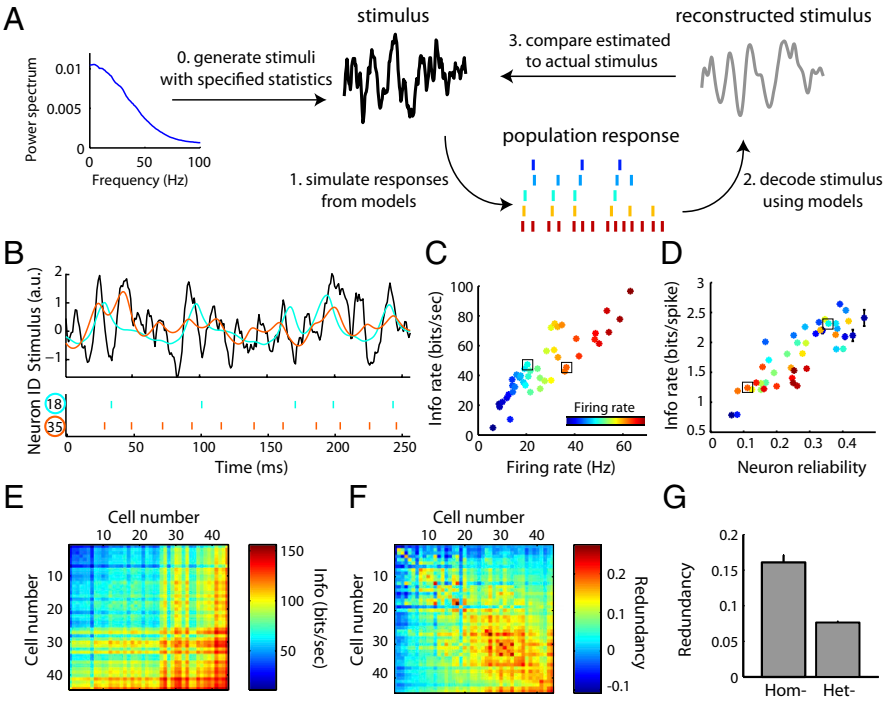

Fig. 2. Using simulated ensemble responses to study stimulus representation in diverse neural populations. (*A*) Schematic of the paradigm used to study how neural populations represent stimuli. Following the generation of noisy stimuli, population spike responses were simulated using the MC models. Bayesian decoding was used to estimate the most probable stimulus given the population response and then compared with the actual stimulus. (*B–D*) Stimulus encoding by single neurons. Stimulus statistics and coloring of neurons same as in Fig. 1. (*B*) Stimulus (*Top*; black), spike trains (*Bottom*), and reconstructions (*Top*; cyan, orange) for two example neurons. These neurons encode the same stimulus differently, as evidenced by their unique spike trains and stimulus reconstructions. (*C*) Quantifying stimulus representation using mutual information (mean ± SEM, $n = 44$ cells) shows that a neuron's information rate is strongly correlated with its firing rate ($r = 0.87$). Boxes indicate neurons shown in *B*. (*D*) As in *C*, but plotted as average information conveyed by single spikes as a function of neuron reliability. (*E–G*) Stimulus encoding by neuron pairs. (*E*) Mutual information for all neuron pairs with neurons ordered along axes by increasing firing rate. Values on (off) diagonal correspond to homogeneous (heterogeneous) pairs. (*F*) As in *E* but plotted as the normalized informational redundancy of the neuron pair. Positive (negative) values indicate population redundancy (synergy) where zero indicates independent stimulus encoding. 90% of pairs were redundant, yet overall redundancy values were small, indicating near-independent encoding. (*G*) Homogeneous pairs (Hom) are significantly more redundant than heterogeneous pairs (Het; $P = 2.5*10^{-16}$, Wilcoxon, $n = 44$ and 946 pairs, respectively).

informative about the stimulus than a single neuron copy alone (Fig. S4). In other words, because spiking is a stochastic process, decoding is improved by considering multiple spike trains from identical model neurons. This allows for "averaging out" the effect of any single neuron's noise. Next, we considered both homogeneous and heterogeneous pairs of neurons, and quantified the informational redundancy of these pairs. This method compares the information of the pair relative to the sum of each neuron's information independently (13), and gives an indication of the efficiency of information representation by the population. For example, do neurons together represent information redundantly (i.e., both neurons communicate identical or partially overlapping messages)? Or do they instead represent information synergistically (i.e., both neurons communicate more information together than both individually)? Although we found that most homogeneous and heterogeneous populations represented information redundantly (Fig. 2 *E* and *F*), homogeneous pairs were twice as redundant as heterogeneous pairs (16% versus 8% informational redundancy; Fig. 2*G*). Given that these neurons do not directly communicate, we note that the appearance of synergism among neurons pairs here is somewhat surprising and is likely due to limitations in our ability to estimate information rates among low firing rate neurons (see *Materials and Methods* for further explanation). Nonetheless, these results demonstrate that although pooling responses over multiple neurons even multiple copies of the same neuron is beneficial, the heterogeneity in intrinsic properties in actual mitral cells is beneficial for efficiently representing sensory information.

Intrinsic Diversity Enables Populations to Generalize Across Stimulus Types. We next investigated the effect of diversity on stimulus coding in larger neuronal populations. In Fig. 3*A* we plot actual and reconstructed stimuli for two example populations: the first, a homogeneous group composed of five copies of the highest firing rate, most informative neuron from Fig. 2*C*; the second, a population composed of neurons with diverse parameters (Fig. 3*D*). Both populations encode stimuli composed of high frequencies with high fidelity (Fig. 3*A*); however, the diverse population is more effective in representing lower-frequency stimuli (Fig. 3*E*) than the homogeneous one (Fig. 3 *B* and *C*). Thus, although the diverse population has 45% fewer spikes than the homogeneous one, the diverse population better uses its allocation of five neurons by representing more of the relevant stimulus space with its (temporal) receptive fields.

To extend this analysis we compared how 250 populations of randomly chosen five-neuron ensembles encoded stimuli with different frequency spectra (e.g., $1/f^{\alpha}$ noise with differing values of α, white noise, etc.; $n = 8$ stimuli total; shown in Fig. S5). These stimuli were chosen to cover a wide range of input frequencies including the range of frequencies these neurons likely receive in vivo (29, 31). We created homogeneous populations, each consisting of five copies of a single MC, and heterogeneous populations generated by randomly selecting five MCs from the recorded set with replacement. To compare population responses across stimulus spectra, we ranked the populations in order of increasing average reconstruction error for each kind of stimulus and compared ranks across different stimuli. Across pairs of stimulus types population ranks were correlated (Fig. 3 *F* and *G*; $r = 0.80$ and 0.71, respectively), meaning that those populations that represent one stimulus well also represent other kinds of stimuli well (termed generalizability). Heterogeneous populations were better than homogeneous ones not just at encoding stimuli on average (Fig. 3*H*), but also at generalizing across different stimuli (specific examples in Fig. 3 *F* and *G* and summary in *I*). Thus, the observed intrinsic diversity helps encode many kinds of stimuli, conferring representational robustness to MC populations.

Populations Optimized for Specific Stimuli Combine Diversity with Homogeneity. Thus far, we have only considered sampling neurons randomly according to a particular rule (homogeneously versus heterogeneously). We next attempted to construct more

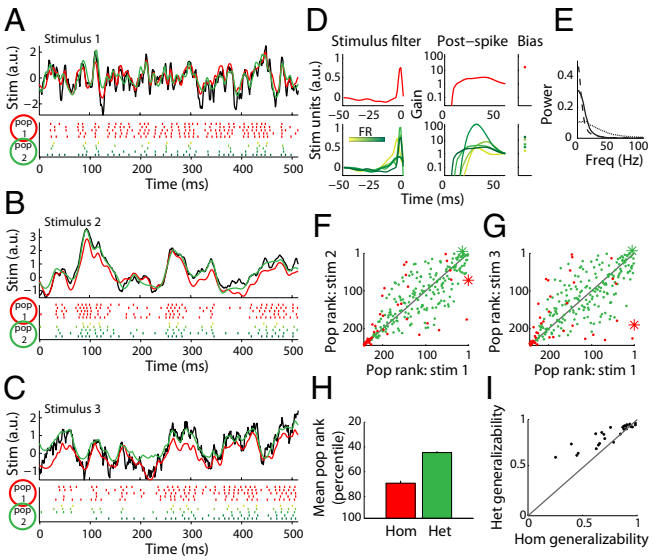

Fig. 3. Populations composed of diverse neurons effectively encode stimuli with very different frequency spectra. (*A–C*) Example stimulus (*Top*; black), rasters (*Bottom*), and reconstructions (*Top*) for a homogeneous population composed of five copies of the most informative neuron (pop 1, red) and a heterogeneous population composed of five neurons with diverse properties (pop 2, green) for three stimuli with different power spectra: stimulus 1, Gaussian white noise (GWN) convolved with an alpha function with $\tau = 3$ ms (*A*); stimulus 2, GWN with alpha function with $\tau = 10$ ms (*B*); stimulus 3, Ornstein–Uhlenbeck process with $\tau = 40$ ms (*C*). Note that although both populations can represent the stimulus in *A* well, only population 2, the diverse population, can also represent the lower frequency stimuli in *B* and *C*. (*D*) Neuron GLM parameters for the populations in *A–C*. *Top* indicates parameters for population 1 and *Bottom* for population 2 (green shades indicate different neurons). (*E*) Power spectra for the three stimuli in *A–C* (dotted, solid, dashed respectively). (*F* and *G*) Relative rankings of stimulus reconstruction accuracy for all homogeneous (hom-, red) and 200 randomly sampled heterogeneous populations (het-, green) for stimuli 1 versus 2 (*F*) or 1 versus 3 (*G*). Populations in *Top Right* indicate those which represent both stimuli accurately. Asterisks indicate populations highlighted in *A–C*. Note that hom populations are among the bottom populations and are further from the unity line than het populations. (H) Average rank of het and hom populations across eight spectrally unique stimuli (*Materials and Methods*). Het populations are consistently ranked higher (more accurate) than hom ones ($P = 0.002$, paired Wilcoxon). (*I*) Plot of generalizability, defined as the correlation of population ranks on stimulus pairs, for hom and het populations across all pairs of eight stimulus types. Each dot corresponds to the generalizability between a pair of stimulus types ($n = 28$ total pairs). Het populations are significantly more likely than hom to generalize to novel kinds of stimuli ($P = 1.5*10^{-4}$, paired Wilcoxon).

optimal groups of neurons for encoding specific stimulus types. We liken this scenario to that of sister MCs associated with a single glomerulus, which receive inputs with a specific temporal structure (26, 32) based on olfactory receptor neuron (ORN) odorant binding kinetics, which differ across glomeruli and ORN subtypes (33, 34). Would the best population for a given stimulus be more diverse than selecting MCs at random from the physiologically based set? Or would the best population be more homogeneous than random, perhaps allowing the responses of unreliable neurons to be improved upon by selecting neurons coding for redundant (i.e., degenerate) stimulus features? To answer these questions, we implemented a greedy search algorithm (35) to build the best population of model MCs to encode a given stimulus by iteratively adding neurons one at a time such that the added neuron maximized the ability of the entire population to represent the stimulus (Fig. 4*A*). Although it is not guaranteed to find the global optimum, it is an efficient and intuitive method of finding neuron groups more informative than those generated through random sampling.

Visualizing the makeup of these greedy-search-selected populations using dimensionality reduction (Fig. S6) reveals that they reflect a balance between diversity—consisting of neurons with different properties, and homogeneity, often including multiple copies of selected neurons (Fig. 4 *B* and *C* and Fig. S7). In addition, the stimulus type dictates the selection of specific neurons and the chosen level of population diversity. For example, the population selected to best encode a white-noise stimulus (Fig. 4*C*) was composed primarily of similar neurons with high firing rates; whereas, diversity in neuron properties was more important for encoding a more naturalistic stimulus with both rapidly and slowly varying temporal components (Fig. 4*B*). Using the greedy search algorithm to select populations for each of the eight stimulus types, we quantified the diversity of these populations and of randomly sampled heterogeneous and homogeneous populations (Fig. 4*D*). Surprisingly, greedy search populations were on average ~25% less diverse than heterogeneous ones when considering either stimulus filter and postspike parameters. Furthermore, quantifying population diversity for MC groups selected to best encode different stimulus types reveals that they have varying levels of diversity (Fig. 4*E* and Fig. S8), suggesting that population diversity should be preferentially tuned to the afferent stimulus distribution.

To ensure that the previous findings are not solely the result of the greedy selection process, we performed additional simulations by randomly constructing populations with differing amounts of diversity and examining the relationship between population diversity and decoding accuracy. As predicted from the greedy search results, we found evidence for a U-shaped relationship between decoding accuracy and population diversity (Fig. 4*F* and Fig. S9), indicating that neural coding is optimized at intermediate levels of diversity. However, population size is also a relevant factor in the importance of population diversity, with diversity being more important to smaller populations than larger ones (Fig. S10). This suggests that heterogeneity will be more important to populations in which the number of neurons devoted to representing a stimulus is relatively small. Furthermore, we found the benefit of neural variability to not be solely dependent upon a single GLM filter dimension (Fig. S11), such as the stimulus filter or bias term.

Discussion

Here we apply the framework of generalized linear models to study how cell-to-cell differences in intrinsic properties of olfactory bulb mitral cells influence stimulus encoding. The statistical modeling approach that we have used accurately captures the neuronal properties determining spiking and avoids overfitting. It also avoids making specific but difficult-to-verify claims about channel densities or properties that can arise from underconstrained Hodgkin–Huxley models (36). We show that diverse populations offer the advantages of more efficient encoding (defined in terms of information per cell or information per spike) and more robust coding of different kinds of stimuli, such as stimuli with wide ranges of spectral properties. This is because neurons encoding partially overlapping (i.e., degenerate) stimulus features can work together to overcome neural spike-generation noise and also encode more stimulus features together than separate. We also show that populations selected to best represent stimuli with specific spectral properties have differing amounts of diversity, which suggests that population diversity should be selectively chosen with respect to the precise stimulus to be encoded. Although variants of this framework have been used to model neural responses previously [including in single neuron modeling competitions (37, 38)] we extend these methods to describe the systematic biological differences among neurons and their impact on population coding. Given the generality of this framework, we

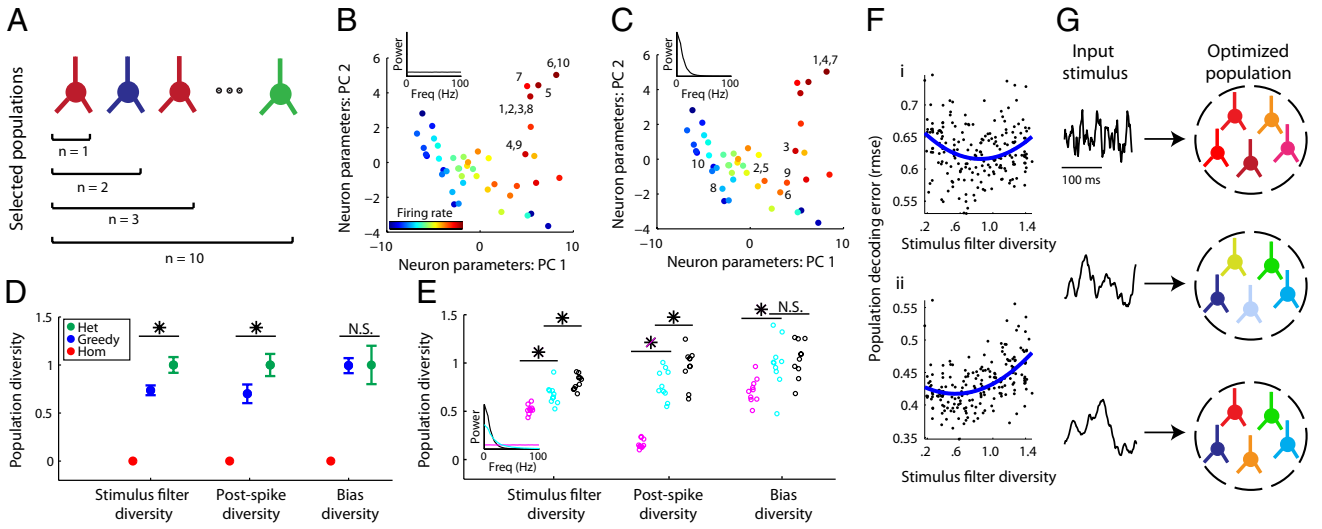

Fig. 4. Populations optimized for stimulus representation combine homogeneity with diversity. (*A*) Cartoon of greedy search algorithm to estimate the population that best represents a particular type of stimulus. Neurons were iteratively added, one at a time, to the current population of neurons such that the neuron chosen maximized the population's reconstruction accuracy. To allow for homogeneity, neurons could be added more than once (e.g., two red neurons). (*B* and *C*) Visualization of the population selected to best represent a white-noise stimulus (*B*) or a low-frequency stimulus (*C*). Graphs show neurons (as dots) projected into a 2D space using principal component analysis (PCs). Population sizes vary from $n = 1$ to $n = 10$, numbers next to dots correspond to algorithm iteration step when each neuron was added. Note that certain neurons are chosen multiple times and that stimulus type dictates the selected population diversity. (*D*) GLM parameter diversity of the greedy-search-selected populations (blue) averaged over eight different choices of stimulus spectra relative to homogeneous (red) and randomly sampled heterogeneous populations (green), $n = 10$ neurons per population. Asterisks indicate where greedy search populations are significantly less diverse than heterogeneous ($P < 0.05$) and population diversity has been normalized to that of randomly sampled heterogeneous. Error bars indicate SEM (blue) and interquartile range (green). (*E*) Greedy-search population diversity for specific stimlus types. Colors indicate different stimulus types corresponding to inset power spectrum (magenta, stimulus as in *B*; cyan, Ornstein–Uhlenbeck process with $\tau = 10$ ms; black, stimulus as in *C*), open circles indicate multiple runs of the greedy search algorithm ($n = 10$ per stimulus type), asterisks indicate significant differences in population diversity between stimulus types. (*F*) Population decoding error as a function of stimulus filter diversity for 200 randomly sampled populations (dots, $n = 5$ neurons per population) for stimulus 1 and 2 as in Fig. 3 (*A* and *B*, respectively). Least-squares fits using a second-order polynomial (blue) show that on average there is an intermediate level of stimulus filter diversity where decoding error is minimized (regression $P < 0.01$). (*G*) Cartoon showing that population diversity should be preferentially selected with respect to the specific incoming stimulus distribution.

believe that this methodology can similarly be extended to describe electrophysiological differences across neuron types and to develop hypotheses about the distinct roles of different neuron types throughout the brain.

One of the key advantages of this approach is that it allows us to use Bayesian stimulus decoding to ask how neuron-to-neuron differences in stimulus filtering and postspike properties influence population coding of arbitrary stimuli. Bayesian decoding is advantageous because it offers an optimal, best-case view of neural encoding, making few assumptions that risk underestimating the complexity of the neural code (18, 23). Although we explored the relationship between stimulus encoding in diverse and homogeneous populations in a previous study (5), performing stimulus reconstruction here allows the identification of the relative importance of variation in specific features of the sets of recorded neurons. This approach also allows us to investigate stimulus encoding in a more general context by simulating responses to arbitrary stimuli. An obvious advantage of simulation approaches is that we are not limited to only analyzing data that we are able to collect during recordings.

Our results make specific, testable predictions on the role of MC intrinsic diversity for encoding olfactory information. First, we show that when populations need to represent a variety of stimulus types, then intrinsic diversity facilitates generalizing representations across stimulus types. Second, when populations need to represent a single kind of stimulus and are allowed to selectively choose their level of variability, populations choose a balance between complete homogeneity and diversity. That is, homogenizing the input received by a population of neurons should lead the population to be less diverse. This *in silico* finding is intriguing because it is consistent with recent experimental findings showing that sister MCs, receiving primary olfactory inputs from the same glomerulus and olfactory receptor subtype, are biophysically more similar to one another than sampling MCs at random (7). Furthermore, our work makes the additional hypothesis that the level of diversity across sister MCs should be adaptive with respect to the unique stimulus distribution that these neurons receive from their olfactory receptor subtype (32–34). Therefore, we predict that the levels of MC intrinsic diversity between sister MCs should be empirically different across glomeruli (Fig. 4*G*).

We note that we made multiple assumptions here for the sake of computational tractability. Because our focus was to study the functional role of MC intrinsic diversity, we chose not to include any of the effects of neural connections such as synapses between neurons in our experiments and simulations. Given that the olfactory bulb possesses extensive lateral circuitry (11), which has been shown to also diversify MC responses (11, 39, 40), we expect that bulbar circuit activity will work in conjunction with intrinsic diversity in vivo to further diversify MC responses (41). Furthermore, when decoding we took the perspective that the best populations were those that resulted in the most faithful reconstruction of the stimulus. However, the biological solution dictating the actual amount of diversity may use alternative criteria for optimality. For example, in vivo olfactory bulb MCs may seek to represent only odor-specific stimulus components or may try to maximize stimulus representation while also minimizing the number of spikes used to transmit the information (42). We chose to avoid assumptions about which features of ORN input are most important for MCs to represent and rather to take the agnostic view that MCs should try to represent the stimulus in its entirety. Our approach, however, can readily be adapted to tasks that require representation of specific stimulus components.

Although these assumptions likely affect the quantitative details of our results, like specifying of the precise balance between diversity and feature similarity, our general finding that a precise stimulus-specific balance exists nevertheless likely holds.

We believe that our results generalize to other neural systems because this circuit motif in which multiple neurons receive highly correlated inputs occurs throughout the brain, including neocortex (43, 44). Thus, we predict that the observed degree of neuronal intrinsic variability plays a substantial role in tuning the output diversity (or redundancy) in these neurons' spiking responses and in improving stimulus encoding. Furthermore, our findings may in part explain the substantial informational redundancy found in neural populations throughout the brain (1, 10). Given that the optimal networks here are neither maximally diverse nor maximally homogeneous, these results suggest similar design principles for other systems. By mixing diversity with neural feature similarity, complex systems can simultaneously maintain efficient functioning while remaining robust to uncertain events.

Materials and Methods

A detailed description of the methods is provided in the *Supporting Information*. In brief, whole-cell patch clamp recordings of mitral cells were obtained in vitro from mouse olfactory bulb slices (5). Spikes were recorded while stimulating neurons with 40 trials of a 2.5-s duration frozen noise stimulus, generated by convolving a white-noise current with an alpha function with $\tau = 3$ ms. Point process models were fit from recordings via previously described methods (18). The models fit from physiological data were used to simulate neuron spike responses to stimuli not presenting during recording. Uncoupled populations were constructed by sampling neurons from the set of model neurons, where all neurons in a population received an identical stimulus. Population responses were decoded using the maximum a posteriori estimator (23) (posterior mode) to reconstruct the time series of stimulus input to the population. Population decoding performance was quantified using mutual information and root mean square error. To approximate the structure of the optimal population for best representing a particular kind stimulus, we implemented a greedy search algorithm (35). Unless otherwise indicated, all error bars indicate SEs and all statistical tests were Wilcoxon rank–sum tests.

ACKNOWLEDGMENTS. We thank B. Doiron, R. Kass, A.L. Kumar, A. Schaefer, C. Shalizi, L. Paninski, and the Urban Laboratory for discussions and manuscript comments. We thank S. D. Burton for collecting the data shown in Fig. S2. Funding was provided by a National Science Foundation Graduate Research Fellowship/RK Mellon Foundation Fellowship (to S.J.T.); National Institutes of Health F32 DC010535 (to R.C.G); and National Institute on Deafness and other Communication Disorders Grant R01DC0005798 (to N.N.U.).

1. Puchalla JL, Schneidman E, Harris RA, Berry MJ (2005) Redundancy in the population code of the retina. *Neuron* 46(3):493–504.
2. Moore CI, Carlen M, Knoblich U, Cardin JA (2010) Neocortical interneurons: from diversity, strength. *Cell* 142(2):189–193.
3. Marder E, Taylor AL (2011) Multiple models to capture the variability in biological neurons and networks. *Nat Neurosci* 14(2):133–138.
4. Altschuler SJ, Wu LF (2010) Cellular heterogeneity: do differences make a difference? *Cell* 141(4):559–563.
5. Padmanabhan K, Urban NN (2010) Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. *Nat Neurosci* 13(10):1276–1282.
6. Angelo K, Margrie TW (2011) Population diversity and function of hyperpolarization-activated current in olfactory bulb mitral cells. *Sci Rep* 1(1):1–50.
7. Angelo K, et al. (2012) A biophysical signature of network affiliation and sensory processing in mitral cells. *Nature* 488(7411):375–378.
8. Graves AR, et al. (2012) Hippocampal pyramidal neurons comprise two distinct cell types that are countermodulated by metabotropic receptors. *Neuron* 76(4):776–789.
9. Marsat G, Maler L (2010) Neural heterogeneity and efficient population codes for communication signals. *J Neurophysiol* 104(5):2543–2555.
10. Schneider DM, Woolley SMN (2010) Discrimination of communication vocalizations by single neurons and groups of neurons in the auditory midbrain. *J Neurophysiol* 103(6):3248–3265.
11. Giridhar S, Doiron B, Urban NN (2011) Timescale-dependent shaping of correlation by olfactory bulb lateral inhibition. *Proc Natl Acad Sci USA* 108(14):5843–5848.
12. Tkačik G, Prentice JS, Balasubramanian V, Schneidman E (2010) Optimal population coding by noisy spiking neurons. *Proc Natl Acad Sci USA* 107(32):14419–14424.
13. Schneidman E, Bialek W, Berry MJ, 2nd (2003) Synergy, redundancy, and independence in population codes. *J Neurosci* 23(37):11539–11553.
14. Edelman GM, Gally JA (2001) Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98(24):13763–13768.
15. Azouz R, Gray CM (1999) Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci* 19(6):2209–2223.
16. Stocks NG (2000) Suprathreshold stochastic resonance in multilevel threshold systems. *Phys Rev Lett* 84:2310–2313.
17. Mejias JF, Longtin A (2012) Optimal heterogeneity for coding in spiking neural networks. *Phys Rev Lett* 108(22):228102.
18. Pillow JW, et al. (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454(7207):995–999.
19. Butts DA, et al. (2007) Temporal precision in the neural code and the timescales of natural vision. *Nature* 449(7158):92–95.
20. Slee SJ, Higgs MH, Fairhall AL, Spain WJ (2005) Two-dimensional time coding in the auditory brainstem. *J Neurosci* 25(43):9978–9988.
21. Narayanan NS, Kimchi EY, Laubach M (2005) Redundancy and synergy of neuronal ensembles in motor cortex. *J Neurosci* 25(17):4207–4216.
22. Kass RE, Ventura V (2001) A spike-train probability model. *Neural Comput* 13(8):1713–1720.
23. Pillow JW, Ahmadian Y, Paninski L (2011) Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains. *Neural Comput* 23(1):1–45.
24. Warland DK, Reinagel P, Meister M (1997) Decoding visual information from a population of retinal ganglion cells. *J Neurophysiol* 78(5):2336–2350.
25. Borst A, Haag J (2001) Effects of mean firing on neural information rate. *J Comput Neurosci* 10(2):213–221.
26. Shusterman R, Smear MC, Koulakov AA, Rinberg D (2011) Precise olfactory responses tile the sniff cycle. *Nat Neurosci* 14(8):1039–1044.
27. Atick JJ, Redlich AN (1992) What does the retina know about natural scenes? *Neural Comput* 4:196–210.
28. Schoppa NE, Westbrook GL (2001) Glomerulus-specific synchronization of mitral cells in the olfactory bulb. *Neuron* 31(4):639–651.
29. Spors H, Wachowiak M, Cohen LB, Friedrich RW (2006) Temporal dynamics and latency patterns of receptor neuron input to the olfactory bulb. *J Neurosci* 26(4):1247–1259.
30. Verhagen JV, Wesson DW, Netoff TI, White JA, Wachowiak M (2007) Sniffing controls an adaptive filter of sensory input to the olfactory bulb. *Nat Neurosci* 10(5):631–639.
31. Khan AG, Sarangi M, Bhalla US (2012) Rats track odour trails accurately using a multilayered strategy with near-optimal sampling. *Nat Commun* 3:703.
32. Carey RM, Verhagen JV, Wesson DW, Pirez N, Wachowiak M (2009) Temporal structure of receptor neuron input to the olfactory bulb imaged in behaving rats. *J Neurophysiol* 101(2):1073–1088.
33. Ghatpande AS, Reisert J (2011) Olfactory receptor neuron responses coding for rapid odour sampling. *J Physiol* 589(Pt 9):2261–2273.
34. Nagel KI, Wilson RI (2011) Biophysical mechanisms underlying olfactory receptor neuron dynamics. *Nat Neurosci* 14(2):208–216.
35. Russell S, Norvig P (2009) *Artificial Intelligence: A Modern Approach* (Prentice Hall, Upper Saddle River, NJ), 3rd Ed.
36. Bhalla US, Bower JM (1993) Exploring parameter space in detailed single neuron models: simulations of the mitral and granule cells of the olfactory bulb. *J Neurophysiol* 69(6):1948–1965.
37. Gerstner W, Naud R (2009) Neuroscience. How good are neuron models? *Science* 326(5951):379–380.
38. Jolivet R, et al. (2008) The quantitative single-neuron modeling competition. *Biol Cybern* 99(4-5):417–426.
39. Dhawale AK, Hagiwara A, Bhalla US, Murthy VN, Albeanu DF (2010) Non-redundant odor coding by sister mitral cells revealed by light addressable glomeruli in the mouse. *Nat Neurosci* 13(11):1404–1412.
40. Arevian AC, Kapoor V, Urban NN (2008) Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. *Nat Neurosci* 11(1):80–87.
41. Chow S-F, Wick SD, Riecke H (2012) Neurogenesis drives stimulus decorrelation in a model of the olfactory bulb. *PLOS Comput Biol* 8(3):e1002398.
42. Weber F, Machens CK, Borst A (2012) Disentangling the functional consequences of the connectivity between optic-flow processing neurons. *Nat Neurosci* 15(3):441–448, S1–S2.
43. Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454(7206):881–885.
44. Kerr JND, Greenberg D, Helmchen F (2005) Imaging input and output of neocortical networks in vivo. *Proc Natl Acad Sci USA* 102(39):14063–14068.

Supporting Information

Tripathy et al. 10.1073/pnas.1221214110

SI Materials and Methods

Neuron Recordings. Whole-cell patch clamp recordings of mitral cells were obtained in vitro from mouse olfactory bulb slices using methods described previously (1). Mitral cells were identified under infrared differential interference contrast optics on the basis of their laminar position in the olfactory bulb and their morphology. All experiments were performed at 35 °C in standard Ringer's solution with excitatory (25 μM 2-amino-5-phosphonopentanoic acid and 10 μM 6-cyano-7-nitroquinoxaline-2,3-dione) and inhibitory (10 μM bicuculline) synaptic activity blocked.

Current-clamp recordings were performed while injecting neurons with a filtered white-noise current stimulus. Noise traces were generated by convolving a 2.5-s white-noise current with an alpha function of the form $t*exp(-t/\tau)$, where $\tau = 3$ ms. We chose this spectral structure as it generates reliable spiking in these neurons and corresponds to the timescale of fast synapses afferent to MCs (2). Each neuron received one of a small number of stimuli generated via this method (most neurons received one of three stimulus templates) and was presented ∼40 stimulus repeats. The amplitude (variance) of the noise used was between 5% and 40% of the direct current (100–800 pA, $\sigma = 20$–80 pA) offset for each cell, with the majority of cells receiving 10–20% of the dc offset. The variance of the noise was selected as previously described (2), to induce reliable firing without large input fluctuations. For all recordings, a 25 or 50 pA hyperpolarizing pulse was injected before stimuli were delivered to measure input resistance and membrane time constant, allowing us to track the stability of recordings over multiple trials. Only neurons whose firing patterns were stable across trials and fired a sufficient number of spikes in each trial (>5 Hz) were used in this study. Upon stimulation most neurons usually underwent a brief nonspiking adaptation period (111 ± 14 ms), which was assessed visually and excluded from the analysis.

Model Fitting. GLM models were fit and simulated using code provided by Jonathan Pillow (Departments of Psychology and Neurobiology, University of Texas at Austin, Austin, TX) (3). Models consisted of a temporal stimulus filter $\boldsymbol{k}$, a postspike history filter $\boldsymbol{h}$,and a constant bias term b. Stimulus and history filters were each represented using 10 spline-like cosine basis functions spaced logarithmically in time. The conditional intensity function of each neuron was modeled as $\lambda(t) = \exp(\boldsymbol{k}\cdot\boldsymbol{x} + \boldsymbol{h}\cdot\boldsymbol{r} + b)$, where $\boldsymbol{x}$ denotes the stimulus and $\boldsymbol{r}$ is the recorded spike response of the neuron. Before fitting, stimuli were downsampled to 1 KHz and standardized by subtracting the steady-state component and dividing by the amplitude of the stimulus noise. LNP models were fit using the spike-triggered average stimulus as the linear filter and estimating the spike-rate nonlinearity using 60 independent histogram bins.

Models were trained using all of the trials from the first 90% of the stimulus presentation and validated using the remaining 10%. In specific, we validated the fit of our models by comparing real and model peri-stimulus time histograms (PSTHs) computed from the test stimulus set (i.e., stimuli not used in the training of the model). We simulated model spike trains using the GLM to probabilistically generate spikes elicited by the test stimulus. PSTHs were computed by summing spikes across trials and smoothing with a Gaussian filter of width $\sigma = 2$ ms. The similarity between real and model PSTHs was reported using Pearson's correlation coefficient. For visualization, MC rasters were randomized across trials.

To assess whether the GLM fitting procedure could also fit neuron responses to multiple stimulus types, we performed an additional set of experiments on mitral cells ($n = 5$ neurons) where each neuron was stimulated with both a high- and low-frequency stimulus (white noise convolved with an alpha function with $\tau = 3$ ms and $\tau = 10$ ms, respectively). We found that the GLM modeling procedure could sufficiently fit neuron responses to each of these stimulus types, indicating that the fitting procedure is not specific to the particular stimulus type used to generate stimulus evoked responses in this study (Fig. S2).

Computation of Neuronal Statistics using GLM Models. We were interested in computing neuronal statistics like average firing rates and trial-to-trial reliability from the fitted GLM models. We computed these by simulating long experiments (∼2 min) of continuous stimulation and computing desired statistics based on these responses. We computed neuron reliability by stimulating each model neuron with multiple trials ($n = 50$) of the same stimulus and calculating reliability as the average zero-lag correlation across trials using a bin size of 5 ms.

To calculate to what extent neurons were driven by intrinsic (history plus bias) versus stimulus components (Fig. S3*B*), we used the model to simulate spike trains while storing the stimulus and intrinsic "currents," which generated the spike trains. Here the stimulus-driven component consists of the convolution of neuron's stimulus filter with the input stimulus; whereas, the intrinsic component is defined as the bias term plus convolution of neuron's spike train with its postspike filter. We calculated the ratio of intrinsic to stimulus inputs as $\langle[stim]_+\rangle / \langle[intrinsic]_+\rangle$ where $[\cdot]_+$ indicates selection of positive values of the currents and $\langle\cdot\rangle$ indicates the mean.

Stimuli Generation for Simulations. We generated zero-mean Gaussian stimuli $\boldsymbol{x}$ with a defined temporal correlation structure and length n by first generating a signal autocorrelogram with the desired spectral structure. This autocorrelogram was used to define a Toeplitz $n*n$ covariance matrix C where the elements of C indicate the pairwise correlations between points of $\boldsymbol{x}$. Correlated stimuli were then generated using the Cholesky decomposition to find a matrix L such that $C = L*L^T$, then multiplying L with a series of uncorrelated normal random variables of length n.

Here we chose eight broadly different stimuli statistics: three stimuli generated via convolving white noise with an alpha function defined as $t*exp(-t/\tau)$, where $\tau = 3$, 5, and 10 ms; three Ornstein–Uhlenbeck processes with $\tau = 10$, 20, and 40 ms, which have flat followed by $1/f^2$-like frequency profiles; a pure white-noise stimulus, with cutoff at 500 Hz; lastly, a "naturalistic" stimulus generated by combining an 8-Hz oscillatory stimulus with an Ornstein–Uhlenbeck process with $\tau = 10$ ms (displayed in Fig. S5).

Decoding. We decoded the population spiking responses using the maximum a posteriori (MAP) estimator (4), which finds the most probable stimulus given a particular population spike response. Stimuli $\boldsymbol{x}$ (typically of length ∼0.5 s, with sampling rate 1 KHz) were decoded from simulated spike responses $\boldsymbol{r}$ by computing the mode of the posterior distribution, $\overset{argmax}{x}\ p(\boldsymbol{x}|\boldsymbol{r})$, where $p(\boldsymbol{x}|\boldsymbol{r}) \propto p(\boldsymbol{r}|\boldsymbol{x})p(\boldsymbol{x})$ via Bayes' rule. Here $p(\boldsymbol{r}|\boldsymbol{x})$ is the likelihood of a response given a stimulus and is given by the set of uncoupled neuron encoding models and $p(\boldsymbol{x})$ is a multivariate Gaussian prior specifying the specific stimulus autocorrelation structure (with covariance matrix, C, used to generate stimuli). In specific, stimuli were decoded using a recently described method (4) that takes advantage of a convenient Gaussian approximation on the posterior distribution $p(\boldsymbol{x}|\boldsymbol{r})$ and its log-concavity to exactly compute the maximum (i.e., the mode) of the posterior

distribution via numerical optimization techniques. This method also provides an estimate of the uncertainty of the stimulus representation (Fig. S4 *F* and *G*). Matlab code for decoding and all other methods related to the simulation and analysis of spike trains generated from GLM models (detailed below) can be found at https://github.com/stripathy/mitral_cell_diversity.

Mutual Information Calculation. We calculated the mutual information (4) of the population response $\boldsymbol{r}$ about the stimulus $\boldsymbol{x}$ as $I(\boldsymbol{x};\boldsymbol{r}) = H(\boldsymbol{x}) - H(\boldsymbol{x}|\boldsymbol{r})$. $H(\boldsymbol{x})$ denotes the entropy of the stimulus and is defined by the multivariate Gaussian stimulus prior $p(\boldsymbol{x})$ and $H(\boldsymbol{x}|\boldsymbol{r})$ denotes the conditional entropy of the stimulus given the response and is estimated by approximating the posterior distribution $p(\boldsymbol{x}|\boldsymbol{r})$ as a multivariate Gaussian $N(\boldsymbol{x_{map}}, C)$, where the covariance matrix, C, is computed as a byproduct of our decoding method. Here we use the fact that the entropy of a Gaussian with covariance matrix C is $\ln\sqrt{(2\pi e)^n|C|}$, where $|\cdot|$ denotes matrix determinant and n is the dimension of the stimulus. Estimates of $I(\boldsymbol{x};\boldsymbol{r})$ were obtained by averaging $H(\boldsymbol{x}|\boldsymbol{r})$ over responses elicited to multiple stimuli realizations ($n = 50$). It is important to note that because this method estimates the entropy of the posterior distribution, it generally provides a better estimate of the mutual information than the commonly used lower-bound estimate of $I(\boldsymbol{x};\boldsymbol{r})$ obtained via the optimal linear estimator (5), especially when the neurons are nonlinear and not well described by an LNP model.

We computed a normalized measure of the redundancy or synergy (6) of a pair of neurons $\boldsymbol{a};\boldsymbol{b}$ relative to each of the neurons independently as $\frac{I(\boldsymbol{x};\boldsymbol{a}) + I(\boldsymbol{x};\boldsymbol{b}) - I(\boldsymbol{x};\boldsymbol{a},\boldsymbol{b})}{I(\boldsymbol{x};\boldsymbol{a},\boldsymbol{b})}$. Positive values indicate informational redundancy and negative values indicate synergy.

To elaborate on our finding of synergistic pairs of neurons (Fig. 2*F*), we note that due to computational constraints we can only decode stimuli of relatively short lengths (0.5 s). Therefore, we will tend to underestimate the information rates of neurons that fire at low firing rates. For example, when performing stimulus decoding to calculate the information rate of a single neuron with a very low firing rate, it may fire zero spikes during the time interval and thus encodes no stimulus information. However, when considering two such neurons, the two will be much more likely to fire at least one spike between them, and thus encode some nonzero stimulus information. In this example, the case of a two-neuron pair would appear synergistic relative to a single neuron alone. Therefore, if we could simulate arbitrarily long stimulus presentations we would expect this apparent synergy effect to disappear.

Calculating Population Stimulus Generalization. To calculate how well heterogeneous and homogeneous populations generalized across stimuli of differing types, we computed the generalizability for each population type. Here generalizability is defined as $corr(ranks_{stim1}, ranks_{stim2})$, or the correlation between population ranks on pairs of stimulus types.

GLM Dimensionality Reduction. We chose to reduce the dimensionality of the space defined by neuronal GLM parameters using principal component analysis (Fig. S6) for visualization and further analysis. Principle components (PCs) were generated by first concatenating waveforms of stimulus, postspike, and bias components across all neurons and standardizing before performing PC analysis. Postspike and bias terms were transformed to units of log(Gain) before concatenating. The first 10 ms of postspike filters were removed and not included in analysis.

Computing Population Diversity. We calculated population diversity as the mean Euclidean distance of GLM parameters computed between all pairs of neurons in a population. We excluded the first 10 ms of the postspike filters across neurons as most neurons were refractory during this period. The average diversity of heterogeneous populations was computed by averaging over 50,000 randomly sampled populations. When reporting the uncertainty in the diversity of randomly sampled populations (Fig. 4*D*), we chose to show a measure of the population variance (interquartile range) as opposed to SEs.

We sampled populations that varied greatly in their amount of diversity (from low to high; Fig. 4*F* and Figs. S9 and S10) through implementing stratified sampling where we first sampled 2 million five-neuron populations and then further subsampled this set to pick populations that varied uniformly in their diversity.

Eliminating Diversity in a Single GLM Dimension. We constructed populations which had diversity eliminated along a single GLM dimension (stimulus, postspike, or bias) by modifying the neuron model parameters from the ones based on the recorded neurons (Fig. S11). For example, to sample neurons where diversity in the stimulus filter had been eliminated, we set the stimulus filter for all neurons that of the mean stimulus filter computed over all neurons. We further ensured that mean of the firing rates across neurons were similar between the original and diversity-reduced populations.

1. Padmanabhan K, Urban NN (2010) Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. *Nat Neurosci* 13(10): 1276–1282.
2. Galán RF, Ermentrout GB, Urban NN (2008) Optimal time scale for spike-time reliability: theory, simulations, and experiments. *J Neurophysiol* 99(1):277–283.
3. Pillow JW, et al. (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454(7207):995–999.
4. Pillow JW, Ahmadian Y, Paninski L (2011) Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains. *Neural Comput* 23(1):1–45.
5. Warland DK, Reinagel P, Meister M (1997) Decoding visual information from a population of retinal ganglion cells. *J Neurophysiol* 78(5):2336–2350.
6. Schneidman E, Bialek W, Berry MJ, 2nd (2003) Synergy, redundancy, and independence in population codes. *J Neurosci* 23(37):11539–11553.

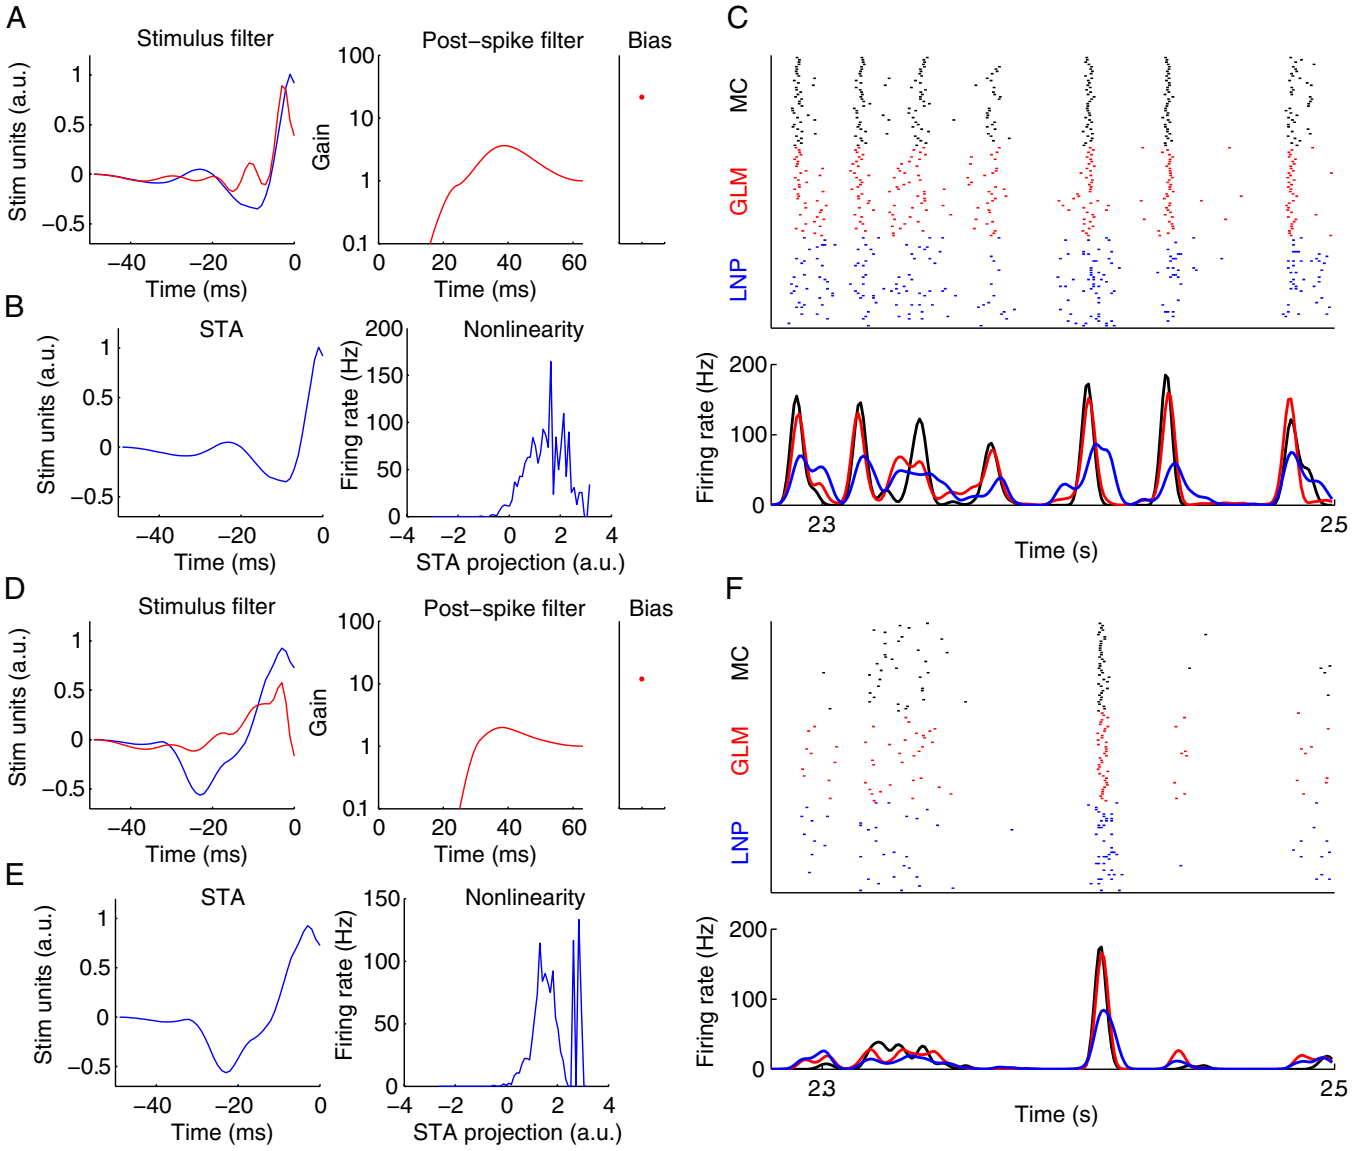

Fig. S1. Comparison of GLM and LNP model parameters and prediction accuracy. (*A*) GLM parameters (red) and spike triggered average current (STA, *Left*, blue) for the neuron in Fig. 1*A*. (*B*) LNP parameters for same neuron. (*C*) Experimental MC, GLM, and LNP rasters (*Top*) and PSTH (*Bottom*). Note that GLM spikes replicate the MC more precisely than the LNP model. (*D*–*F*) As in *A*–*C* but for neuron in Fig. 1*B*.

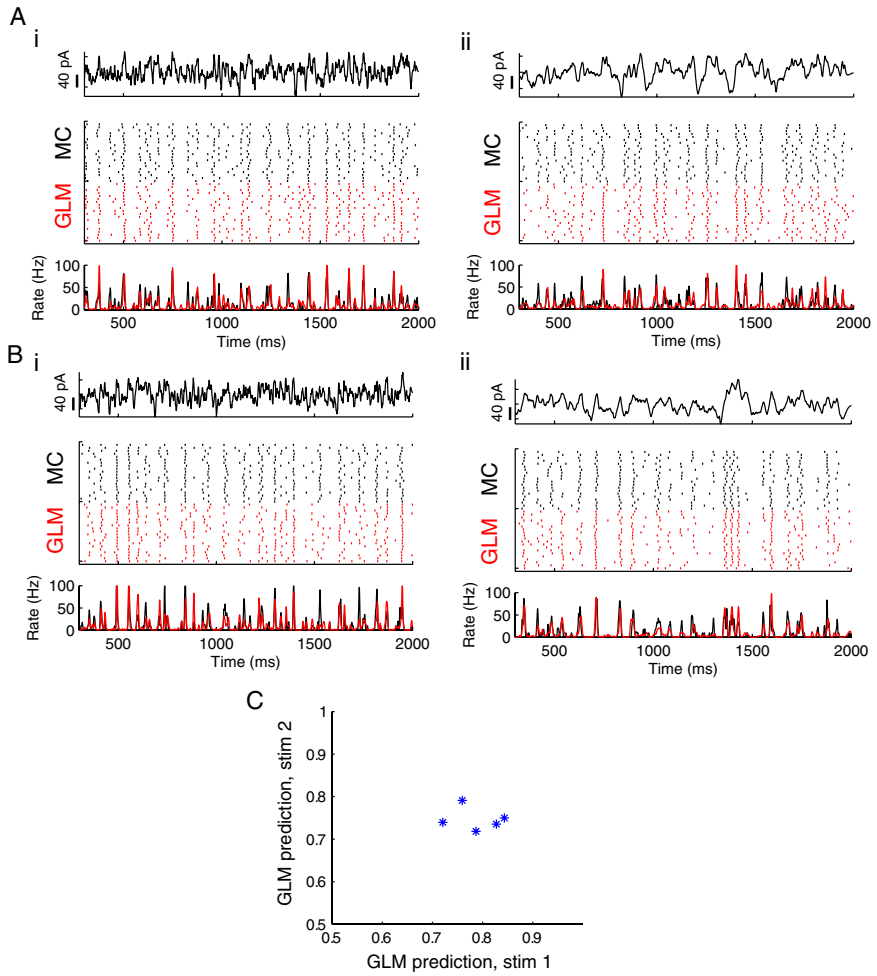

Fig. S2. GLM modeling procedure is effective at fitting neuron responses to different stimulus types and temporal correlation structures. (*A*) Stimulus and spike rasters for a recorded MC (black) and simulated model responses (GLM, red) for a high-frequency stimulus (*i*, white noise convolved with alpha function with τ = 3 ms, stim 1) or a low-frequency stimulus (*ii*, white noise convolved with alpha function with τ = 10 ms, stim 2). Entire stimulus and spike response shown to illustrate temporal correlation difference between stimuli in *i* and *ii*. (*B*) As in *A*, but for a different mitral cell. (*C*) GLM models accurately predict neuron responses to each of these two stimulus types (n = 5 MCs). Prediction accuracy computed as the correlation coefficient between MC and model PSTH.

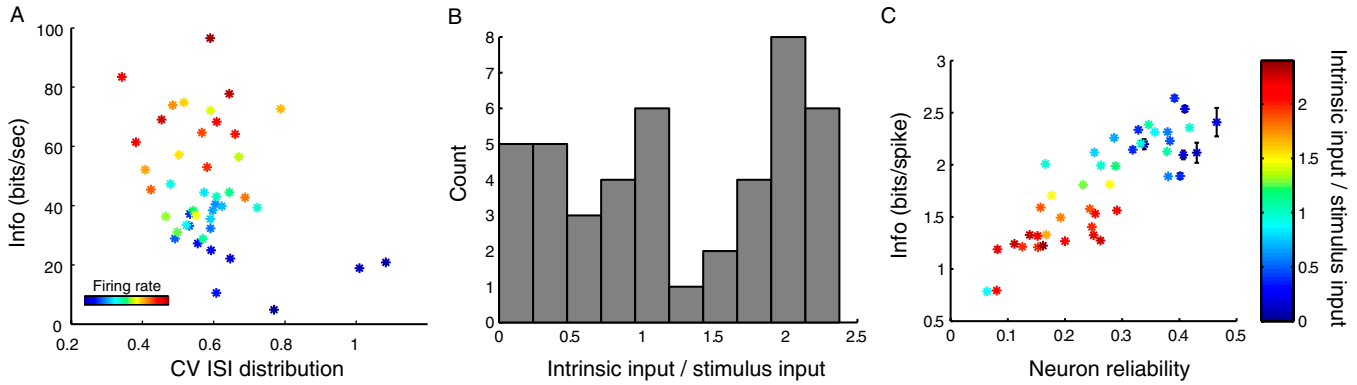

Fig. S3. Relationship between neuronal information rate and other spiking characteristics. (*A*) Neuron information rate as a function of coefficient of variation (CV) of the interspike interval distribution (ISI) (CV ISI distribution). Neurons with higher information rates tend to have lower CVs, indicating that they fire more regularly (i.e., less Poisson-like). (*B*) Histogram of the relative ratio of neuron spiking due to stimulus driven or intrinsic components (bias plus history term) as indicated from the GLM. Neurons vary from being driven entirely by stimulus to being driven primarily by intrinsic components. (*C*) Same as Fig. 2*D* but neurons are colored by ratio of intrinsic to stimulus input (as in *B*). Neurons that are primarily stimulus driven are both more reliable and encode more information per spike.

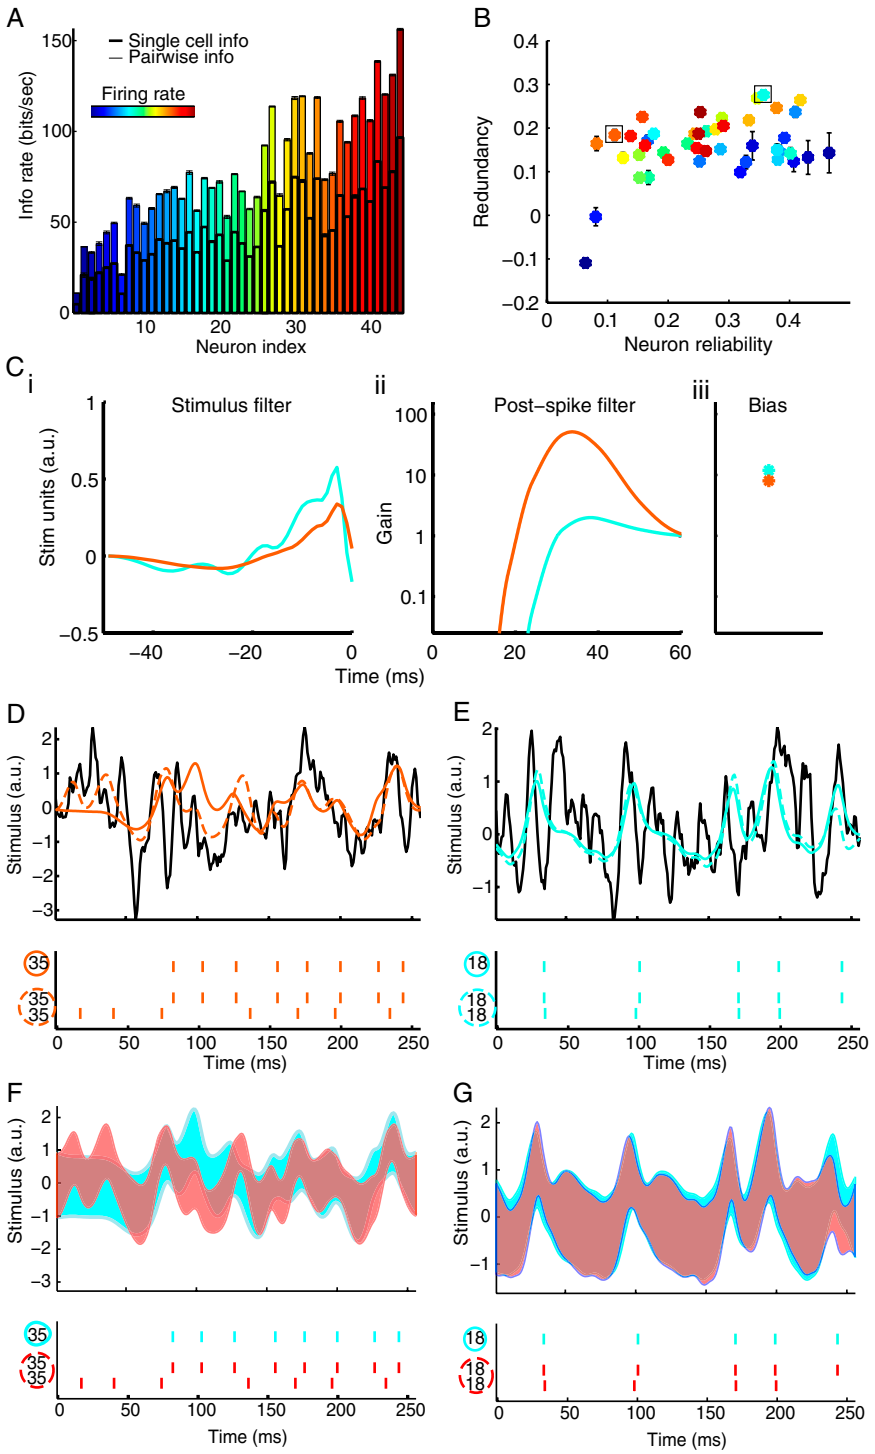

Fig. S4. Two copies of the same neuron are more informative than a single neuron alone. (*A*) Stimulus-response mutual information for single neurons (thick line) and populations of two copies of the same cell (thin line). In all cases, two neurons convey more information than a single neuron alone. (*B*) Informational redundancy of homogeneous pairs versus single neuron populations plotted as a function of single neuron reliability. Boxes show neurons 18 and 35, which are also highlighted in Fig. 2*B*. (*C*) GLM parameters for boxed neurons in *B*. These two neurons were chosen as an example because they have very different GLM parameters which cause them to spike differently in response to the same stimulus. Neuron 18 (cyan) fires primarily as a result of the stimulus whereas neuron 35 (orange) has a large postspike rebound current causing it to fire relatively periodically. (*D* and *E*) Example stimulus and reconstructions for neurons 35 and 18. Note that for neuron 35, an unreliable neuron, there is an improved stimulus representation when another copy of 35 is added; this is effect is less pronounced for neuron 18, which fires more reliably. (*F* and *G*) As in *D* and *E* but showing uncertainty in stimulus reconstruction. Blue shaded trace indicates neuron 18 and red trace indicates neuron 35. Stimulus not shown for clarity. Error bars represent 1 SD of the Gaussian estimate of the decoded posterior distribution. For both neurons, the stimulus representation is less uncertain with neuron pairs versus single neurons, highlighting the benefit of redundancy or pooling representations over neurons.

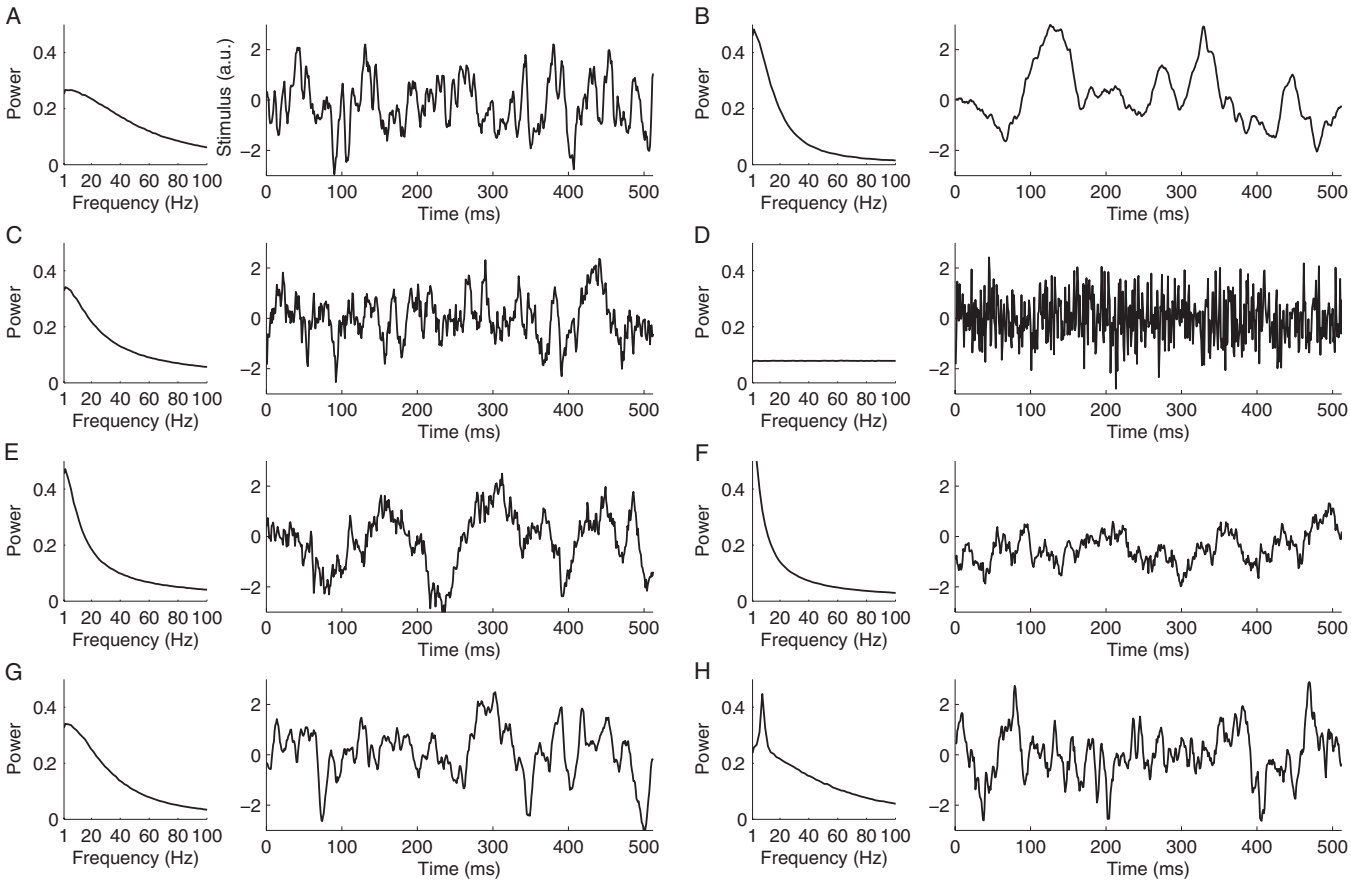

Fig. S5. Examples of the eight stimulus statistics used in this study. (*A*) White noise convolved with alpha function with $\tau = 3$ ms. (*B*) White noise convolved with alpha function with $\tau = 10$ ms. (*C*) Ornstein–Ulhenbeck process with $\tau = 10$ ms. (*D*) Pure white-noise stimulus (max frequency = 500 Hz). (*E*) Ornstein–Ulhenbeck process with $\tau = 20$ ms. (*F*) Ornstein–Ulhenbeck process with $\tau = 40$ ms. (*G*) White noise convolved with alpha function with $\tau = 5$ ms. (*H*) Naturalistic sniffing stimulus generated by combining 8-Hz oscillation with Ornstein–Ulhenbeck process with $\tau = 10$ ms.

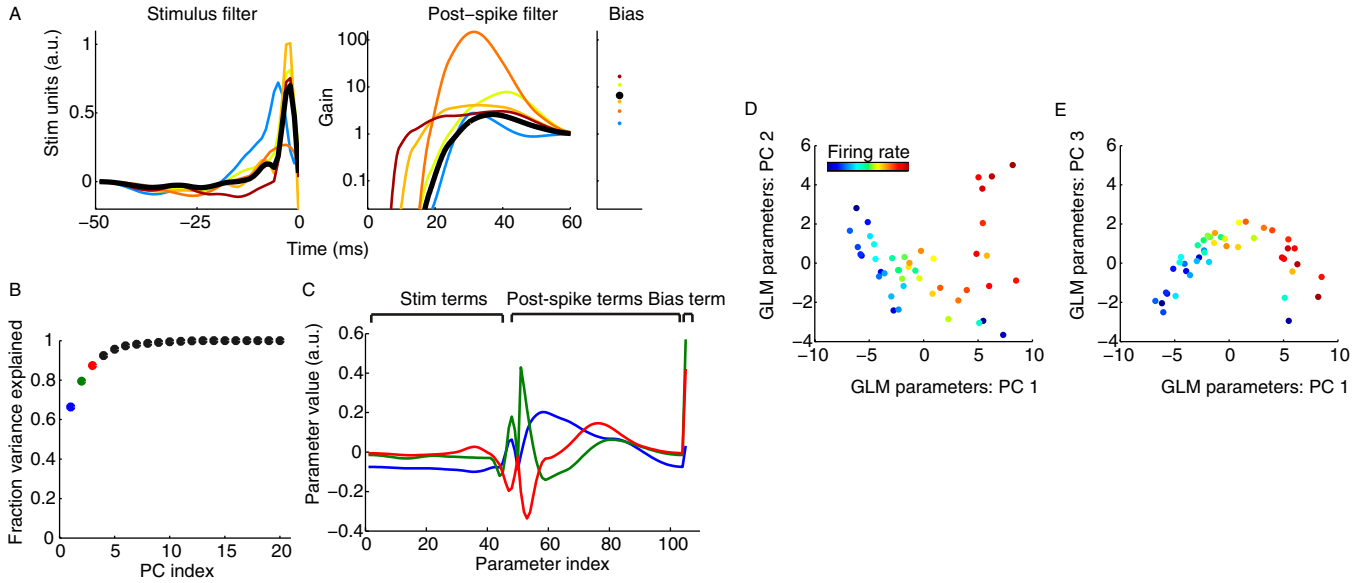

Fig. S6. Decomposition of neuronal GLM parameter space into a small number of principal components. (*A*) GLM parameters for an example five-neuron population (colors) and mean GLM parameters across all neurons (black). (*B*) Percent variance explained for each subsequent PC. The first three principal components explain 85% of variance among GLM parameters. (*C*) Visualization of the first three PCs computed from GLM parameters. (*D* and *E*) Projection of neurons (dots) into space defined by PCs 1 and 2 (*D*) or PCs 1 and 3 (*E*). This analysis plots neurons such that those with similar GLM parameters are plotted close to one another. The computed PCs largely reflect differences among postspike and bias terms, and to a lesser extent stimulus filters. Neurons with negative PC1 tend to have low firing rates, low-amplitude stimulus filters (relative to the mean across neurons) and longer refractory periods with less of a tendency to burst 20–40 ms following a spike and vice versa for positive PC1. Neurons with high PC2 tend to have high baseline excitability, very short refractory periods, and increased amplitude stimulus filters relative to the mean.

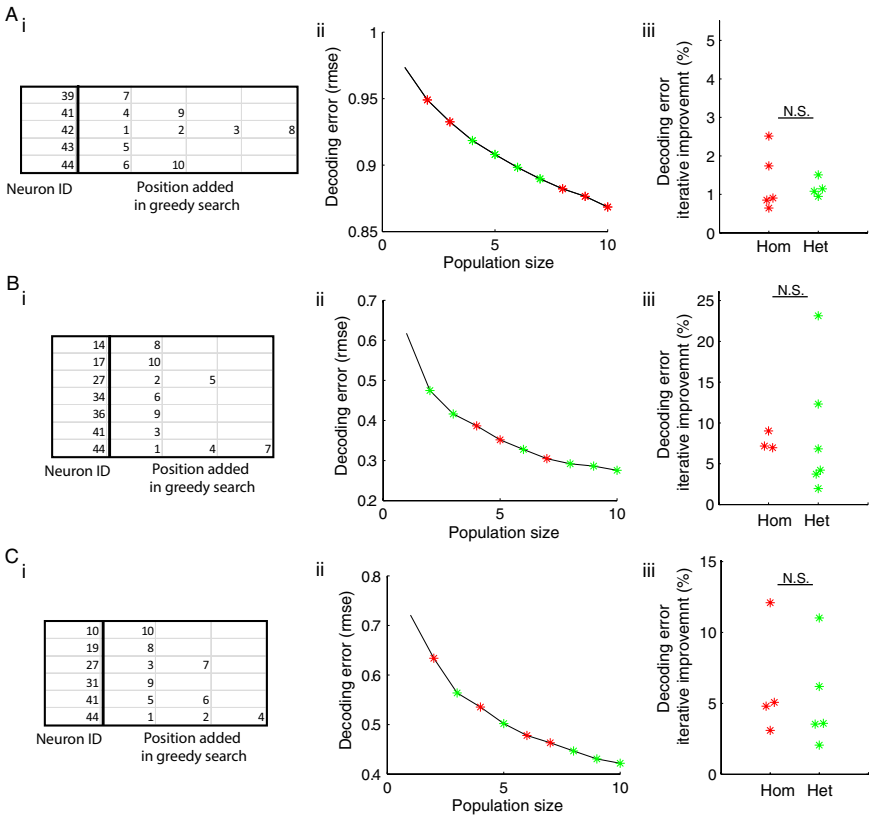

Fig. S7. Selection of homogeneous neurons during greedy search does not preferentially occur after coding improvements have saturated. (*A, i*) Detailed analysis of neurons added and their selected position during greedy search algorithm. Neuron population selected to best encode a white-noise stimulus. (*ii*) Decoding error as a function of population size during greedy search optimization procedure. Colored asterisks indicate whether neuron added at *n*th iteration is a homogeneous neuron (red, i.e., a copy of selected neuron is already in the population) or heterogeneous (green). Same population as shown in *i*. (*iii*) Same data as *ii*, plotted as iterative improvement in decoding error when adding an additional neuron, broken down by if neuron added is a homogeneous neuron (red) or heterogeneous neuron (green). There is no significant difference ($P > 0.05$) in decoding error improvement between homogeneous and heterogeneous groups. (*B* and *C*) As in *A*, but for a low-frequency stimulus (*B*, white noise convolved with alpha function with $\tau = 10$ ms) or a high-frequency stimulus (*C*, white noise convolved with alpha function with $\tau = 3$ ms).

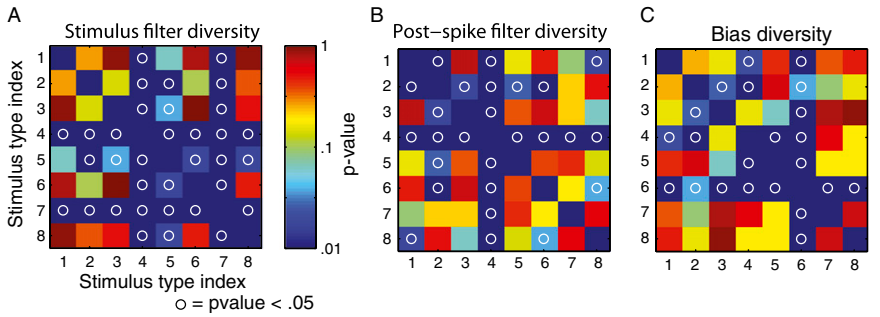

Fig. S8. Companion figure to Fig. 4*E*, showing that populations selected through the greedy search procedure to best represent different stimulus types tend to have different amounts of GLM parameter diversity. (*A–C*) Pairwise comparisons between the amounts of population diversity resulting when populations are optimized to best represent different stimulus types. Population diversity values are first computed by running the greedy search algorithm multiple times for each stimulus type ($n = 10$). Colored squares indicate *P* value for statistical test comparing population diversity between each pair of stimulus types. Circles between stimulus pairs indicates that populations are significantly different in terms of their diversity (Wilcoxon, $P < 0.05$). Stimuli are numbered as in Fig. S5 (1): White noise convolved with alpha function with $\tau = 3$ ms (2). White noise convolved with alpha function with $\tau = 10$ ms (3). Ornstein–Ulhenbeck process with $\tau = 10$ ms (4). Pure white-noise stimulus (max frequency = 500 Hz) (5). Ornstein–Ulhenbeck process with $\tau = 20$ ms (6). Ornstein–Ulhenbeck process with $\tau = 40$ ms. (7) White noise convolved with alpha function with $\tau = 5$ ms. (8) Naturalistic sniffing stimulus generated by combining 8-Hz oscillation with Ornstein–Ulhenbeck process with $\tau = 10$ ms.

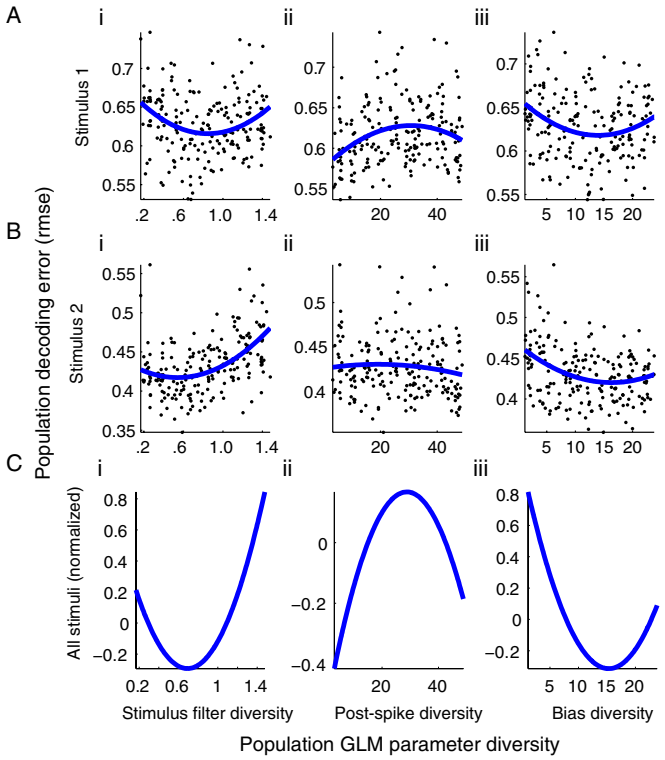

Fig. S9. Evidence for a U-shaped relationship between population diversity and decoding error. (*A* and *B*) Decoding error for five-neuron populations (black dots) as a function of population diversity along stimulus filters (*i*); postspike filters (*ii*); and bias parameters (*iii*) for stimulus 1 (*A*, as in Fig. 3*A*) or stimulus 2 (*B*, as in Fig. 3*B*). (*i–iii*) Two-hundred heterogeneous populations were drawn such that populations of varying diversity (from superdiverse through subdiverse) were sampled with equal probability (*Materials and Methods*; $n = 5$ neurons per population). Blue line shows fit of a quadratic polynomial, used to test for expected U-shaped relationship. In all cases, the regression coefficient associated with the quadratic term of the polynomial fit was positive and significant ($P < 0.01$), except for *A, ii* and *B, ii*, indicating that reconstruction error is minimized at an intermediate values of stimulus filter and bias diversity. The reason why there does not appear to be a concave-up U-shaped relationship for postspike filters is due to sampling confounds: low postspike diversity populations tend to have higher firing rates than high postspike diversity populations. (*C*) As in *A*, but showing U curves averaged across each of the eight stimuli. In this case, the decoding error was first normalized to *z* scores before performing the regression, allowing comparison across stimuli.

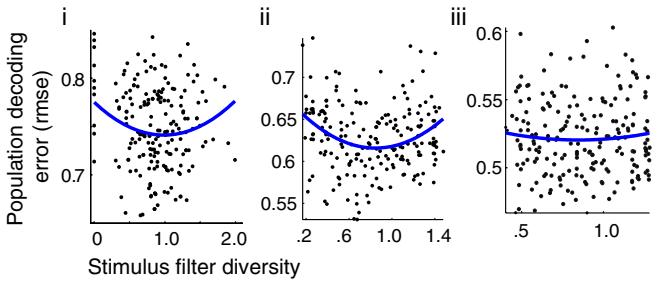

Fig. S10. Lack of a substantial U-shaped relationship between population diversity and decoding error for larger sized populations. Decoding error for neuron populations (black dots) of size $n = 2$ neurons (*i*, data reanalyzed from Fig. 2); 5 neurons (*ii*, same as Fig. 4*F*); or 10 neurons (*iii*) per population as a function of population diversity along the GLM stimulus filter dimension for stimulus 1 (as in Fig. 3*A*). (*i–iii*) Two-hundred heterogeneous populations were drawn such that populations of varying diversity (from superdiverse through subdiverse) were sampled with equal probability (*Materials and Methods*). Blue line shows fit of a quadratic polynomial, used to test for expected U-shaped relationship. The simplest explanation for why the observed U-shaped curve effect gets weaker with more neurons per population is that each of these larger populations have saturated in their ability to represent the stimulus. In this case, it matters less whether the populations are diverse (or not) because there are enough neurons in each population to effectively represent the stimulus.

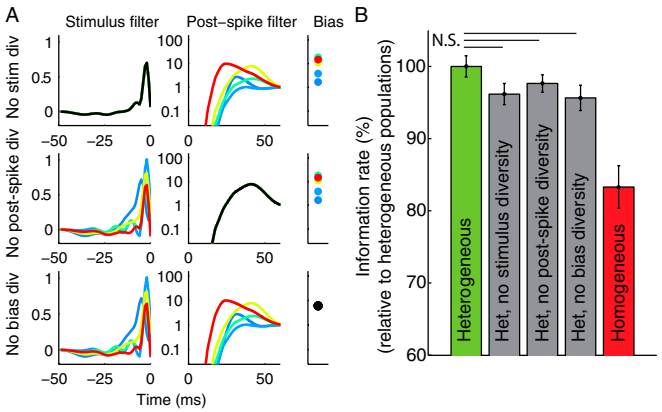

Fig. S11. Benefit of neuron variability does not depend on a single GLM model dimension. (*A*) Example five-neuron populations where population variability in a single GLM dimension (stimulus, postspike, and bias) has been eliminated (*Top*, *Middle*, and *Bottom*, respectively). (*B*) Mean information rates for heterogeneous, homogeneous, and stimulus, postspike, and bias-reduced diversity populations. Information rates computed relative to heterogeneous populations ($n = 5$ neurons per population, 200 populations per condition, stimulus is high-frequency stimulus, white noise convolved with alpha function with $\tau =$ 3 ms). None of the reduced diversity populations were significantly different from random heterogeneous populations ($P > 0.05$, Wilcoxon, N.S.) indicating that the coding benefits of diversity do not rely upon a single GLM dimension. This figure suggests that the representational advantage of neuron variability is not specifically tied to any one of the three GLM dimensions.