Modeling Maximum number of addressable targets per second.

This supplementary note describes simulations to estimate the maximum number of light-evoked spikes that can be driven per second under the experimental conditions used in this study. Simulations are based on empirical data, thus representing the current state of the art using our approach. Further advances in opsins, SLMs, and lasers will change these numbers. Estimates presented below refer to L2/3 pyramidal cells of the mouse barrel cortex, and may not hold true in other classes of neurons or at different depths. A key point is that the constraints of two photon imaging and GCaMP6s (the tools we used in this study to empirically verify photo-activation) make it impractical to experimentally verify these simulated upper bounds. These constraints include the slow temporal dynamics of calcium signals compared to action potentials, the low sensitivity of GCaMP6s to detecting individual action potentials under real world conditions, and the speed of volumetric two photon imaging. Another point is that the total heating of the brain during stimulation places an additional upper bound depending on what amount of heating is deemed acceptable. This stimulation considers the total number of light-evoked spikes that can be driven in one second, which result in less than 2° of heating (Supplementary Figure 22).

The number of cells that can be activated during one second is calculated as the product of the number of cells that can be activated simultaneously (i.e., per SLM frame) by the rate in which different targets can be addressed (i.e., the SLM refresh rate). Both of these values are affected by the specifications of the SLM. In this section, we model two SLMs with different resolution and operating speed: the Meadowlark 512L (512x512 pixels) and the Hamamatsu X10468-03 (800x600 pixels), as well as a theoretical 'Ideal' SLM with a nematic crystal that switches instantaneously and suffers no losses in diffraction efficiency. We show that with a single 5 ms laser pulse of 0.2mW/μm^2 (60mW) we can spike the majority of ST-ChroME expressing neurons (Fig 1l, 3d). With 4.1W of power available at the objective, we would expect to be able to deliver this level of power to over 60 targets simultaneously using the 'Ideal' SLM. Real SLMs suffer from losses in diffraction efficiency. We simulate these losses computationally for holograms containing between 1 and 200 targets placed near the zero order (Supplementary Fig 23a). Using the diffraction efficiency value determined above for both SLMs (calculated for phase masks targeting 50 spots), we expect to have enough power to stimulate 55 or 56 targets simultaneously depending on the SLM used. This is in strong agreement with our demonstrations of being able to drive 50 cells *in vivo* (Fig 7).

In the case of the 'Ideal' SLM that can instantly switch phase masks, the limiting factor that determines the upper bound for the number of addressable targets is only the exposure time needed to activate a cell. For example, a 5ms stimulation (which activated >88% of ST-ChroME cells) would allow a maximal operation of 200Hz leading to ~12,000 addressable targets in a second (68 Targets x 200 Hz x 88% success rate). Each of those targets could be a unique cell with a single evoked spike, or the same cells revisited in arbitrary patterns over a second. Real SLMs, however, take time to transition from one phase pattern to another. The two SLMs modeled take 3.4 and 90 ms respectively to switch patterns, including both the rising and falling phase, according to their manufacturer's specifications (personal communications from Meadowlark Inc, and published materials Hamamatsu). We empirically confirmed these values in the case of the Meadowlark 512L (Supplementary Figure 7). If we were to cease illumination during the transition time between one hologram and the next while maintaining a 5ms stimulation duration, we would slow down the stimulation rate to 119 or 10 Hz respectively; and, including the 12% failure rate, would, in principle, be able to spike ~5800 (55 Targets x 119 Hz x 0.88) or ~520 (56 Targets x 10 Hz x 0.88) targets per second.

A simple way to stimulate faster than 200 Hz with a single SLM would be to use a shorter illumination time. However shorter illumination times run the risk of failing to spike cells. We reasoned that an estimate for the spike probability with short stimulus durations is the fraction of cells whose latency to spike is less than a given hypothetical duration. In other words, if a cell spikes in a certain amount of time, no additional illumination is needed to ensure that cell spikes. As channels take time to close and cells can spike even after the light turns off, this is necessarily an underestimate of the spike probability. Using this logic, we model the expected probability of spiking of ST-ChroME expressing neurons with arbitrarily short pulses using the empirical distribution of latency to spike collected from ST-ChroME-expressing L2/3 pyramidal neurons *in vivo* and *ex vivo* (Supplementary Figure 23b). Furthermore, this model accounts for the fraction of cells (~12%) that could not be activated (Fig 1l).

Modeling shows that while shorter stimulation times allow the stimulation rate to be increased, they generally do not improve the maximum number of spikeable targets (Supplementary Figure 23c-d). The stimulation duration that maximizes the number of evoked spikes per second is consistently around 5 ms ('Ideal' SLM 4.7ms ~11000 Targets, Meadowlark 512L 5.3 ms ~5300 Targets, and Hamamatsu 6.9 ms ~510 Targets, Supplementary Figure 23d). Note that given our underestimates built into the spike probability model, these values are below our initial estimates that assumed a fixed 5ms stimuli with a measured 88% success rate. Nonetheless, depending on the specific application, stimulating faster may be more desirable than achieving the maximum number of targets.

An additional strategy to improve stimulation rate is to illuminate continuously, including during the time in which the SLM phase pattern is being updated (Supplementary Figure 23e). However, it is possible this approach risks degraded diffraction efficiency or off-target illumination. We characterized the illumination pattern during high speed phase transitions (using the Meadowlark 512L), via a stroboscopic imaging technique (see methods, Supplementary Figure 7). By assembling images acquired at various steps in the hologram transition, we were able to reconstruct and visualize the pattern of illumination at high spatial and temporal resolution. We show that during transitions, the diffraction efficiency is significantly lower, but light is directed to the zero order where it is blocked and does not appear as off-target illumination (Supplementary Figure 7, Supplementary Movie 1).

We approximate this loss of diffraction efficiency as a decrease in power over the course of the transition time. Thus, short stimuli are not power efficient, but can operate at much faster stimulation rates (Supplementary Figure 23f). This creates a tradeoff between stimulation speed and total number of evoked spikes, which the experimenter can select to suit the demands of a specific experiment (Supplementary Figure 23 c-d). The net effect is a ~12% increase in total evoked spikes (best stimuli for Meadowlark 512L is 5.45 ms exposure with continuous illumination, resulting in ~5800 evoked spikes, Supplementary Figure 23d). While in theory this approach could be pushed to arbitrarily fast frame rates at the expense of simultaneously addressable targets, the meadowlark SLM's drivers have a maximum operating speed of 350 Hz, and we have not attempted beyond 300Hz stimulation. This continuous illumination strategy will result in more energy being delivered to the system as a whole. However, any light lost during transitions will be directed to the zero order, where it is blocked and will not reach the animal (see Supplementary Figure 22 for estimates of induced heat).

Next, we examined how efficiently each of the SLMs are able to direct light to different locations in space. All SLMs become less efficient as you direct light to the limit of their range. Theoretically, SLMs with higher pixel counts should be able to direct light more efficiently to wider areas. We measure the diffraction efficiency for a 1000 individual holograms throughout a 550 x 550 x 150 µm volume (as described[7]) for both SLMs and fit a smooth function (X and Y axes are averaged for display purposes). We then simulate many holograms with 50 targets randomly distributed within a range and calculate

the overall diffraction efficiency (Supplementary Figure 23g). Given this relationship, we can calculate how many targets could be spiked within bounding boxes of different sizes (Supplementary Figure 23h). While both SLMs exhibit decreases in efficiency when sampling from a larger range, at the maximum field of view we image (550 x 550 x 150 µm) we would still be able to spike >3400 unique targets per second with the Meadowlark 512L, or >370 unique targets with the Hamamatsu SLM.

Finally, we acknowledge that in many cases it may be difficult to find thousands of opsin positive cells in a field of view, and potentially even more challenging to detect one added spike in each of those cells with conventional imaging approaches. However, we note that this approach works equally well to target each cell only once in a second as well as to revisit the same cell or cells many times, in a pattern limited by the frame rate. As with many techniques, as SLM and laser technology improves so too will the total number of addressable cells. Improvements in the speed and diffraction efficiency of SLMs will help stimulation systems approach the theoretical 'Ideal' SLM; whereas more powerful lasers will increase the maximum number of simultaneous targets for any SLM, but will be limited by the heat tolerance of the brain.