Title
Sparse thalamocortical convergence

Journal
Current Biology, 31(10)

ISSN
0960-9822

Author
Ringach, Dario L

Publication Date
2021-05-01

DOI
10.1016/j.cub.2021.02.032

Peer reviewed

Published in final edited form as:

Curr Biol. 2021 May 24; 31(10): 2199–2202.e2. doi:10.1016/j.cub.2021.02.032.

Sparse thalamocortical convergence

Dario L. Ringach

Departments of Psychology and Neurobiology, David Geffen School of Medicine, University of California, Los Angeles Los Angeles, CA 90095, USA.

Summary

How many thalamic neurons converge onto a cortical cell? This is an important question because the organization of thalamocortical projections can influence the cortical architecture [1, 2]. Here we estimate the degree of thalamocortical convergence in primary visual cortex by taking advantage of the cortical expansion – neurons within a restricted volume in primary visual cortex have overlapping receptive fields driven by a smaller set of inputs from the lateral geniculate nucleus [3–5]. Under these conditions, the measurements of cortical receptive fields in a population can be used to infer the receptive fields of their geniculate inputs and the weights of their projections using non-negative matrix factorization [6]. The analysis reveals sparse connectivity [7], where a handful (~2 − 6) of thalamic inputs account for 90% of the total synaptic weight to a cortical neuron. Together with previous findings [8] these results paint a picture consistent with the idea that convergence of a few inputs partly determine the retinotopy and tuning properties of cortical cells [8–13].

eTOC Blurb

Ringach shows that the receptive fields of a population of neurons in primary visual cortex can be used to reconstruct their thalamic inputs and the weight of their projections. The analyses reveal sparse, thalamocortical connectivity, where a handful of neurons account for a substantial part of the cortical response.

Graphical Abstract

dario@ucla.edu, Twitter: @DarioRingach.

Author contributions

D. L. R. conceived of the experiment, collected the data, conducted the analyses, and wrote the manuscript.

Declaration of interests

The author declares no competing interests.

Inclusion and diversity

We worked to ensure sex balance in the selection of non-human subjects. The author of this paper self-identifies as Latinx.

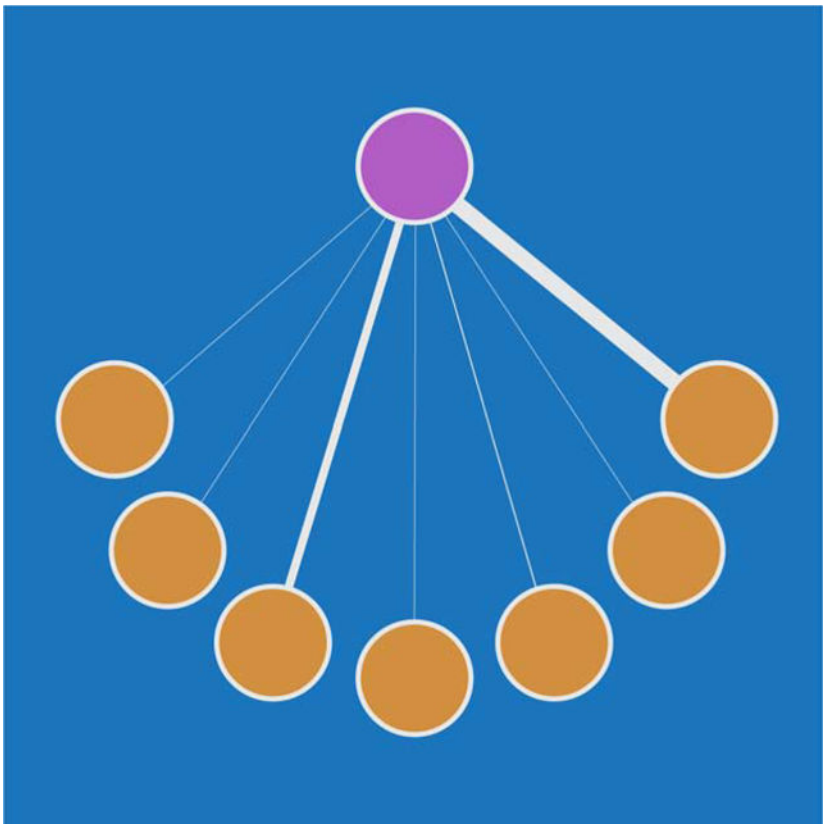

Results and Discussion

We use two-photon, calcium imaging in awake mice to record the responses of neurons in primary visual cortex to a small, contrast-reversing 2 × 2 checkerboard presented at different locations on a grid covering a large part of the visual field (Figure 1A). An individual 2 × 2 checker measures 5.6 × 5.6 deg^2 which is just under the optimal stimulus size for thalamic neurons (cf. Figure 3I in [14]). The average response of neurons to stimulation at each location yields the classical receptive field (Figure 1B, top). Individual receptive fields can be reasonably fit by anisotropic, two-dimensional Gaussians (Figure 1B, bottom). Note that these measurements simply represent the locations on the visual field where visual stimulation by a small stimulus leads a neuron to respond – there is no attempt to map ON and OFF subregions separately.

Let us assume the receptive fields of the population $\{r_i(x,y)\}_{i=1}^{n}$ can be expressed as a non-negative, linear combination of a smaller set of geniculate receptive fields, $r_i(x,y) \cong \sum_{i=1}^{m} w_{i,j} h_j(x,y)$ (Figure 2A). Here, $\{h_k(x,y)\}_{k=1}^{m}$ are the receptive fields of putative thalamic neurons ($m < n$), and the synaptic weights between thalamic and cortical neurons are non-negative, $w_{i,j} \geq 0$. In matrix form, we can write $R = WH (W, H \geq 0)$ where R $(n \times p)$ are the cortical receptive fields (with p being the number of different grid locations), $W(n \times m)$ are the synaptic weights, and $H(m \times p)$ are the thalamic receptive fields. For a given number of geniculate neurons, we can use non-negative matrix factorization [6,

15] to estimate both the thalamic receptive fields H and their projections w Note that no other constraints other than non-negativity are imposed. In particular, the shape and size of the thalamic receptive fields (represented by the matrix H) are an outcome of the analysis. I validated and assessed the robustness of this approach using computer simulations before applying it to experimental data (Supplemental Information).

How many geniculate receptive fields are needed to account for the retinotopy of a given population? Given an experimental dataset R containing the measured receptive fields, we can first examine how the quality of the factorized approximation improves as we increase the putative number of geniculate receptive fields, m (Figure 2B). The dependence of the mean-squared error (MSE) with the number of geniculate receptive fields falls off exponentially (Figure 2B, solid curve). We define the optimal as the smallest number that achieves at least a 90% reduction in MSE (Figure 2B, red point – see Figure S1F for a robustness analysis of this choice). In this example, m=18 inputs are sufficient to explain the structure of n=101 cortical receptive fields. The reconstructed thalamic receptive fields tile a compact region of the visual space (Figure 2C–F, left panels). Moreover, the receptive fields of individual cortical neurons in the population are well described by the linear combination of a small number of geniculate inputs (Figure 2C–F). For each cortical cell, we calculate the fraction of the total synaptic input accounted for by the largest k inputs (Figure 3A). The effective number of inputs is defined as the smallest number of inputs that contribute at least 90% of the total synaptic weight. The distribution of the effective number of inputs across all n = 425 neurons in our experiments shows that 4.69 ± 2.19 (1SD) geniculate inputs are sufficient to account for the retinotopy of the cortical population (Figure 3B).

Discussion

At first sight the result appears in conflict with a more direct estimate of 80 ± 10 thalamic inputs into layer 4 neurons [7]. In this previous study, the authors measured the unitary EPSC and its integral (the charge contribution) of individual geniculate neurons during cortical silencing. To estimate the number of inputs a cell receives they asked how many times one would need to resample from the empirical charge-contribution distribution for the accumulated values to reach the total charge produced by visual stimulation. This method yielded an estimate of 80 ± 10 thalamic inputs, which is an order of magnitude larger than obtained in Figure 3B.

Can these data be reconciled? One salient feature of the empirical charge-contribution distribution is that most values are rather small and only a handful are large (cf. Figure 6B in [7]). When there is a large inequality between the contributions of the inputs, the specific way one defines the effective number of inputs can yield disparate results. To demonstrate this point, consider a hypothetical example where a cell receives a total of 100 inputs with a contribution of 0.1% and a single input contributing the remaining 90% of the charge. Applying the resampling method [7] to this case results in an estimate of 127 inputs (the reason for the curious overestimate is that the probability of not observing the large input after 101 draws is relatively large, ~exp(−1) = 0.37). In contrast, if we ask what is the minimum number of inputs that account for 90% of the total charge the answer is, clearly, just one. Finally, if we ask how many non-zero inputs the cell receives, irrespective of the

magnitude of their contributions, the answer is 101. Thus, different definitions yield wieldy different numbers.

The preceding discussion suggests we may be able to reconcile the present estimate of thalamic convergence with the one reported previously [7]. To evaluate quantitatively if this scenario is feasible, let us assume a symmetric Dirichlet distribution of contributions with concentration parameter $\alpha = 0.03$ and a total number of $N = 69$ inputs (the Dirichlet distribution ensures the sum of all contributions is 100%). For this particular selection of values, the minimum number of inputs required to reach 90% of the total is ~4.6 (Figure 4A) while the resampling method yields ~80 inputs (Figure 4B), thereby offering a potential explanation for the observed differences. The distribution of charge contribution is the marginal of the Dirichlet distribution, which turns out to be a Beta distribution with shape parameters $(\alpha, (N-1)\alpha)=(0.03,2.04)$. As expected, this distribution is sparse – small values appear frequently while large values are very rare (Figure 4C). Note that the largest contribution observed among the 23 measurements reported by Lien and Scanziani was 6% (see Figure 6 in ref [7]). The probability that any one realization from a Beta distribution with the above shape parameters will be smaller than this value is $p = 0.9456$. This means that the probability that the maximum of 23 such observations would not exceed 6% is $p^{23} \sim 0.2764$. Therefore, there is a reasonable chance that inputs with larger contributions exist but were missed in the previous study due to limited sampling. In other words, the findings from the two studies can be reconciled assuming a sparse distribution of charge contributions and applying the definitions and estimation methods used in each study. Both studies reinforce the conclusion that thalamocortical connectivity is sparse and suggest that fitting a Beta distribution to synaptic contribution data might provide a better yardstick to compare results.

Although the present findings can be reconciled with prior data, it is worth mentioning experimental factors that could bias our estimates towards lower numbers. First, the selected spatial and temporal parameters of the stimuli is expected to drive most thalamic neurons, including sustained ON/OFF and transient OFF classes, but to a lesser degree, directionally selective neurons (as the stimulus has no net motion) and very few of the so-called “slow” or “contrast suppressed” neurons [14]. Thus, it is possible that the stimulus fails to engage all different cell thalamic classes equally. If multiple cell classes converge onto single cortical neurons, this can cause an underestimate of the actual number of inputs. Note this a limitation of the stimulus and not the technique. Second, ON and OFF-center thalamic inputs with largely overlapping receptive fields could potentially be lumped into a single receptive field in the present analysis. This can be avoided by separately measuring ON and OFF responses and independently estimating ON and OFF inputs. Third, as the stimuli are small, it is possible that weak, sub-threshold inputs are not detectable in our measurements [16, 17]. The present estimates should be interpreted with these limitations in mind. Note that alternative, anatomical approaches to the question, such as single-cell viral tracing of the inputs of layer 4 cells to label all its thalamic inputs can yield important connectivity data, but do not yet allow us to estimate the strength of synaptic connections, which is necessary to compute the effective number of inputs. The number of cells labeled in such a way only provides a weak, upper bound on the number of inputs.

These findings indicate that, at least when it comes to explaining the structure of population receptive fields, the estimated thalamocortical connectivity is sparse in the sense that a handful of inputs contribute a large fraction of the response [18] – a result we modeled mathematically and showed to be consistent with prior data [7]. In the mouse, we further know that the similarity in tuning for orientation and spatial frequency between a pair of neurons increases as the overlap between their receptive fields increases [8]. Altogether, these results paint a picture consistent with the notion that the selection of a small number of inputs by a cortical cell determines, at least in part, both its retinotopy and tuning [9, 11, 12].

STAR Methods

Lead contact

Further information and requests for resources and reagents should be directed to the Lead Contact, Dario Ringach (dario@ucla.edu)

Materials availability

This study did not generate new unique reagents.

Data and Code Availability

The dataset for this study, including raw kernels and their Gaussian fits, has been deposited at https://figshare.com/articles/dataset/Sparse_Thalamocortical_Convergence/13864970

Experimental Model and Subject Details

All procedures were approved by UCLA's Office of Animal Research Oversight (the Institutional Animal Care and Use Committee) and were in accord with guidelines set by the US National Institutes of Health. A total of 5 mice, both male (3) and female (2), aged P35–56, were used. In two mice we collected data from two separate populations, yielding a total of seven experimental datasets. All these animals were from a TRE-GCaMP6s line G6s2 (Jackson Lab), where GCaMP6s is regulated by the tetracycline-responsive regulatory element (tetO). Mice were housed in groups of 2–3, in reversed light cycle.

Method Details

Surgery—Carprofen and buprenorphine analgesia were administered pre-operatively. Mice were then anesthetized with isoflurane (4–5% induction; 1.5–2% surgery). Core body temperature was maintained at 37.5C using a feedback heating system. Eyes were coated with a thin layer of ophthalmic ointment to prevent desiccation. Anesthetized mice were mounted in a stereotaxic apparatus. Blunt ear bars were placed in the external auditory meatus to immobilize the head. A portion of the scalp overlying the two hemispheres of the cortex (approximately 8mm by 6mm) was then removed to expose the underlying skull. After the skull is exposed, it was dried and covered by a thin layer of Vetbond. After the Vetbond dries (15 min) it provides a stable and solid surface to affix an aluminum bracket (a head holder) with dental acrylic. The bracket is then affixed to the skull and the margins sealed with Vetbond and dental acrylic to prevent infections.

Imaging and eye movements—We began imaging sessions 5–7 days after surgery. We used a resonant, two-photon microscope (Neurolabware, Los Angeles, CA) controlled by Scanbox acquisition software and electronics (Scanbox, Los Angeles, CA). The light source was a Coherent Chameleon Ultra II laser (Coherent Inc, Santa Clara, CA) running at 920nm. The objective was an x16 water immersion lens (Nikon, 0.8NA, 3mm working distance). The microscope frame rate was 15.6Hz (512 lines with a resonant mirror at 8kHz). A camera synchronized to the frame rate of the microscope imaged the eye and pupil during data collection. These data were subsequently analyzed to determine the center and size of the pupil within the image plane. The distribution of eye movements was computed, yielding a mode and scatter that was larger along the horizontal axis. Our analyses yielded similar results whether performed on the entire dataset or by based on data segments where the eye position was within 1SD of the mode.

Visual stimulation—We used a Samsung CHG90 monitor positioned 30 cm in front of the animal. The screen was calibrated using a Spectrascan PR-655 spectro-radiometer (Jadak, Syracuse, NY), and the result used to generate the appropriate gamma corrections for the red, green and blue components via a GeForce RTX 2080 Ti graphics card. Visual stimuli were generated by a Processing sketch using OpenGL shaders (see http://processing.org). Stimuli were presented only on the right visual hemifield, which covered 100° × 55°, as recordings took place from the monocular zone of primary visual cortex on the left hemisphere. Two by two checkerboards, switching contrast every 100 msec, were presented for 0.5 sec at a random location on an 18 × 10 grid. Stimulus onset asynchrony was 1 sec. Stimulus locations were randomized in blocks and we presented a total of 6 blocks. Transistor-transistor logic (TTL) pulses generated by the stimulus computer signaled the onset of stimulation at each location and were sampled by the microscope and time-stamped with the frame and line number that being scanned at that time.

Visual stimulation—We used a Samsung CHG90 monitor positioned 30 cm in front of the animal. The screen was calibrated using a Spectrascan PR-655 spectro-radiometer (Jadak, Syracuse, NY), and the result used to generate the appropriate gamma corrections for the red, green and blue components via a GeForce RTX 2080 Ti graphics card. Visual stimuli were generated by a Processing sketch using OpenGL shaders (see http://processing.org). Stimuli were presented only on the right visual hemifield, which covered 100° × 55°, as recordings took place from the monocular zone of primary visual cortex on the left hemisphere. Two by two checkerboards, switching contrast every 100 msec, were presented for 0.5 sec at a random location on an 18 × 10 grid. Stimulus onset asynchrony was 1 sec. Stimulus locations were randomized in blocks and we presented a total of 6 blocks. Transistor-transistor logic (TTL) pulses generated by the stimulus computer signaled the onset of stimulation at each location and were sampled by the microscope and time-stamped with the frame and line number that being scanned at that time.

Data and Robustness Analysis—We used Matlab to compute the non-negative matrix factorization using the alternating least-squares algorithm. To analyze the robustness of the method, we simulated an input consisting of LGN receptive fields with centers located on the vertices of a noisy hexagonal grid (Figure S1A). Average grid spacing was 2σ, where σ

is the width of the Gaussian defining the LGN receptive fields. We then simulated cortical neuron sampling from inputs near the center of the grid. For each neuron, we centered its receptive field at (x_0, y_0) where both variables are distributed uniformly $U(-1.5,1.5)$. This introduces scatter in the receptive field centers. We randomly oriented an ellipse at the center location with a minor axis of 1.2 and a major axis equal to $1.5+U(0,2)$. This introduced variability in the axes of elongation and the aspect ratio of the receptive fields. The cortical cell then sampled from all the LGN cells with centers within the interior of the ellipse with a synaptic weight that decayed exponentially from (x_0, y_0) and with a space constant equal to 3. The level sets of some of the receptive fields generated by this procedure are shown in Figure S1A. For each cell, we computed the number of inputs as the minimum number required to account for 90% of the total synaptic input. The mean squared error fell exponentially initially and it increases when the model is overfit (Figure S1B). Non-negative matrix factorization for an optimal selection of $m=12$ yielded LGN receptive fields with centers that aligned with those simulated (Figure S1C). The confusion matrix shows that over 90% of the time the number of inputs were recovered exactly (Figure S1D). The distribution of the effective number of inputs was nearly identical to the simulated one (Figure S1E). Importantly, the performance of the algorithm is robust to the selection of the number of receptive fields in the geniculate (Figure S1F). There is a substantial range of values over which the bias is near zero and the standard deviation of the estimates is low. Here, we define bias as the mean of $\bar{n}_i - n_i$, where $\bar{n}_i$ is the estimated number of inputs to a cell and n_i is the actual, simulated value; the precision is defined as the standard deviation of $\bar{n}_i - n_i$. The shaded are in Figure S1F shows that over a reasonable range of values the bias and the standard deviation are low. Thus, the estimates appear robust to the selection of m.

Quantification and Statistical Analysis—The experiments were conducted in a total of 5 mice yielding a pooled dataset of $n = 425$ cells (Figure 3B). We use standard deviation (SD) as a measure of dispersion of the estimated number of inputs, while also showing the entire distribution of values (Figure 3B). Figure 2B shows an estimate of the number of neurons in the LGN required to explain the population receptive fields of $n = 101$ cells in a single experiment. Such data were fit with an exponential function $A + B\exp(-Cm)$ using the 'fit' function in Matlab. Non-negative matrix factorization was computed using the 'nnmf' function in Matlab. The experimental design did not involve assigning subjects to different groups and evaluating the outcome of different conditions. Thus, no statistical comparison between groups was needed.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The author wants to thank Elaine Tring for performing the cranial window implants. This study was supported by EB022915 and NS116471 to D. L. R.

References

1. Kremkow J, and Alonso JM (2018). Thalamocortical Circuits and Functional Architecture. Annu Rev Vis Sci 4, 263–285. [PubMed: 29856937]
2. Mazade R, and Alonso JM (2017). Thalamocortical processing in vision. Vis Neurosci 34, E007. [PubMed: 28965507]
3. Smith SL, and Häusser M.(2010). Parallel processing of visual space by neighboring neurons in mouse visual cortex. Nature neuroscience 13, 1144–1149. [PubMed: 20711183]
4. Olshausen BA, and Field DJ (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? Vision Res 37, 3311–3325. [PubMed: 9425546]
5. Olshausen BA, and Field DJ (2005). How close are we to understanding v1? Neural Comput 17, 1665–1699. [PubMed: 15969914]
6. Lee DD, and Seung HS (1999). Learning the parts of objects by non-negative matrix factorization. Nature 401, 788–791. [PubMed: 10548103]
7. Lien AD, and Scanziani M.(2018). Cortical direction selectivity emerges at convergence of thalamic synapses. Nature 558, 80–86. [PubMed: 29795349]
8. Jimenez LO, Tring E, Trachtenberg JT, and Ringach DL (2018). Local tuning biases in mouse primary visual cortex. J Neurophysiol.
9. Ringach DL (2011). You get what you get and you don't get upset. nature neuroscience 14, 123–124. [PubMed: 21270775]
10. Jin J, Wang Y, Swadlow HA, and Alonso JM (2011). Population receptive fields of ON and OFF thalamic inputs to an orientation column in visual cortex. Nature neuroscience 14, 232–238. [PubMed: 21217765]
11. Ringach DL (2004). Haphazard wiring of simple receptive fields and orientation columns in visual cortex. Journal of neurophysiology 92, 468–476. [PubMed: 14999045]
12. Soodak RE (1987). The retinal ganglion cell mosaic defines orientation columns in striate cortex. Proc Natl Acad Sci U S A 84, 3936–3940. [PubMed: 3108884]
13. Paik SB, and Ringach DL (2011). Retinal origin of orientation maps in visual cortex. Nat Neurosci 14, 919–925. [PubMed: 21623365]
14. Piscopo DM, El-Danaf RN, Huberman AD, and Niell CM (2013). Diverse visual features encoded in mouse lateral geniculate nucleus. The Journal of neuroscience 33, 4642–4656. [PubMed: 23486939]
15. Lee DD, and Seung HS (2001). Algorithms for Non-negative Matrix Factorization. In Advances in Neural Information Processing Systems 13, Leen TK, Dietterich TG and Tresp V, eds. (MIT Press), pp. 556–562.
16. Carandini M, and Ferster D.(2000). Membrane potential and firing rate in cat primary visual cortex. J Neurosci 20, 470–484. [PubMed: 10627623]
17. Ringach DL, and Malone BJ (2007). The operating point of the cortex: neurons as large deviation detectors. J Neurosci 27, 7673–7683. [PubMed: 17634362]
18. Hurley N, and Rickard S.(2009). Comparing measures of sparsity. IEEE Trans Information Theory 55, 4723–4741.

Highlights

- Population receptive fields in V1 are used to reconstruct inputs from the LGN.
- Estimated connectivity shows sparse thalamocortical convergence.
- A handful of inputs account for a large fraction of the cortical responses.
- Clarifies the notion of “effective number of inputs”.

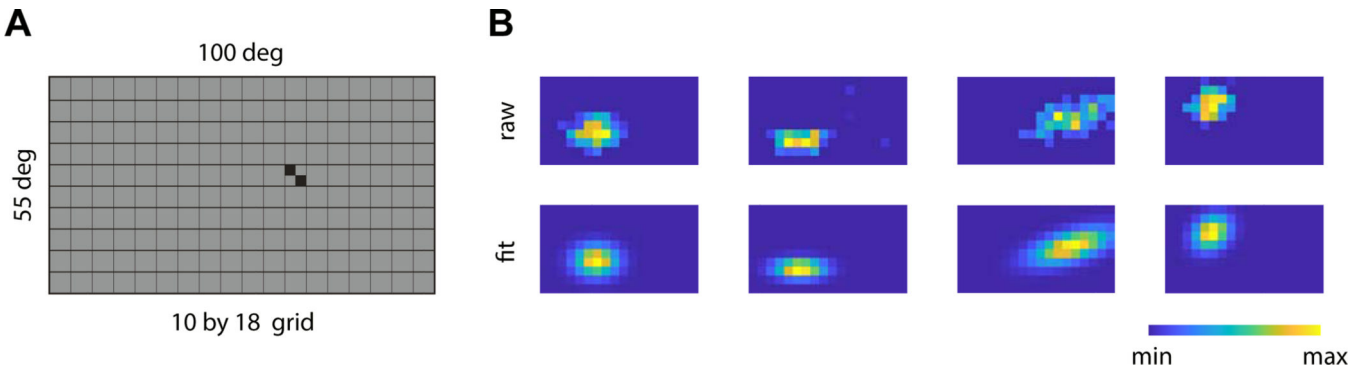

Figure 1. Measuring population receptive fields in mouse primary visual cortex.

A. The stimulus consists of flickering 2×2 checkerboards on a 10 by 18 grid covering 100 by 55 deg of the contralateral hemifield. The average response to stimulation each location is represented as an image depicting the receptive fields of the neurons. **B.** Examples of four (raw) receptive fields (top) and their Gaussian fits (bottom) from neurons in mouse V1.

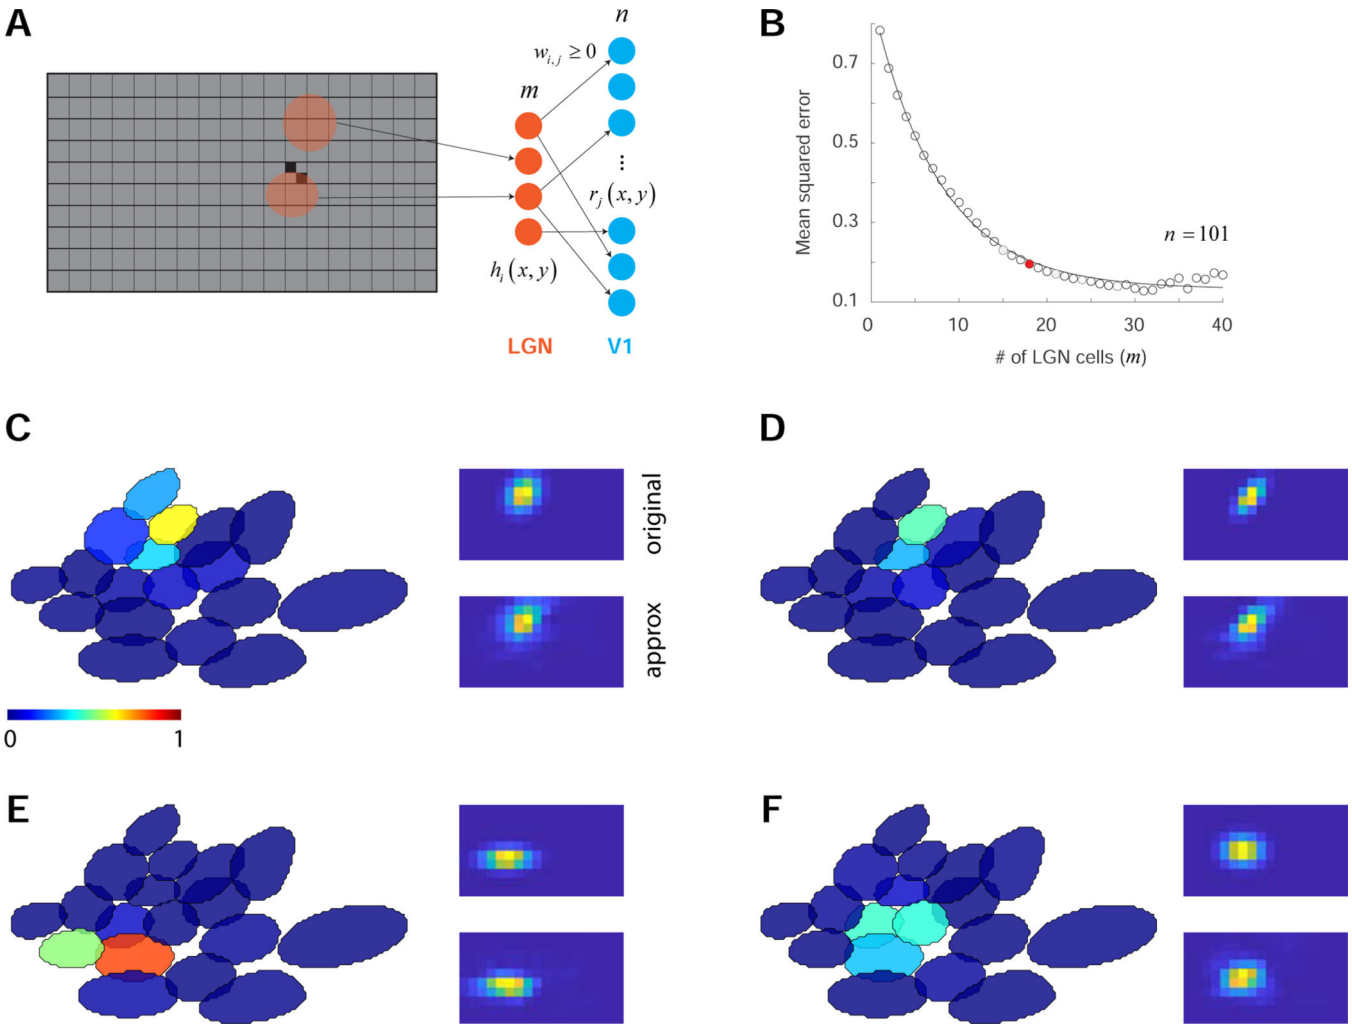

Figure 2. Modeling population receptive fields by non-negative matrix factorization.

A. We model the population of V1 receptive fields as a non-negative, linear combination of LGN inputs. Shaded areas represent the receptive fields of putative LGN neurons. **B.** The mean-squared error of the non-negative factorization falls off with the number of putative LGN inputs. In this experiment, a total of 101 receptive fields could be explained with only 18 inputs (red point). **C-F.** Examples of the reconstruction of V1 receptive fields from the linear combination of LGN inputs from the same population as in **B**. In each case, the left panels show the 1σ level-sets of the LGN receptive fields estimated by the factorization. The color code indicates the fraction of the total synaptic weight of their projections to a cortical cell. The right panels show the measured receptive fields along with its factorized representation. See also Figure S1.

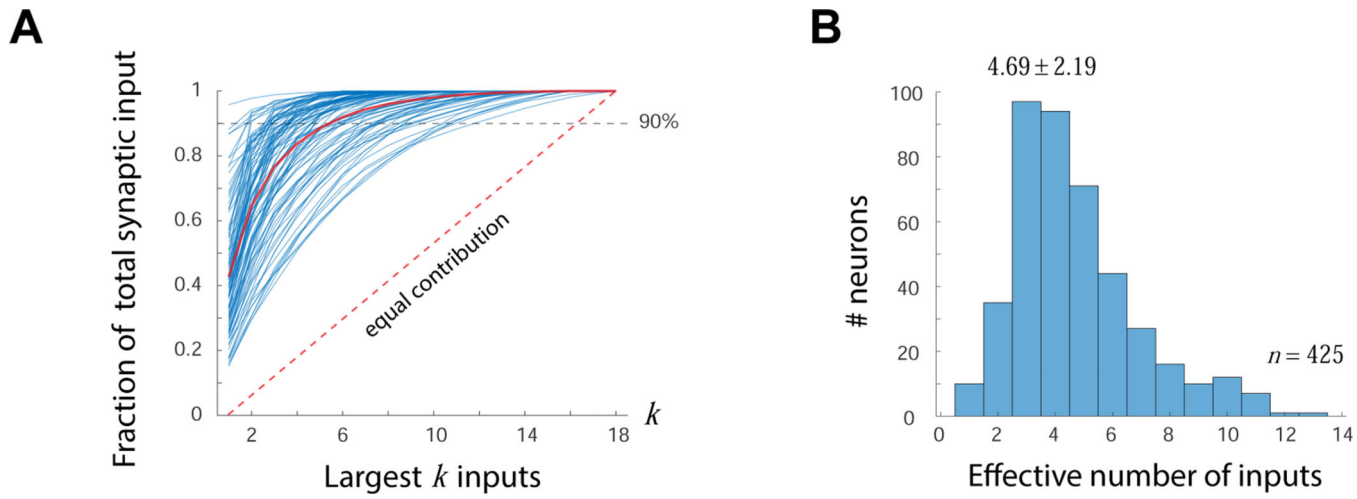

Figure 3. Sparse connectivity of thalamocortical projections.

A. Fraction of total synaptic weight explained by the largest k inputs. Blue curves represent individual cells. The red curve is the average behavior across population. For each cell, the effective number of inputs is defined as the minimum number required to exceed 90% of the total input (dashed horizontal line). For reference, note that if cells had inputs with equal contributions the result would be the red, dashed line. **B.** Distribution of the effective number of total inputs across all neurons in all the experiments. See also Figure S1.

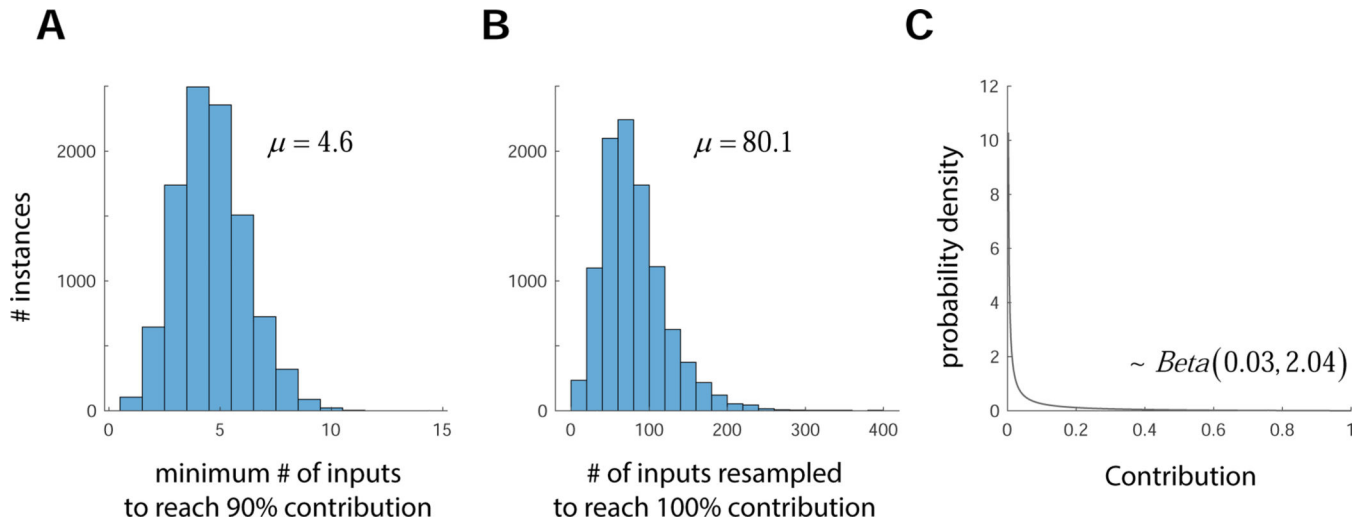

Figure 4. A statistical model reconciles different estimates of the number of inputs.

By assuming a specific multivariate Dirichlet distribution of contributions we can simultaneously match the minimum number of inputs required to explain 90% of the response in the present study (**A**) and the number of inputs estimated by resampling (**B**) as done by Lien and Scanziani [7]. The distribution of contributions under such a model is a Beta distribution (**C**) (the marginal of the Dirichlet distribution) where small values are frequent, while large values are rare.