Supporting Information

Gavornik et al. 10.1073/pnas.0901835106

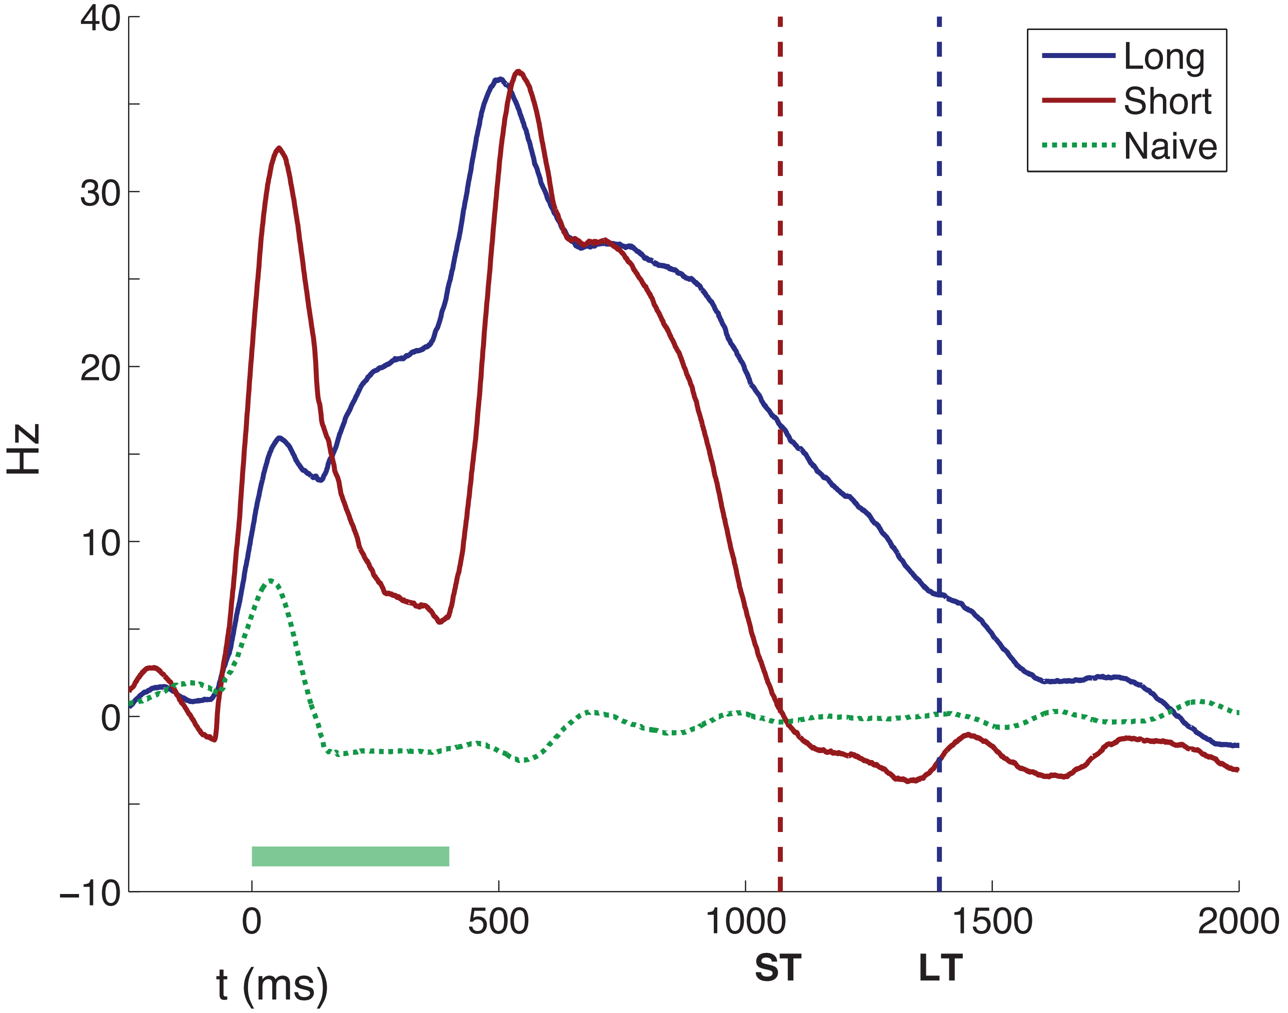

Fig. S1. Example responses of single V1 neurons to visual stimulation. Plot of the average evoked response of a neuron recorded in a naïve animal (dotted green) and the after training responses of neurons trained with short (red) and long (blue) reward times. In the naïve animal, neurons respond briefly during the period of stimulation (green bar). During training, left eye and right eye stimulations are paired with rewards delivered after a short (ST) or long (LT) delay period, respectively (dashed vertical lines). After training, neuronal responses evoked by a given stimulus can persist until the reward time paired with that stimulus. The plots show the difference between the dominant and non-dominant eye responses per neuron smoothed with a Gaussian kernel (SD 50 ms); for detailed methods and plots of average population responses, see Shuler MG, Bear MF [(2006) Reward timing in the primary visual cortex. *Science* 311:1606–1609].

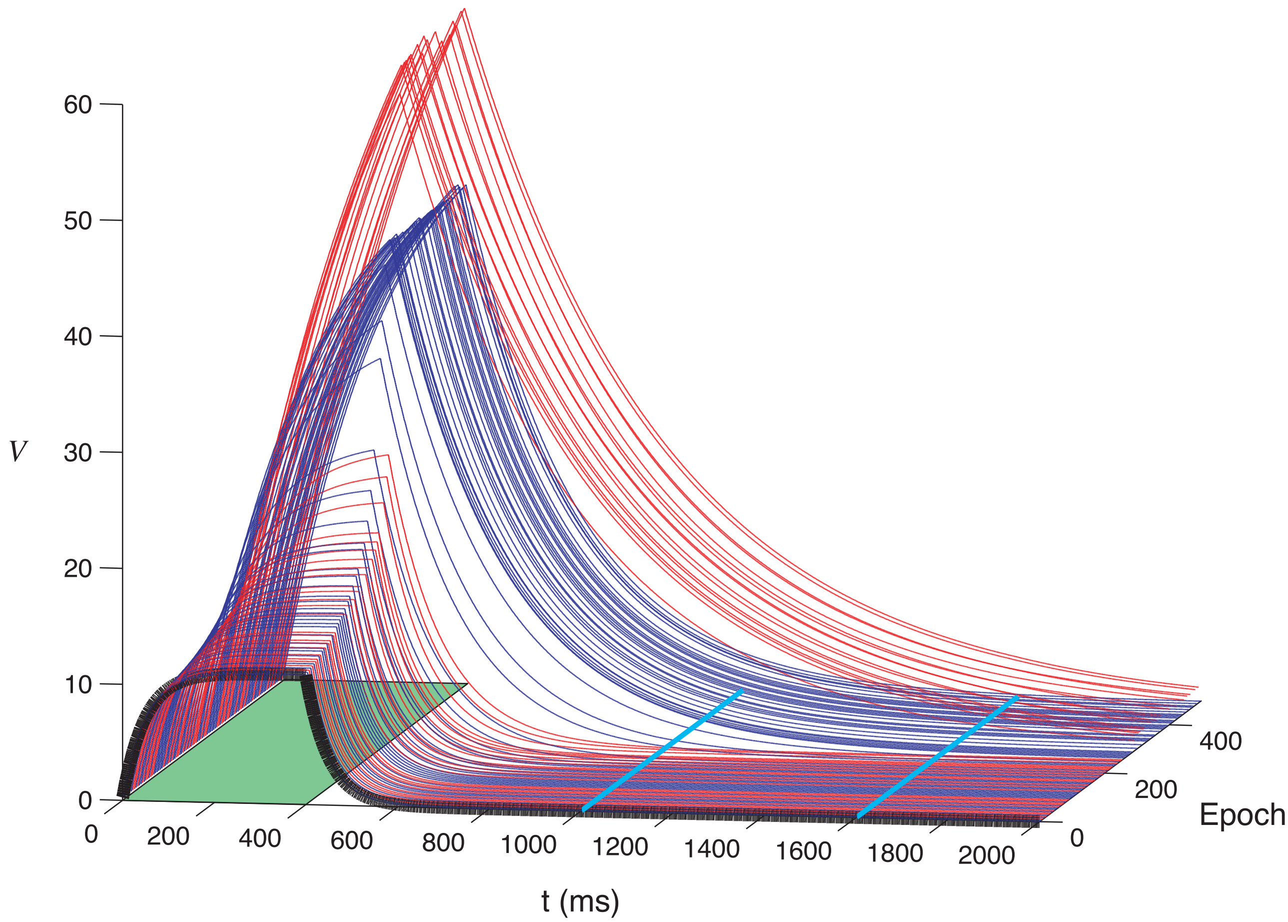

Fig. S2. Training in the rate based model. This plot shows the response of the neurons responsive to left (blue lines) or right eye (red lines) stimulation during each epoch of a training session. The stimulus is active during the period indicated by the green patch and cyan lines show reward times. The response of the naïve network, indicated by thick black line, decays quickly back to zero following stimulation. As training progresses, the responses to both inputs increase until the appropriate activity level is reached at the time of reward at which time learning stops.

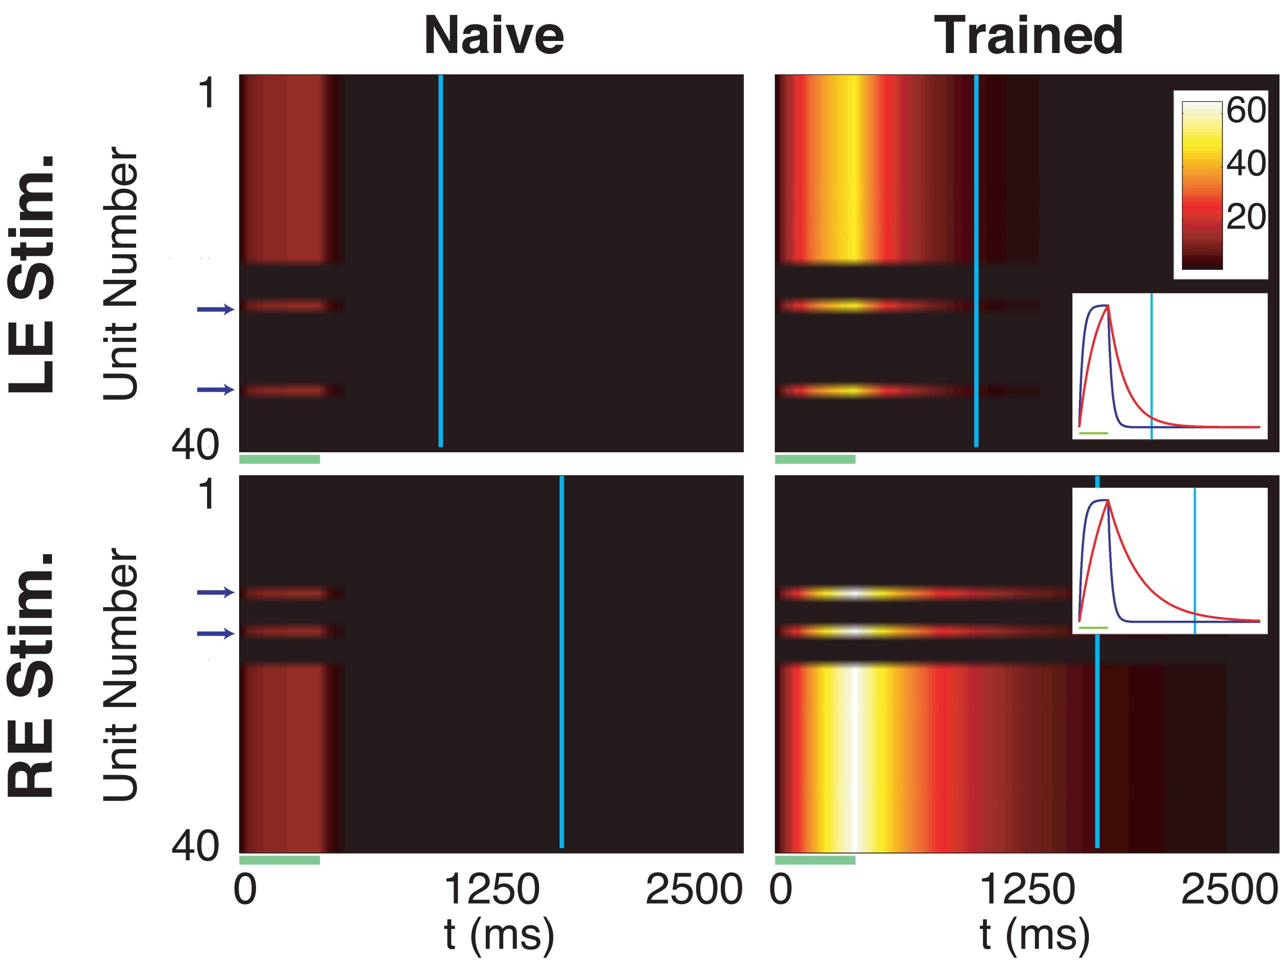

Fig. S3. Network trained with overlapping inputs. Naïve network structure is the same as in Fig. 3, but the input patterns are set so that several neurons respond binocularly (binocular neurons marked with arrows). Neurons in the trained network respond as with monocular training, except that binocular neurons have learned appropriate responses for both reward times. See Fig. S4 for plot of a binocular neuron's activity elicited by both inputs.

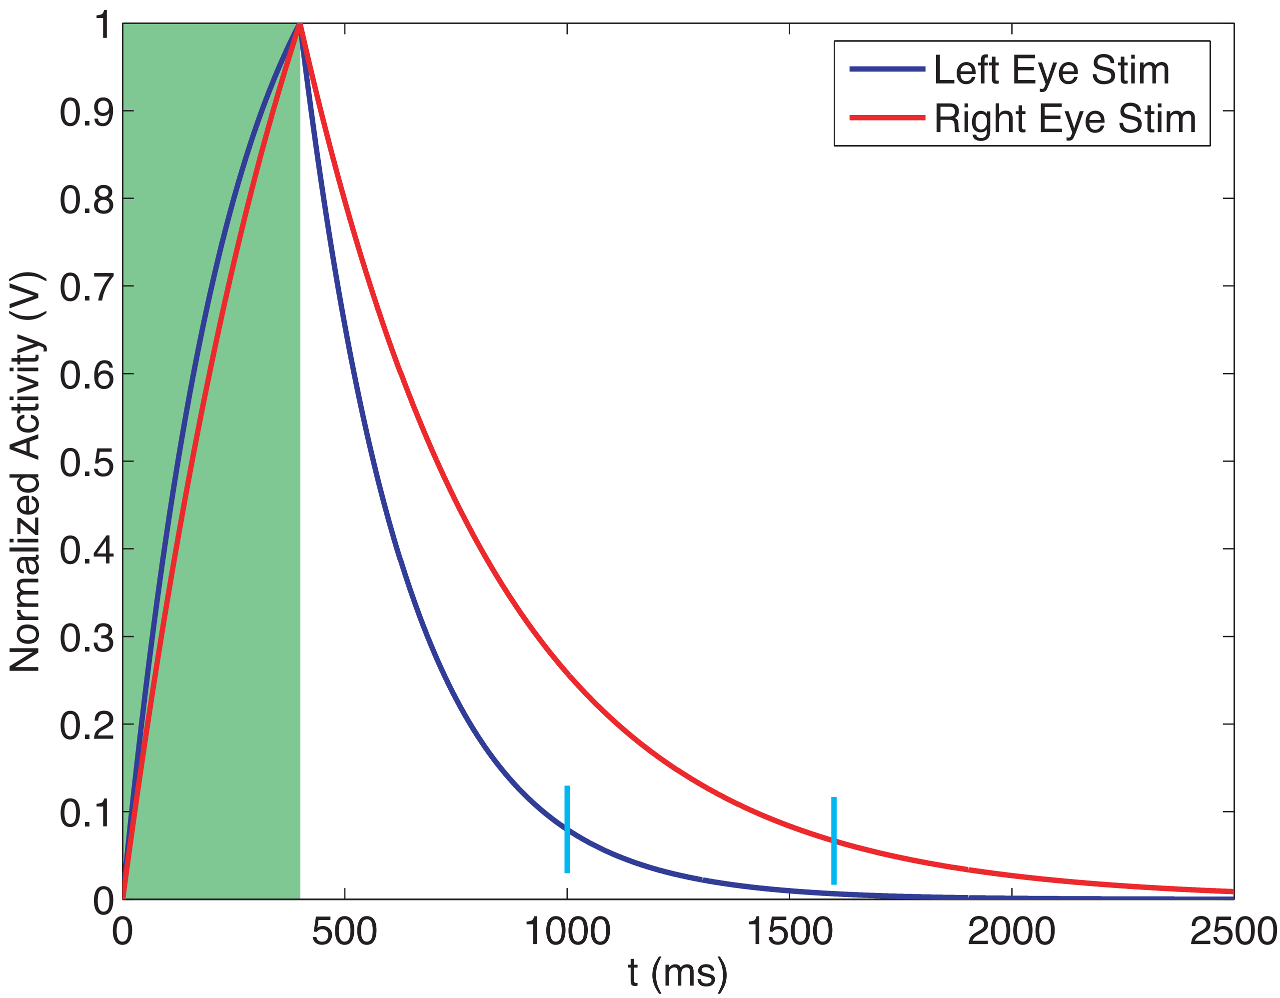

Fig. S4. Response of a binocular neuron to left and right eye stimulation. This plot demonstrates the normalized activity of a single binocularly responsive neuron to both left (blue line) and right (red line) eye stimulation. The stimulus is active during the period indicated by the green patch. The reward times used during training for both stimuli are shown by ticks at 1,000 and 1,600 ms. The binocular neuron participates in computation of both temporal representations and responds with different time constants depending on which input is presented to the network. The neuron shown is one of the binocularly responsive neurons from Fig. S3.

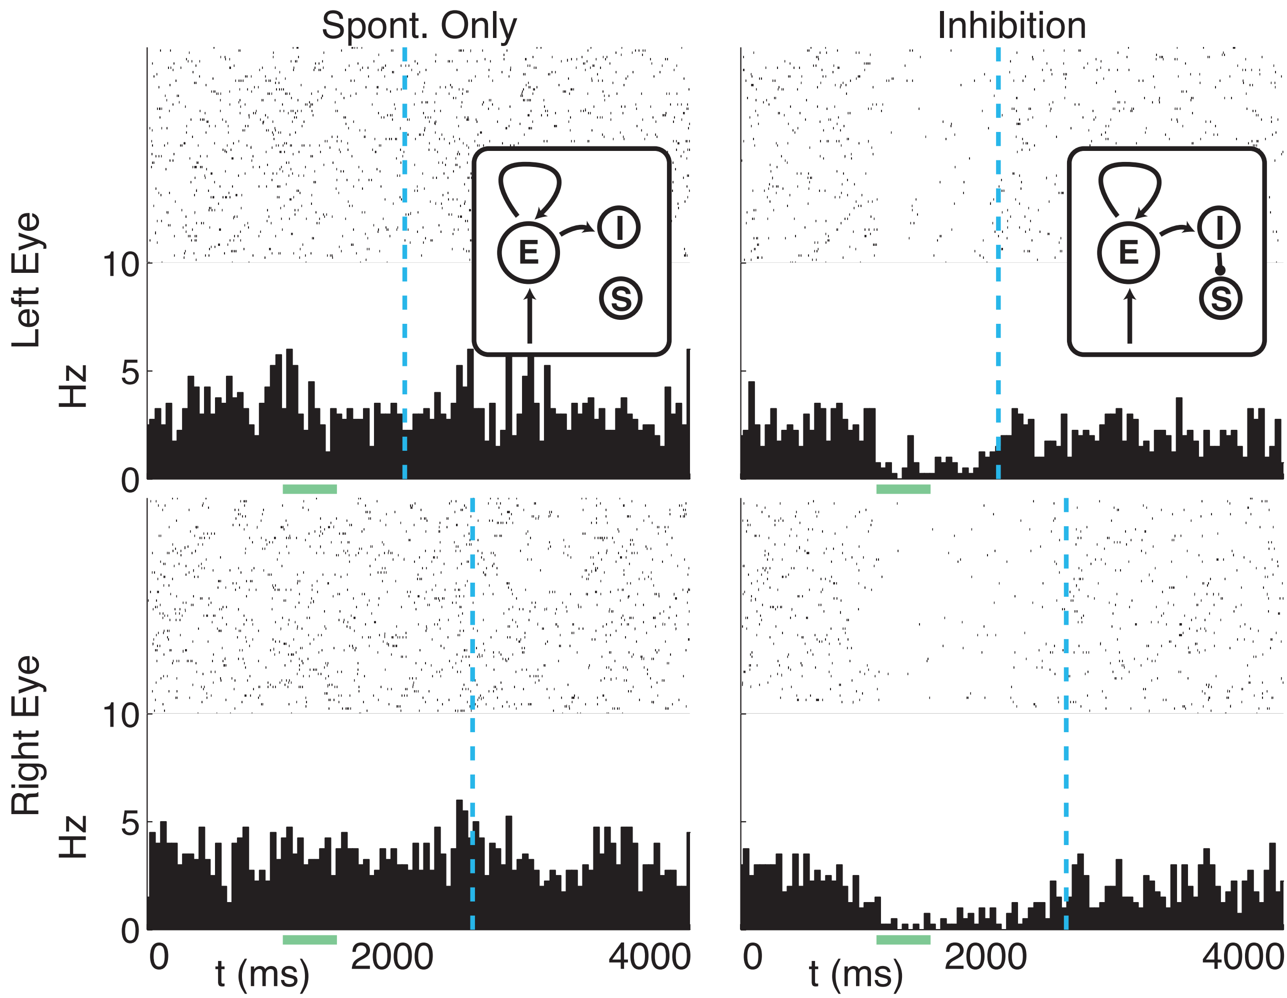

Fig. S5. Sustained decrease. This plot demonstrates how a population of neurons trained to show sustained increase in firing rates can engender a sustained decrease when embedded in a network including inhibitory neurons. As shown in the *Inset*, recurrent layer neurons (E) drive inter-neurons (I) that inhibit the activity of a population of spontaneously spiking neurons (S). The rasters and histograms in the left column show spiking activity of the neurons in S without inhibition; this activity is not affected by the feed-forward stimulation of E (indicated by green bars). Including inhibitory connection between I and S (right column) results in a decreased firing rate of S that persists until the time of reward (dashed line) for both left and right-eye stimulation (top and bottom rows). Here, reward is still encoded by recurrent excitatory weights and the sustained decrease is a derivative form of representation.

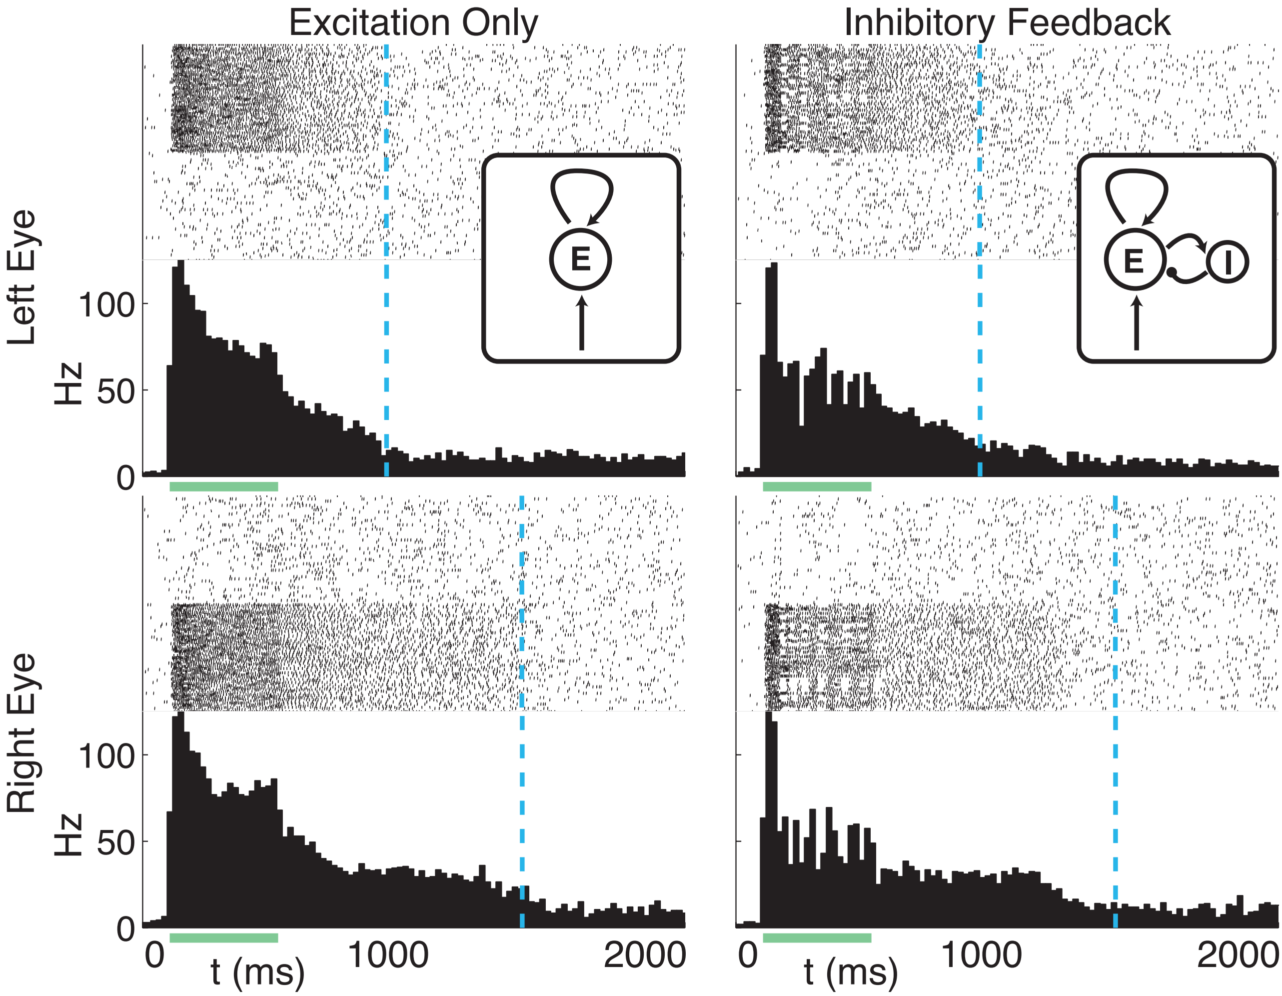

Fig. S6. Inhibition can decreases evoked firing rates. This plot demonstrates one way homeostatic mechanisms could be used by the brain to limit maximum firing rates of neurons in trained networks. Neurons in the recurrent layer (E in *Inset*) were trained to respond to left and right eye stimulation with two different reward times (dashed lines). In the left column, neurons in E receive only excitatory feed-forward and lateral inputs. In the right column, these same neurons excite a smaller population of inhibitory neurons (I) that project back into the recurrent layer. The inhibition decreases the magnitude of evoked response with little impact on temporal representations. Connections between layers E and I were established randomly and weights were set heuristically to demonstrate this result. Stimulation is active during the periods indicated by green bars.

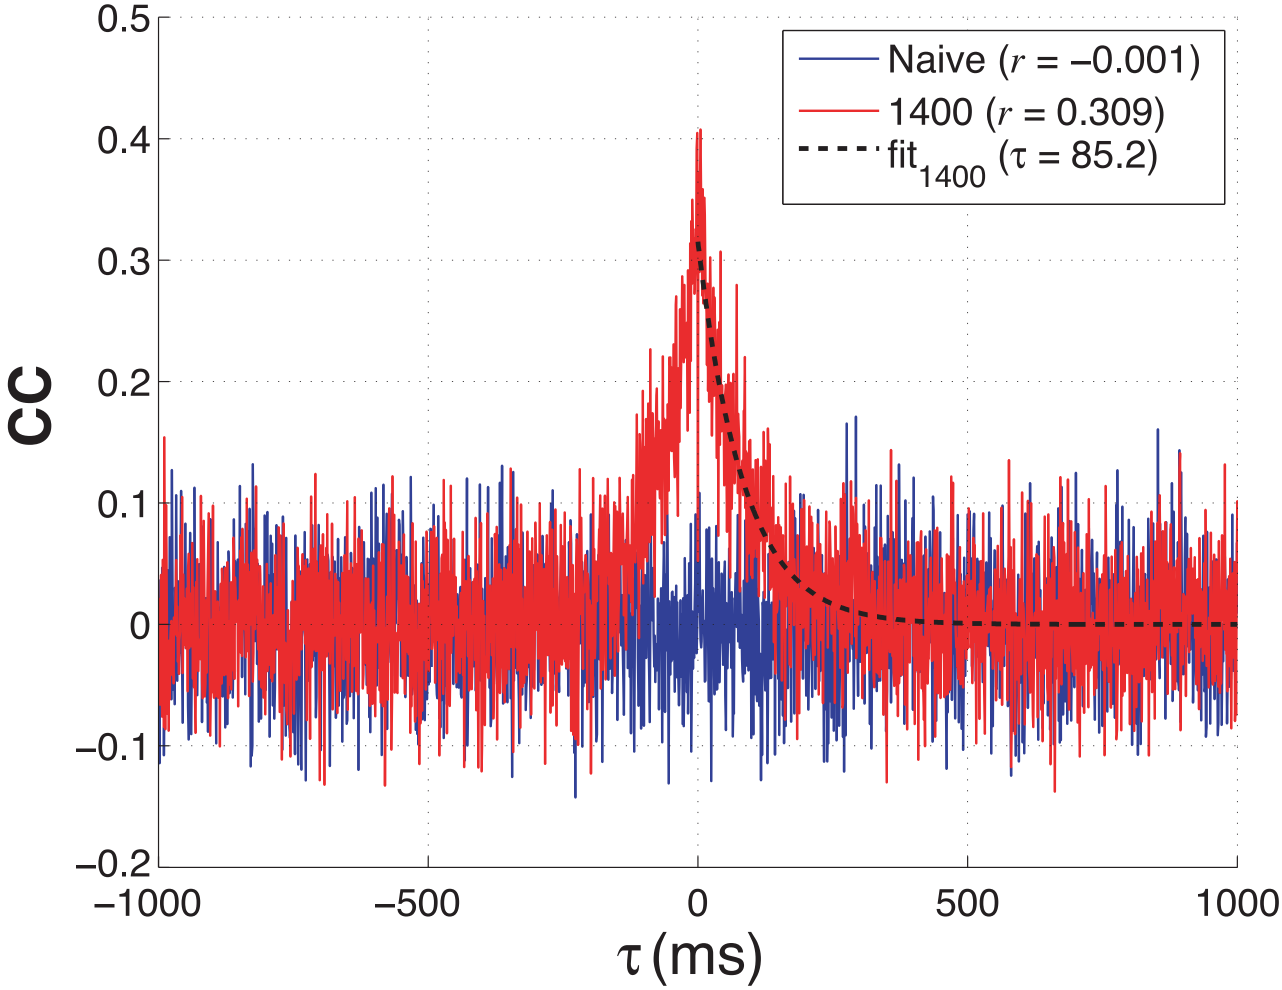

Fig. S7. Noise Correlations. Spontaneously active neurons in a network trained with a reward time of 1,400 ms have higher correlation coefficients ($r = 0.309$) than in a naïve network ($r = -0.001$). Plots of the cross-correlograms show a flat temporal profile in the naïve case (black) and a region of increased correlation between neurons in the trained network (gray). An exponential fit (dashed line) shows the region of increased correlation. Note correlation values are calculated across grouped neural populations (see in *SI Appendix*).

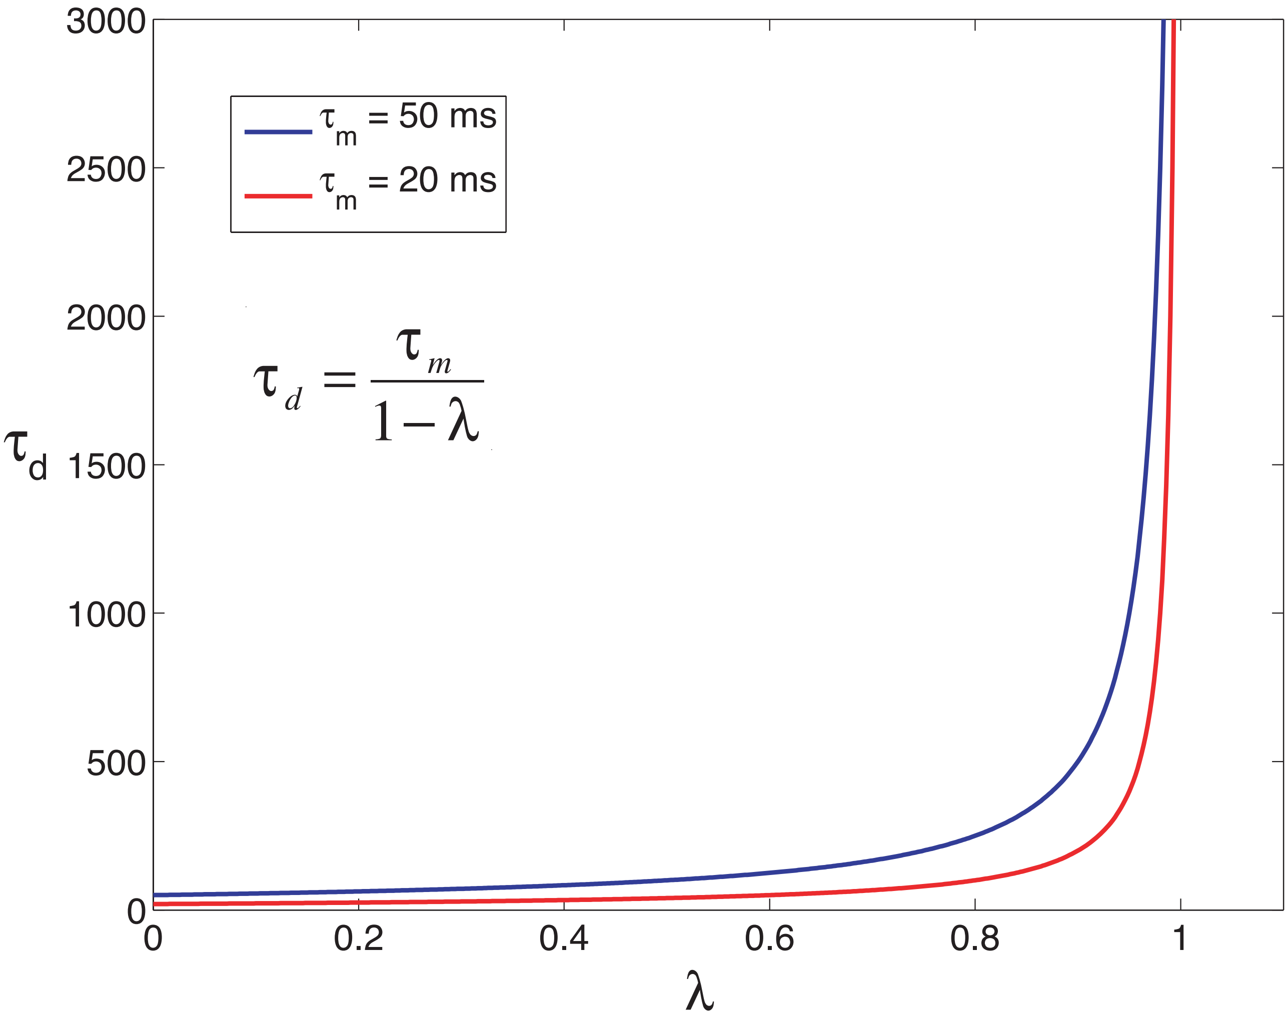

Fig. S8. Dependence of τ_d on λ. The value of λ (Eq. **S4** in *SI Appendix*) required to set the correct weight matrix to encode a particular time (Eq. **S3** in *SI Appendix*) limits the effective interval range that our approach can learn robustly. This plot shows the dependence of the network time constant (in milliseconds) on the eigenvalue of the weight matrix for two different values of the intrinsic neuron decay time constant. The steep rise of the curve in the temporal region of interest for our training task (1–2 s) means that recurrent synaptic weights must be learned to a high degree of precision. This makes robust learning inherently difficult. Note that the relatively faster dynamics associated with the shorter τ_m (red line) result in a steeper curve than with the longer value of τ_m (blue line). An analogous plot can also be generated numerically for the spiking neuron model.

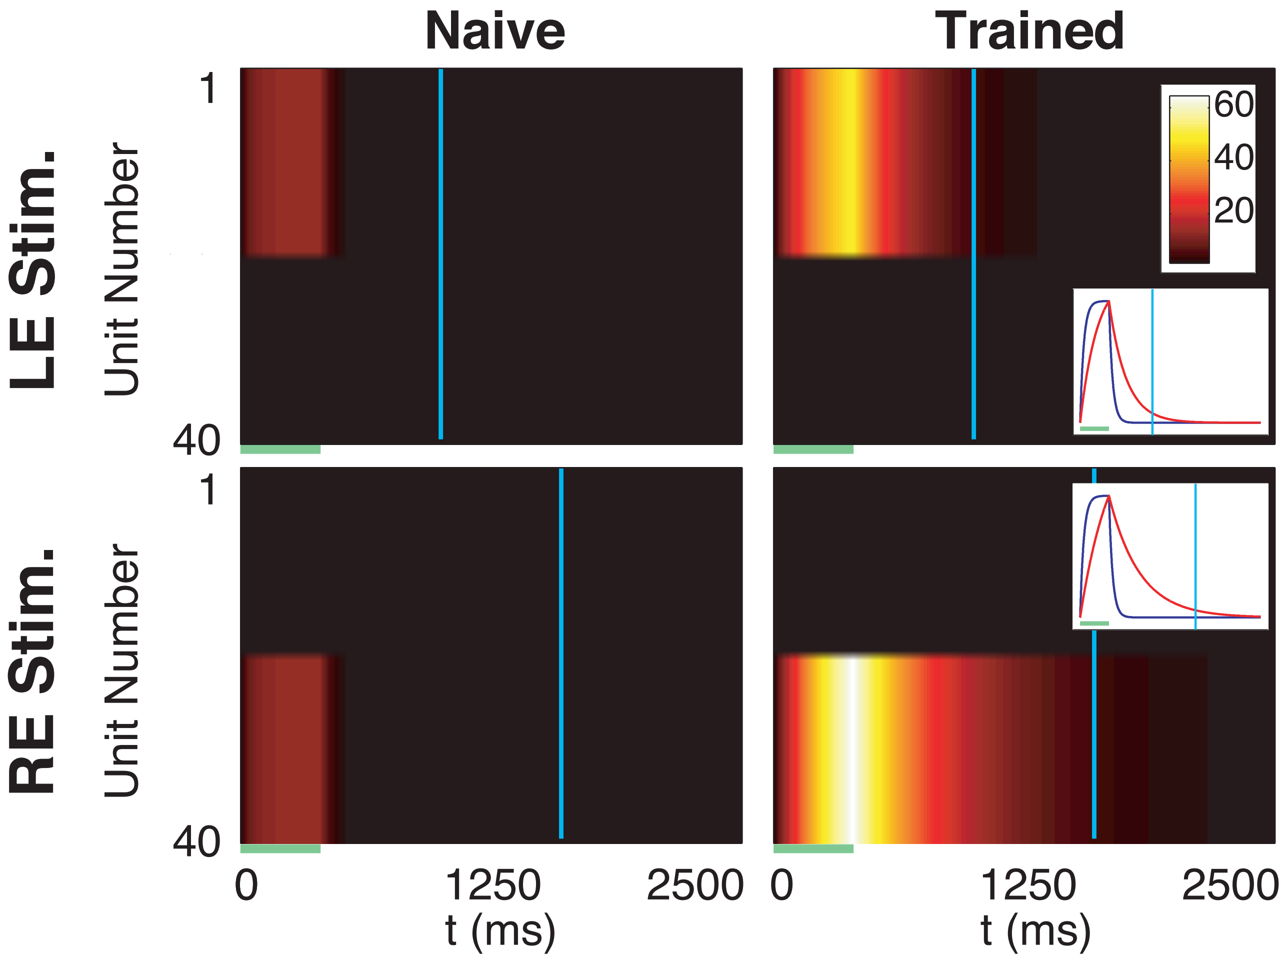

Fig. S9. Training with local inhibition. This plot shows the results of training the rate-based model with local reward inhibition. The results are virtually identical to training using inhibition based on average activity in the entire recurrent layer (see Fig. 3).

Other Supporting Information Files

SI Appendix