A Bayesian Solution to Count the Number of Molecules within a Diffraction Limited Spot

Alexander Hillsley[1], Johannes Stein[2,3], Paul W. Tillberg[1], David L. Stern[1], Jan Funke[1*]

[1] HHMI Janelia Research Campus, Ashburn, VA,
[2] Wyss Institute of Biologically Inspired Engineering, Boston, MA
[3] Department of Genetics, Harvard Medical School, Boston, MA

April 18, 2024

Abstract

We address the problem of inferring the number of independently blinking fluorescent light emitters, when only their combined intensity contributions can be observed at each timepoint. This problem occurs regularly in light microscopy of objects that are smaller than the diffraction limit, where one wishes to count the number of fluorescently labelled subunits. Our proposed solution directly models the photo-physics of the system, as well as the blinking kinetics of the fluorescent emitters as a fully differentiable hidden Markov model. Given a trace of intensity over time, our model jointly estimates the parameters of the intensity distribution per emitter, their blinking rates, as well as a posterior distribution of the total number of fluorescent emitters. We show that our model is consistently more accurate and increases the range of countable subunits by a factor of two compared to current state-of-the-art methods, which count based on autocorrelation and blinking frequency, Furthermore, we demonstrate that our model can be used to investigate the effect of blinking kinetics on counting ability, and therefore can inform experimental conditions that will maximize counting accuracy.

*Corresponding Author

1 Introduction

Molecular counting aims to determine the absolute number of subunits associated in an object, a quantity that is often essential to understanding the underlying biology of a system. The activation of T cells, for example, is sensitive to single ligands, and quantifying the molecular count of ligand-receptor interactions informs understanding of system sensitivity and activation of downstream signalling pathways [1]. Furthermore, molecular counting enables identification of oligomeric states by counting the subunits in a cluster, allowing differentiation of monomers, dimers, trimers, and higher order oligomers. Several biological processes depend on oligomer quantity: TGF-β signaling, for example, depends on the oligomeric state of Smad [2, 3]. Similarly, the oligomeric state of G-protein coupled receptors influences downstream GPCR signaling [4, 5].

Often, as in the aforementioned examples, the subunits of interest are separated by only a few nanometers and thus can not be quantified through standard fluorescence microscopy techniques. In contrast to larger structures (whole cells or organelles), which can be discerned in fluorescence microscopy, individual molecules in diffraction limited spots below a resolution of approximately 200 nm can not be separated with light microscopy. Even super-resolution techniques [6, 7] cannot discern subunits that are closer

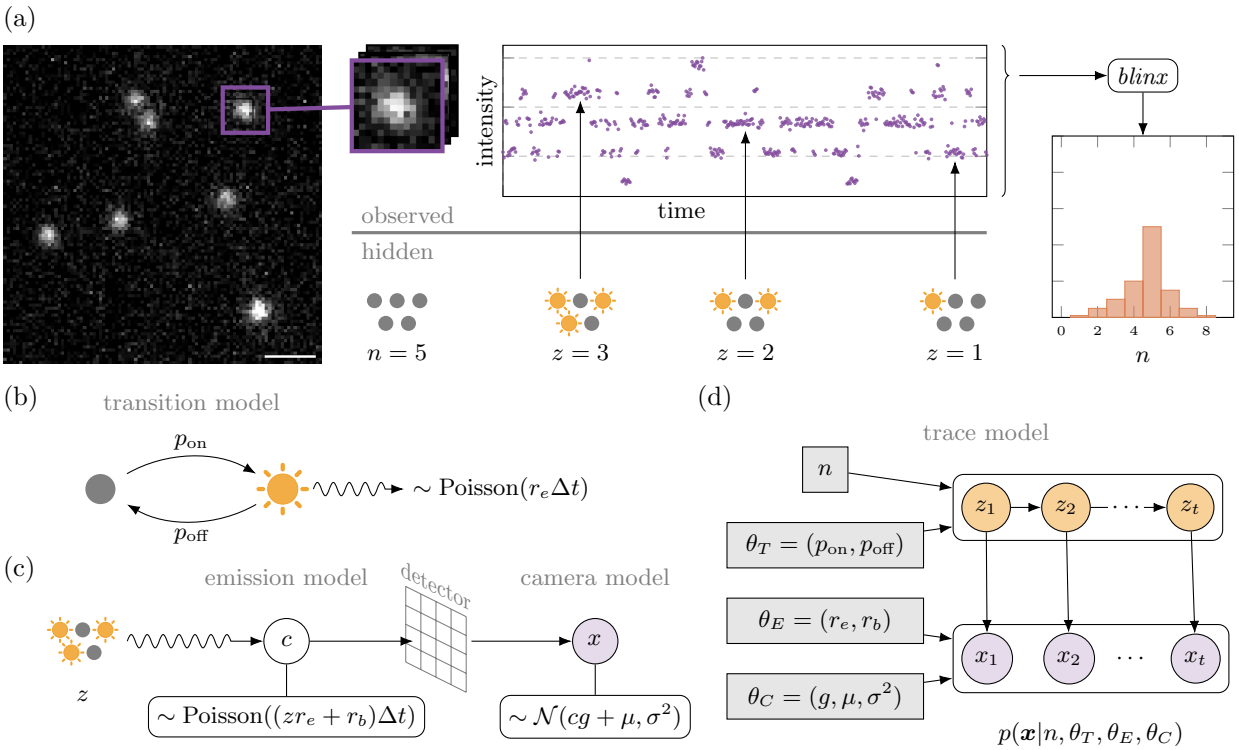

Figure 1: **(a)** Overview of the *blinx* method (scale bar: 1 μm). *blinx* is fit on a series of observed intensity measurements, produced by a hidden state of independent stochastically blinking emitters. **(b)** p_{on} and p_{off} (θ_T) parameterize the blinking kinetics of each individual emitter **(c)** The distribution of observed intensities is effected by both poisson distributed shot noise (θ_E), and normally distributed readout noise (θ_C). **(d)** *blinx* is based on a hidden Markov model conditioned on the total count n and seven fittable parameters.

than 10 nm apart [8].

We present a Bayesian solution to estimate the number of fluorescently labelled subunits in a single diffraction limited spot. Our solution is based on a probabilistic model that incorporates the photophysics of blinking fluorescent emitters and models their dynamics over time as a Markov chain. Given a trace of the combined intensity of all subunits contained in a spot over time (Fig. 1a), this model can be used to obtain the most likely number of emitters contained in that spot: the molecular count.

Perhaps the simplest method of molecular counting is to correlate the combined fluorescent intensity of a spot with the number of subunits, *i.e.*, the more subunits located within a spot, the brighter the spot is expected to be [9, 10]. This method works well for qualitative measurements, but, due to the noise in intensities measured from any single fluorophore, this approach lacks the precision required for accurate estimates of molecular counts.

Molecular counting methods that incorporate temporal fluctuations of intensity are more accurate than correlation approaches. Methods such as fluorescence correlation spectroscopy (FCS) [11, 12, 13] and balanced super-resolution optical fluctuation imaging (bSOFI) [14] fit higher order statistics to fluctuations in fluorescent intensity over time to estimate the molecular count. In other methods, temporal variations in intensity are induced rather than just observed; for example, counting distinct bleaching

steps can provide estimates of the number of fluorescent emitters [15, 16, 17, 18].

Blinking fluorophores, such as those used in PALM [19, 20] and STORM [21], can be used to facilitate solving the counting problem [22, 23]. The known blinking behavior allows for more precise models than the random fluctuations measured in FCS and bSOFI, while the repeated transitions in intensity provide more information than the irreversible switches that occur in bleaching-based counting. A major limitation of all of these methods is the effect of photo-bleaching, because this limits the amount of time a single fluorophore can be observed, and makes it difficult to distinguish between blinking and photobleaching.

In contrast, DNA-PAINT [24] provides blinking fluorescence that is functionally immune to photobleaching due to the continuous replenishment of fluorophores from solution [25]. Localization based FCS (lbFCS+) [26, 27] combines the structured blinking of DNA-PAINT with the principles of FCS, fitting the autocorrelation function of intensity over time to produce a count. As such, it is able to accurately count up to six molecules in a single diffraction limited spot. Quantitative DNA-PAINT (qPAINT) [28] estimates the molecular count based on the frequency of blinking events. For example, if a blinking rate of one event per second is calibrated to one molecule, an observation of ten events per second corresponds to a count of ten molecules. However, both of these methods are limited in that they fit summary statistics, rather than the data directly. In contrast, our solution fits the model to each frame of the measurement in sequence, fully utilizing both the temporal and intensity information available.

Furthermore, all prior methods provide a single estimate of molecular count, whereas a probabilistic method would provide valuable information for downstream analysis. Bayesian approaches can estimate a likelihood for each possible condition and have previously been used to infer the number of FRET conformational states [29, 30], the number of photobleaching steps [18], and the molecular count from observed intensity [23].

We propose *blinx*, a Bayesian model to estimate the molecular count directly from a fluctuating sequence of fluorescent intensity measurements over time. Based on a fully differentiable hidden Markov model, *blinx* fits seven parameters that describe the photo-physics and kinetics of the system, directly to the sequence of measurements and estimates a likelihood for each possible count. These likelihoods can be compared directly, producing a posterior distribution of counts given the observation sequence. We first run *blinx* as a forward model, to simulate different experimental conditions and determine their effects on counting ability. We demonstrate a substantial improvement in counting accuracy compared to both lbFCS+ and qPAINT and an ability to count across different kinetic regimes. Finally, we prove the counting ability of *blinx* experimentally by validating the estimated count with ground truth measured by super-resolution DNA-PAINT imaging.

2 Method

Starting from the photo-physics of the system, we derive a probabilistic model that accounts for both the observed intensity, and the temporal fluctuation kinetics. Given an intensity trace, the output of our method is the posterior distribution over molecular counts, providing not only the most likely count, but also a comparison to all other possible counts.

2.1 Model

We consider the case of multiple fluorescent emitters within a single diffraction limited spot. As such, we can only measure their combined intensity contributions x_t at a timepoint t. We assume that each emitter blinks stochastically and independently of the others, yielding a fluctuating trace $\boldsymbol{x} = (x_1, \ldots, x_T)$ of total intensity measurements over T frames (Fig. 1a). Our goal is to infer the number of emitters n, given $\boldsymbol{x}$.

Formally, we are modelling a posterior distribution of the number of emitters n given the intensity trace $\boldsymbol{x}$, *i.e.*:

$$p(n|\boldsymbol{x}) \propto p(n)p(\boldsymbol{x}|n).$$

Assuming a uniform prior on the number of emitters, $p(n)$ can be ignored and it remains to model the like-

lihood $p(\boldsymbol{x}|n)$. We model this likelihood as a hidden Markov model where the hidden state, denoted by z_t, represents the number of emitters that are active and emitting photons at time t (see Fig. 1d). The hidden Markov model is conditioned on a set of parameters θ, including θ_C camera parameters, θ_E emission parameters, and θ_T transition parameters,

$$p(\boldsymbol{x}|n,\theta) = \prod_{t=1}^{T} p(x_t|z_t,\theta)p(z_t|z_{t-1},\theta,n), \tag{1}$$

where we define $p(z_1|z_0,\theta,n) = p(z|\theta,n)$, *i.e.*, the initial distribution of the hidden state z is the stationary distribution.

Following this framework, the likelihood is the product of two probability distributions: an intensity model $p(x_t|z_t,\theta)$ and a transition model $p(z_t|z_{t-1},\theta,n)$.

2.1.1 Intensity Model

The distribution of observed intensity x_t given the number of currently active emitters z_t can be directly derived from the photo-physics of our microscope. In the following section we omit the subscript t for clarity.

Working backwards through the light path, the measured intensity value x is a function of the number of photons c detected (Fig. 1c). The detector itself contributes noise to the system, known as the readout noise, which we model as $p(x|c,\theta_C)$. The number of photons hitting the detector, in turn, depends on the number of active emitters z and is distributed following a shot noise model $p(c|z,\theta_E)$.

The probability of observing an intensity x given z can then be obtained by marginalizing over c, *i.e.*:

$$p(x|z,\theta) = \sum_c \underbrace{p(x|c,\theta_C)}_{\text{readout}} \underbrace{p(c|z,\theta_E)}_{\text{shot}}. \tag{2}$$

Readout Noise The readout noise from the detector is often assumed to be negligible (*i.e.* when using an EMCCD camera). However, when using an sCMOS camera, its contribution is significant [31]. To accommodate both systems, here we include the readout noise in our model. However, the readout noise contribution can be negated by setting $\mu = 0$, and $\sigma^2 << 1$. Given c photons, the readout intensity of the camera is normally distributed [31]

$$p(x|c,\theta_C) = \mathcal{N}[cg+\mu,\sigma^2](x), \tag{3}$$

where $\theta_C = (g,\mu,\sigma^2)$ are the camera's gain, offset, and variance, respectively.

In the following, it will be beneficial to express the readout noise as a zero mean normal distribution that is independent of the number of photons detected. To that end, we transform our measurement x into photon space by accounting for the camera gain and offset:

$$\tilde{x} = \frac{x-\mu}{g} \quad x = \tilde{x}g + \mu. \tag{4}$$

We can now express our readout noise distribution in terms of the transformed measurement $\tilde{x}$, which we further shift by $-c$ to obtain a zero-mean distribution:

$$p(\tilde{x}|c,\theta_C) = \mathcal{N}[c,\frac{\sigma^2}{g^2}](\tilde{x}) \tag{5}$$

$$= \mathcal{N}[0,\frac{\sigma^2}{g^2}](\tilde{x}-c). \tag{6}$$

Shot Noise Due to shot noise, the number of photons detected over a time interval Δt is Poisson distributed [32] (Fig. 1c). Our model accounts for two sources of emitted photons: first, the z active emitters produce photons at a rate zr_e, and second, out of plane emitters produce a relatively constant number of background photons r_b. Combined, these sources produce the expected number of photons $\lambda = (zr_e + r_b)\Delta t$ per frame.

$$p(c|z,\theta_E) = \text{Poisson}[\underbrace{(zr_e+r_b)\Delta t}_{\lambda}](c). \tag{7}$$

For large values of λ, the Poisson distribution approaches a normal distribution with both mean and variance of λ and therefore we approximate our shot noise distribution as:

$$p(c|z,\theta_E) \approx \mathcal{N}[\lambda,\lambda](c). \tag{8}$$

Combined Intensity Model Taking together the readout noise transform (Eq. 6) and shot noise approximation (Eq. 8), the intensity distribution (Eq. 2) can now be rewritten as

$$p(\tilde{x}|z, \theta_C, \theta_E) = \sum_c \mathcal{N}[0, \frac{\sigma^2}{g^2}](\tilde{x} - c)\mathcal{N}[\lambda, \lambda](c). \quad (9)$$

Similar to [31], we note that this expression resembles a convolution, *i.e.*, the sum of two independent random variables (the photon count c and a zero-mean, photon-independent camera readout noise, (Eq. 6)). Since the two distributions are normal, we can rewrite the above as a single normal distribution with summed means and variances:

$$p(\tilde{x}|z, \theta_C, \theta_E) = \mathcal{N}[\lambda, \lambda + \frac{\sigma^2}{g^2}](\tilde{x}). \quad (10)$$

2.1.2 Transition Model

To model the step-like temporal fluctuations in intensity observed in the trace $\boldsymbol{x}$, a distribution is needed that describes the change in the number of active emitters z over time. To accomplish this, we assume that the process is Markovian and the number of active emitters z_t at time t depends only on the number of active emitters at the previous time point z_{t-1}. On the individual emitter level, we define p_{on} as the probability of an inactive emitter at time $t-1$ to become active at time t. Conversely, we define p_{off} as the probability of an active emitter at $t-1$ to become inactive at time t (Fig. 1b).

For the case of a single emitter (*i.e.*, $n = 1$), the probability of observing this emitter as active (*i.e.*, $z_t = 1$) if the emitter was inactive in the previous frame, is p_{on}, or $(1 - p_{\text{off}})$ if it was previously active (and similarly for the probability of observing the emitter as inactive). For multiple emitters, however, a transition from z_{t-1} to z_t can be caused by multiple events. Consider, for example the case of two emitters ($n = 2$). A transition from $z_{t-1} = 1$ to $z_t = 1$ can either be caused either by no change occurring at all or by one emitter activating while the other deactivates.

In the general case ($n \geq 1$), we consider all possible events that can lead to an observed change in z. To that end, we marginalize over the number a of active emitters that deactivated from $t-1$ to t:

$$p(z_t = z|z_{t-1}, \theta_T, n) = \sum_{a=0}^{z_{t-1}} \binom{a}{z_{t-1}} p_{\text{off}} \binom{z - z_{t-1} + a}{n - z_{t-1}} p_{\text{on}}, \quad (11)$$

We assume that all emitters share the same p_{on} and p_{off}, and that both remain constant over time.

2.1.3 Inference

The likelihood $p(\boldsymbol{x}|\theta, n)$ of observing $\boldsymbol{x}$, see (Eq. 1), depends on a total of seven parameters θ and the total number of emitters n. However, we are interested in the posterior distribution $p(\boldsymbol{x}|n)$. Accounting for the priors on θ in addition to the likelihood, this is known as the model evidence and defined as:

$$p(\boldsymbol{x}|n) = \int p(\boldsymbol{x}|\theta, n)p(\theta|n)d\theta.$$

Unfortunately, no closed form solution exists for this integral. Although no closed for solution exists for this integral, multiple approximations exist including variational methods [30] and the Laplace approximation [33]. We use the later, which is the maximum likelihood solution corrected by the Occam factor:

$$\ln p(\boldsymbol{x}|n) \approx \ln p(\boldsymbol{x}|\theta_{\text{MAP}}, n) + \underbrace{O(\boldsymbol{x}, \theta_{\text{MAP}})}_{\text{Occam factor}}, \quad (12)$$

where the Occam factor $O(\boldsymbol{x}, \theta_{\text{MAP}})$ is defined as

$$O(\boldsymbol{x}, \theta_{\text{MAP}}) = \ln p(\theta_{\text{MAP}}) + \frac{M}{2}\ln(2\pi) - \frac{1}{2}\ln|A|$$
$$A = -\nabla\nabla \ln p(\theta_{\text{MAP}}|\boldsymbol{x}).$$

Here, M is the dimensionality of θ, which is constant. Because we are interested only in comparing likelihoods this term can safely be ignored, resulting in:

$$\ln p(\boldsymbol{x}|n) \approx \ln p(\boldsymbol{x}|\theta_{\text{MAP}}, n) + \ln p(\theta_{\text{MAP}}) - \frac{1}{2}\ln|A|. \quad (13)$$

Therefore, to estimate the posterior distribution over n we first need to find the maximum likelihood parameters θ_{MAP} for each value of n.

This model is fully differentiable, allowing us to estimate the maximum likelihood parameters through gradient ascent:

$$\theta_{k+1} = \theta_k + \nu \frac{\partial p(\boldsymbol{x}|\theta, n)}{\partial \theta}\bigg|_{\theta_k},$$

where ν is the step size, dynamically calculated by the Adam optimizer, and the gradients of θ are taken with respect to the likelihood: $p(\boldsymbol{x}|n, \theta)$. The maximum likelihood parameters θ_{MAP} are then obtained,

$$\theta_{\mathrm{MAP}} = \arg\max_{\theta} p(\boldsymbol{x}|\theta, n), \tag{14}$$

thus allowing us to estimate the posterior distribution $p(\boldsymbol{x}|n)$. Additionally, we can obtain an estimate of the the maximum likelihood molecular count:

$$n^* = \arg\max_{n} \ln p(\boldsymbol{x}|n). \tag{15}$$

3 Results

3.1 Simulated Experiments

Running *blinx* as a forward model to simulate experiment, we generated traces for counts $n = 1$ to 30 (Fig. 2a,b). We determined experimentally relevant camera parameters from a calibration of the sCMOS camera (θ_C: $g = 2.17$, $\mu = 4791$, $\sigma^2 = 774$). We determined kinetic (θ_T: $p_{\mathrm{on}} = 0.036$, $p_{\mathrm{off}} = 0.028$), and emission parameters (θ_E: $r_e = 2.79$, $r_b = 6.77$) from an initial experiment, and these values closely matched those reported in the literature [27]. We simulated traces for 4000 frames, corresponding to an imaging time of 14 minutes.

For inference, we placed tight empirically determined, Gaussian priors on the camera parameters. We also placed flat uniform priors on the kinetic parameters, allowing the model to fit blinking rates without external bias, and loose priors on the emission parameters (r_b, r_e) to prevent over-estimation of count (see Section 4). We experimentally determined the prior on r_b by averaging the intensity of the pixels immediately outside the region of interest. We estimated the prior on r_e by first fitting a count of $n = 1$ to the trace. The distribution of all observed z-states is unimodal. Therefore, whichever z-state is the most common, the second most common state will be $z \pm 1$. We observed that fitting $n = 1$ captures this specific transition, and provides a good estimate of photon emission rate of a single emitter r_e.

blinx accurately fit simulated traces with counts up to $n = 30$ (Fig. 2c). For $n = 10$ and below, *blinx* estimated the true count with a substantially higher likelihood than all other possibilities. Above $n = 10$ the posterior broadened as other possibilities become more likely. The most likely estimate was the true count in 177/300 (59%) of traces and was within 1 of the true count in 267/300 (89%) of all traces. Importantly, in many cases *blinx* is able to fit the true count despite the fact that not every state is occupied *i.e.* the model is not just counting the number of peaks in the intensity histogram (Fig. 2a,b). This is a tripling in performance compared to lbFCS+, which counted accurately up to $n = 6$ and was within 1 of the true count for 64/300 (21%) of traces but failed to estimate higher counts. Upon further analysis, the main limitation of lbFCS+ is the estimation of the intensity of a single emitter, corresponding to r_e in *blinx*. lbFCS+ relies on a histogram based method to determine this value, while *blinx* is able to jointly optimize this parameter with the others. Providing this value to lbFCS+ led to a substantial recovery of performance. With r_e provided, lbFCS+ identified 200/300 (67%) of traces within 1 of the true count.

Both *blinx* and lbFCS+ count without prior knowledge of the kinetic parameters. In contrast, qPAINT is heavily dependant on a kinetic calibration and therefore not comparable in this context. A direct comparison to qPAINT is shown in Section 3.3

A longer trace provides more information to the model and increases the counting accuracy. We simulated and fit traces of lengths ranging from 100 to 10,000 frames, with the same parameters as the previous section (Fig. 2d). As expected, accuracy increased and variance of the posterior distribution decreased, as trace length increased. Interestingly for short trace lengths, *blinx* consistently underestimated

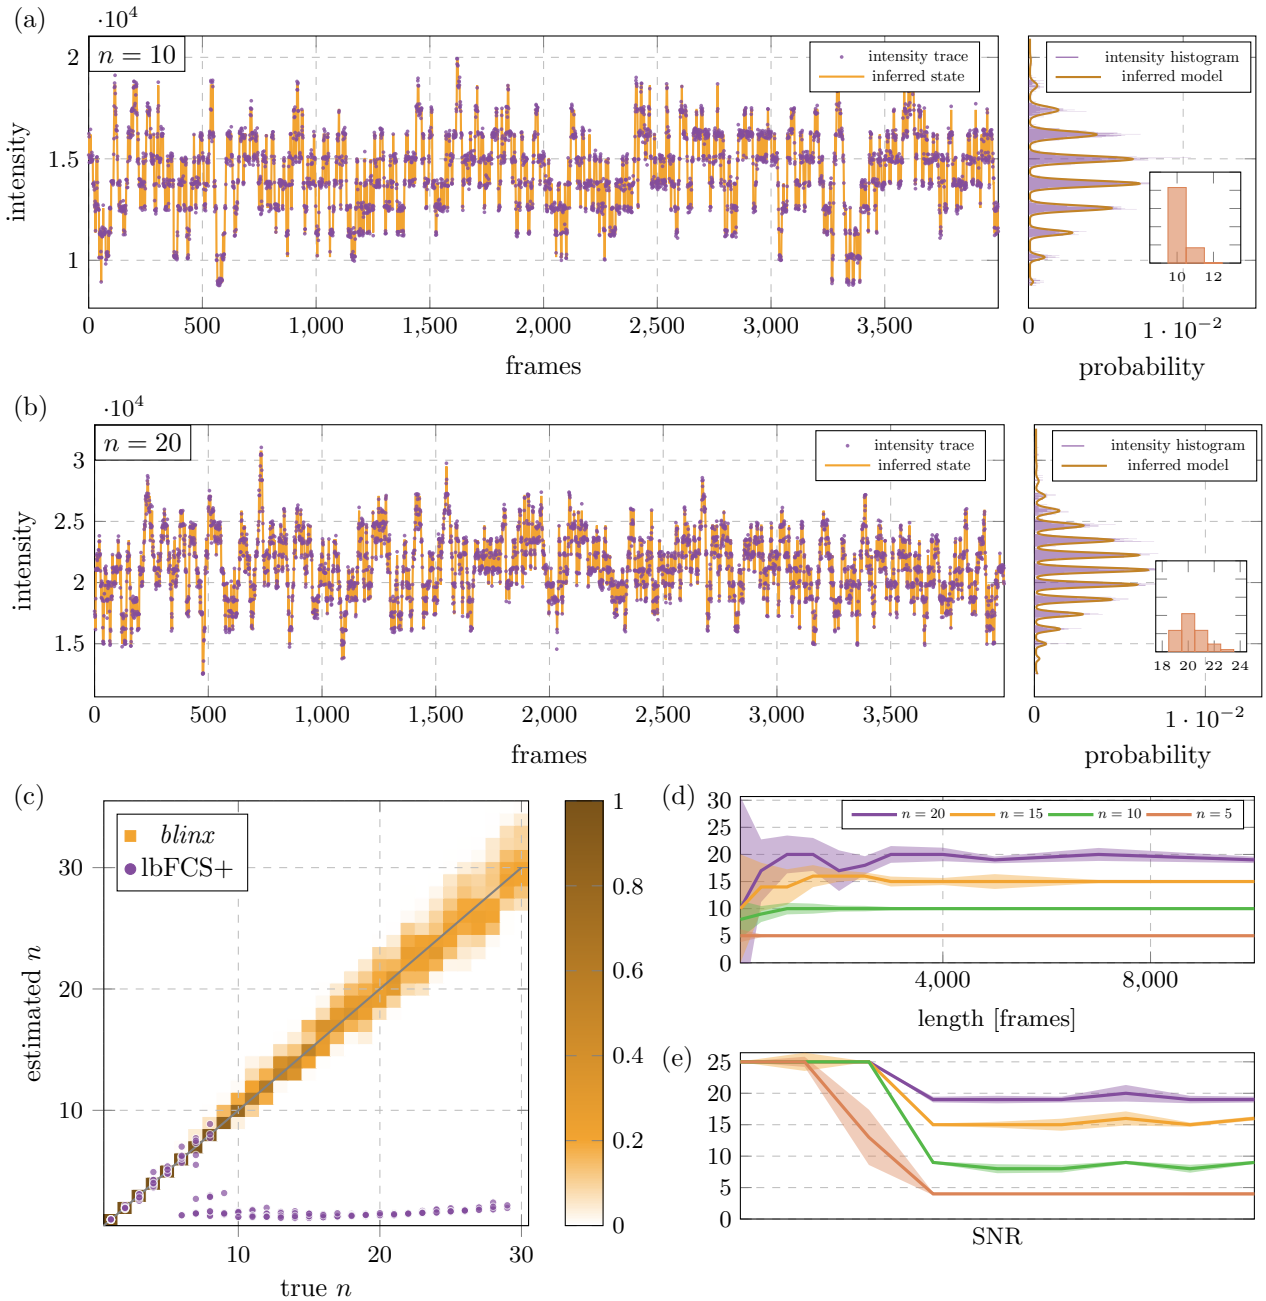

Figure 2: **(a, b)** Traces simulated from $n = 10/20$ emitters, and a density plot of measured intensities. The inferred state and inferred model demonstrate accurate fitting of both the intensity and transition distributions. Inset: posterior distribution estimated by *blinx*. **(c)** The *blinx* posterior is able to accurately estimate significantly higher molecular counts than the current state of the art, lbFCS+. *(caption continued on next page)*

Figure 2: *(continued)* **(d)** As trace length increases, the variance of the *blinx* posterior decreases, with limited improvement above 4000 frames. **(e)** At low signal-to-noise, *blinx* severely over estimates the count, estimating the maximum count tested as the most likely.

the true count. This is likely because only a subset of possible states were observed in this short time.

We determined the robustness of our model and the minimum signal-to-noise ratio (SNR) needed to achieve an accurate count. We artificially increased the variance of the readout noise σ^2, to simulate traces with SNRs ranging from 9 to 1 (where 9 corresponds to the parameters previously used). Because noise in our model is a function of intensity, quantifying the SNR of a trace is not trivial. For simplicity, here we define SNR as the difference in intensity between the first two states (z_0 and z_1) divided by the standard deviation of the difference. In effect this is the higher bound of SNR for a given trace. *blinx* shows accurate counting for SNRs ≥ 4 (Fig. 2e). Interestingly, for SNR ≤ 4, *blinx* estimates the count as 25, the highest n tested, no matter the true count. This is due to the model adding states to compensate for the wide distribution of intensities (see Section 4).

3.2 Effect of kinetic parameters

The performance of our model depends on the true kinetic parameters of the system (p_{on} and p_{off}). To find a regime that maximizes our models counting ability, we simulated and fit traces from $n = 1$ to 20, with a range of kinetic parameters, while holding all other parameters constant. We ensured that this range covered experimentally relevant kinetics from the literature including: qPAINT: (p_{on}, p_{off}) = (0.006, 0.2) [28] and lbFCS+ (0.02, 0.02) [27]. We summarized the resulting posteriors by calculating the expected squared error over all true counts: $\sum w(n^* - n)^2$, where w is the probability of the estimated count n (Fig. 3a).

The result can be divided into two clear regimes: $p_{\text{on}} < p_{\text{off}}$ and $p_{\text{on}} \geq p_{\text{off}}$. In the first regime, where $p_{\text{on}} < p_{\text{off}}$ we see significant limitations to our models counting ability. Furthermore, when p_{on} is substantially lower than p_{off}, our model loses almost all ability to count and estimates $n = 1$ or 2, regardless of the true n^*. In contrast, in the second regime, where $p_{\text{on}} \geq p_{\text{off}}$ we observe accurate counting up to $n = 20$. A clear difference between these two regimes can be seen in the distribution of their hidden states z (Fig. 3a, insets). When $p_{\text{on}} < p_{\text{off}}$, the distribution is shifted towards 0, and a majority of the frames occupy the lowest two z-states (Fig. 3c). In this regime, the true count becomes indistinguishable for all methods, without prior information. When $p_{\text{on}} \geq p_{\text{off}}$, this distribution is centered, or even shifted towards n and a larger fraction of the states are visited (Fig. 3d). This provides ample information for the model to infer the correct n.

In DNA-PAINT, the blinking rate is determined by the kinetics of single-stranded DNA binding, which in turn depends on experimental conditions such as temperature and concentration, and the specific DNA sequence [34]. As a result, the kinetic parameters p_{on} and p_{off} of our model can be tuned by adjusting the temperature and imager concentration (Fig. 3b). We imaged DNA origami with a known count of $n = 1$ at 25°C and an imager concentration of 10 nM and measured $p_{\text{on}} = 0.028$ and $p_{\text{off}} = 0.072$ (Fig. 3b), which places these conditions firmly in the poor counting accuracy regime of $p_{\text{on}} < p_{\text{off}}$ (Fig. 3a). To move to a more favorable kinetic regime for experiments, we increased the imager concentration and decreased the temperature. Increasing imager concentration to 30 nM raised p_{on} to 0.071 and decreasing temperature to 13°C decreased p_{off} to 0.037 (Fig. 3b). As expected, the effects of temperature and concentration were largely independent of one another [34]. We also observed a substantial decrease in SNR with decreasing temperature and with increasing concentration. This was an expected side effect of increasing concentration, because increasing imager concentration increases r_b. But the effect of temperature on SNR was surprising. We hypothesize that this is due to the stabilization of partial binding between imager and docker strands at low temperatures. Balancing the increase in counting accuracy and the decrease in

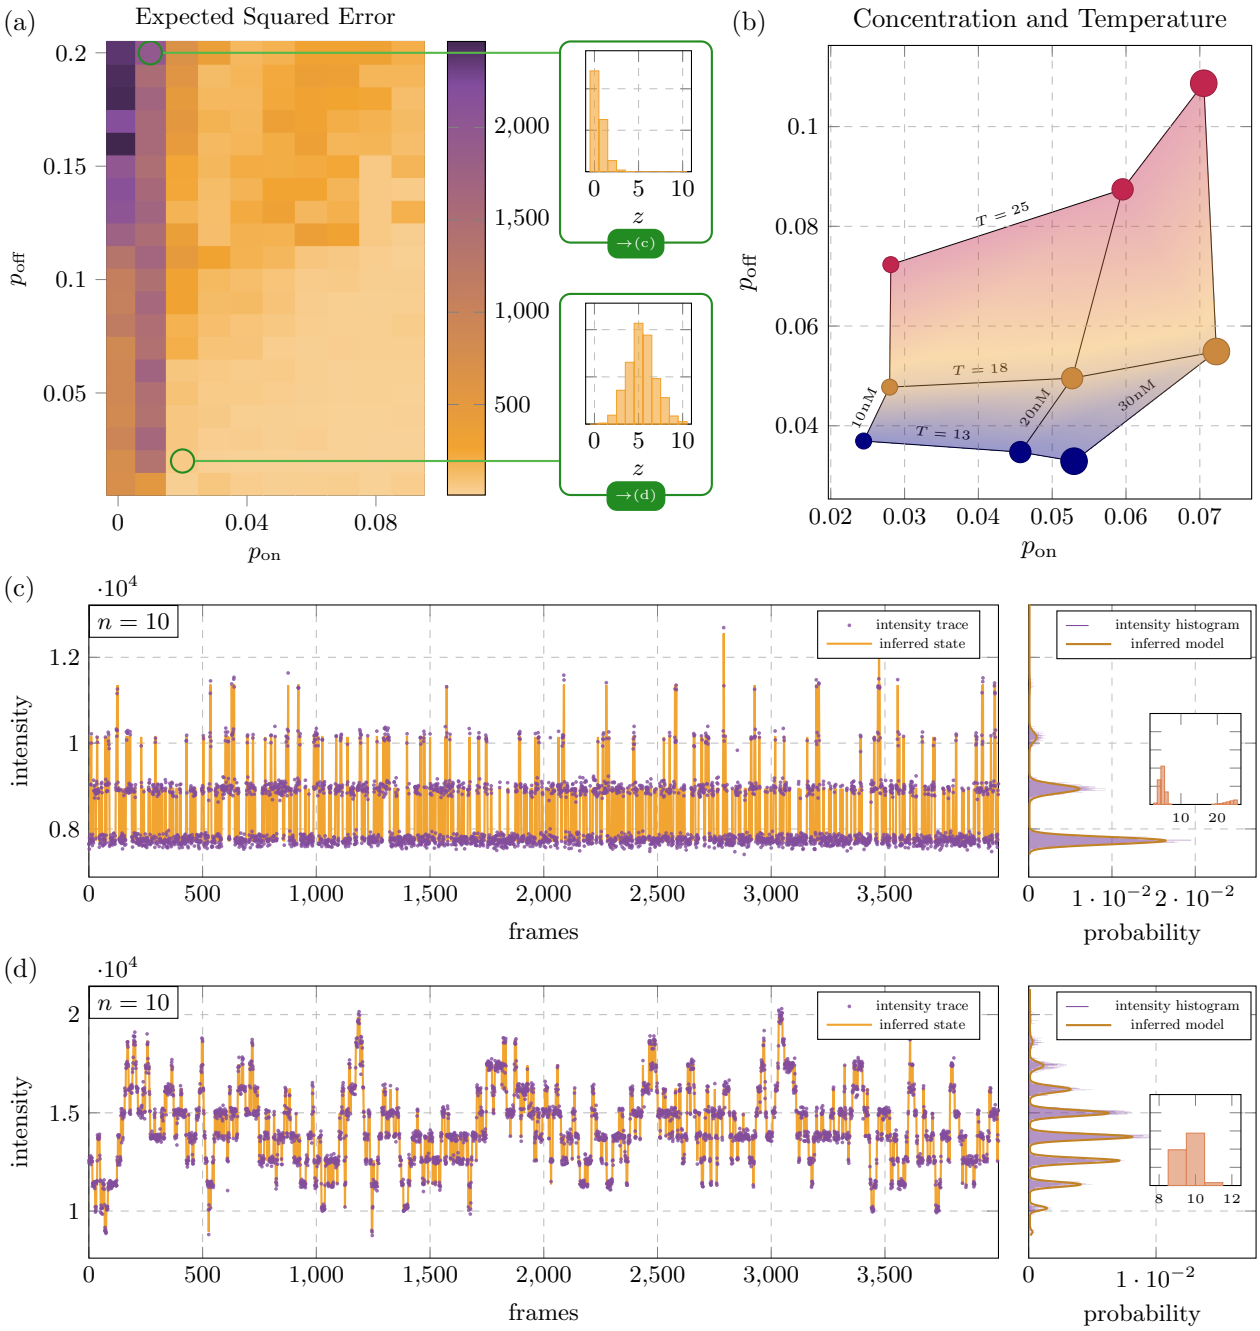

Figure 3: **(a)** The counting ability of *blinx* varies significantly with the blinking kinetics of the emitters, showing increased model accuracy when $p_{\text{on}} > p_{\text{off}}$. Top inset: when $p_{\text{on}} < p_{\text{off}}$, the distribution of observed z states is shifted towards 0. Bottom inset: for $p_{\text{on}} = p_{\text{off}}$, the distribution of z states is centered at $1/2\ n$. **(b)** In DNA-PAINT p_{on} and p_{off} are governed by DNA binding kinetics and therefore are highly sensitive to changes in temperature and concentration, allowing experimental tuning of these parameters. *(caption continued on next page)*

Figure 3: *(continued)* **(c)** Trace and *blinx* fit, p_{on} =0.01, p_{off} =0.2, and true count of n=10, representative of (a) top inset. **(d)** Trace and *blinx* fit, p_{on} =0.02, p_{off} =0.02, and true count of n=10, representative of (a) bottom inset.

SNR, we identified imaging conditions of 20 nM and 13°C as optimal on our microscope.

3.3 qPAINT Kinetic Regime

qPAINT relies on the accurate measurement of the average dark time between blinking events [28] and is therefore optimized for an entirely different kinetic regime, where blinking events are short and infrequent (Fig. 4a,b). This regime presents a challenge to *blinx* (Fig. 3a top inset). If the only states ever observed are $z = 0$ or 1, there is insufficient information to estimate count without prior knowledge. qPAINT faces the same limitation and relies on a calibration of the blinking kinetics of a single binding site. *blinx*, being a Bayesian model can incorporate a similar calibration as priors of the kinetic parameters. With this addition, and the tightening of the priors on r_e and r_b, the counting accuracy of *blinx* is restored (Fig. 4c). *blinx* estimated the true count as the most likely in 133/300 (44%) of traces and estimated within 1 of the true count in 269/300 (90%) of all traces, comparable to the favorable kinetic regime (Fig. 2c). However, obtaining these priors requires additional experimental steps, and is often not trivial. Thus, a tradeoff exists. Without any calibration, *blinx* is able to accurately count in a subset of kinetic conditions where $p_{\text{on}} > p_{\text{off}}$. But with an experimental calibration, the counting ability of *blinx* is expanded to a wider range of kinetic conditions.

Due to the stochastic nature of blinking, multiple emitters can be active at any given time-point, which becomes increasingly likely at higher counts. This is not compatible with the qPAINT requirement of well separated, single emitter blinking events. As a result at higher n the measured average dark time is longer than the true average dark time and results in qPAINT underestimating molecular count at higher n (especially noticeable above $n = 20$, see Fig. 4c). In *blinx*, the activity of multiple emitters at a single timepoint is explicitly modeled in both the intensity and transition models previously derived ((Eq. 7) and (Eq. 11)). As a result, *blinx* avoids underestimating, and provides an accurate measurement of the molecular count up to $n = 30$.

3.4 Experimental Counting

To experimentally validate the counting performance of *blinx*, we used DNA-Origami, which provides nano-scale control over the number and location of emitters [35]. DNA-Origamis were designed containing 1 and 4 DNA-PAINT docker strands, spaced in a grid 20 nm apart. This distance was chosen so that the true number of docker strands could be visually confirmed through super-resolution post-processing. Incorporation efficiency is roughly 80 percent for each docker site [36], so only a fraction of the origamis were expected to contain all 4 dockers. Origamis were first imaged at 13°C with low laser power (1.5%) and 20 nM imager concentration to collect traces (4,000 frames at 200 ms exposure time) for counting with *blinx* (Fig. 5a,b). The system was then allowed to warm to 25°C a buffer exchange was performed, and new imager was added at 10 nM. The origamis were imaged again at high laser power (40%), and post-processed with Picasso [24] to obtain super-resolution ground truth (Fig. 5c). Only origamis that had a visual count matching the designed count (1 or 4) were selected for analysis with *blinx*.

Of the 131 traces with a known count of 1, *blinx* correctly counted 112 (85%) and the average of the posterior distributions is shown in Fig. 5d. Possible counts up to $n = 8$ were modelled, but the likelihoods for counts above $n = 3$ were negligible for all traces. For the traces with a known count of 4, 71/110 (65%) were correctly identified as 4, and 103/110 (94%) were identified as between 3 and 5. The average posterior distribution is shown in Fig. 5e. Importantly, traces were not preprocessed or filtered before undergoing *blinx* counting analysis.

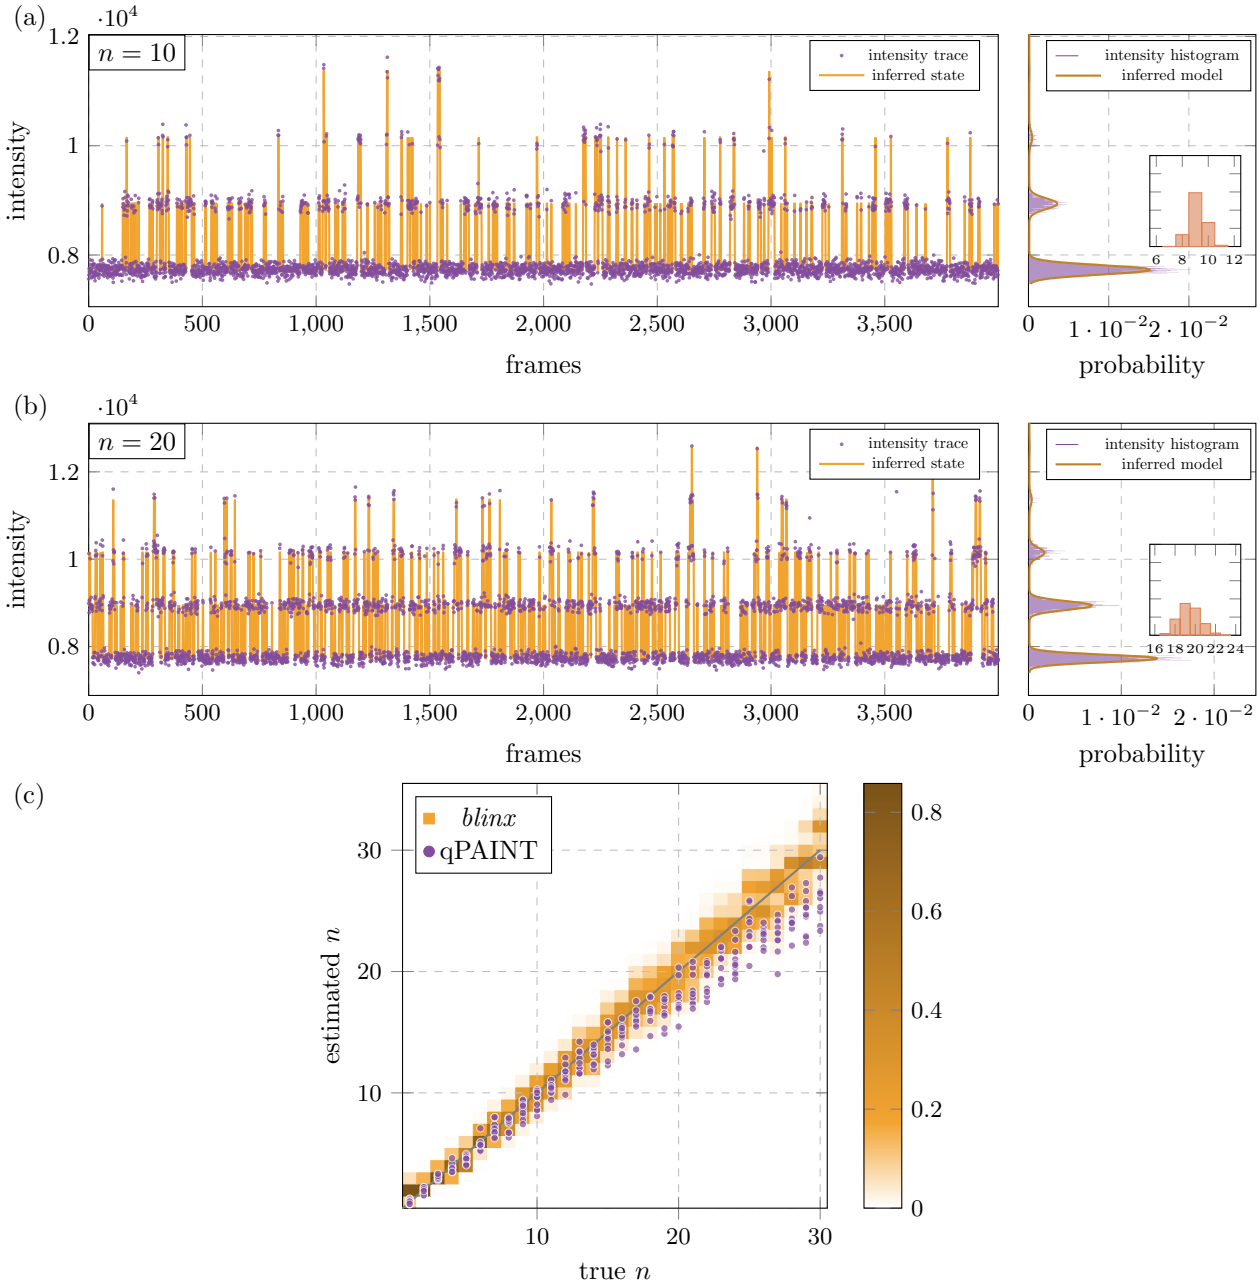

Figure 4: **(a, b)** simulated traces and *blinx* fits, p_{on} =0.006, p_{off} =0.2, and $n = 10/20$ emitters. *blinx* was fit with strong priors on all parameters. **(c)** At higher molecular counts ($n > 15$) qPAINT begins to underestimate the count, while the posterior of *blinx* remains accurate.

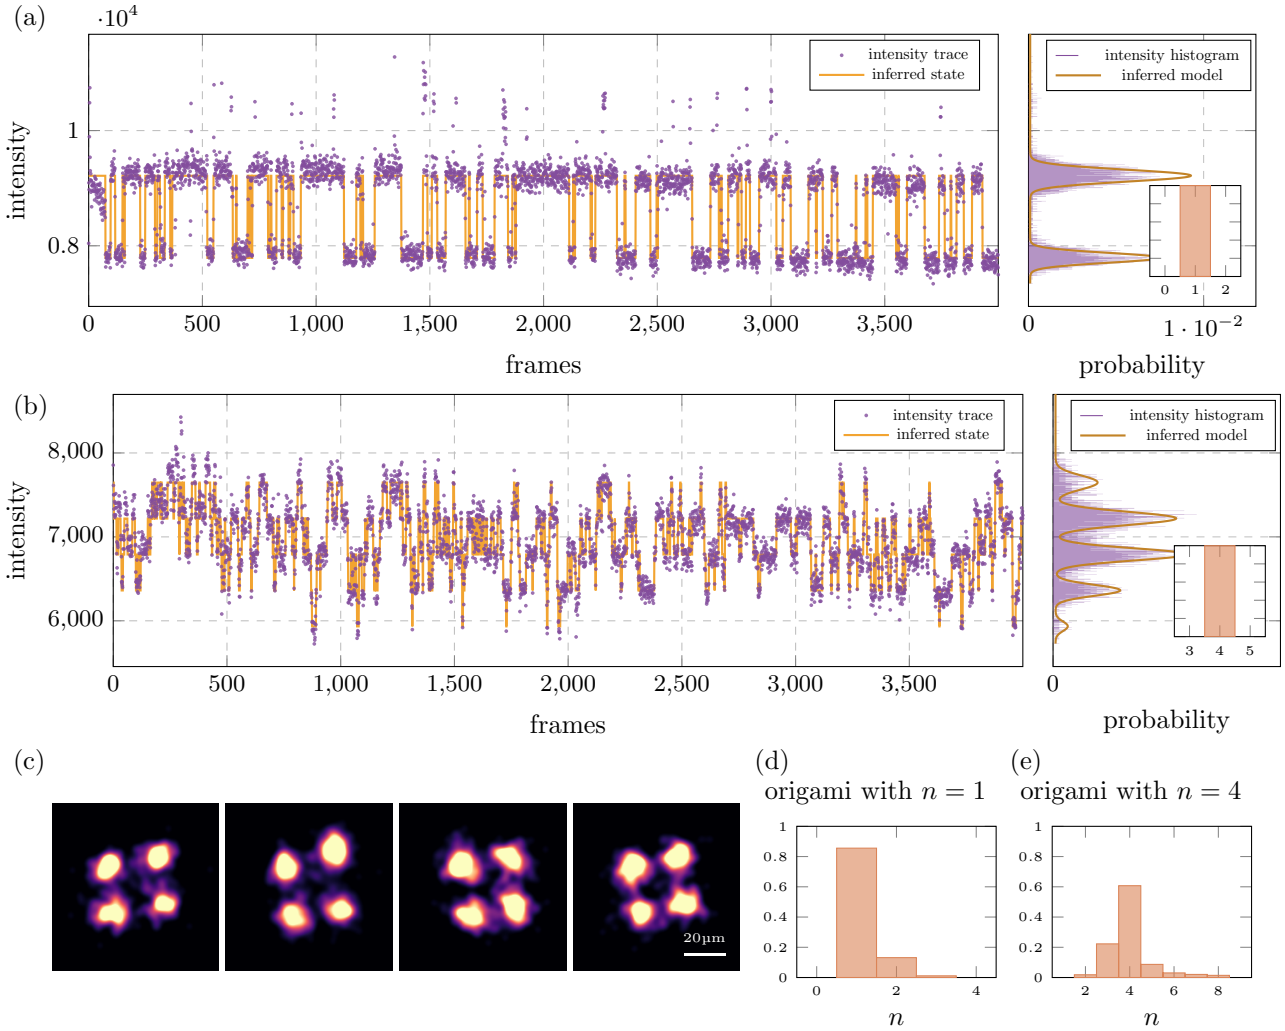

Figure 5: **(a, b)** Experimental traces with super-resolution confirmed counts of one (a) and four (b), and *blinx* fits **(c)** Super-resolution visualization of origamis made with four binding sites. **(d, e)** Average *blinx* posterior distributions of 131 traces with a known n of one (d) and 110 traces with a known n of four (e)

4 Discussion

Although we used DNA-PAINT for the experiments presented here, our model is not limited to DNA-PAINT, and is compatible with any system that generates repeated fluorescent blinking that reports the underlying molecular count. We anticipate that *blinx* will be adaptable to any experimental molecular counting pipeline, to provide accurate probabilistic estimates up to at least $n = 30$.

Regardless of the emitter chosen, several assumptions are essential to *blinx*: Firstly, we assume that all subunits behave independently. In the intensity model, this means that observed intensity will scale linearly with the number of active emitters z. In the transition model, this means that the blinking kinetics of one spot have no effect on the blinking kinetics of any other. In practice, subunits in extremely close proximity could potentially interact and violate this assumption. Other methods like qPAINT make this same assumption and have been shown to work well in experimental systems [37, 38]. Second, we assume that all subunits within a spot have identical properties (p_{on}, p_{off}, r_e). Rather than model each subunit individually, we model the summed behavior of all subunits combined. In other words, we assume that each subunit is interchangeable with any other within the same spot. In some applications, where non-uniform behavior among subunits may occur, *e.g.*, if one unit is less accessible to diffusion of imager than the others [39]. In these cases, we would expect a decrease in model performance. Finally, we also assume that all properties remain constant over time. Experimentally, drift in emission properties (r_e, r_b) can result from an unstable focus and should be corrected before any processing with *blinx*. Changes in kinetic properties over time are also sometimes observed and could be caused by photo-bleaching, damage to emitters or changes in temperature. Unlike the other two assumptions, this assumption is not fundamental to the structure of *blinx* and slight extensions of our model could account for these time dependant parameters. In fact, we anticipate that intentional temporal changes in parameters could be used to gain more information from the system and further increase model performance.

While *blinx* has been designed to be applicable to as many different systems as possible, there are a few limitations that should be considered. The primary limitation of *blinx* is the tendency to overestimate the molecular count in traces with low SNR. This is caused by the distribution at the heart of the intensity model. The model expects intensities to be observed in evenly spaced, normally distributed peaks (see histograms in Fig. 5a,b). When intensity is observed between two of these peaks, it can sometimes be cheaper (in terms of likelihood) to add a new z-state to the model, rather than account for that intensity with the given distribution. In the extreme limit, an infinite number of states could be used to perfectly explain any intensity trace. This limitation can be avoided by specifying a prior on r_e, the photon emission rate of a single fluorophore, which effectively increases the cost of adding additional states.

Another limitation is that the intensity model does not capture all the noise in the system. As seen in Fig. 5a, there are a occasional frames where the measured intensity is greater than expected and fits the model distribution poorly. To compensate for those outliers we incorporated a baseline outlier probability into the intensity model, *i.e.*, a fixed uniform likelihood for observing any intensity within the recorded range of values. To further reduce the influence of outliers on counting accuracy, we excluded the highest 0.5% of intensities values when calculating the likelihood (the specific value is adjustable as a hyperparameter).

The Bayesian framework of *blinx* opens this method up to many future opportunities. For example, the simple modification of setting $p_{\text{on}} = 0$ could support the counting of photobleaching events. Additionally, by incorporating dynamic priors, our model could also capture changing conditions over time, potentially even increasing the counting ability. Further, with minor modifications to the transition distribution to account for photobleaching, this method could be extended to other stochastically blinking emitters, such as those used in PALM or STORM, opening the door for molecular counting in living samples.

5 Code Availability

Code for the model *blinx* is available at https://github.com/funkelab/blinx.

Code and data for the experiments described in this manuscript can be found at https://github.com/funkelab/blinx_experiments.

6 Experimental Methods

DNA Origami DNA origami with 20 nm spaced docking strands, were designed with the Picasso design module [24] and assembled following the procedure in [24], with single-strand DNA oligos purchased from IDT. A repetitive 11 nt docker sequence (CTCCTCCTCCT) was added to the 3' end of select staples as the docker sequence. Such repetitive sequences have been shown to increase p_{on} [39]. DNA origamis samples were prepared for imaging in ibidi μ-Slide 8-well glass bottom slides, again following the procedure described in [24].

DNA-PAINT DNA-PAINT imaging was performed following the standard procedure detailed in [24]. Imaging solution contained buffer B (5 mM Tris-HCl, 10 mM $MgCl_2$ 1 mM EDTA, pH 8) and an oxygen scavenging system consisting of PCA, PCD, and Trolox. Super-resolution images were acquired at an with a 7 nt imager (GAGGAGG-Cy3B, Biosyn) at a concentration of 10 nM, while counting analysis images were acquired with a 8nt imager (GAGGAGGA-Cy3B, Biosyn) at a concentration of 20 nM. All images were post-processed in the Picasso-localize module to detect spots and correct drift over the time course.

Microscopy TIRF imaging was performed on a Zeiss Elyra 7 microscope, equipped with a pco.edge 4.2 sCMOS camera (pixel size 6.5 μm), and an α Plan-Apochromat 63x/1.46 TIRF oil immersion objective. An Okolab H101-Cryo-BL system provided temperature control. It was found that this temperature control system was not compatible with super-resolution experiments, due to vibrations caused by the pump, and transferred to the objective through the cooling jacket. Therefore for super-resolution experiments, the temperature control system was turned off, and the sample temperature allowed to increase to room temp (25°C) before imaging.

7 Acknowledgments

We thank Brian English, Chris Obara, Lorena Benedetti, and Jennifer Lippencott-Schwartz for their expertise in super-resolution microscopy and for sharing use of their microscopes. We thank James Liu for his advice on DNA-PAINT and on troubleshooting imaging systems. We thank Diane Adjavon, Caroline Malin-Mayor, Manan Lalit, Magdalena Schneider, and Cedric Allier for their help in model development and python implementation of the model. We thank Shalin Mehta for insight into modeling and calibrating the photo-physics of these systems. We thank Mark Aronson for his help reviewing and editing this manuscript. This work was supported by the Howard Hughes Medical Institute.

8 Author Contributions

A.H. and J.F. derived the model and wrote all code. A.H. performed all simulations and microscopy experiments. J.S. provided training on DNA-PAINT and DNA-Origami experiments and analysis. P.T. and D.S. provided valuable insights and guidance for both method development and experiments. A.H. and J.F. wrote the manuscript. A.H., J.F., P.T. and D.S., conceived the project. All authors revised the manuscript and have given approval to the final version.

References

[1] Darrell J. Irvine, Marco A. Purbhoo, Michelle Krogsgaard, and Mark M. Davis. Direct observation of ligand recognition by t cells. 419(6909):845–849, 2002.

[2] Gareth J. Inman and Caroline S. Hill. Stoichiometry of active smad-transcription factor complexes on dna. 277(52):51008–51016, 2002.

[3] Aristidis Moustakas and Carl-Henrik Heldin. From mono- to oligo-smads: The heart of the matter in tgf-b signal transduction. 16(15):1867–1871, 2002.

[4] James H. Felce, Simon J. Davis, and David Klenerman. Single-molecule analysis of g protein-coupled receptor stoichiometry: Approaches and limitations. 39(2):96–108, 2018.

[5] Gerda E. Breitwieser. G protein–coupled receptor oligomerization. 94(1):17–27, 2004.

[6] Eric Betzig, George H. Patterson, Rachid Sougrat, O. Wolf Lindwasser, Scott Olenych, Juan S. Bonifacino, Michael W. Davidson, Jennifer Lippincott-Schwartz, and Harald F. Hess. Imaging intracellular fluorescent proteins at nanometer resolution. *Science*, 313(5793):1642–1645, 2006.

[7] Michael J. Rust, Mark Bates, and Xiaowei Zhuang. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM). *Nature Methods*, 3(10):793–796, 2006.

[8] Jessica Valli, Adrian Garcia-Burgos, Liam M. Rooney, Beatriz Vale de Melo e Oliveira, Rory R. Duncan, and Colin Rickman. Seeing beyond the limit: A guide to choosing the right super-resolution microscopy technique. *Journal of Biological Chemistry*, 297(1):100791, 2021.

[9] Jürgen J. Schmied, Andreas Gietl, Phil Holzmeister, Carsten Forthmann, Christian Steinhauer, Thorben Dammeyer, and Philip Tinnefeld. Fluorescence and super-resolution standards based on dna origami. 9(12):1133–1134, 2012.

[10] Pavel Tolar, Hae Won Sohn, and Susan K. Pierce. The initiation of antigen-induced b cell antigen receptor signaling viewed in living cells by fluorescence resonance energy transfer. 6(11):1168–1176, 2005.

[11] Shotaro Otsuka, Jeremy O. B. Tempkin, Wanlu Zhang, Antonio Z. Politi, Arina Rybina, M. Julius Hossain, Moritz Kueblbeck, Andrea Callegari, Birgit Koch, Natalia Rosalia Morero, Andrej Sali, and Jan Ellenberg. A quantitative map of nuclear pore assembly reveals two distinct mechanisms. 613(7944):575–581, 2023.

[12] Malte Wachsmuth, Christian Conrad, Jutta Bulkescher, Birgit Koch, Robert Mahen, Mayumi Isokane, Rainer Pepperkok, and Jan Ellenberg. High-throughput fluorescence correlation spectroscopy enables analysis of proteome dynamics in living cells. 33(4):384–389, 2015.

[13] Antonio Z. Politi, Yin Cai, Nike Walther, M. Julius Hossain, Birgit Koch, Malte Wachsmuth, and Jan Ellenberg. Quantitative mapping of fluorescently tagged cellular proteins using FCS-calibrated four-dimensional imaging. 13(6):1445–1464, 2018.

[14] Stefan Geissbuehler, Noelia L. Bocchio, Claudio Dellagiacoma, Corinne Berclaz, Marcel Leutenegger, and Theo Lasser. Mapping molecular statistics with balanced super-resolution optical fluctuation imaging (bSOFI). 1(1):4, 2012.

[15] Maximilian H. Ulbrich and Ehud Y. Isacoff. Subunit counting in membrane-bound proteins. 4(4):319–321, 2007.

[16] Ankur Jain, Ruijie Liu, Biswarathan Ramani, Edwin Arauz, Yuji Ishitsuka, Kaushik Ragunathan, Jeehae Park, Jie Chen, Yang K. Xiang, and Taekjip Ha. Probing cellular protein complexes using single-molecule pull-down. 473(7348):484–488, 2011.

[17] Johan Hummert, Klaus Yserentant, Theresa Fink, Jonas Euchner, Yin Xin Ho, Stanimir Asenov Tashev, and Dirk-Peter Herten. Photobleaching step analysis for robust determination of protein complex stoichiometries. 32(21):ar35, 2021.

[18] Jon Garry, Yuchong Li, Brandon Shew, Claudiu C. Gradinaru, and Andrew D. Rutenberg. Bayesian counting of photobleaching steps with physical priors. 152(2):024110, 2020.

[19] Prabuddha Sengupta, Tijana Jovanovic-Talisman, Dunja Skoko, Malte Renz, Sarah L. Veatch, and Jennifer Lippincott-Schwartz. Probing protein heterogeneity in the plasma membrane using PALM and pair correlation analysis. 8(11):969–975, 2022.

[20] Sang-Hyuk Lee, Jae Yen Shin, Antony Lee, and Carlos Bustamante. Counting single photoactivatable fluorescent molecules by photoactivated localization microscopy (PALM). 109(43):17436–17441, 2012.

[21] Lekha Patel, David Williamson, Dylan M. Owen, and Edward A. K. Cohen. Blinking statistics

and molecular counting in direct stochastic reconstruction microscopy (dSTORM). *Bioinformatics*, 37(17):2730–2737, 2021.

[22] Geoffrey C. Rollins, Jae Yen Shin, Carlos Bustamante, and Steve Pressé. Stochastic approach to the molecular counting problem in superresolution microscopy. *Proceedings of the National Academy of Sciences*, 112(2), 2015.

[23] Daniel Nino, Nafiseh Rafiei, Yong Wang, Anton Zilman, and Joshua N. Milstein. Molecular counting with localization microscopy: A bayesian estimate based on fluorophore statistics. 112(9):1777–1785, 2017.

[24] Joerg Schnitzbauer, Maximilian T. Strauss, Thomas Schlichthaerle, Florian Schueder, and Ralf Jungmann. Super-resolution microscopy with dna-paint. *Nature Protocols*, 12(6):1198–1228, 2017.

[25] Florian Stehr, Johannes Stein, Julian Bauer, Christian Niederauer, Ralf Jungmann, Kristina Ganzinger, and Petra Schwille. Tracking single particles for hours via continuous DNA-mediated fluorophore exchange. 12(1):4432, 2021.

[26] Johannes Stein, Florian Stehr, Patrick Schueler, Philipp Blumhardt, Florian Schueder, Jonas Mücksch, Ralf Jungmann, and Petra Schwille. Toward absolute molecular numbers in DNA-PAINT. 19(11):8182–8190, 2019.

[27] Johannes Stein, Florian Stehr, Ralf Jungmann, and Petra Schwille. Calibration-free counting of low molecular copy numbers in single dna-paint localization clusters. *Biophysical Reports*, 1(2):100032, 2021.

[28] Ralf Jungmann, Maier S. Avendaño, Mingjie Dai, Johannes B. Woehrstein, Sarit S. Agasti, Zachary Feiger, Avital Rodal, and Peng Yin. Quantitative super-resolution imaging with qpaint. *Nature Methods*, 13(5):439–442, 2016.

[29] Jason Hon and Ruben L. Gonzalez. Bayesian-estimated hierarchical HMMs enable robust analysis of single-molecule kinetic heterogeneity. 116(10):1790–1802, 2019.

[30] Jonathan E. Bronson, Jingyi Fei, Jake M. Hofman, Ruben L. Gonzalez, and Chris H. Wiggins. Learning rates and states from biophysical time series: A bayesian approach to model selection and single-molecule FRET data. 97(12):3196–3205, 2009.

[31] Fang Huang, Tobias M. P. Hartwich, Felix E. Rivera-Molina, Yu Lin, Whitney C. Duim, Jane J. Long, Pradeep D. Uchil, Jordan R. Myers, Michelle A. Baird, Walther Mothes, Michael W. Davidson, Derek Toomre, and Joerg Bewersdorf. Video-rate nanoscopy using sCMOS camera–specific single-molecule localization algorithms. 10(7):653–658, 2013.

[32] Shalin B. Mehta, Molly McQuilken, Patrick J. La Riviere, Patricia Occhipinti, Amitabh Verma, Rudolf Oldenbourg, Amy S. Gladfelter, and Tomomi Tani. Dissection of molecular assembly dynamics by tracking orientation and position of single molecules in live cells. 113(42):E6352–E6361, 2016.

[33] Christopher M. Bishop. *Pattern recognition and machine learning*. Information science and statistics. Springer, 2006.

[34] Ralf Jungmann, Christian Steinhauer, Max Scheible, Anton Kuzyk, Philip Tinnefeld, and Friedrich C. Simmel. Single-molecule kinetics and super-resolution microscopy by fluorescence imaging of transient binding on DNA origami. 10(11):4756–4761, 2010.

[35] Paul W. K. Rothemund. Folding DNA to create nanoscale shapes and patterns. 440(7082):297–302, 2006.

[36] Maximilian T. Strauss, Florian Schueder, Daniel Haas, Philipp C. Nickels, and Ralf Jungmann. Quantifying absolute addressability in DNA origami with molecular resolution. 9(1):1600, 2018.

[37] Lisa S. Fischer, Christoph Klingner, Thomas Schlichthaerle, Maximilian T. Strauss, Ralph Böttcher, Reinhard Fässler, Ralf Jungmann, and Carsten Grashoff. Quantitative single-protein imaging reveals molecular complex formation of integrin, talin, and kindlin during cell adhesion. 12(1):919.

[38] Izzy Jayasinghe, Alexander H. Clowsley, Ruisheng Lin, Tobias Lutz, Carl Harrison, Ellen Green, David Baddeley, Lorenzo Di Michele, and Christian Soeller. True molecular scale visualization of variable clustering properties of ryanodine receptors. 22(2):557–567, 2018.

[39] Fehmi Civitci, Julia Shangguan, Ting Zheng, Kai Tao, Matthew Rames, John Kenison, Ying Zhang, Lei Wu, Carey Phelps, Sadik Esener, and Xiaolin Nan. Fast and multiplexed superresolution imaging with DNA-PAINT-ERS. 11(1):4339, 2018.