Neuron, Volume *91*

Supplemental Information

Bottom-Up and Top-Down Input Augment the Variability of Cortical Neurons

Camille Gómez-Laberge, Alexandra Smolyanskaya, Jonathan J. Nassi, Gabriel Kreiman, and Richard T. Born

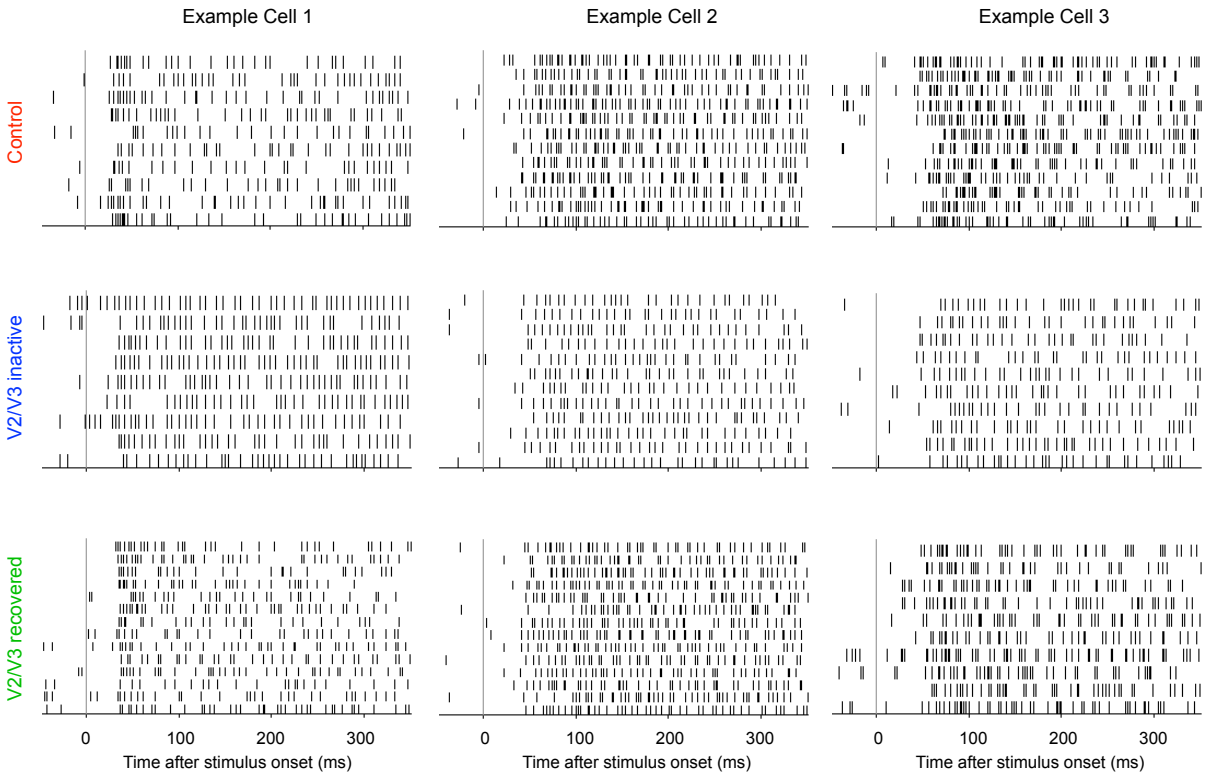

Figure S1, related to Figure 1. Input Inactivation Regularizes the Spike Trains of Cortical Neurons

Spike raster from three example neurons (columns) in visual area MT during a 350 ms random dot stimulus of preferred direction of motion during control (top row), V2/V3 inactivation (middle row), and after V2/V3 recovery (bottom row). Consistent with the analysis presented in Figure 1, input inactivation leads to more regular spike trains due to a disproportionate elimination of the shortest interspike intervals.

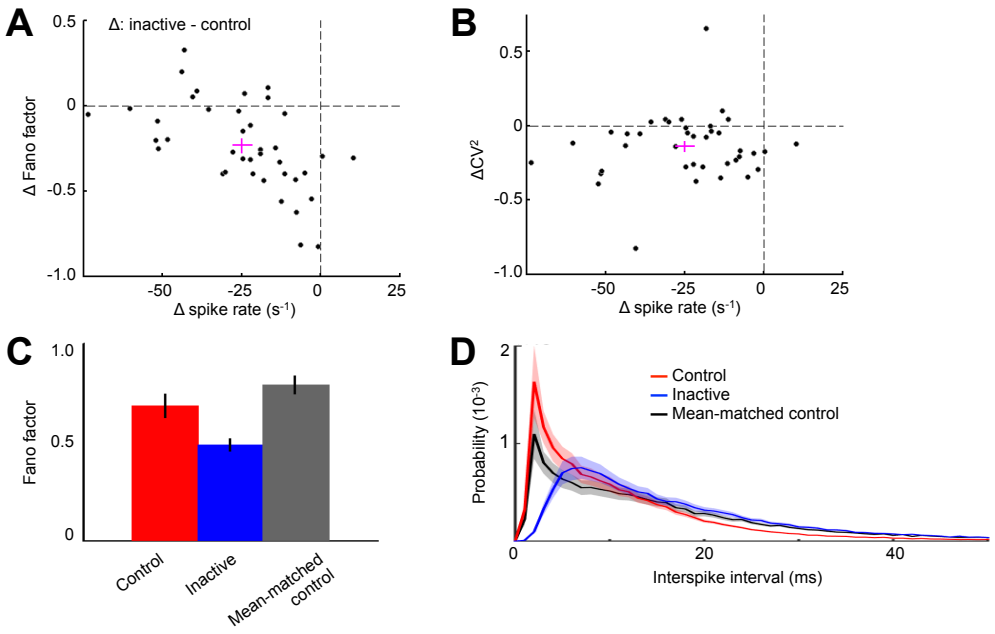

Figure S2, related to Figure 2. Variability Reduction is not an Artifact of Decreasing Spike Rate

(A) Scatter plot of changes in Fano factor versus changes in spike rate per neuron (cross represents mean ± SEM; n = 39 cells). Fano factor and spike rate change in opposite direction ($r = –0.56$, $p < 0.001$). Thus, the largest changes in Fano factor occurred in cells with the smallest changes in spike rate. (B) Same as in (A) except for CV^2. A relationship between CV^2 and spike rate was not detected ($r = 0.12$, $p > 0.05$). (C) Fano factor for control, inactive, and rate-matched control data (by random spike deletion). Artificially reducing the spike rate of the control data did not lead to detectable changes in the Fano factor. Bars represent mean ± SEM. (D) Interspike interval histogram for the rate-matched control data (black) compared to the original control (red) and inactivation (blue) data. The rate-matched histogram retains the exponential shape and, therefore, the irregularity observed in the original control data. Lines and shading represent mean ± SEM.

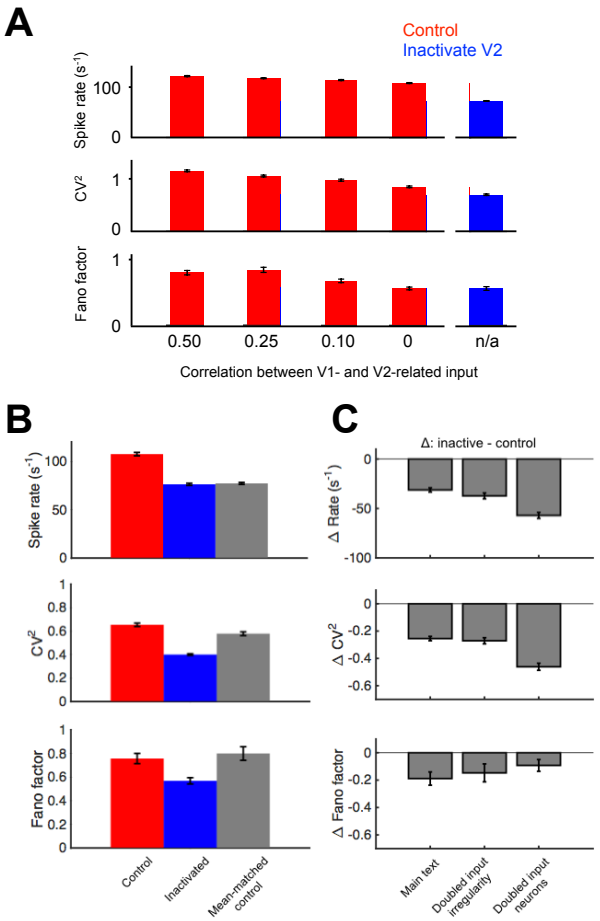

Figure S3, related to Figure 3. Supplemental Simulations of the Within-pool Correlation Model

(A) The experimental effect of input inactivation was reproducible provided that at least some weak correlation existed between V1 and V2. Spike rate, CV^2 and Fano factor for control (red) and V2 inactivated (blue) conditions. Each column measures the effect of inactivation for a different amount of correlation between V1 and V2 inputs. Correlation between V1 and V2 was implemented by copying a fraction of the template spike train (**μ**[t] in Experimental Procedures) that synchronized each group. (B) Rate-matched control simulations. Spike rate, CV^2 and Fano factor for control (red), inactivated (blue) and mean rate-matched control (gray) data. The mean-matched control CV^2 slightly decreased from control levels, but its Fano factor was indistinguishable from control levels. Both measures remained significantly larger than those from the inactivation model. (C) Model robustness. Inactivation-induced changes in spike rate, CV^2 and Fano factor from the simulation in the main text and from additional simulations with either larger input interval irregularity or larger ensembles of input neurons. In all cases, the inactivation-induced effects were detected (t test, $p < 0.01$). Bars represent mean ± SEM.

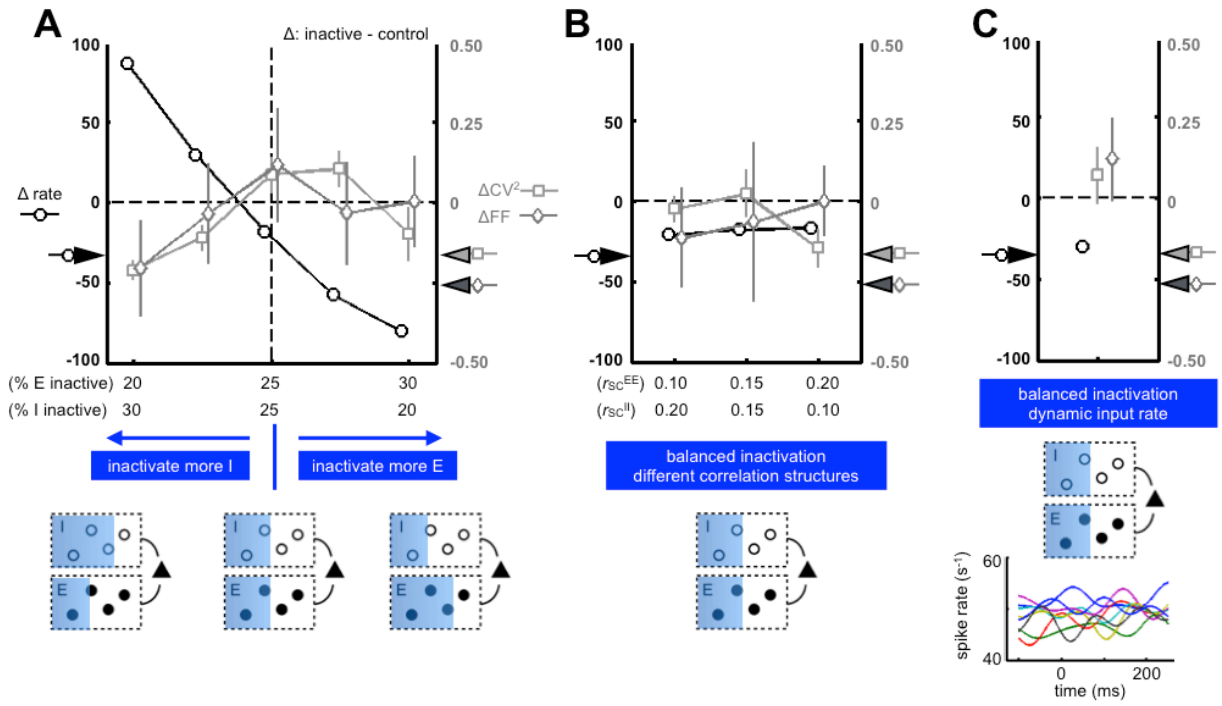

Figure S4, related to Figure 3. Alternatives Lacking Within-pool Correlation did not Reproduce Inactivation Effects

(A) Inactivation-induced changes in spike rate (circles read at left ordinate) and variability (squares/diamonds read at right ordinate) as a function of remaining E–I neuron proportion. A schematic of the single ensemble model and the inactivated neurons is provided at bottom. Horizontal arrows indicate the population mean of the experimental MT data. For most simulations, direct inactivation of different E–I proportions strongly affected spike rate but not variability. The extreme case of disproportionately silencing most of the inhibitory cells (far left) led to a clear reduction in variability, but it was caused by refractoriness due to large increases in spike rate and, therefore, inconsistent with the experiments. (B) Same as balanced inactivation in (A) except inputs were correlated either evenly (center), with stronger inhibitory correlations (left), or with stronger excitatory correlations (right). (C) Same as balanced inactivation in (A) except inputs were independently driven by rapidly fluctuating input rates such as produced by an inhomogeneous Poisson process. These models, lacking within-pool correlation, did not reproduce the experimental data. Error bars represent SEM.

SUPPLEMENTAL EXPERIMENTAL PROCEDURES

Analysis of neuronal response variability. Spiking irregularity (i.e., within-trial variability) was measured by the squared coefficient of variation (CV^2) of the interspike intervals $\mathbf{T}_i$, for i = 1, … , M – 1, between the M spikes observed during one trial

$$\mathrm{CV}^2 = \frac{\mathrm{Var}[\mathbf{T}]}{\mathrm{E}[\mathbf{T}]^2}, \qquad [1]$$

where $\mathrm{Var}[\mathbf{T}]$ indicates the variance, and $\mathrm{E}[\mathbf{T}]$ indicates the expected value.

The average CV^2 over repeated trials of each stimulus condition is reported in the text. The trial-to-trial variability of the spike count **X** was measured by the Fano factor (FF)

$$\mathrm{FF} = \frac{\mathrm{Var}[\mathbf{X}]}{\mathrm{E}[\mathbf{X}]}, \qquad [2]$$

where the expected value and variance were estimated across trials for each stimulus condition. For long observation times, the CV^2 and FF are equal for a class of stationary stochastic point processes called *renewals* (e.g., a Poisson process is a special case of a renewal). For sustained responses, visual cortical neurons are generally well described by this theory and, thus, allow us to interpret changes of within- or across-trial variability accordingly.

Stimulus responses were subject to recommended criteria to ensure the reliability of the CV^2 and FF estimates (Nawrot et al., 2008). All estimates for each stimulus condition were calculated from spike trains of exactly 250 ms duration: 50–300 ms before and after stimulus onset for fixation and response data, respectively (see horizontal gray bar

in Figure 1D). A neuron's response to a given stimulus condition was included only if a minimum of 5 spikes per response and a minimum of 8 repeated trials were obtained. This led to the inclusion of 39 single MT neurons (out of 43 held for entire recovery block), each contributing responses to an average 11 repeated trials of 13 stimulus conditions per inactivation block (4,392 total trials per block). The same criteria held for the V1 data, except that its larger stimulus set only permitted 5 repeated trials per condition. Consequently, we admitted these data acknowledging that a larger FF bias toward unity could conceal inactivation effects. In V1, all 36 single neurons encountered were included, each contributing responses to exactly 5 repeated trials of an average 17 stimulus conditions per inactivation block (3,060 total trials per block). Population mean and standard error were estimated across neuron-conditions.

The dynamics of trial-to-trial variability were estimated for MT (Figure 1D) and V1 (Figure 4C) to examine their relationship to evoked changes in spike rate. These estimates were obtained using a centered, 70 ms sliding window that advanced in 17.5 ms steps. The FF is robust to low spike counts; however, the obtained quantities in these curves will be more biased toward unity than those calculated from the static 250 ms window due to the smaller counting window (Nawrot et al., 2008). They nevertheless convey complementary temporal data on stimulus- and inactivation-evoked effects with respect to changes in spike rate response.

Interspike interval histograms were obtained for each neuron-condition by pooling all observed intervals across repeats (Figure 1E). The population mean and SEM of the histograms were taken over neuron-conditions. The mean histogram (discounting a 2

ms refractory period) for each inactivation condition constrained a gamma probability density function f for an interval t > 0, given shape k > 0 and scale θ > 0 parameters

$$f(t; k, \theta) = \frac{t^{k-1} e^{-\frac{t}{\theta}}}{\theta^k \Gamma(k)}, \quad [3]$$

where Γ(k) is the gamma function. Shape parameters greater than 1 indicate intervals that are more regular than those obtained by a Poisson process. The mean ± SEM for both parameters across inactivation blocks are reported (Figure 1E, inset).

Bottom-up input inactivation simulations. The output spike rate and variability statistics in Figure 3 were calculated from input ensembles consisting of 50 excitatory and 50 inhibitory neurons, each firing at a constant spike rate of $\lambda = 100$ spikes s^{-1}. Their spike trains were realized using the gamma distributed intervals of equation [3] (Experimental Procedures) measured in seconds with parameters $k = 2$ and scale $\theta = (\lambda k)^{-1} = 0.005$. Inactivation set λ to zero for half of the excitatory and inhibitory neurons.

The “independent model” (i.e., with uncorrelated input) realized i.i.d. spike trains for each neuron. This model led to unsatisfactory results because the input spike rates and high irregularity were substantially attenuated in the output before inactivation. This is because observing N excitatory inputs is equal in probability to observing N inhibitory inputs, independent of time scale. Therefore, an action potential is only generated when random fluctuations permit the unlikely accrual of excitation, which leads to (i) rate attenuation because the majority of excitatory synapses are annihilated by inhibitory ones, and (ii) interval regularization because accumulating N randomly unchecked

excitatory events to overcome the sub-threshold deficit effectively constructs a gamma process with shape parameter $k \propto N$ (with unit proportionality for perfect accumulation).

The "uniform correlation" model correlated all excitatory and inhibitory neurons to one template $\boldsymbol{\mu}$ (see Experimental Procedures), such that excitation and inhibition fluctuated in unison. This model led to similar results as the independent model, since the majority of excitatory synapses were canceled in a timely manner by a comparable number of inhibitory synapses. Only the random fluctuations in total spike count across neurons contributed to suprathreshold depolarization.

The "within-pool correlation" model synchronized all excitatory neurons to one template $\boldsymbol{\mu}_E$ and all inhibitory neurons to another i.i.d. template $\boldsymbol{\mu}_I$. This case led to more realistic results because the output spike rate was preserved as was the interval irregularity. Only in this case did inactivation reduce the output rate, CV and Fano factor, as measured in the MT data. We found that the results also held when the excitatory ensemble was divided into two partially correlated parts (e.g., V1 and V2). However, the inactivation effects shrank as the ensembles become more independent. In the limiting case, independent components no longer reproduced the inactivation-induced decrease in Fano factor (Figure S3A).

Control analyses for simulated spike rate changes. As for the experimental data, we performed supplementary simulations to ensure that the inactivation-induced variability reductions were not artifacts of rate change. This was done by running a control model with λ matched to the output spike rate in the inactivation model of the main text (where $\lambda = 100\ s^{-1}$). Figure S3B shows that the mean-matched control data remained

significantly more variable than the inactivation data. Furthermore, differences between control and mean-matched Fano factor were not detected, although CV did decrease slightly from control levels.

Top-down input inactivation simulations. These simulations were set up like the within-pool correlation model described above, except an imbalanced number of excitatory and inhibitory neurons were silenced. The total number of inactivated neurons always added up to 50% of the control model population. Only the relative change in E–I balance seems to modulate the inactivation-induced effect on rate. For example, we obtained similar results when forcing the LGN input source to provide balanced input and the V2 input source to provide I > E input (therefore starting off with an imbalanced control model). In this case, when V2 inactivation silenced half of the total population, a balanced input ensemble remained.

Model robustness. To evaluate the robustness of our input inactivation results, we performed additional simulations that individually varied each input ensemble variable: spike rate λ, interval irregularity k, ensemble synchrony b, ensemble size $N_E + N_I$, and the E-I imbalance. In the main text, we showed that inactivation effects were sensitive to ensemble synchrony b (Figure 3E) and that E–I imbalance mainly affected inactivation rate (Figure 4G). The above supplementary results on spike rate artifacts showed that our model was robust to variations of spike rate λ (Figure S3B). Consequently, we ran our simulations after doubling the input interspike interval irregularity $k \rightarrow k/2$ and once more after doubling the number of input neurons. In both cases, the inactivation results were qualitatively consistent with those of the model in the main text (Figure S3C).

Alternative models. We considered alternative models that did not employ the correlation structure of the within-pool model described above and in the main text. These models all started off with the uniform correlation model where uniform correlations exist between all input neurons, and (i) an imbalance in the number of E and I neurons takes place (Figure S4A), (ii) an imbalance in the strength of the correlations within the E and I pools takes place (Figure S4B), or (iii) a rapidly fluctuating input rate parameter drives each input neuron (Figure S4C). None of these alternative explanations could reproduce the experimental data, further supporting our conclusion that the crucial feature of the model in the main text is the within-pool correlation of the E and I input.

Effect of spike-count correlation in the model. The five-step algorithm for generating synchronized ensembles in Experimental Procedures section allowed us to vary spike count correlation (r_{SC}) for any given amount of synchrony dictated by the asynchrony parameter b. This algorithm led to $r_{SC} \approx 0$, for the tested asynchrony parameter range $0.1 \leq b \leq 1000$. In step 4 of the algorithm, by replacing the underlined the phrases "the same number of spikes as **x**" with "a spike count within N spikes of **x**," one can increase r_{SC} by increasing N. We found that varying $r_{SC} \geq 0$ did not alter the effect of temporal correlations on input inactivation.