Classification:

Biological sciences (Neuroscience)

Title:

A synthetic likelihood solution to the silent synapse estimation problem

Authors:

Michael Lynn[a], Kevin F.H. Lee[a], Cary Soares[a], Richard Naud[a,e] & Jean-Claude Béïque[a,b,c,d]

Affiliations:

[a] Department of Cellular and Molecular Medicine, [b]Canadian Partnership for Stroke Recovery, [c] Centre for Neural Dynamics, [d] Brain and Mind Research Institute,

University of Ottawa, Ottawa, Canada K1H 8M5

[e]Department of Physics, STEM Complex, room 336, 150 Louis Pasteur Pvt., University of Ottawa, K1N 6N5, Ottawa, Canada.

Corresponding author:

Jean-Claude Beique

University of Ottawa

451 Smyth Road, RGN 3501N

Ottawa, Ontario, Canada
K1H 8M5

Tel: 1-613-562-5800 ext. 4968
jbeique@uottawa.ca

Keywords: silent synapses, plasticity, statistical inference, synthetic likelihood

Abstract

Functional features of populations of synapses are typically inferred from random electrophysiological sampling of small subsets of synapses. Are these samples unbiased? Here, we developed a biophysically constrained statistical framework for addressing this question and applied it to assess the performance of a widely used method based on a failure-rate analysis to quantify the occurrence of silent (AMPAR-lacking) synapses. We thus simulated this method *in silico* and found that it is characterized by strong and systematic biases, poor reliability and weak statistical power. Key limitations were validated by whole-cell recordings from hippocampal neurons. To address these shortcomings, we employed our experimentally constrained simulations to develop a synthetic likelihood estimator that infers silent synapse fraction with no bias and low variance. Together, this generalizable implementation highlights how computational models of experimental methodologies can substantially improve inference.

Significance statement

The plasticity potential of neural networks is influenced by the presence of silent (AMPA-lacking) synapses. Silent synapses are more prevalent during early postnatal development, yet are found in adults where drugs of abuse are thought to regulate their occurrence. Current methods estimate this global fraction by recording from small subsets of electrically stimulated synapses. We show that a widely used failure rate-based calculation returns highly variable, positively biased estimates of silent synapse fraction. We developed an alternate estimator employing simulated likelihood functions which exhibits dramatically decreased bias and variance. This method allows the reliable interrogation of fine-grained changes in the functional make-up of synaptic populations, and exemplifies how computational models of methodologies can be adapted into statistical inference tools.

Introduction

Activity-dependent synaptic plasticity has captured global research attention for well over four decades due to its hypothesized role in learning and memory (1). Silent synapses contain only NMDA receptors, not AMPA receptors, and represent preferential sites for receptor insertion during long-term potentiation (LTP; 2, 3). Thus, their relative proportions in neural circuits is believed to be a fundamental determinant of the plasticity potential of neural networks. Silent synapses are more prevalent in developing networks (4, 5) yet continue to play a role in adult brain circuits for instance in newborn neurons (6), and mature synapses can at times undergo 'silencing' by removal of AMPARs (7). Furthermore, addictive drugs such as cocaine and morphine have been reported to induce *de novo* formation of silent synapse in reward-related pathways of adult animals (8-11). More generally, theoretical work has shown that high fractions of 'potential' synapses (*e.g.* silent) are associated with optimal information storage over a range of conditions (12-14). Together, whereas these observations point to the fundamental role of silent synapses in both developing and mature brain circuits, they also emphasize the corollary challenge of quantifying with precision their occurrence in synaptic populations with experimentally realistic throughput.

While sound methods exist to demonstrate the presence of individual silent synapses (Fig. 1Ai-ii) (2, 3), these binary single-synapse classification methods can only estimate the fraction of silent synapses in a population by pooling across repeated experimental attempts in a minimum stimulation paradigm (4, 5). In contrast, multi-synapse classification methods, such as the failure-rate analysis (FRA) method, have been developed to provide an estimate of the fraction of silent synapses for every attempt: here an experimenter records from a small, but unknown, number of synapses simultaneously, and the FRA formalism estimates the proportion of these which are silent by means of a comparison of failure rates over many sweeps (Fig. 1Bi-ii) (methodology reviewed in 15). A number of studies have used this method to quantify changes in silent synapse fraction, particularly in rodent models of drug addiction, but also during critical periods of plasticity in the visual cortex and hippocampus (8-11; 16-21, 32). Yet, although FRA seeks to scale up the throughput of silent synapse estimation, its scalability hinges on its reliability and accuracy (especially compared with traditional binary classification methods), which have not been investigated with targeted experiments and statistical analyses.

Here, we first employ a combination of biophysically constrained numerical simulations and electrophysiological recordings to examine the performance of existing methods for silent synapse detection and quantification. Our analysis reveals substantial and general limits on the reliability and validity of the silent synapse fraction estimates obtained using the FRA formalism. To overcome these inherent limitations, we build on approaches developed in the dynamical systems field and employed our computational model to generate an approximate maximum-likelihood estimator, which provides dramatic improvements in estimator bias and variance. A power analysis reveals that the enhanced estimator can resolve fine-grained changes in silent fraction with experimentally feasible sample sizes. Our estimator does not require a novel set of observations, and can thus be deployed on existing FRA datasets. These tools and findings make large scale interrogation of changes in silent synapses possible and reliable.

Results

While minimum stimulation experiments have been critically important for demonstrating the existence of silent synapses, they do not formally allow estimation of the true silent synapse fraction (s_t) contained in a synaptic population. A simple experimental technique, the failure rate analysis (FRA), has been developed to provide a quantitative estimate ($\hat{s}_{FRA}$) of the true underlying silent fraction, s_t. In this method, the rate of synaptic failures at two holding potentials (-60 mV and +40 mV) are used as inputs to a mathematical equation (see Supplemental Methods for derivation), which returns an estimated fraction of silent synapses. Despite this formalism, it is unclear how biophysical variables such as stochastic neurotransmitter release and the inherent variability in the number of synapses stimulated shape the statistical power and discriminability of this method. Here, we first use a set of increasingly constrained and experimentally realistic numerical simulations and statistical tools to examine the performance and general usefulness of this estimator.

A. Nonsilent synaptic populations and FRA estimator performance

As a first step, we directly determined the extent to which the stochasticity of synaptic release affect the behavior of the FRA estimator for a typical electrophysiological recording. We thus numerically simulated the simplest scenario where a single, non-silent synapse with probability of release Pr = 0.5 is electrically stimulated (Fig. 1Ci). Henceforth we will refer to non-silent synapses as active synapses (i.e., synapses containing both AMPA and NMDARs). We simulated Bernoulli-distributed trials (failures and successes) in 50-sweep epochs. The FRA technique, applied to a single, active synapse (i.e., $s_t = 0$) should return $\hat{s}_{FRA} \approx 0$. Fig. 1Cii depicts one example simulation where an estimate of $\hat{s}_{FRA} \approx 28.6$ was returned. The distribution of $\hat{s}_{FRA}$ across multiple experiments (Fig. 1D) was strikingly broad (S.D. = 31.3%) and non-Gaussian, showing high kurtosis (Fisher's kurtosis = 3.00) and a prominent left tail (Fisher-Pearson coefficient of skewness = -1.01). These attributes therefore make the FRA estimator unsuitable for traditional parametric tests that assume normality. Furthermore, 45.3% of the $\hat{s}_{FRA}$ distribution falls below zero, which are biologically irrelevant estimates. Since negative silent synapse values are conspicuously absent from published datasets using FRA (8-11, 16-21), it is possible that these negative values returned from the equation were arbitrarily set to a biologically relevant value of 0% silent (i.e., systematically 'zeroed'). This procedure would artificially inflate the mean silent synapse estimate.

We next asked whether the FRA estimator's accuracy could be improved by simply increasing the number of sweeps from the recommended value of 50 (15) (Fig. 1E). Increasing the number of sweeps per epoch to 500 still produced a remarkably variable estimate (Fig. 1E inset; 500 sweeps: S.D. = 9.1%). Alternatively, recordings from multiple synapses may improve the reliability of $\hat{s}_{FRA}$, due to the intuitive averaging out of stochastic binary events. We thus numerically simulated FRA experiments on small numbers of synapses, first analytically calculating a fixed Pr for all synapses such that the mean failure rate was 50% (as per standard experimental practice for this approach; Fig. 1F). Surprisingly, $\hat{s}_{FRA}$ variability was virtually identical over the entire range of active synapse numbers (Fig. 1G, inset). Therefore, in the simple case of non silent populations, the variability of the FRA estimates is not markedly improved when monitoring multiple synapses.

Figure 1: The FRA technique returns high-variance estimates under a variety of conditions with no silent synapses present. **Ai:** Single-synapse classification techniques rely on electrical (left) or optical (right) stimulation of a single synapse. **Aii**: The responses (here averaged electrical traces), are classified as successes and failures for the -70mV condition (left) and the 50mV condition (right). **Bi:** In contrast, the multi-synapse failure rate analysis method relies on stimulating multiple synapses and (**Bii**) comparing failure rates at two potentials to estimate silent fraction. **Ci:** Depiction of a single active synapse with Pr = 0.5 being measured by FRA, yielding simulated experiments (**Cii**) which each return one estimate. **D**: The distribution for FRA-estimated silent fractions was strikingly broad for one active synapse. **E**: FRA estimate distributions narrow slightly as the number of sweeps per voltage condition increase (inset, colors indicate sweep num.). **F:** Depiction of multiple active synapses measured by FRA, where Pr was scaled to give a failure rate of 0.5. **G:** The standard deviation of the estimate distribution was not dependent on the number of active synapses stimulated (inset, colors indicate sweep num.).

We considered the generalizability of our findings by asking how $\hat{s}_{FRA}$ estimates varied over a wide Pr range and ensemble synapse number (Fig. S1 and Supplemental Information). Our analysis revealed that the FRA estimator variance, surprisingly, is inversely proportional to synapse determinism (i.e., as $Pr \rightarrow 0^+$ and as $Pr \rightarrow 1^-$), and does not significantly decrease with any combination of Pr and synapse number. Taken together, these findings outline the inability of the FRA method to return reliable (i.e., low-variance) estimates of silent synapse fraction in simple conditions with no silent synapses present ($s_t = 0\%$), in most experimentally plausible conditions.

B. Silent synapse-containing synaptic populations and FRA estimator performance

Under conditions where silent synapses are present ($s_t > 0\%$), the quantitative accuracy of the FRA method is naturally dependent on the particular arrangement of synapses sampled. It is not intuitively obvious whether traditional electrical stimulation selects an unbiased subsample of a larger synaptic population, either in terms of silent fraction or of individual synapse properties such as release probability. Therefore, before assessing the accuracy of the FRA equation, we first developed an obligatory intermediate method to simulate experimentally feasible 'draws' of synapse sets, with varying release probabilities, from larger mixed populations (Fig. 2A).

Briefly, we started with initial sets of 100 synapses with some true fraction silent (s_t). Release probabilities were uniformly distributed ($Pr \sim U(0,1)$). We next simulated as faithfully as possible a typical FRA electrophysiological experiment. Thus, the initial set of activated synapses was subjected to multiple rounds of stochastic synapse elimination (20% per round) corresponding to a decrease in electrical stimulation intensity that an experimenter typically performs during an FRA experiment. The rounds of elimination were ended when the failure rate of active synapses reached between 20% and 80%, as dictated by described experimental guidelines (15). Synapse sets were discarded if they did not reach this criteria before the elimination of active synapses. The simulation thus returned stochastic, small subpopulations of silent and active synapses, each with an associated release probability, and with the failure rate range as the main constraint.

To assess how the subpopulation sampling varied with s_t, we performed Monte Carlo simulations with 10 000 draws per condition. We found, surprisingly, that the number of active synapses in each subpopulation reaching the selection criterion was very small - typically between 1 and 5 (Fig. 2Bi). As one may intuit, large numbers of synapses in a draw are rare because they require very low Pr for every synapse to achieve the desired failure rate (Fig. 2Bii). This experimentally-derived selection criterion therefore leads to a bias towards sampling synapses with low Pr. The number of silent synapses in the draws scaled with s_t, as expected (Fig. 2Ci). However, the sampled subsets were intriguingly enriched in silent synapses compared to the ground-truth population (Fig. 2D). This bias arises from the fact that the failure-rate criterion only depends on active synapse monitoring: therefore, the last elimination round ineluctably eliminates an active, rather than a silent, synapse. The number of active synapses is therefore under-represented in the sample, which leads to an overestimation of the silent fraction. Given that the small size of the sampled synaptic population (1-5), this systematic bias is consequential. This appears to be a general feature of the FRA method when stimulation amplitude is adjusted to reach a particular failure rate.

Figure 2: Experimentally bounded simulations reveal a fundamental bias inherent in bulk electrical stimulation. **A:** Schematic depicting simulation of synapse set choice according to failure-rate criteria. **Bi:** Probability distribution of number of active synapses obtained during each bulk electrical stimulation experiment. **Bii:** Release probability of active synapses as a function of number of active synapses. Small dots denote mean release probability of each simulated subsample, while large dots denote the mean across all subsamples. **C:** Probability distribution of number of silent synapses obtained during each bulk electrical stimulation experiment. **D:** Sampled fraction silent is shown against the true fraction silent. **Ei**: Depiction of returned samples of silent/active synapses. **Eii**: Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction procedure. **Eiii**: Distribution of FRA-estimated silent fraction for two ground truth silent fraction values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii**: FRA estimator standard deviation for non-zeroed and zeroed distributions.

We next considered the performance of the FRA estimator under this experimentally realistic scenario. Individual simulations returned highly variable synapse sets (Fig. 2Ei) and inaccurate estimates ($\hat{s}_{FRA}$) (Fig. 2Eii). Over multiple simulated experiments, the estimated silent synapse fraction showed distinctly non-normal distributions (Fig. 2Eiii). We next undertook a formal analysis of the FRA estimator's bias and variance analysis. The bias of $\hat{s}_{FRA}$ increased with the true silent fraction, s_t, and this occurred irrespective of whether negative estimates were erroneously set to zero ('zeroing') or not (Fig. 2Fi). The standard deviation of the estimate was high across all silent fractions (Fig. 2Fii), reaching a maximum of 40% with zeroing applied. We conclude that the FRA estimator displays a high positive bias and large variance, and thus is a flawed estimation technique.

While we drew the release probability of individual synapses from a uniform distribution following previous work assessing Pr from iGluSnFR-mediated single-synapse release events (22), we additionally asked whether different release probability distributions of a synaptic population would influence the FRA estimator's performance. First, we considered the gamma distribution (parameters: $\lambda = 5.8; n = 3$) reported from staining-destaining experiments (Fig. 1G in ref. 23). This distribution produced negligible changes in both the synapse selection bias, and in the bias and variance of the estimator (Fig. S3). We additionally simulated the case with extremely small release probabilities ($\lambda = 5.8$, $n = 1$) (Fig. S4). Although the mean number of active synapses rises from ~3 to ~10 (Fig. S4Bi), the estimator with zeroing applied still shows comparable bias (27% bias at 50% silent, compared to 35% bias at 50% silent for a uniform distribution) and variance (peak standard deviation of 30%) (Fig. S4Fi-ii). These results demonstrate the significant propensity of the FRA estimator towards large positive biases and large variances over a wide range of possible synaptic release probability distributions.

C. Electrophysiological verification of numerically simulated FRA estimate inaccuracies

We next experimentally assessed the propensity of the FRA approach to provide variable estimates of silent synapse fraction by performing voltage-clamp recordings from CA1 neurons from p14-15 mice while electrically stimulating Schaffer's collateral synapses (Fig. 3A; see Methods). As per recommended practice for typical FRA experiments, we tuned the stimulation intensity such that we obtained clear monosynaptic currents with ~50% failure rates (Fig. 3B,C). We performed FRA experiments in triplicate for each recording from individual CA1 neurons (Fig. 3D), and we verified that the evoked EPSC amplitudes at both -70mV and +50mV were stable between these successive iterations (Fig. S5A), confirming that we were stimulating the same subsets of synapses during each FRA iteration.

Since we performed the FRA experiments in triplicate, our recordings returned three independent $\hat{s}_{FRA}$ estimates of the same synaptic population. We found that these estimates were highly variable within the same neuron, and often returned biologically irrelevant negative values (Fig. S5B), as predicted from our simulations (Fig. 3E). Together, these results demonstrate that the estimated silent synapse fraction returned by the FRA method on a fixed population naturally fluctuate in agreement with results from our numerical simulations. This inherent lack of precision can be understood as resulting from stochastic neurotransmitter release combined with the fundamental properties of the FRA estimator. Because we cannot directly measure s_t, we are unable to formally compare the bias and variance of the estimator with experimental ground truth, but these results nonetheless demonstrate the strong agreement between our computational model and experimental findings in documenting the fundamental limits of the FRA formalism.

Figure 3: Electrophysiological recordings at the CA3-CA1 synapses agree with simulated FRA distributions. **A:** Experimental schematic illustrating recording from CA1 neurons while stimulating CA3 afferents. **B:** Representative averaged traces showing successes and failures at Vm = +40mV (top) and Vm = -70mV (bottom). **C:** Exemplar experiment showing failure rates and estimated silent fraction. **D:** Triplicate FRA experiment on a single neuron returns highly variable estimates consistent with biophysical limitations of the FRA method. **E:** Three independent FRA estimates (triplicates) for each neuron (circles) and the mean of these estimates (bar). **F:** Histogram of all triplicate estimates (blue) overlaid with simulated histogram from null distribution (0% silent, black). Distributions are significantly different (p = 0.049, Kolmogorov-Smirnov two-sided test).

D. An improved estimation technique based on synthetic likelihood functions

In principle, it should nonetheless be possible to extract reliable information about the true state of a synaptic population from unreliable and biased estimates returned by the FRA method. For instance, if we knew the distributions of failure rates at both hyperpolarized and depolarized potentials for each s_t, we could write expressions for likelihood functions and construct a maximum-likelihood estimator (MLE). The logic of this approach is to first determine the analytical function establishing the probability of an observation given some model parameters, and then to find the model parameter (here the silent synapse fraction) that maximizes this likelihood function. This type of estimator is of particular interest here since MLE minimizes the variance, at least asymptotically (37).

In the present case, however, the likelihood functions are analytically intractable, as they result from a complex combination of the sampling protocol and the stochastic properties of synapses. Some simplifying assumptions could be made to reach a closed-form likelihood function, but this would require assuming that release probability distributions are independent of sampled synapse number, and that the sampled subsets themselves reflect unbiased samples from the population (both assumptions are shown by our model to be false). Alternatively, popular approaches for noisy dynamical systems sidestep analytical solutions and rather construct synthetic likelihood functions from Monte Carlo simulations (33-35). Here, the dynamical system is iteratively simulated and summary statistics are extracted, allowing to likelihood-based estimation from a stochastic model. These likelihood function are approximations only in the sense that they are based on a finite number of stochastic simulations, but will converge to the true likelihood function. We use this formalism to construct an estimator, FRA-MLE, from synthetic likelihood functions (Fig. S6A-B) where the summary statistic is the biased FRA estimate itself ($\hat{s}_{FRA}$) (See Supplemental Methods for more details). Since our noisy sampling model already incorporates technical choices made by experimenters (e.g. the range of acceptable failure rates), this allows us the flexibility to construct estimators, $\hat{s}_{FRA-MLE}(\boldsymbol{\theta})$ which are a function of an experimental parameter vector $\boldsymbol{\theta}$ containing methodological variables and synaptic properties.

In contrast with the large bias found for the FRA estimator, we found that $\hat{s}_{FRA-MLE}$ is an entirely unbiased estimator (Fig. 4A). This critically fixes the inherent $\hat{s}_{FRA}$'s overestimation of silent fraction (Fig. 2Fi), by yielding accurate estimates across the range of possible silent synapse values. Also in contrast with the large variance found for the FRA estimator (Fig 2Fii), we found that although the variance of $\hat{s}_{FRA-MLE}$ increases with ground truth silent fraction, it remains lower than that of $\hat{s}_{FRA}$ (Fig. 4B; Fig. S6C). Therefore, the new FRA-MLE method circumvents the inherent sampling bias and high variance of the original $\hat{s}_{FRA}$.

As a final test, we applied FRA-MLE to the experimental values collected from hippocampus (Fig. S7). We used FRA-MLE to arrive at a maximally likely estimate, $\hat{s}_{FRA-MLE}$, across all experimental observations. We found that the joint likelihood function was sharply tuned, peaking at a maximal likelihood estimate of 8.7% silent for our experiments (Fig. S7B). This was substantially lower than the estimate obtained by simply taking the mean of all experimental FRA values (16.8%), in agreement with our findings of a pervasive positive bias for FRA. We additionally verified our estimator and numerical simulation approach as follows: We employed our statistical sampling model (*i.e.,* successive rounds of synapse elimination) to construct a hypothetical experimental distribution of FRA estimates for the maximally likely silent fraction $\hat{s}_{FRA-MLE} = 8.7\,\%$ obtained for our experiments. We found that this

Figure 4: Statistical properties of FRA-MLE estimator compared with other methods. **A:** FRA -MLE (blue) and FRA (green) estimator bias. **B**: FRA-MLE (blue) and FRA (green) estimator standard deviation. **C:** Minimum samples required for each estimator to discriminate silent fractions from a null population where no silent silent synapse is present. Inset depicts a zoomed view of 15-50% silent ($\alpha = 0.05$, $\beta = 0.2$ for all cases).

model-generated distribution was not significantly different from the experimental distribution (Fig. S7C; p = 0.90; KS-test, two-tailed). Together, these findings reiterate the general propensity of FRA to return positively biased estimates, and demonstrate the ability of FRA-MLE to extract not just locally optimal solutions, but solutions which are also in good quantitative correspondence with experimental distributions.

E. A power analysis quantifies performance improvements of the FRA-MLE estimator

To achieve a high-throughput characterization, one needs to take two sets of n samples from either of two separate conditions, and conclude with an experimentally realistic sample size whether the two conditions do not have the same proportion of silent synapses (*i.e.* reject the null hypothesis). We used the framework of power analyses to ask whether our FRA-MLE estimator compares favourably with other methods (FRA or binary classification) in terms of the minimum sample size, n_{min}, required to discriminate silent-containing populations from an active-only population. Depending on the estimator, calculating n_{min} is done by either analytically or numerically calculating the minimum sample size to achieve a threshold statistical power (*i.e.* a threshold false negative rate; see Supplemental Methods). More specifically, since the FRA-MLE method calculates an approximate likelihood function explicitly, we could use the log-likelihood ratio (LLR) to test against our null hypothesis (38), an approach that is not possible with the FRA estimator. We can compare against the binary method (Fig. 1A) either using a LLR as done for the FRA-MLE estimator or with a paired sample test as done with the FRA method. Thus we compare four statistical protocols: FRA estimates using a paired test, FRA-MLE estimates using LLR, binary classification using a paired test and binary classification using LLR.

Figure 4C shows n_{min} against the true silent synapse fraction. As is expected, we found that more samples are required to discriminate smaller differences in silent synapse fractions from a fully active synapse population. Importantly, we found a dramatic improvement in statistical power afforded by the FRA-MLE technique (Fig. 4C) compared to FRA. While the FRA estimator required $n_{min} = 36$ to discriminate populations containing less than 15% silent, FRA-MLE required only $n_{min} = 8$. Averaged across all silent fractions, FRA-MLE achieved an identical statistical power to the FRA estimator with a 3.82-fold reduction in sample size. Binary classification required 1.88-fold more samples than FRA, although when LLR-based hypothesis testing was incorporated, the binary technique now required fewer samples than FRA such that its power almost matched that of the FRA-MLE method. Together, this power analysis shows that the FRA-MLE estimator renders more fine-grained silent fractions detectable with fewer samples than comparable methods.

Discussion

It has been long understood that the plasticity potential of a network is in part determined by the fraction of synapses which exist in a silent state, and the proportion of silent synapses is regulated both throughout development (4-5) and in response to diseases or drugs (8-20). Despite a growing appreciation of the overall importance of silent synapses in mature neural networks, a systematic analysis of quantification methods for this parameter is, to our knowledge, lacking. Here, using statistical inference methods alongside electrophysiological recordings, we have demonstrated the significant bias and variance inherent to the popular failure-rate analysis method of silent synapse quantification. Through a set of

numerically simulated likelihood functions relying on minimal assumptions, we have furthermore proposed and validated an alternate maximum-likelihood estimator which corrects for these flaws. This approach illustrates the advantages of a model-based statistical inference approach.

Our computational model is based on two minimal constraints – first, that the range of failure rates for an acceptable experiment is between 20% and 80%, and second, that release probabilities across a synaptic population are drawn from a uniform distribution (22). It is formally possible that specific constellations of central synapses in the brain may deviate from this assumption. For instance, a population dominated by very low release probability synapses would inevitably result in the sampling of a greater number of synapses, which could in principle increase the statistical power of the method and potentially decrease bias. To address this possibility, we considered the release probabilities derived from staining-destaining experiments in hippocampus that did not follow a uniform distribution (Fig. S3; ref. 23) and found that it induced minimal changes in the mean estimates of sampled synapses. These considerations in effect highlight a strength of the MLE model-based approach in that experimenters can ensure that the statistical model is precisely calibrated to their experiment's methodology (e.g., here, to a specified release probability distribution).

We compared electrical minimum stimulation, historically considered the gold standard for silent synapse detection (2-3) with the FRA method. We found that minimum stimulation required a greater sample size (typically >2x larger) for the same degree of discriminability between populations. A sample size of 100 to discriminate between populations differing by 15% silent may appear restrictive, but methods such as 2P glutamate uncaging can provide a statistically equivalent alternative to electrical minimum stimulation for a binary classification of synapses (i.e., active vs silent), while offering the potential for much faster mapping (i.e., several synapses per neuron) with an additional spatial component. Analysis of the electrophysiological responses induced by localized 2P glutamate uncaging can classify synapses as silent (depolarized response) or active (hyperpolarized and depolarized response (24-25, 29-30). This approach further circumvents the FRA assumption that synaptic release probability is invariant between hyperpolarized and depolarized holding potential - a problematic assumption that is not always valid (e.g., depolarization-mediated suppression from excitation; ref. 41).

Recent work has drawn attention to the importance of silent synapses, suggesting their presence in mature networks and outlining potential regulations by drugs of abuse. This renewed interest emphasizes the need to consider the quantitative accuracy of current methods available to quantify silent synapses. We have provided several lines of evidence that argue that the widely-used FRA technique, in its current form, is a poor tool when fine discriminability is required. Our findings demonstrate an unsuspected inherent high variability and positive bias obtained with this approach, and define constraints on the granularity of the changes that can be expected to be detected with experimentally realistic sample sizes. We have provided an enhanced FRA-MLE technique based on numerical simulations which provide satisfactory correction to two inherent problems of the FRA formalism: high bias and variance. This principled framework should allow for deeper investigation into subcellular and fine-grained differences in silent synapse populations which may underlie network function and dysfunction in disease states. We have further shown that the incorporation of a statistical model, as well as the use of a high-throughput method such as two-photon uncaging, can render the class of single-synapse measurements highly competitive with even the corrected FRA-MLE technique – with the additional strength of subcellular or input-specific interrogation. An understanding of the dynamic network-level changes in silent synapse

fractions will ultimately aid in understanding how the plasticity potential of distributed networks of neurons is established and modified.

Methods

A Python package which implements the optimal estimator, FRA-MLE, is available at https://github.com/micllynn/SilentMLE.

Failure rate analysis equation and methodology

In FRA, experimenters measure synaptic failure rates during electrical stimulation while voltage-clamping a neuron at hyperpolarized and depolarized membrane potentials. The resulting failure rates, $\hat{F}_h$ (hyperpolarized) and $\hat{F}_d$ (depolarized), are utilized as inputs to estimate the fraction silent synapses of some recorded population:

$$\hat{s}_{FRA} = 1 - \frac{\ln \hat{F}_h}{\ln \hat{F}_d} \cdot 100\,\% \tag{1}$$

(2; See Supplemental Methods for full derivation and experimental procedures). The recommended number of sweeps to produce each failure rate estimate is 50-100, and $\hat{F}_h$ should be kept around 0.5 (15).

Numerical simulations

All numerical simulations were performed using custom scripts employing standard functions of technical computing libraries Numpy (39) and Scipy (40) in a Python 3.6 environment. See Supplemental Methods for more details.

Electrophysiology

Hippocampal brain slices from P14-P15 mice (C57/Bl6) were made, and whole-cell recordings were carried out from CA1 neurons using previously described methods while stimulating Shaffer collateral afferents with a monopolar stimulating electrode. See Supplemental Methods for more details.

References

1. Bliss TV, Lomo T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol* 232(2):331–356.

2. Liao D, Hessler NA, Malinow R (1995) Activation of postsynaptically silent synapses during pairing-induced LTP in CA1 region of hippocampal slice. *Nature* 375(6530):400–404.

3. Isaac JTR, Nicoll RA, Malenka RC (1995) Evidence for silent synapses: implications for the expression of LTP. *Neuron* 15(2):427–434.

4. Durand GM, Kovalchuk Y, Konnerth A (1996) Long-term potentiation and functional synapse induction in developing hippocampus. *Nature* 381(6577):71–75.

5. Isaac JT, Crair MC, Nicoll RA, Malenka RC (1997) Silent synapses during development of thalamocortical inputs. *Neuron* 18(2):269–280.

6. Chancey JH, Poulsen DJ, Wadiche JI, Overstreet-Wadiche L (2014) Hilar mossy cells provide the first glutamatergic synapses to adult-born dentate granule cells. *J Neurosci* 34(6):2349–2354.

7. Abrahamsson T, Gustafsson B, Hanse E (2007) Reversible synaptic depression in developing rat CA3 CA1 synapses explained by a novel cycle of AMPA silencing-unsilencing. *J Neurophysiol* 98(5):2604–2611.

8. Huang YH, et al. (2009) In vivo cocaine experience generates silent synapses. *Neuron* 63(1):40–47.

9. Brown TE, et al. (2011) A silent synapse-based mechanism for cocaine-induced locomotor sensitization. *J Neurosci* 31(22):8163–8174.

10. Graziane NM, et al. (2016) Opposing mechanisms mediate morphine- and cocaine-induced generation of silent synapses. *Nat Neurosci* 19(7):915–925.

11. Ma Y-Y, et al. (2016) Re-silencing of silent synapses unmasks anti-relapse effects of environmental enrichment. *Proc Natl Acad Sci USA* 113(18):5089–5094.

12. Brunel N, Hakim V, Isope P, Nadal J-P, Barbour B (2004) Optimal Information Storage and the Distribution of Synaptic Weights. *Neuron* 43(5):745–757.

13. Brunel N (2016) Is cortical connectivity optimized for storing information? *Nat Neurosci* 19(5):749–755.

14. Clopath C, Brunel N (2013) Optimal Properties of Analog Perceptrons with Excitatory Weights. *PLoS Comput Biol*.

15. Graziane N, Dong Y (2016) Measurement of Silent Synapses. *Electrophysiological Analysis of Synaptic Transmission*, Neuromethods. (Springer New York, New York, NY), pp 217–224.

16. Neumann PA, et al. (2016) Cocaine-Induced Synaptic Alterations in Thalamus to Nucleus Accumbens Projection. *Neuropsychopharmacology* 41(9):2399–2410.

17. Huang X, et al. (2015) Progressive maturation of silent synapses governs the duration of a critical period. *Proc Natl Acad Sci USA* 112(24):E3131–40.

18. Ma Y-Y, et al. (2014) Bidirectional modulation of incubation of cocaine craving by silent synapse-based remodeling of prefrontal cortex to accumbens projections. *Neuron* 83(6):1453–1467.

19. Lee BR, et al. (2013) Maturation of silent synapses in amygdala-accumbens projection contributes to incubation of cocaine craving. *Nat Neurosci* 16(11):1644–1651.

20. Koya E, et al. (2012) Silent synapses in selectively activated nucleus accumbens neurons following cocaine sensitization. *Nat Neurosci* 15(11):1556–1562.

21. Rumpel S, Hatt H, Gottmann K (1998) Silent synapses in the developing rat visual cortex: evidence for postsynaptic expression of synaptic plasticity. *J Neurosci* 18(21):8863–8874.

22. Soares C, Lee KFH, Béïque J-C (2017) Metaplasticity at CA1 Synapses by Homeostatic Control of Presynaptic Release Dynamics. *Cell Rep* 21(5):1293–1303.

23. Branco T, Staras K, Darcy KJ, Goda Y (2008) Local Dendritic Activity Sets Release Probability at Hippocampal Synapses. *Neuron* 59(3):475–485.

24. Beique J-C, Lin D-T, Kang M-G, Aizawa H, Takamiya K, Huganir, RL (2006) Synapse-specific regulation of AMPA receptor function by PSD-95. *Proc Natl Acad Sci USA* 103(51):19535-19540.

25. Soares C, Lee KFH, Nassrallah W, Béïque J-C (2013) Differential subcellular targeting of glutamate receptor subtypes during homeostatic synaptic plasticity. *J Neurosci* 33(33):13547–13559.

26. Losonczy A, Makara JK, Magee JC (2008) Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature* 452(7186):436–441.

27. Branco T, Clark BA, Hausser M (2010) Dendritic Discrimination of Temporal Input Sequences in Cortical Neurons. *Science* 329(5999):1671–1675.

28. Zito K, Scheuss V, Knott G, Hill T, Svoboda K (2009) Rapid Functional Maturation of Nascent Dendritic Spines. *Neuron* 61(2):247–258.

29. Lee KFH, Soares C, Thivierge J-P, Béïque J-C (2016) Correlated Synaptic Inputs Drive Dendritic Calcium Amplification and Cooperative Plasticity during Clustered Synapse Development. *Neuron* 89(4):784–799.

30. Busetto G, Higley MJ, Sabatini BL (2008) Developmental presence and disappearance of postsynaptically silent synapses on dendritic spines of rat layer 2/3 pyramidal neurons. *J Physiol* 586(6):1519–1527.

31. Ashby MC, Isaac JTR (2011) Maturation of a Recurrent Excitatory Neocortical Circuit by Experience-Dependent Unsilencing of Newly Formed Dendritic Spines. *Neuron* 70(3):510–521.

32. Beroun A, et al. (2018) Generation of silent synapses in dentate gyrus correlates with development of alcohol addiction. *Neuropsychopharmacology 43:*1989 – 1999.

33. Wood SN (2010) Statistical inference for noisy nonlinear ecological dynamic systems. *Nature* 466(7310):1102–1104.

34. Sunnåker M, et al. (2013) Approximate Bayesian Computation. *PLoS Comput Biol* 9(1):e1002803–10.

35. Diggle PJ, Gratton RJ (1984) Monte Carlo Methods of Inference for Implicit Statistical Models. *J. Royal Stat. Soc Ser B* (Methodological) 46(2):193–212.

36. Geddes SD, et al. (2016) Target-specific modulation of the descending prefrontal cortex inputs to the dorsal raphe nucleus by cannabinoids. *Proc Natl Acad Sci USA* 113(19):201522754–5434.

37. Cramér, H. (1945) *Mathematical methods of statistics* (Almqvist & Wiksells, Uppsala).

38. Kass RE, Eden UT, Brown EN. (2014) *Analysis of Neural Data* (Springer, New York).

39. van der Walt S, Colbert SC, Varoquaux G (2011) The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering* 13(2):22–30.

40. Jones E, Oliphant E, Peterson P, et al. (2001-) *S*cipy: Open Source Scientific Tools for Python. http://www.scipy.org/ [Accessed 2019-09-24].

41. Kreitzer AC, Regehr WG. (2001) Retrograde Inhibition of Presynaptic Calcium Influx by Endogenous Cannabinoids at Excitatory Synapses onto Purkinje Cells. *Neuron* 29:717-727.

Figure captions

Figure 1: The FRA technique returns high-variance estimates under a variety of conditions with no silent synapses present. **Ai:** Single-synapse classification techniques rely on electrical (left) or optical (right) stimulation of a single synapse. **Aii**: The responses (here averaged electrical traces), are classified as successes and failures for the -70mV condition (left) and the 50mV condition (right). **Bi:** In contrast, the multi-synapse failure rate analysis method relies on stimulating multiple synapses and (**Bii**) comparing failure rates at two potentials to estimate silent fraction. **Ci:** Depiction of a single active synapse with Pr = 0.5 being measured by FRA, yielding simulated experiments (**Cii**) which each return one estimate. **D**: The distribution for FRA-estimated silent fractions was strikingly broad for one active synapse. **E**: FRA estimate distributions narrow slightly as the number of sweeps per voltage condition increase (inset, colors indicate sweep num.). **F:** Depiction of multiple active synapses measured by FRA, where Pr was scaled to give a failure rate of 0.5. **G:** The standard deviation of the estimate distribution was not dependent on the number of active synapses stimulated (inset, colors indicate sweep num.).

Figure 2: Experimentally bounded simulations reveal a fundamental bias inherent in bulk electrical stimulation. **A:** Schematic depicting simulation of synapse set choice according to failure-rate criteria. **Bi:** Probability distribution of number of active synapses obtained during each bulk electrical stimulation experiment. **Bii:** Release probability of active synapses as a function of number of active synapses. Small dots denote mean release probability of each simulated subsample, while large dots denote the mean across all subsamples. **C:** Probability distribution of number of silent synapses obtained during each bulk electrical stimulation experiment. **D:** Sampled fraction silent is shown against the true fraction silent. **Ei**: Depiction of returned samples of silent/active synapses. **Eii**: Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction procedure. **Eiii**: Distribution of FRA-estimated silent fraction for two ground truth silent fraction values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii**: FRA estimator standard deviation for non-zeroed and zeroed distributions.

Figure 3: Electrophysiological recordings at the CA3-CA1 synapses agree with simulated FRA distributions. **A:** Experimental schematic illustrating recording from CA1 neurons while stimulating CA3 afferents. **B:** Representative averaged traces showing successes and failures at Vm = +40mV (top) and Vm = -70mV (bottom). **C:** Exemplar experiment showing failure rates and estimated silent fraction. **D:** Triplicate FRA experiment on a single neuron returns highly variable estimates consistent with biophysical limitations of the FRA method. **E:** Three independent FRA estimates (triplicates) for each neuron (circles) and the mean of these estimates (bar). **F:** Histogram of all triplicate estimates (blue) overlaid with simulated histogram from null distribution (0% silent, black). Distributions are significantly different (p = 0.049, Kolmogorov-Smirnov two-sided test).

Figure 4: Statistical properties of FRA-MLE estimator compared with other methods. **A:** FRA -MLE (blue) and FRA (green) estimator bias. **B**: FRA-MLE (blue) and FRA (green) estimator standard deviation. **C:** Minimum samples required for each estimator to discriminate silent fractions from a null population where no silent silent synapse is present. Inset depicts a zoomed view of 15-50% silent ($\alpha = 0.05$, $\beta = 0.2$ for all cases).

<u>Supplemental Figure 1</u>: The FRA equation returns similar results over the parameter space of release probabilities and synapse numbers. **A:** Depiction of simulation with one active synapse. **Bi:** Representative simulations for failures and successes at Vm = -70mV for Pr = 0.05, Pr = 0.5 and Pr = 0.95. **Bii:** Failure rate distributions have standard deviations which are parabolically related to release probability. **Ci:** Representative simulations and silent fraction estimates for full FRA experiments with Pr = 0.05, Pr = 0.5 and Pr = 0.95. **Cii:** In contrast to failure rate, the distributions for estimated silent fraction exhibit similar standard deviations for various release probability values. **Ciii:** Same as **Cii**, but plotting the distributions as cumulative probability densities. **D:** Depiction of simulation with multiple active synapses. **Ei:** Estimate broadness as a function of release probability for various numbers of active synapses. **Eii:** Estimate broadness as a function of failure rate for various numbers of active synapses.

<u>Supplemental Figure 2:</u> The FRA equation is buffered against changes in failure rate. **A:** Representative traces of FRA experiments where the hyperpolarized and depolarized failure rates differ by 10%. **B:** Graphical depiction of the returned FRA estimates, demonstrating log(failure rate) plotted against failure rate.

Supplemental Figure 3: Experimentally bounded simulations with a gamma release probability distribution (23). **A:** Schematic depicting simulation of synapse set choice according to failure-rate criteria. **Bi:** Probability distribution of number of active synapses obtained during each bulk electrical stimulation experiment. **Bii:** Release probability of active synapses as a function of number of active synapses. Small dots denote mean release probability of each simulated subsample, while large dots denote the mean across all subsamples. **C:** Probability distribution of number of silent synapses obtained during each bulk electrical stimulation experiment. **D:** Release probability for all synapses at the start of the simulation was drawn from a uniform distribution. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction procedure. **Eiii:** Distribution of FRA-estimated silent fraction for two ground truth silent fraction values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions.

Supplemental Figure 5: Triplicate experiments at the CA3-CA1 synapse in hippocampus demonstrate high recording stability for both AMPA-R and NMDA-R currents (**A**); and the estimate spread over the triplicate experiments is high and compares well to the numerically simulated estimate spread for three independent iterations of FRA experiments on the same synapse ensemble (**B**).

Supplemental Figure 6: A maximum likelihood estimator for silent synapse estimation. **A**: A numerically simulated likelihood function for a representative single FRA observation, demonstrating the wide spread of ground truths consistent with this observation. **B:** When 20 observations are considered jointly, the likelihood function narrows to a more precise estimate. **Ci**: Estimate distribution for the FRA and FRA-MLE method given a ground truth of 0% silent. **Cii**: Estimate distribution for the FRA and FRA-MLE method given a ground truth of 25% silent.

Supplemental Figure 7: The FRA-MLE method applied to experimental data collected at the CA3-CA1 synapse. **A:** Likelihood function associated with a single experimental observation. **B:** Likelihood function associated with experimental observations from all cells considered jointly, yielding a maximum likelihood estimate of 8.7% silent. **C:** Distribution of experimental silent estimates overlaid with the maximum-likelihood-estimated distribution, simulated numerically for silent fraction = 8.7%. The distributions are not statistically different (p = 0.90; KS-test, two-tailed)

Supplemental figure captions

Supplemental Figure 1: The FRA equation returns similar results over the parameter space of release probabilities and synapse numbers. **A:** Depiction of simulation with one active synapse. **Bi:** Representative simulations for failures and successes at Vm = -70mV for Pr = 0.05, Pr = 0.5 and Pr = 0.95. **Bii:** Failure rate distributions have standard deviations which are parabolically related to release probability. **Ci:** Representative simulations and silent fraction estimates for full FRA experiments with Pr = 0.05, Pr = 0.5 and Pr = 0.95. **Cii:** In contrast to failure rate, the distributions for estimated silent fraction exhibit similar standard deviations for various release probability values. **Ciii:** Same as **Cii**, but plotting the distributions as cumulative probability densities. **D:** Depiction of simulation with multiple active synapses. **Ei:** Estimate broadness as a function of release probability for various numbers of active synapses. **Eii:** Estimate broadness as a function of failure rate for various numbers of active synapses.

Supplemental Figure 2: The FRA equation is buffered against changes in failure rate. **A:** Representative traces of FRA experiments where the hyperpolarized and depolarized failure rates differ by 10%. **B:** Graphical depiction of the returned FRA estimates, demonstrating log(failure rate) plotted against failure rate.

Supplemental Figure 3: Experimentally bounded simulations with a gamma release probability distribution (23). **A:** Schematic depicting simulation of synapse set choice according to failure-rate criteria. **Bi:** Probability distribution of number of active synapses obtained during each bulk electrical stimulation experiment. **Bii:** Release probability of active synapses as a function of number of active synapses. Small dots denote mean release probability of each simulated subsample, while large dots denote the mean across all subsamples. **C:** Probability distribution of number of silent synapses obtained during each bulk electrical stimulation experiment. **D:** Release probability for all synapses at the start of the simulation was drawn from a uniform distribution. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction procedure. **Eiii:** Distribution of FRA-estimated silent fraction for two ground truth silent fraction values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions.

Supplemental Figure 4: Experimentally bounded simulations with a gamma release probability distribution biased towards low values ($\lambda = 5.8$, $n = 1$). **A:** Schematic depicting simulation of synapse set choice according to failure-rate criteria. **Bi:** Probability distribution of number of active synapses obtained during each bulk electrical stimulation experiment. **Bii:** Release probability of active synapses as a function of number of active synapses. Small dots denote mean release probability of each simulated subsample, while large dots denote the mean across all subsamples. **C:** Probability distribution of number of silent synapses obtained during each bulk electrical stimulation experiment. **D:** Release probability for all synapses at the start of the simulation was drawn from a uniform distribution. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction procedure. **Eiii:** Distribution of FRA-estimated silent fraction for two ground truth silent fraction values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions.

Supplemental Figure 5: Triplicate experiments at the CA3-CA1 synapse in hippocampus demonstrate high recording stability for both AMPA-R and NMDA-R currents (**A**); and the estimate spread over the

triplicate experiments is high and compares well to the numerically simulated estimate spread for three independent iterations of FRA experiments on the same synapse ensemble (**B**).

Supplemental Figure 6: A maximum likelihood estimator for silent synapse estimation. **A**: A numerically simulated likelihood function for a representative single FRA observation, demonstrating the wide spread of ground truths consistent with this observation. **B:** When 20 observations are considered jointly, the likelihood function narrows to a more precise estimate. **Ci**: Estimate distribution for the FRA and FRA-MLE method given a ground truth of 0% silent. **Cii**: Estimate distribution for the FRA and FRA-MLE method given a ground truth of 25% silent.

Supplemental Figure 7: The FRA-MLE method applied to experimental data collected at the CA3-CA1 synapse. **A:** Likelihood function associated with a single experimental observation. **B:** Likelihood function associated with experimental observations from all cells considered jointly, yielding a maximum likelihood estimate of 8.7% silent. **C:** Distribution of experimental silent estimates overlaid with the maximum-likelihood-estimated distribution, simulated numerically for silent fraction = 8.7%. The distributions are not statistically different (p = 0.90; KS-test, two-tailed)

Supplemental Information

Derivation of FRA equation: Let n_s represent the number of silent synapses in some ensemble, n_a the number of active synapses, F_h the failure rate at hyperpolarized potentials (n_a synapses active) and F_d the failure rate at depolarized potentials ($n_a + n_s$ synapses active since NMDA-receptor-containing silent synapses can only conduct under depolarized settings). Then:

$$F_h = (1 - \text{Pr})^{n_a} \rightarrow n_a = log_{1-\text{Pr}}(F_h)$$
$$F_d = (1 - \text{Pr})^{n_a+n_s} \rightarrow n_a + n_s = log_{1-\text{Pr}}(F_d)$$

The silent fraction is $s = \dfrac{n_s}{n_a + n_s}$, so:

$$s = 1 - \frac{n_a}{n_s + n_a}$$
$$s = 1 - \frac{log_{1-\text{Pr}}(F_h)}{log_{1-\text{Pr}}(F_d)}$$
$$s = 1 - \frac{\log F_h}{\log F_d}$$

This can be written as a percentage: $s_{\text{pc}} = 1 - \dfrac{\log F_h}{\log F_d} \cdot 100\,\%$.

Behavior of FRA in simple nonsilent contact scenarios: We further explored how $\hat{s}_{FRA}$ varied with probability of release and the number of synapses in the ensemble (Fig. S1). Pr has been shown to vary widely between synapses (22-23) and one might intuitively expect that high-Pr synapses might return more reliable $\hat{s}_{FRA}$ estimates since they release neurotransmitter more deterministically. We first explored this intuition with one active synapse, where Pr of 0.05, 0.5 and 0.95 were simulated. The distributions of failure rate per se do indeed, as one might expect, become narrower as $\text{Pr} \rightarrow 0^+$ and as $\text{Pr} \rightarrow 1^-$ (Fig. S1A-B). However, at these two limits, the standard deviation of $\hat{s}_{FRA}$ paradoxically increases (Fig. S1C) – the estimates become less reliable at very high or very low release probabilities. It is possible that there may be some combinations of synapse numbers and Pr exhibiting low estimate variability. However, by simulating a large number of such combinations, we find that there is a lower bound of ~25% on the standard deviation of $\hat{s}_{FRA}$ (Fig. S1D-Ei). In fact, although the variance of the estimator depends on both the Pr and the number of synapses, the variance of the estimator as a function the failure rate is independent of the number of synapses (Fig. S1Eii). These counterintuitive behaviors can be understood by appreciating the logarithm-transform taking place in the equation for the FRA estimator (Fig. S2). First, under high-Pr conditions, as failure rates approach 0, the natural logarithm of failure rates decreases nonlinearly, which amplifies failure rate variability under high-Pr conditions (Fig. S2B, left). Second, under low-Pr conditions, as failure rates approach 1, the natural logarithm approaches 0 (Fig. S2B, left). Taking the fraction of two stochastic numbers approaching 0, as in the FRA equation, amplifies the variability present. Therefore, due to the fundamental structure of the equation, the variability of the FRA estimates is not markedly improved neither by sampling from more synapses nor by sampling deterministic synapses.

Supplemental Methods

Simulations with active synapses only *(Fig. 1, S1-S2)*. We simulated single active (nonsilent) synapses or groups of active synapses. For each synapse i we assigned a release probability by drawing from a uniform distribution $\text{Pr}_i \sim U(0,1)$.

We numerically stimulated stochastic neurotransmitter release at hyperpolarized voltages (V_h; only active synapses) and depolarized voltages (V_d; active and silent synapses). 50 trials were simulated at each of V_h and V_d. Trials were modelled as Bernoulli trials with failure rate $F = \prod_{i=1}^{m}(1 - \text{Pr}_i)$ over m synapses considered, where $m = \begin{cases} n_a & \text{if } V = V_h \\ n_a + n_s & \text{if } V = V_d \end{cases}$. This yielded $\hat{F}_h$ and $\hat{F}_d$, which are estimates of the true failure rate over 50 trials at V_h and V_d respectively. The FRA calculation was performed as reported in equation (1) and Methods, with $\hat{F}_h$ and $\hat{F}_d$ as the equation inputs. Simulated currents (for purposes of illustration only) were generated by modeling the total current as a sum of all successfully releasing synapses, where each success is treated as generating a stochastic unitary current with an (arbitrary) mean amplitude of 10pA.

Simulations with active and silent synapses *(Fig. 2, S3-S4)*. In this section, we describe our simulations of minimum-simulation electrophysiological experiments, where a larger set of synapses with defined release probabilities and some fraction silent are iteratively reduced into some experimentally feasible subset on which to perform experiments. This iterative reduction models a gradual reduction in electrical stimulation intensity that an electrophysiologist would perform.

Assume the true fraction of silent synapses in a population is s_t. For each replicate, an initial set of silent and active synapse numbers $\{n_s, n_a\}$ were each chosen from a binomial distribution with an initial population size of $n_t = 100$ synapses and s_t silent probability (e.g. the distribution of n_s silent synapses was $p(n_s = k) = \binom{n_t}{k} s_t^k (1 - s_t)^{n_t - k}$). Release probabilities for each synapse were I.I.D and were drawn from either a uniform or gamma distribution as noted in the text, producing a vector of release probabilities $\mathbf{Pr_a} = [Pr_1, Pr_2, \dots, Pr_{n_a}]$ for active synapses and $\mathbf{Pr_s} = [Pr_1, Pr_2, \dots, Pr_{n_s}]$ for silent synapses. The synapse set $\{n_s, n_a\}$ and associated release probability set $\{\mathbf{Pr_s}, \mathbf{Pr_a}\}$ were then subjected to iterative rounds of reduction, which yielded an experimentally feasible subset.

Iterative reduction of the synapse sets proceeded until the failure rate criterion was met: $0.2 < F_{-70} < 0.8$. Here, $F_{-70} = \prod_{i=1}^{n_a}(1 - \text{Pr}_i)$ where $Pr_i \in \mathbf{Pr_a}$. F is analytically calculated at each iteration, rather than estimated using Monte Carlo simulations, for purposes of computational efficiency. Iterative reduction consisted of the following sequential steps: 1) Using $p_{elim} = 0.2$ as the elimination probability, stochastically eliminate $p_{elim}(n_s + n_a)$ synapses along with their associated release probabilities in $\{\mathbf{Pr_s}, \mathbf{Pr_a}\}$. 2) After each elimination, calculate F_{-70}. 3). If F_{-70} falls within criterion range, then store $\{n_s, n_a\}$ and $\{\mathbf{Pr_s}, \mathbf{Pr_a}\}$. If F_{-70} does not fall within criterion range, loop back to 1) and repeat.

Together, this iterative decrease yielded Monte Carlo-simulated subsamples (replicates) of active and silent synapses along with associated release probabilities $\{\mathbf{Pr_s}, \mathbf{Pr_a}\}$. For each replicate, we used $\{\mathbf{Pr_s}, \mathbf{Pr_a}\}$ to simulate 50 Bernoulli trials at each of $V_m = -70mV$ (n_a synapses) and $V_m = +40mV$ ($n_a + n_s$ synapses) similar to above, yielding stochastic failure rate estimates, $\hat{F}_{-70}$ and $\hat{F}_{+40}$. These

estimates were used as input to the FRA estimator, to return one silent fraction estimate, $\hat{s}_{FRA}$, per Monte Carlo simulation.

For each ground truth silent value, s_t, we obtained ~20 000 replicates which each yielded an estimate, $\hat{s}_{FRA}$. These numerically simulated distributions formed the basis of our bias and variance analysis the FRA estimator. For each s_t, we can then compute:

$$\text{Bias}(\hat{s}_{FRA}, s_t) = \mathbf{E}(\hat{s}_{FRA}) - s_t$$
$$\text{Var}(\hat{s}_{FRA}) = \mathbf{E}[(\hat{s}_{FRA} - \mathbf{E}(\hat{s}_{FRA}))^2].$$

<u>Experimental measurement of silent synapse estimate variability in hippocampus (Fig. 3)</u>. Hippocampal slices (300um thick) were made from P7-14 wild-type (C67/Blk) mice using a previously described technique (36). In brief, mice were anesthetized using isofluorane, decapitated, and the brain was extracted and placed in ice-cold choline-based cutting solution for the duration of the cut. Cutting solution was comprised of 119 mM choline-Cl, 2.5 mM KCl, 1 mM CaCl2, 4.3 mM MgSO4-7H2O, 1 mM NaH2PO4, 1.3 mM sodium L-ascorbate, 26.2 mM NaHCO3, and 11 mM glucose, and was equilibrated with 95% O2 and 5% CO2. 300uM coronal sections were made of the hippocampus. Slices were recovered in a chamber containing standard Ringer's solution of the following composition: 119 mM NaCl, 2.5 mM CaCl2, 1.3 mM MgSO4- 7H2O, 1 mM NaH2PO4, 26.2 mM NaHCO3, and 11 mM glucose, at a temperature of 37 °C, continuously bubbled with a mixture of 95% O2 and 5% CO2. Slices were recovered for 1 h in the recovery chamber and equilibrated to a temperature of ~25 °C until recordings were performed.

Neurons were visualized in the CA1 subfield using an upright microscope (Examiner D1 equipped with differential interference contrast (DIC) (Zeiss; 40×/0.75 N.A. objective). Using 4-6mOhm fire-polished glass electrodes (Sutter Corporation), CA1 subfield neurons were recorded in the whole-cell patch-clamp configuration at room temperature (20-22C) with a cesium-based internal comprised of: [Get cesium internal recipe and put here]. As an external solution, Ringer's solution was continuously circulated. Whole-cell recordings were carried out using an Axon Multiclamp 700B amplifier, sampled at 10 kHz, digitized with an Axon Digidata 1440A (or 1550) digitizer, and filtered at 2 kHz. An electrical stimulation electrode was placed apposed to the apical dendritic arbor to activate Schaffer collaterals. Stimulation intensity was tuned such that ~50% synaptic failures and successes each were recorded.

The FRA protocol consisted of 100 sweeps of electrical stimulation (0.1ms pulse width; 10-15 seconds sweep-to-sweep time). 50 sweeps were performed at $V_m = -70mV$ and 50 sweeps were performed at $V_m = +40mV$. This FRA protocol was performed in triplicate on single neurons, with FRA iterations immediately following one another. Recordings were discarded when the access resistance changed by >30%. Liquid junction potential was not compensated for. The data was analyzed offline with the aid of a custom Python module written by the author and available online at: <u>https://github.com/micllynn/synappy</u>. Synaptic failures and successes were determined in an automated manner by setting a threshold for the baseline-subtracted electrically evoked current to exceed to be considered a success. The FRA calculation was performed as indicated in equation (1), by comparing the failure rate at -70mV and +50mV. Mean estimate spread was calculated by subtracting the minimum FRA-estimated silent synapse fraction from the maximum FRA-estimated silent synapse fraction within the same cell.

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services.

<u>A maximum likelihood estimator derived from synthetic likelihood functions</u> (Fig. 4, S6, S7). For a jointly observed pair of failure rates at hyperpolarized and depolarized potentials, $\hat{F}_h$ and $\hat{F}_d$, we wish to estimate the most likely s_t. If we had an analytically tractable likelihood function, we could use likelihood-based inference techniques (e.g. maximum likelihood estimation) to find our most likely s_t. However, two factors make it difficult to analytically derive a likelihood function: first, the number of active and silent synapses in each sampled subset are only probabilistically related to s_t; and second, each synapse has release probabilities Pr which are a complex function of the starting population's Pr and a bias imposed by the number of other (active) synapses present in that particular subset.

To circumvent these obstacles, we use a statistical inference approach inspired by recent treatments of noisy dynamical systems (33). While each individual simulation of synaptic release trains contains a large amount of information, we can extract summary statistics for each simulation and construct a synthetic likelihood function yielding the probability of obtaining these summary statistics at each underlying value of s_t. For convenience, we choose the biased estimator $\hat{s}_{FRA}$ as the summary statistic as it contains information about both $\hat{F}_h$ and $\hat{F}_d$. Our synthetic likelihood function can therefore be written in terms of the summary statistic: $p(\hat{s}_{FRA} = d \mid s_t, Pr) = \dfrac{n_d^{s_t}}{n_{tot}^{s_t}}$, where d is a given value of observed data, $n_d^{s_t}$ is the count of Monte Carlo-simulated replicates using the parameters s_t, Pr which satisfy $\hat{s}_{FRA} = d$ (i.e. the observation is obtained), and $n_{tot}^{s_t}$ is the total number of Monte Carlo-simulated replicates using the parameters s_t, Pr. As above, s_t refers to the true silent fraction and Pr denotes a synaptic release probability distribution.

When more than one piece of data is acquired, is convenient to write a joint log-likelihood function over a given vector of observations $\mathbf{D} = (d_1, d_2, \ldots, d_n)$. We can express this as $\ell(s_t \mid \mathbf{D}) = \sum_{i=1}^{n} \log[p(d_i \mid s_t)]$. We can then construct an MLE estimator for s_t given some vector of measured data $\mathbf{D}$, expressed as $\hat{s}_{FRA-MLE}(\mathbf{D}) = \underset{s_t \in [0,1]}{\operatorname{argmax}}(\ell(s_t \mid \mathbf{D}))$. This estimator's bias and variance was evaluated in an identical way as the FRA estimator (above), and the MLE estimator was found to exhibit no bias and significantly lower variance.

Two caveats to this approach: First, note that instead of constructing $p(\hat{s}_{FRA} \mid s_t)$ point-by-point with some step size $\delta s_t \to 0$, we could conceivably model the likelihood function as a continuous probability density over all values of s_t, by estimating parameters of a normal distribution similar to (33). Here, one could compute $\mu_{s_t} = \mathbf{E}[\hat{s}_{FRA}]$ and $\sigma_{s_t}^2 = \mathbf{E}[(\hat{s}_{FRA} - \mathbf{E}[\hat{s}_{FRA}])^2]$ for each s_t. However, we have established that the distribution of $\hat{s}_{FRA}$ for each s_t has nonzero skewness and kurtosis, making it ineligible for normal distribution fitting.

Second, while we chose one summary statistic $\hat{s}_{FRA}$ to construct a synthetic likelihood function, it is also possible to directly use the failure rates $\hat{F}_h$ and $\hat{F}_d$ as two summary statistics. This would appear to be advantageous because we are using two summary statistics instead of one. However, this does not provide better estimation capabilities because the distributions for $\hat{F}_h$ show no relationship with s_t in our simulations. Therefore, we chose a single summary statistic $\hat{s}_{FRA}$.

<u>Power analysis of estimators for s_t</u> (Fig. 4). We employed a power analysis approach to compare the statistical power of different techniques for silent synapse detection. Parameters for power analyses were $\alpha = 0.05$ (false positive rate) and $\beta = 0.2$ (false negative rate) unless otherwise noted.

Given two populations of synapses, P_1 and P_2, having distinct silent fractions $s_t^{(1)}$ and $s_t^{(2)}$ silent fractions, respectively. We want to determine n_{min}, the minimum number of samples required to discriminate these populations with some estimation technique. To do this, we wish to express statistical power, p, as a function of the number of samples: $p(n)$. We can then find an n_{min} such that $p(n > n_{min}) < \beta$. Our approach to explore $p(n)$ was to rely on Monte Carlo sampling methods paired with an efficient binary search algorithm to search the space of possible $n \in [0,2^{11}]$ (an arbitrary maximum limit on n to search). Here, at each iteration, a sample size guess n_{guess} is made. After Monte Carlo simulations and an evaluation of statistical power (whether $p(n_{guess}) < \beta$ or $p(n_{guess}) > \beta$). n_{guess} is adjusted up (if insufficient statistical power), or down (if sufficient statistical power). The amount that n_{guess} is adjusted by is halved at each time point in order to converge on the true value for n_{min} in precisely $t = 11$ iterations (i.e. the algorithm has order $O(log_2(n))$ where n is the initial guess.)

The algorithm is described in the Python source code. Briefly:

0) Using the Monte Carlo methodology described above, generate numerical probability distributions for estimates of $s_t^{(1)}$ and $s_t^{(2)}$. These distributions are denoted $\mathbf{p^{(1)}}(\hat{\mathbf{s}}_{\mathbf{FRA}})$ and $\mathbf{p^{(2)}}(\hat{\mathbf{s}}_{\mathbf{FRA}})$.

1) Set $n_{guess} = 2^{11} = 2048$. Draw $k = 10000$ independent replicates, each composed of n_{guess} paired experimental samples sampled from $\mathbf{p^{(1)}}(\hat{\mathbf{s}}_{\mathbf{FRA}})$ and $\mathbf{p^{(2)}}(\hat{\mathbf{s}}_{\mathbf{FRA}})$. This yields corresponding replicate matrices $\mathbf{D}_1$ and $\mathbf{D}_2$, which have dimension (k, n_{guess}).

2) *Evaluation of statistical power*. For each one of the k columns in $\mathbf{D}_1$ and $\mathbf{D}_2$, use a Wilcoxon Rank-Sum statistical test and record whether the replicate is classified as a positive with $p < \alpha$. Since the populations have distinct silent fractions, all positives are true positives so we can record the number of true positive n_{TP}. Compute the false-negative rate as $\text{FNR} = 1 - n_{TP}/k$.

3a) If $\text{FNR} < \beta$: Update $n_{guess}(t + 1) \leftarrow n_{guess}(t) - \dfrac{|(n_{guess}(t) - n_{guess}(t-1)|}{2}$. Loop back to (1).

3b) If $\text{FNR} > \beta$: Update $n_{guess}(t + 1) \leftarrow n_{guess}(t) + \dfrac{|(n_{guess}(t) - n_{guess}(t-1)|}{2}$. Loop back to (1).

4) Repeat until $|n_{guess}(t) - n_{guess}(t-1)| = 1$. Then $n_{min} = \min\{n_{guess}(t), n_{guess}(t-1)\}$.

In Fig. 4C, we test each of the four estimators (*fra, binary, fra-mle, binary(llr)*) by testing the null hypothesis that P_1 is silent. So we fix $s_t^{(1)} = 0$ while $s_t^{(2)} \in [0,1)$ and compute $n_{min}(s_t^{(2)})$ for each estimator:

i) For $\hat{s}_{FRA}$ (*fra*) the computation of $n_{min}(s_t^{(2)})$ proceeds numerically as described in the algorithm above.

ii) For the binary sampling case (*binary*), the algorithm was identical to above, except instead of drawing estimates from a Monte Carlo-simulated distribution, we simply considered each draw of n_{guess}

samples as constituting Bernoulli trials with $p = s_t^{(1)}$ or $s_t^{(2)}$. In step 2, a Chi-Squared test was used instead of a Rank-Sum test as the data is categorical (silent or active binarization).

iii) For $\hat{s}_{FRA-MLE}$ (*fra-mle*), we employ a log-likelihood (LLR) approach. We first calculate likelihood functions as described above. In all cases the null hypothesis is $H_0 : s_t = 0$ and the alternate hypothesis is $H_A : s_t = \hat{s}_{FRA-MLE}$. We apply Wilks' theorem, calculating a test statistic $D = -2[\ell(s_t = \hat{s}_{FRA-MLE}) - \ell(s_t = 0)]$ where $D \sim \chi^2(\text{df} = 1)$. This allows us to employ a Chi-Squared test on D in place of a Wilcoxon Rank-Sum test (step 2 of the algorithm above). All other components of the algorithm are identical.

iv) For a binary sampler incorporating a model (*binary(llr)*), we derive an analytical expression for minimum sample size, n_{min}, with a LLR testing approach (starting from the binomial distribution) :
$n_{min}(s_t) = \frac{log(\beta)}{1 - s_t}$.