

SCIENTIFIC REPORTS



OPEN

Spontaneous activity emerging from an inferred network model captures complex spatio-temporal dynamics of spike data

Cristiano Capone^{1,2}, Guido Gigante^{1,2}  & Paolo Del Giudice³

Inference methods are widely used to recover effective models from observed data. However, few studies attempted to investigate the dynamics of inferred models in neuroscience, and none, to our knowledge, at the network level. We introduce a principled modification of a widely used generalized linear model (GLM), and learn its structural and dynamic parameters from *in-vitro* spike data. The spontaneous activity of the new model captures prominent features of the non-stationary and non-linear dynamics displayed by the biological network, where the reference GLM largely fails, and also reflects fine-grained spatio-temporal dynamical features. Two ingredients were key for success. The first is a saturating transfer function: beyond its biological plausibility, it limits the neuron's information transfer, improving robustness against endogenous and external noise. The second is a super-Poisson spikes generative mechanism; it accounts for the undersampling of the network, and allows the model neuron to flexibly incorporate the observed activity fluctuations.

Dynamic models, in neuroscience and in general, embody in a mathematical form the causal relationships between variables deemed essential to describe the system of interest; of course, the value of the model is measured both by its ability to match observations, and by its predictive power. Typically, along this route, parameters appearing in the model are assigned through a mix of insight from experiments and trial-and-error and, in turn, the study of the model dynamics helps understanding the relevance of each parameter in determining the dynamic regimes accessible to the system. Another approach, initially quite detached from dynamic modeling, is rooted in the domain of statistical inference and is focused on data-driven model building; a seminal example in neuroscience was offered by the application of maximum-entropy inference of Ising-like models to multi-electrode recordings of neural activity^{1–4}; in this case, as well as in the later extension to kinetic Ising-like models, the Montecarlo/Glauber dynamics of the model is only meant as a means to sample the probability distribution of interest and it is not claimed to offer a detailed model of the actual dynamics at work in the system^{5–7}. Recently, Generalized Linear Models (GLM) (which incorporate kinetic Ising models as a special case) have been recognized as flexible and powerful inference models^{8,9}. In time, efforts have been made to make contact between the two approaches, e.g., in the case of neuroscience, by endowing the coupling structure of the inference model with features motivated by biological plausibility^{10–12}; it has also been recognized that a GLM is close to a stochastic version of the spike-response model. Recently, the repertoire of the driven dynamics of GLM models of single neurons has been explored¹³; however, to our knowledge, a largely open issue is to endow network inference models with predictive power in terms of the system dynamics. Our approach to this problem is to explore the free, spontaneous dynamics of the inferred model in its relation with the one of the biological system generating the data and, in doing this, to identify the role of different elements of the inference model in determining the spontaneous dynamics of the neuronal network. Indeed, an obvious but persisting problem in the application of inference models to neuroscience has been how to assess the meaning and value of the inferred parameters; a recurring example is offered by the inferred synaptic couplings^{5,7,10}: in the literature, a cautionary remark is usually included, recognizing that the inferred couplings (whether in the form of synaptic efficacies of more complicated synaptic kernels) are to be meant as ‘effective’, leaving of course open the problem of what exactly this

¹Physics department, “Sapienza” University, Rome, Italy. ²INFN, Sezione di Roma, Rome, Italy. ³Istituto Superiore di Sanità, Rome, Italy. Cristiano Capone and Guido Gigante contributed equally. Correspondence and requests for materials should be addressed to P.D.G. (email: paolo.delgiudice@roma1.infn.it)

means (especially in the face of the dramatic subsampling of the underlying biological network). The value of the inferred model cannot be assessed directly by a detailed correspondence between its elements and corresponding elements of the biological system; using the output of a GLM to decode input stimuli is an interesting recent approach¹⁴. Our attitude is that spontaneous activity is a good testing ground to assess whether the inferred model does indeed capture essential features of the biological system, and we choose a case in which the spontaneous activity of the biological network is highly non-stationary and irregular: a cultured neuronal network generating spontaneously a wide spectrum of activity fluctuations, from population bursts, to neuronal avalanches and noisy oscillations^{15–22} (and references therein). Such richness is hard to reproduce with a GLM, and to our knowledge no models proposed so far were able to cope with it. For this reason we chose population bursting activity as a challenging benchmark to test our model. In our GLM approach we introduced novel and quite simple ingredients inspired by biological observation, such as activity dependent negative feedback over different time-scales for the single neuron¹⁵ (spike frequency adaptation, SFA), a bounded transfer function, and a generative probabilistic mechanism for spike generation with super-Poisson fluctuations. We show that these ingredients are crucial to endow the system with non-stationary spontaneous activity and also to account for detailed dynamic features such as temporal correlations and spatio-temporal evolution of network bursts.

Results

Bounded firing rate, super-Poisson spike generation, and their impact on the spontaneous activity of the inferred model. The model we introduce is an extension of GLMs as widely employed for inferring the structure of neuronal networks. The main points of departure from more standard GLMs lie in the choice of the non-linear transfer function and the probabilistic generative model; besides we introduce, at the single neuron level, an activity-dependent self-inhibition current, mimicking spike-frequency adaptation (SFA) effects widely observed in real neurons (see Methods for details).

Figure 1, panel A, gives a schematic view of the model. The input current $H_i(t)$ to neuron i is the sum of synaptic currents (mediated by kernels $k_{ij}(\Delta)$), constant external current h_i , and five activity-dependent SFA currents $c_i^{(x)}(t)$, $x = 1 \dots 5$, with different time-scales. $H_i(t)$ determines, through a non-linear transfer-function $f[H_i]$ the expected firing rate $\lambda_i(t)$ for the spike count at the next time-bin $t + dt$ ($dt = 10$ ms). The transfer function and generative model most commonly employed in the literature are, respectively, the exponential function and the Poisson distribution; in the following we will label this choice ‘Exp-Poisson’; our model, on the other hand, makes use of a sigmoid function and of a Negative-Binomial distribution, and will be referred to in the following as ‘Sig-NegBin’ model.

Parameters of synaptic kernels, transfer function, SFA currents, and external currents are fitted maximizing the likelihood using the adaptive iRprop algorithm²³. Both models can then be simulated to generate new spike trains according to two basic simulation modalities.

The first one (‘driven’ mode) feeds the model with the actual spike trains observed in the data and at each time t uses the model’s output to make a 1-step prediction of the network activity at step $t + dt$; thus the generated activity never re-enters the network dynamics; the probabilistic distance between the 1-step prediction and the actual recorded network activity is ultimately what is minimized during the optimization procedure. In the second modality (‘free’ mode) the models are left free to evolve and the generated spike trains are fed back into the network, driving its dynamics at subsequent times.

We trained both the Sig-NegBin and the Exp-Poisson model on data from multi-electrode recordings of *ex-vivo* cultured cortical neurons¹⁷, that show a clear non-stationary dynamics (Fig. 1B, top row) in the rastergrams and in the whole network activity (red line), with and irregular bursts of synchronous activity interleaved by periods of low, noisy activity (‘inter-burst intervals’ or IBIs). At the end of the training, both models show, in driven mode, an activity that follows quite closely the experimental spike trains (Fig. 1B, middle and bottom rows; red line: data; blue line and rastergrams: simulation). However, when the simulation is switched to free mode at time $t = 30$ s, the behavior of the two models clearly diverges: whilst the Sig-NegBin is able to autonomously sustain a highly non-stationary activity, the Exp-Poisson only exhibits stationary fluctuations. Note that, in free mode, the activity of the Sig-NegBin model does not closely follow the activity of the real network anymore (top row, last 10 s); this is not surprising: the choice of a probabilistic model as a GLM stems from the assumption that the network dynamics is, to a large extent, influenced by self-induced random fluctuations and, thus, non-deterministic; it is therefore expected that, at least after a transient phase, the model and the data will drift apart. Yet the activity spontaneously generated by the model clearly resembles the one seen in the data in a statistical way, showing irregular bursts interspersed with interval of relative quiescence.

Such resemblance is made more quantitative in Fig. 1C. Bursts of network global activity are detected through the algorithm developed in¹⁵, both for the data (top row) and the autonomous activity of the Sig-NegBin model (bottom row). The comparison between burst amplitudes (total spike counts, left), bursts durations (center) and IBI durations (right) histograms shows a semi-quantitative agreement between the dynamics of the biological network and the inferred model (mean and standard deviation of each distribution are reported with red and green lines respectively; mean \pm standard deviation, data vs simulation: burst amplitude 500 ± 440 vs 420 ± 300 ; burst duration (0.35 ± 0.24) s vs (0.28 ± 0.19) s; IBI (3.4 ± 2.9) s vs (2.7 ± 2.4) s). The main discrepancy between model and data is visible in the IBI distribution; ‘doublets’ of bursts in the data, clearly separated by the rest of the IBI distribution, are absent in the model. Based on previous modeling work on similar data²⁴, the discrepancy is likely due to elements such as short-term synaptic facilitation and depression, that we did not include in our minimal model.

We remark that the analogous distributions for the Exp-Poisson model are not reported because of the extreme sparsity of burst events in that case. Moreover we observe that, despite the major difference in the exhibited dynamic behavior, the inferred synaptic kernels for our model and the Exp-Poisson model are correlated ($c_{\text{Pearson}} = 0.49$ with a $P < 0.05$; correlation between the total areas under each synaptic kernel, as a measure of

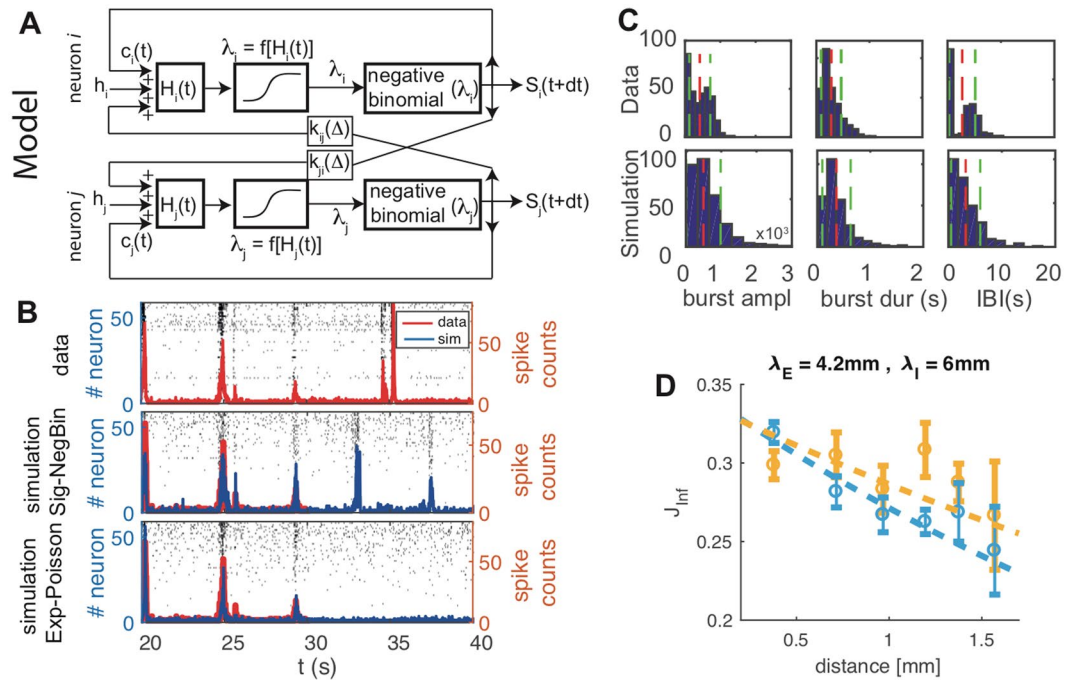


Figure 1. Spontaneous activity of the inferred model reproduces non-stationary network dynamics. **(A)** Schematic representation of the model. The current $H_i(t)$ felt by neuron i sums spikes emitted by other neurons j in the recent past ($t - \Delta$) through synaptic kernels $k_{ij}(\Delta)$, a constant external current h_i , and spike-frequency adaptations signals that integrate, over different time-scales, the spiking activity S_i of neuron i itself. This current, fed into a non-linear transfer-function $f[H_i]$, determines the expected number $\lambda_i(t)$ of spikes $S_i(t + dt)$ that the neuron will probabilistically generate at the next time bin. **(B)** Rastergrams of *ex-vivo* cultured cortical neurons show a strongly non-stationary spontaneous dynamics (top row; red line: whole network activity). The proposed Sig-NegBin model (central row, blue line) closely follows the data used to drive its dynamics (red line, t between 20 and 30 s); when left free to run autonomously at $t = 30$ s, the model displays large irregular bursts of activity, qualitatively mimicking the behavior observed in the recordings. The standard Exp-Poisson model (bottom row, blue line), on the other hand, whilst still able to follow the data in driven mode, fails to produce any large activity fluctuation, settling instead into a noisy, stationary state. **(C)** Comparison between burst amplitudes (total spike counts, left), bursts durations (center) and inter-burst-intervals (IBI) durations (right) histograms both for the data (top row) and the autonomous activity of the Sig-NegBin model (bottom row). Mean and standard deviation of each distribution are reported with red and green lines respectively. Results for the standard Exp-Poisson model are not shown because of the extreme sparsity of burst events in that case. **(D)** Average area under the inferred synaptic kernels (absolute value) as a function of the distance separating the electrodes; cyan and yellow circles are for excitatory and inhibitory kernels respectively; dashed lines: exponential fits; the strength (area under the kernel) of inferred excitatory synapses decays faster than that of the inhibitory ones.

‘synaptic strength’). The inferred synaptic kernels also show a correlation with the distance separating the electrodes where the activity of the pre- and post-synaptic neurons were recorded (Fig. 1D; cyan and yellow circles are for excitatory and inhibitory kernels respectively; dashed lines: exponential fits). It is worth noting that the connections’ decay with distance is faster for excitatory than for inhibitory inferred synapses; this is in partial agreement with findings in²⁵. We also remark that the inferred network does not comply with the ‘Dale’s law’: the latter was not enforced as a constraint, and each neuron forms in general both excitatory and inhibitory synapses.

The choice of a sigmoidal, saturating transfer function, besides seeming just natural in a GLM meant to model neural activities, turns out to be key to achieve the above results. We repeated the inference procedure with a hybrid Sig-NegBin model, in which the sigmoidal transfer function is replaced with an exponential and the spikes are still generated according to a negative binomial distribution (Fig. 2A). As above, when during the first 10 s the model is simulated (blue line and rastergrams) in driven mode, it shows a good agreement with the driving data (red line); when the driving is turned off, instead, the model settles in a quiescent state, absent any large bursts of activity. In free mode, the saturation of the sigmoidal function $f[\cdot]$ (Fig. 2B, top-right panel) allows a wide distribution of input currents H_i to the single neuron (Fig. 2B, bottom panel) to give rise to a distribution of firing rates $\lambda_i = f[H_i]$ (Fig. 2B, top-left panel) that is strongly bimodal (notice the logarithmic scale). On the other hand, the fast growth of the exponential function (Fig. 2C, top-right panel) pushes the inference procedure towards somewhat narrower, more conservative distribution of currents (Fig. 2C, bottom panel), resulting in a unimodal distribution of firing rates. The comparison of the two models in driven and free modes suggests that the exponential transfer function makes the system very selectively susceptible to the fluctuations experienced during learning,

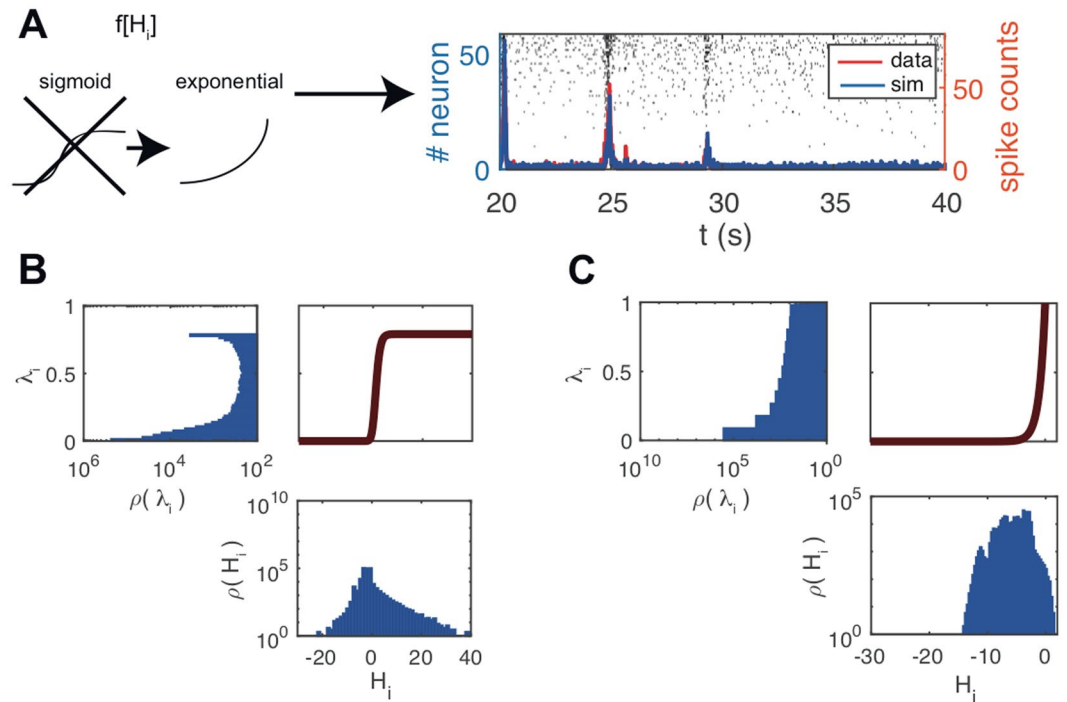


Figure 2. Bounded spike rate is necessary to generate non-stationary spontaneous activity in the inferred model. **(A)** Population activity (blue line) and rastergram generated by the model when the bounded transfer function is replaced by an exponential and inference procedure is carried out from scratch. Red line represents the data driving the dynamics for the first 10 s ($20 \leq t \leq 30$ s); afterwards, when left free to generate its spontaneous activity, the model is unable to sustain any non-stationary bursting activity and settles in a quiescent state. **(B,C)** Transfer function $f[\cdot]$ (respectively sigmoid and exponential) (top-right panel); distribution of input currents to the single neuron $\rho(H_i)$ (bottom panel); distribution of firing rates $\rho(\lambda_i)$, $\lambda_i = f[H_i]$ (top-left panel; note that the axes are flipped, with the vertical scale being shared with the transfer function plot). The sigmoid transfer function allows for a bimodal distribution of firing rates.

while the sigmoid transfer function allows for susceptibility to a wider range of fluctuations, which we believe is the root of the ability to spontaneously generate network bursts.

A negative binomial spike generative model naturally accounts for the effects of subsampling. The recorded neurons from the cultured network constitute of course a dramatic subsampling of the network (Fig. 3A, leftmost sketch), though, in fact, the activity of several neurons will be collected by each electrode. The effects of the unobserved neurons on the behavior of the observed part of the network could be in principle very complex²⁶, yet we argue that a Negative Binomial distribution for the spike counts generation, instead of a Poisson distribution as commonly adopted, captures a relevant part of those effects and is in fact critical to account for the non-stationarity of the spontaneous activity of the network.

One way to approach the problem is to assume that at each time t the generated spike count of a generic neuron $S_i(t)$ is determined by both the current $H_i(t)$, accounting for the recurrent activity of other observed neurons and external currents (see Methods), and the currents due to unobserved, ‘hidden’ neurons, h_i^{hid} (Fig. 3A), such that the observed spike count probability distribution would be obtained by integrating over the (unknown) distribution of h_i^{hid} :

$$p(S_i|H_i) = \int dh_i^{\text{hid}} p(S_i|H_i, h_i^{\text{hid}}) p(h_i^{\text{hid}}). \quad (1)$$

We will still assume that $p(S_i(t)|H_i(t), h_i^{\text{hid}})$ is a Poisson distribution, with mean given by $\lambda_i(t) = f[H_i(t) + h_i^{\text{hid}}]$.

It is intuitive that $p(S_i(t)|H_i(t))$ will now generate super-Poissonian fluctuations; and it turns out that larger fluctuations are critical for the ability of the Sig-NegBin model to spontaneously generate bursting activity, as shown above (Fig. 3B upper panel; red line: data; blue line: Sig-NegBin model in free mode). In fact, an inferred Sig-Poisson model shows very sparse bursting (Fig. 3B, middle panel). But it is not just that a broader spike generation distribution favors larger fluctuations in the network activity. We simulated the inferred Sig-NegBin model replacing the Negative Binomial with a Poisson spike generator (Fig. 3B, bottom panel); such ‘Poisson replay’ of the inferred Sig-NegBin model generates frequent bursts, each lasting for long time on average. Thus the inferred synaptic couplings of the Sig-NegBin model, unlike the ones for the Sig-Poisson, appear to be compatible with a bistable network dynamics, between UP and DOWN states, as also suggested by the results for the sigmoid transfer function (Fig. 2B). Therefore, the main role of a super-Poisson spike generation distribution seems to be to

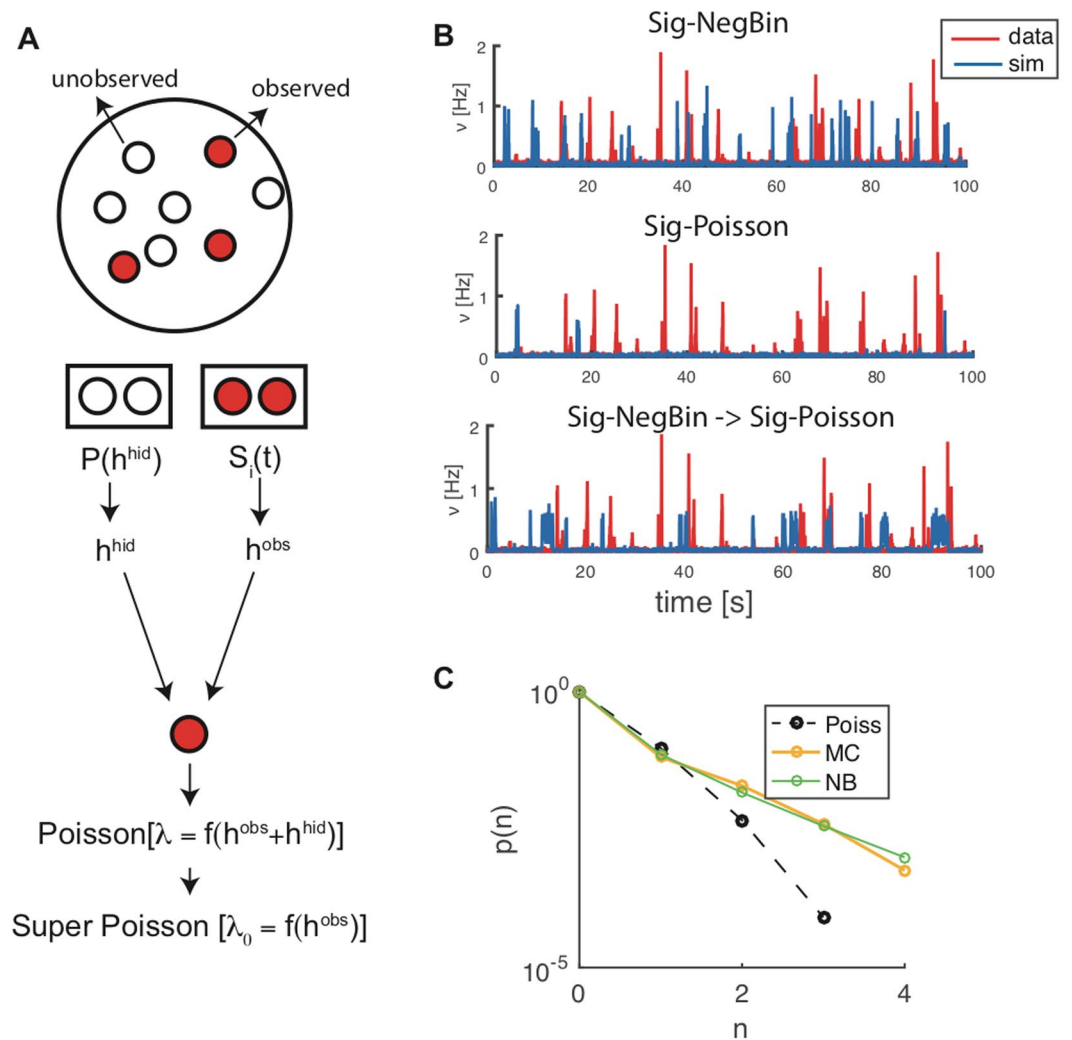


Figure 3. Rationale and implications of choosing a negative binomial as generative model of spike counts. **(A)** Sketch showing the effects of the unobserved neurons on the fluctuations in the observed part of the network. **(B)** Spontaneous bursting activity generated by different models (red line: data; blue line: model in free mode). Upper panel: Sig-NegBin model; middle panel: Sig-Poisson; bottom panel: Sig-NegBin model after replacing the Negative Binomial with a Poisson spike generator. **(C)** Comparison between the spike counts generated by a negative binomial distribution ($r_{\text{NB}} = 0.22$) with a mean value equal to the average spike count of the neuron (green line), the Monte Carlo sampled $p(S_i|H_i)$ with best matching values for the first two moments of spike counts (orange line) and Poisson distribution for the same average spike count (black dashed line), reported as a reference.

destabilize the UP state, making the inferred model once again more robust to the variability of the free network's activity and thus able to self-sustain a highly non-stationary activity.

Taking together the results in Figs 2 and 3, the suggested picture is that on the one hand the sigmoid gain function allows the inference procedure to explore and use ranges of couplings that are effectively forbidden for an exponential gain function, and are essential to spontaneously generate large sudden increase of activity (burst onset); on the other hand, the NegBin generator plays the dynamic role of efficiently destabilizing, with large fluctuations of the spike counts, the high-activity states which are spontaneously generated by the large recurrent excitation. The two ingredients are robustly coupled to spontaneously produce bursting activity.

Albeit the NegBin distribution represents just one of the forms that $p(S_i(t)|H_i(t))$ can take, it has been suggested to adequately model fluctuations in observed neural activity^{27,28}, which makes it a natural candidate. Yet, a NegBin would result from Eq. 1 assuming an exponential transfer function and a log-Gamma distribution for h_i^{hid} : the first condition is patently in contradiction with our choice of a saturating transfer function and the second one is not expected to hold in general. We nevertheless tested the adequacy of the NegBin assumption for the Sig-NegBin model as follows. We assumed that $p(h_i^{\text{hid}})$ in Eq. 1 is a scaled and shifted version of the distribution of the observed currents H ; in turn we estimated such distribution from the inferred synaptic couplings and the observed experimental spike counts, for the neuron with maximal average activity, for which we expected the differences in the spike count distribution would both show up more clearly, and matter more for the dynamics. From this, we performed a Monte Carlo estimate of $p(S_i|H_i)$ for different values of the mean and the standard

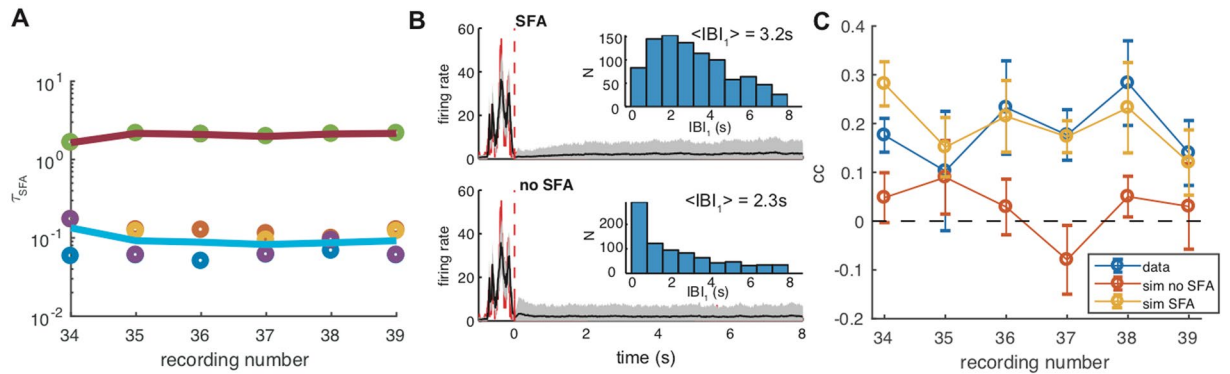


Figure 4. SFA timescales, and effect of SFA currents on bursts statistics. **(A)** Inferred τ_{SFA} for different recordings from the same preparation; the five inferred time-scales cluster around two values for all recordings, one about 100 ms, the other around 2 s. **(B)** Comparison between the spontaneous model dynamics with and without SFA. In black, and gray-shading, are represented the population firing rate and its variance over 500 simulations. Red line is the population activity from data during a chosen large burst. Insets: histograms of the time to the subsequent burst. SFA induces a short-term refractoriness for burst generation. **(C)** Correlation between the burst amplitude and the time since the preceding burst. SFA determines positive correlations in the model, as those seen in the data; notice that models are separately inferred for the different recordings, and it is seen that inferred models closely mirror the fine structure of experimental correlations.

deviation defining $p(h_i^{hid})$. Figure 3C shows the distribution of spike counts for a NegBin distribution that is very close to the one used in the Sig-NegBin model ($r_{NB} = 0.22$), for a mean value equal to the average spike counts of the chosen neuron (green line); the sampled $p(S_i|H_i)$ with best matching values for the first two moments is reported in orange; the two distributions are very similar, (in black we report for reference the Poisson distribution for the same average spike count, to provide a scale for the discrepancies between the other two distributions). Thus spike counts compatible with a NegBin distribution are naturally accounted for by a plausible assumption on the effect of sub-sampling underlying Eq. 1, with marked deviations from a Poisson distribution. We repeated the above procedure sampling h_i^{hid} from a Gaussian distribution, finding similar results; this suggests that, provided suitable values for the mean and variance are chosen, the precise shape of the $p(h_i^{hid})$ has a minor effect. As a further check, we sampled in driven mode the sequence of $\lambda_i(t)$ generated by both the Exp-Poisson and the Sig-NegBin models for all times and all neurons i ; we then estimated the unconditioned distribution $p(s)$ for the spike count s for each model by convolving the sampled distribution $p(\lambda)$ with $p_{Poisson}(s|\lambda)$ and $p_{NB}(s|\lambda)$ respectively ($\int p(\lambda)p(s|\lambda)d\lambda$). Both models predict a higher than observed probability for larger spike counts at the expense of probabilities for lower values; at our $dt = 10$ ms, 4 spikes correspond to a spike frequency of 250 Hz, a value that probably engages refractoriness mechanisms in real neurons; such mechanisms are not accounted for by neither models, since the Poisson and the NegBin distributions do not force a hard maximum value for the spike count. The likelihood ratio on the data, though, clearly favors the Sig-NegBin over the Exp-Poisson as the model best accounting for the observed spike counts; having the Sig-NegBin slightly less parameters (see Methods), this result would be made even stronger applying corrections for different model complexities.

The inferred model captures detailed temporal and spatial aspects of the neural dynamics.

Figure 4A illustrates the inferred time-scales for the SFA component. Although we allowed for five independent time-scales to be inferred (but constrained to be equal for all neurons), it is interesting to note that the inferred values consistently aggregate around just two values, at about 100 ms and 2 s respectively, for different temporal segments of the same network's activity (Fig. 4A). This result is consistent with previous work on the same data aiming to infer the main time-scales of the bursting dynamics using a completely different approach¹⁵.

To assess how relevant the adaptation mechanism is for the model dynamics, we performed inference with and without SFA. Then we simulated the model in driven mode until just after the end of a burst of large amplitude, where adaptation effects are expected to be more pronounced, leaving the model free to spontaneously evolve afterwards. In Fig. 4B the SFA and non-SFA dynamics are reported (top and bottom, respectively); black line is the model activity averaged over 500 simulations of the dynamics, while gray shading marks the plus/minus one standard deviation range; red line is the population activity from data. For each realization we collected the next IBI, that is the time interval to the next burst spontaneously generated by the model. The IBI histogram (inset) attains a maximum very close to zero when SFA is not present; such maximum shifts toward higher values with SFA; the two distributions also have different averages (2.3 s vs 3.2 s) and coefficients of variation (0.98 vs 0.65). Thus SFA significantly increases the average interval to the next burst, reducing at the same time its variability.

This difference in the low-IBI distribution can be traced back to the slow increase (resp. decrease) of the average post-burst activity with and without SFA, visible in the plots of Fig. 4. We also remark that single bursts are not visible across the shaded regions in such plots because bursts, across the 500 simulations are broadly distributed in time so that the corresponding peaks are not visible in the plots.

It is noteworthy that our GLM model is able to reproduce such a detailed feature, comparably with model networks of adapting integrate-and-fire neurons.

We also found in data, consistently with the results reported in²⁹, a positive correlation between burst amplitude and the length of the previous IBI, an effect that could be possibly attributed to some adaptation mechanism, such as SFA. We estimated this correlation from simulations of the models inferred, with and without SFA, on different recordings; Fig. 4C compares the result of the two models with that found in their experimental counterparts. It is seen that the model without SFA is unable to recover correlations that are significantly different from zero, while the model with SFA produces correlations that are consistent with that measured in the data.

We also addressed the ability of the model to account for the temporal structure and spatial organization of the network activity. We first considered the average order of activation of different spatial locations within a burst: time-rank is the index of the time bin when a neuron spikes first since the burst onset. By evaluating the average rank for both each neuron in the model and each electrode in the data, we found them strongly correlated ($c = 0.9$). In Fig. 5A we report the spatial distribution of average time-ranks for data (left) and simulations (right).

Looking farther into the burst temporal structure, we asked to what extent the time-ranks in the model and in the data are comparable on a single burst basis. For each burst in the data, we selected and compare the simulated burst with the closest time-rank pattern. Figure 5B,C show two examples of such comparison. To provide a global comparison, capturing the correlation in time-ranks between data and simulations, we proceeded as follows: independently for each neuron, we shuffled the vector of its time-ranks across all the simulated bursts, thereby destroying spatial correlations in the time-ranks, while preserving the average ranks (Fig. 5A, right).

For each burst in the data we took the closest simulated burst and the closest shuffled simulated one, computed the corresponding euclidean distances (d_{shuff}^{min} and d_{sim}^{min}) and the distribution of their differences Fig. 5D (blue histogram). It is seen that, for the large majority of bursts, the simulated burst is closer to the data burst than the surrogate. To quantitatively support this indication, for each burst in the data we took the closest from two different realizations of the simulation shuffling and evaluated the respective euclidean distances between the data and the two shuffles. The distribution of differences between such distances is reported in Fig. 5D (blue histogram), and is used as a reference to compare the significance of the first distribution. The two histograms in Fig. 5D are statistically different (2-sample t-Student returns $p = 10^{-5}$), confirming that the simulated bursts consistently reproduce the order of burst activation more reliably than surrogate bursts. We therefore can conclude that inferred model captures most of the spatial development of neural activity at the single burst level.

To inquire into the relationship between the inferred synaptic structure and the ensuing network dynamics, we also asked whether for a neuron with low average time-rank (early spiking), the efficacy of its outgoing synapses correlates with the time-ranks of its post-synaptic neurons. Figure 5E shows indeed a high negative correlation between the efficacy of outgoing synapses and the time-rank of post-synaptic neurons ($c = -0.8$).

Discussion

Several criteria can be identified to evaluate a model, yet the ability to reproduce the behavior of the system under analysis is certainly among the most relevant. We have seen that the addition of few computational ingredients can enable a probabilistic, generative network to autonomously sustain a rich, highly non-stationary dynamics, as observed in the experimental data employed to infer the model's parameters.

In recent years, the field of statistical inference applied to neuronal dynamics has mainly focused on devising models and procedures that could reliably recover the real or effective synaptic structure of a network of neurons, under different dynamical and stimulation conditions^{3,30,31}, even though, more often than not, an experimental ground truth was not readily available.

Although, quite surprisingly, the study of the dynamics of the inferred models has been largely neglected, we are not the first to address the issue. In¹¹ the authors demonstrated, on an in-vitro population of ON and OFF parasol ganglion cells, the ability of a GLM to accurately reproduce the dynamics of the network¹⁴; studied the response properties of lateral intraparietal area neurons at the single trial, single cell level; the capability of GLM to capture a broad range of single neuron response behaviors was analyzed in¹³. In all these works, however, the focus was on the response of neurons to stimuli of different spatio-temporal complexity; even where network interactions were accounted for, they proved to be, for the overall dynamics displayed by the ensemble, an important, yet not decisive correction. To our knowledge, no published study to date has focused on the autonomous dynamics of GLM networks inferred from neuronal data.

Important progress has also been made on the ability of GLMs to predict single neuron spiking in an ensemble of neurons, with accuracy higher than that provided by methods based on the instantaneous state of the ensemble itself⁹ and the PSTH¹¹. In the lingo of the present paper, though, these attempts were based on simulations of the inferred network in driven mode, with at most the activity of a single neuron free to reenter just that single neuron's dynamics. To our knowledge, our work for the first time has attempted a direct, microscopic comparison between the activity autonomously generated by an inferred GLM network and the one observed in the biological data. The free mode of simulation opens the way, in principle, to multi-time-step, whole network activity predictions. Being the model probabilistic and the data noisy, it is of course expected that the ability of the model, in free simulation mode after having been driven by the data for some time, to follow the experimentally observed spike trains should quickly deteriorate. In fact, this is what we found with our model: a burst becomes practically unforeseeable going some 100–200 ms in the past. This result is in contrast with the impressive findings reported in³², in which the authors, on recordings similar to the ones used in this work, through a new model-free method based on state-space reconstruction, were able to predict the occurrence of a burst even 1 s in advance. We applied the state-reconstruction method on our data with negative results, in agreement with what we found with our model.

We presented several semi-quantitative comparisons between the spontaneous spike trains generated by the model network and the data from which the model's parameters have been inferred. Our model clearly outperforms typically used GLM inference models, with exponential gain functions and Poisson spike generative model.

A natural question arises as to how representative the chosen experimental system is of neuronal network dynamics in general and the capability of the presented framework to generalize to other kinds of non-stationary

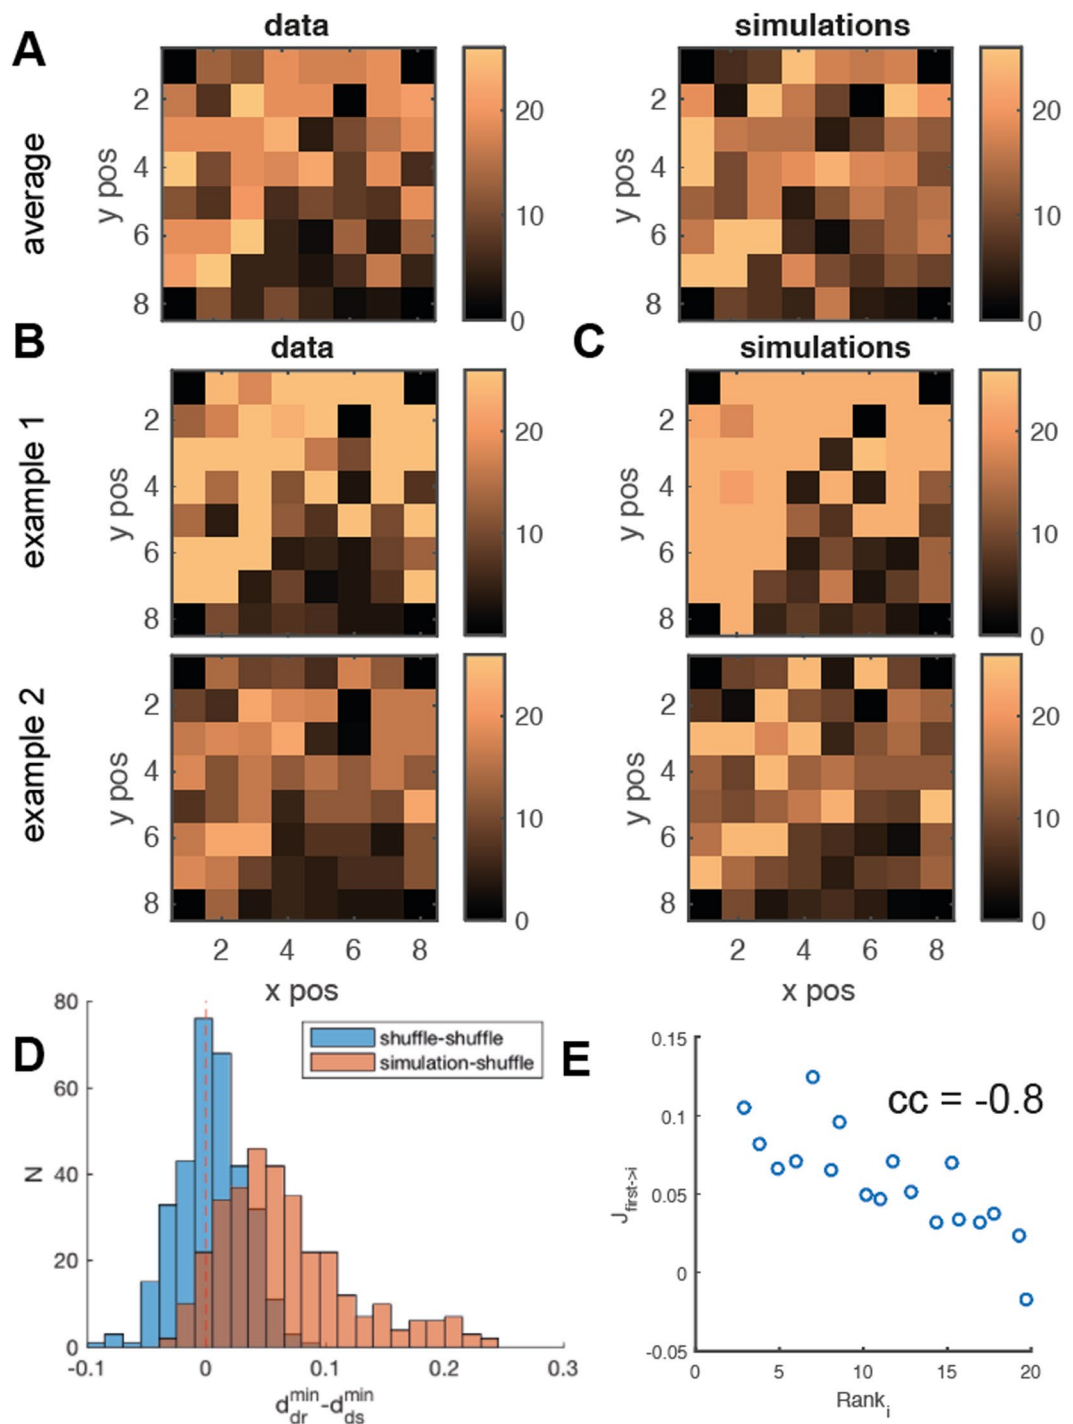


Figure 5. Spatio-temporal burst dynamics, data vs model. **(A)** Spatial distribution of the average time-rank of neuron activations during bursts for data and model. Averages are computed over all bursts in the data and a simulation of comparable length. **(B)** Spatial distribution of time-rank for 2 single bursts in data. **(C)** Spatial distribution of time-rank for the 2 bursts from simulation, most similar to those shown in B. **(D)** For each burst in the data, we computed its distances to the closest bursts from simulation and from shuffled simulation (see text). We show the distribution of the differences in such distances. Simulated bursts are typically closer to experimental ones w.r.t. shuffled simulations (red histogram). Similarly for each data burst we evaluated the closest burst from two different shufflings of simulation (blue histogram). The latter is used as a reference for our measure. The 2-sample t-Student returns $p = 10^{-5}$. **(E)** Scatterplot of post-synaptic efficacies of the neuron with smallest time-rank vs the rank of its post-synaptic targets. The synaptic structure correlates with the early development of the burst.

neural activities. Our primary concern in selecting an *ex-vivo* bursting culture, known to exhibit complex spontaneous spatio-temporal behavior^{15–21}, has been to study a system that, absent any external stimuli, would show a similar non-stationary collective behavior. As discussed above, the capability of GLMs to mimic the response properties of single neurons or a network to stimuli has already been well investigated. On the other hand it will be interesting to study the implications of the elements we introduced in the present work in terms of flexibility and robustness of inference in driven conditions.

Given our goal to establish the capability of a GLM to reproduce a complex, non-linear dynamics, we found two ingredients - a non-symmetric sigmoidal transfer function and a Negative Binomial spike-generation model - that proved to be minimal: absent one of them, the model's performances were drastically impaired. Even if we do not claim these specific ingredients to be necessary in general, we are persuaded they represent instances of mechanisms that are both biologically plausible and computationally desirable for the model. Although a sigmoidal, saturating transfer function would appear to be a natural choice for a model meant to reproduce neural activities, surprisingly this option has never been explored, to our knowledge, to perform network inference on neuronal recordings. While the use of an exponential transfer function in GLMs is grounded in general statistical requirements³³, in our case an asymmetric sigmoid appears to be the single most important factor explaining the success of the proposed model. As already noted, the saturation of the sigmoidal function allows naturally for a bimodal distribution of firing rates and thus makes the model's behavior more robust in the face of the intrinsic fluctuations of network activity. A certain degree of asymmetry also resulted from inference, allowing for differential sensitivity to high and low input levels.

It is interesting to note how this finding contrast with the success of non-saturating transfer functions in deep learning literature, where the introduction of rectified linear units, in place of the standard logistic ones, has represented one of the major breakthroughs of recent years³⁴. Such units exhibit 'intensity equivariance', that is the ability to readily generalize to data points that differ only for a scale factor; whilst such property is clearly valuable when dealing with data such as natural images and sounds, when applying machine learning techniques to very noise and sparse data, such as in the case studied here, bounded transfer functions are probably beneficial exactly for the opposite reason: they filter out most of the incoming information to gain a poorer but more robust coding in the output.

We have seen that, if the role of a sigmoid gain function seems to facilitate a bistable behavior of the network, the negative binomial's role is to efficiently destabilize, with large fluctuations of the spike counts, the high-activity state. This hints at the possibility that the present inference model may also be suited to account for UP and DOWN states observed in slices (and it could possibly counteract the effect observed in³⁵). Recent experimental evidence has emerged supporting the Negative Binomial distribution as a candidate for the spike counts variability in real neurons²⁸; moreover, a negative binomial spike generation has already been adopted in a model combining GLM with sparseness priors for the connectivity³⁶. Although, then, our second ingredient already finds experimental and theoretical support in the literature, we provide a new hypothesis on why the super-Poisson statistics arises, as the effect of input fluctuations generated by the activity of neurons that have not been recorded. And our hypothesis directly hints to the Negative Binomial as just one possible way to model such effect, where over-dispersed spike counts are, instead, a necessary signature of it; it is this more general feature, in our opinion, that proved to be so important for the success of the proposed model. Interestingly, our findings resonate with the recently proposed role of fluctuating unobserved variables in the emergence of criticality in a wide range of systems³⁷.

Implicit in our approach is the assumption that the recorded neurons have been chosen at random from the whole network; this is probably not the case in many instances, where a local set of units is sampled instead; such non-random sub-sampling can potentially lead to strong deviations from expected behavior²⁶; this of course could produce systematic effects in the inference. Besides, our model cannot incorporate the effects of slow changes in single neuron excitability, as observed experimentally in^{38,39}, and analyzed in a statistical model in²⁸.

Methods

Experimental data. Electro-physiological recordings are fully described in¹⁷, and were taken from cortical neurons from newborn rats within 24 hours after birth, following standard procedures. Data are freely available from S. Marom (<http://marom.net.technion.ac.il/>). Throughout the paper we made use of 6 consecutive recordings, from 60 electrodes, lasting about 40 minutes each (in Fig. 4A,C, numeric labels are as in the original data files). We treated spikes from each electrode as belonging to a same neuron (no spike sorting); in all our computations, we neglected electrodes (typically less than two) with an average firing rate less than 0.1 Hz.

Model. The input current $H_i(t)$ to neuron i is:

$$H_i(t) = h_i + \sum_{\Delta} \sum_j^{N_n} k_{ij}(\Delta) S_j(t - \Delta) - \sum_{x=1}^5 g_{\text{SFA}}^{(x)} c_i^{(x)}(t). \quad (2)$$

N_n is the number of neurons, h_i is a constant external current; $k_{ij}(\Delta)$ are synaptic kernels of the same form as in¹¹:

$$k_{ij}(\Delta) = \sum_l w_{ij}^{(l)} \text{rc}^{(l)}(\Delta) \quad (3)$$

where $\text{rc}^{(l)}(\Delta)$ are 'raised cosines':

$$\text{rc}^{(l)}(\Delta) = \frac{1}{2} [1 + \cos(a \log(\Delta + \delta) - \phi^{(l)})] \quad (4)$$

with fixed parameters a , δ , and $\phi^{(l)}$.

The five activity-dependent SFA adaptation currents are described by $c_i^{(x)}$, $x = 1 \dots 5$, evolving according to:

$$c_i^{(x)}(t + dt) = c_i^{(x)}(t) \left(1 - \frac{dt}{\tau_{\text{SFA}}^{(x)}} \right) + S_i(t) \frac{dt}{\tau_{\text{SFA}}^{(x)}}. \quad (5)$$

In the Exp-Poisson model the SFA currents are not included, to be consistent with its original formulation. On the other hand, in the Sig-NegBin model we set the diagonal $w_{ii}^{(l)} = 0$, to avoid that two mechanisms with similar computational roles (SFA and the kernels through which the past spikes of the single neuron affect the future activity of the neuron itself) could compete with each other, making the results more difficult to interpret.

The current-to-rate transfer function is:

$$f_{\text{exp}}[H] = \exp(H) \quad (6)$$

for the Exp-Poisson model, and:

$$f_{\text{sigmoid}}[H] = \frac{\lambda_{\infty}}{(1 + \exp(-H))^{\gamma}} \quad (7)$$

for our Sig-NegBin model, with λ_{∞} setting the asymptotic firing rate for high input, and γ governing the asymmetry of the sigmoid for low and high input.

The average firing rate at time $t + dt$ is given by:

$$\lambda_i(t + dt) = f[H_i(t)]. \quad (8)$$

For each time-bin at t , given $\lambda_i(t + dt)$, the model probabilistically generates spikes according to:

$$p_{\text{Poisson}}(S_i(t + dt) | \lambda_i(t + dt)) = \exp(-\lambda_i) \frac{\lambda_i^{S_i}}{S_i!} \quad (9)$$

for the Exp-Poisson model, and:

$$p_{\text{NB}}(S_i(t + dt) | \lambda_i(t + dt)) = \frac{\Gamma(S_i + r_{\text{NB}})}{S_i! \Gamma(r_{\text{NB}})} \left(\frac{\lambda_i}{\lambda_i + r_{\text{NB}}} \right)^{S_i} \left(\frac{r_{\text{NB}}}{\lambda_i + r_{\text{NB}}} \right)^{r_{\text{NB}}} \quad (10)$$

for our Sig-NegBin model.

The model's parameters are fitted maximizing the likelihood:

$$\text{LL} = \sum_t \sum_i \log(p(S_i(t + dt) | \lambda_i(t + dt))) \quad (11)$$

using the adaptive iRprop algorithm²³.

For the Exp-Poisson model the inferred parameters are the $w_{ij}^{(l)}$, including the diagonal terms $w_{ii}^{(l)}$ (SFA currents are in this case neglected by keeping $g_{\text{SFA}}^{(x)} = 0$). Following¹¹, the set of $\text{rc}^{(l)}$ functions is different for diagonal and off-diagonal terms. For the off-diagonal terms the 4 rc functions have parameters: $a = 1.153$, $\delta = 0.2560$ ms, $\phi_s = \frac{\pi}{2}(-1, 0, 1, 2)$; for the diagonal terms there are 6 functions with parameters: $a = 2.974$, $\delta = 0.3477$ ms, $\phi_s = \frac{\pi}{2}(-2, 3, 4, 5, 6, 7)$. For all the rcs , $\text{rc}(\Delta) = 0$ for $\Delta > 150$ ms. The total number of parameters for the models is thus: $4N_n(N_n - 1) + 6N_n$.

For the Sig-NegBin model the inferred parameters are the $w_{ij}^{(l)}$ (only off-diagonal $i \neq j$ terms, with the same $\text{rc}^{(l)}$ functions above), $\tau_{\text{SFA}}^{(x)}$, $g_{\text{SFA}}^{(x)}$, λ_{∞} , and γ . The total number of parameters for this models is thus: $4N_n(N_n - 1) + 10 + 2$. Therefore, the Exp-Poisson model has $6N_n - 12$ parameters more than the Sig-NegBin; with $N_n = 60$ (assuming no electrode is neglected in the recordings, see above), this amounts to a difference of 348 parameters (about +2.5%).

The r_{NB} parameter, instead, was chosen *a priori* as follows. For a NegBin variable $X \sim \text{NegBin}(\lambda, r_{\text{NB}})$, we have $\langle X \rangle = \lambda$ and $\text{Var}[X] = \lambda \left(1 + \frac{\lambda}{r_{\text{NB}}} \right)$. For very low, stationary firing rates the single neuron behavior is often empirically well approximated by a Poisson process⁴⁰⁻⁴²; notably, during the inter-burst intervals in our data, the single neuron inter-spike interval distribution turns out to be close to exponential, as expected for a Poisson process. Therefore, in order for the NegBin distribution to be clearly non-Poissonian only for higher firing rates, we took as a reference the median $m = 0.037$ (among the neurons) of the average spike count (over the time record for each neuron); we then set the 'λ-threshold' to be about $5m$ (a value that, in the data, clearly separates the down states from the bursts), and we chose to impose, at that frequency, a variance-to-mean ratio of 2 (for Poisson, this ratio is always 1 for every value of the mean): $1 + \frac{5m}{r_{\text{NB}}} \equiv 2$. This simple calculation gives $r_{\text{NB}} \simeq 0.2$, a value close to the one we used throughout the paper. We also checked the dependence of the presented results on the value of r_{NB} . The capacity of the network to spontaneously generate bursts is not affected on the entire range we explored; we found a mild dependence on r_{NB} of the average inter-burst interval; in this sense the choice of r_{NB} represents an additional handle for the model to better match the burst statistics observed in the data.

References

1. Nguyen, H. C., Zecchina, R. & Berg, J. Inverse statistical problems: from the inverse Ising problem to data science. *Advances in Physics* **66**, 197–261, <https://doi.org/10.1080/00018732.2017.1341604> (2017).
2. Tkačik, G. *et al.* Thermodynamics and signatures of criticality in a network of neurons. *Proceedings of the National Academy of Sciences* **112**, 11508–11513, <https://doi.org/10.1073/pnas.1514188112> (2015).
3. Roudi, Y., Dunn, B. & Hertz, J. Multi-neuronal activity and functional connectivity in cell assemblies. *Current Opinion in Neurobiology* **32**, 38–44, <https://doi.org/10.1016/j.conb.2014.10.011> (2015).
4. Schneidman, E., Berry, M. J. II, Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007 (2006).
5. John, H., Roudi, Y. & Tyrcha, J. Ising models for inferring network structure from spike data. In Quiroga, R. Q. & Panzeri, S. (eds) *Principles of Neural Coding*, chap. 27, 527–546 (CRC Press, 2013).
6. Roudi, Y. & Hertz, J. Mean field theory for nonequilibrium network reconstruction. *Phys. Rev. Lett.* **106**, 048702, <https://doi.org/10.1103/PhysRevLett.106.048702> (2011).
7. Cocco, S., Leibler, S. & Monasson, R. Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics methods. *Proceedings of the National Academy of Sciences* **106**, 14058–14062, <https://doi.org/10.1073/pnas.0906705106> (2009).
8. Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P. & Brown, E. N. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology* **93**, 1074–1089 (2005).
9. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nature Neuroscience* **13**, 105–111 (2010).
10. Capone, C., Filosa, C., Gigante, G., Ricci-Tersenghi, F. & Del Giudice, P. Inferring synaptic structure in presence of neural interaction time scales. *PLoS One* **10**, e0118412 (2015).
11. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
12. Stevenson, I. H. *et al.* Bayesian inference of functional connectivity and network structure from spikes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **17**, 203–213 (2009).
13. Weber, A. I. & Pillow, J. W. Capturing the dynamical repertoire of single neurons with generalized linear models. *Neural Computation* **29**, 3260–3289 (2017).
14. Park, I. M., Meister, M. L., Huk, A. C. & Pillow, J. W. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience* **17**, 1395–1403 (2014).
15. Gigante, G., Deco, G., Marom, S. & Del Giudice, P. Network events on multiple space and time scales in cultured neural networks and in a stochastic rate model. *PLoS Computational Biology* **11**, e1004547 (2015).
16. Baltz, T., Herzog, A. & Voigt, T. Slow oscillating population activity in developing cortical networks: models and experimental results. *J Neurophysiol* **106**, 1500–1514, <https://doi.org/10.1152/jn.00889.2010> (2011).
17. Eytan, D. & Marom, S. Dynamics and effective topology underlying synchronization in networks of cortical neurons. *The Journal of Neuroscience* **26**, 8465–8476 (2006).
18. Giugliano, M., Darbon, P., Arsiero, M., Lüscher, H.-R. & Streit, J. Single-neuron discharge properties and network activity in dissociated cultures of neocortex. *Journal of Neurophysiology* **92**, 977–996 (2004).
19. Gritsun, T., le Feber, J., Stegenga, J. & Rutten, W. L. Experimental analysis and computational modeling of interburst intervals in spontaneous activity of cortical neuronal culture. *Biological Cybernetics* **105**, 197–210 (2011).
20. Park, I. M., Xu, D., DeMarse, T. B. & Principe, J. C. Modeling of synchronized burst in dissociated cortical tissue: An exploration of parameter space. In *Neural Networks, 2006. IJCNN'06. International Joint Conference on*, 581–586 (IEEE, 2006).
21. Wagenaar, D. A., Pine, J. & Potter, S. M. An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience* **7**, 11 (2006).
22. Plenz, D. & Schuster, H. G. *Criticality in neural systems*. (Wiley-VCH New York, NY, 2014).
23. Igel, C. & Hüsken, M. Empirical evaluation of the improved rprop learning algorithms. *Neurocomputing* **50**, 105–123 (2003).
24. Masquelier, T. & Deco, G. Network bursting dynamics in excitatory cortical neuron cultures results from the combination of different adaptive mechanisms. *PLoS One* **8**, e75824 (2013).
25. Peyrache, A. *et al.* Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep. *Proceedings of the National Academy of Sciences* **109**, 1731–1736 (2012).
26. Levina, A. & Priesemann, V. Subsampling scaling. *Nature Communications* **8**, 15140 (2017).
27. Onken, A., Grünewälder, S., Munk, M. H. & Obermayer, K. Analyzing short-term noise dependencies of spike counts in macaque prefrontal cortex using copulas and the flashlight transformation. *PLoS Computational Biology* **5**, e1000577 (2009).
28. Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nature Neuroscience* **17**, 858–865 (2014).
29. Lombardi, F., Herrmann, H. J., Plenz, D. & de Arcangelis, L. Temporal correlations in neuronal avalanche occurrence. *Scientific Reports* **6**, 24690 (2016).
30. Tyrcha, J., Roudi, Y., Marsili, M. & Hertz, J. The effect of nonstationarity on models inferred from neural data. *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P03005 (2013).
31. Nghiem, T.-A., Telenczuk, B., Marre, O., Destexhe, A. & Ferrari, U. Maximum entropy models reveal the correlation structure in cortical neural activity during wakefulness and sleep. *arXiv preprint arXiv:1801.01853* (2018).
32. Tajima, S., Mita, T., Bakkum, D. J., Takahashi, H. & Toyozumi, T. Locally embedded presages of global network bursts. *Proceedings of the National Academy of Sciences* **114**, 9517–9522 (2017).
33. Lindsey, J. K. *Applying generalized linear models* (Springer Science & Business Media, 2000).
34. Nair, V. & Hinton, G. E. Rectified linear units improve restricted Boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, 807–814 (2010).
35. Rostami, V., Mana, P. P., Grün, S. & Helias, M. Bistability, non-ergodicity, and inhibition in pairwise maximum-entropy models. *PLoS Computational Biology* **13**, e1005762 (2017).
36. Linderman, S., Adams, R. P. & Pillow, J. W. Bayesian latent structure discovery from multi-neuron recordings. In *Advances in Neural Information Processing Systems, 2002–2010* (2016).
37. Schwab, D. J., Nemenman, I. & Mehta, P. Zipf's law and criticality in multivariate data without fine-tuning. *Physical Review Letters* **113**, 068102 (2014).
38. Gal, A. *et al.* Dynamics of excitability over extended timescales in cultured cortical neurons. *Journal of Neuroscience* **30**, 16332–16342 (2010).
39. Gal, A. & Marom, S. Self-organized criticality in single-neuron excitability. *Physical Review E* **88**, 062717 (2013).
40. Eggermont, J. J. The correlative brain. In *The correlative brain*, 267–281 (Springer, 1990).
41. Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience* **18**, 3870–3896 (1998).
42. Averbeck, B. B. Poisson or not Poisson: differences in spike train statistics between parietal cortical areas. *Neuron* **62**, 310–311 (2009).

Acknowledgements

We are grateful to Shimon Marom for having shared his data, and for several illuminating discussions along the way. We thank Gabriel Baglietto for a critical reading of the manuscript. The work was supported by the EU Horizon 2020 Research and Innovation Programme under HBP SGA2 (grant no. 785907 to P.D.G.).

Author Contributions

C.C., G.G. and P.D.G. designed research. C.C. and G.G. developed and implemented algorithms. C.C., G.G. and P.D.G. analyzed results and wrote the paper.

Additional Information

Competing Interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018