POPULATION CODING USING FAMILIARITY-CONTINGENT NOISE

Gerard J. Rinkus

Brandeis University, 415 South St. Waltham, MA 02453, USA grinkus@brandeis.edu

Many prior neural models of decision-making use a global arousal measure, perhaps reflecting norepinephrine levels, to titrate randomness into the choice process. The value (expected reward), V, of each possible choice (hypothesis) is computed. Then the V distribution is converted to a probability distribution, ρ, as a function of arousal level; i.e., higher arousal → more randomness added → less likely that the highest-V choice wins; lower arousal (i.e., more focused attention) → less randomness added → more likely that the highest-V choice wins. In the main, these prior models have used localist representations (codes) of choice; i.e., one coding unit per choice, whether that unit be a single cell or a distinct population of cells. Our proposed model departs from earlier work in two ways. 1) Instead of arousal/attention, it uses a global measure of *familiarity*, G, i.e., the degree of match between the expected and actual inputs, to titrate randomness. 2) It uses a sparse distributed code, i.e., each choice's code is a *set* of Q cells and any given cell participates in many codes. Instead of expected reward, we define a cell's V as the degree of match between its receptive field and its current input pattern, i.e., a *local* degree of *evidence*.

The figure's top row shows hypothetical V values over a representational field with 24 cells grouped into six WTA clusters. It contrasts two cases: *unfamiliarity* (all cells have weak local evidence, $V \approx 0$) and *perfectly familiarity* (each cluster has a cell with $V=1$). We call the set of $Q=6$ cells with the maximum V, $\hat{V}$, in the cluster (black bars), the *most favored* code, or $\hat{V}$ code. Note, the $\hat{V}$ code is the same in both cases. But, the average, G, of the $\hat{V}$ code differs greatly, ~0.1 for unfamiliar case, 1 for familiar. Normatively, when unfamiliarity is detected, a new code having little overlap with any previously assigned code should be assigned. Our model achieves this by making the V-to-ρ map be a constant function (green line). Choosing six winners from the uniform distributions (bottom left) yields the minimal expected overlap between the final code (bottom row) and the $\hat{V}$ code (code separation). Conversely, when perfect familiarity is detected ($G=1$), the model should reactivate the code that represented the current (familiar) condition in the past, i.e., the $\hat{V}$ code. Thus, the V-to-ρ map becomes highly expansive (green sigmoid), yielding the highly peaked distributions (lower right). This maximizes the probability that the $\hat{V}$ cell in each cluster wins, and thus, that the $\hat{V}$ code, as a whole, gets reactivated (code completion). More generally, morphing the V-to-ρ map smoothly based on G confers the property that similar inputs map to similar codes.