ORIGINAL ARTICLE

Open Access

Blind assessment of localisation microscope image resolution

Eric J Rees[1*], Miklos Erdelyi[1,2], Dorothea Pinotsi[1], Alex Knight[2], Daniel Metcalf[2] and Clemens F Kaminski[1]

Abstract

Background: This paper analyses the resolution achieved in localisation microscopy experiments. The resolution is an essential metric for the correct interpretation of super-resolution images, but it varies between specimens due to different localisation precisions and densities.

Methods: By analysing localisation microscopy as a statistical method of Density Estimation, we present a method that produces a blind estimate of the resolution in a super-resolved image. This estimate is derived directly from the raw image data without the need for comparisons with known calibration specimens. It is corroborated with simulated and experimental data.

Results and discussion: Localisation microscopy has a resolution limit equal to 2σ, where σ is the r.m.s. localisation precision, evaluated as an average Thompson precision, Cramer Rao bound, or otherwise. Further, for a limited-sampling case in which there is only one localisation per fluorophore, the expected resolution of an optimised super-resolution image is worsened to approximately 3σ, due to smoothing processes that are necessarily involved in visualising the specimen with limited data. This 2σ or 3σ resolution can be estimated for any localisation microscopy specimen, and this metric can corroborate or replace empirical estimates of resolution. Other quantifiable resolution losses arise from sparse labelling, fluorescent label size, and motion blur.

Keywords: Density estimation, Localisation microscopy, Resolution, Super-resolution

Background

Localisation microscopes can image biological samples, under physiological conditions or even *in vivo*, with an effective resolution better than 20 nm (Betzig et al. 2006, Huang et al. 2008, Schroff et al. 2008). This capability is providing striking new insights into cell biology on this previously inaccessible scale, such as the intracellular layout of the actin cytoskeleton, morphology of membrane complexes, virus particle assembly, and the progression of protein aggregation in neurodegenerative diseases (Hell 2007, Heilemann et al. 2008, Kaminski et al. 2011). In fact the method has the scope to study all manner of intracellular pathology and protein behaviour in unprecedented detail.

Briefly, localisation microscopy techniques such as PALM, GSDIM and (*d*)STORM, work as follows. A traditional fluorescence microscope is used to take video images of the specimen in question. The specimen is labelled with photoswitchable fluorophores, which highlight the protein or tissue of interest. The user applies controlled illumination and solution chemistry conditions so that these fluorophores photoswitch stochastically, which produces a "blinking" fluorescence. This blinking fluorescence, combined with a suitable camera exposure time, enables the measurement of random, bright spots of light which correspond to sparse subsets of the fluorescent labels. The fluorescence images of individual fluorophores are thus separated, and the position of each one can be precisely determined computationally from its image(s), once a sufficient quantity of data has been captured. Commonly the analysis is done by a "Segmentation and Sparse Gaussian Fitting" algorithm (Wolter et al. 2010). Finally, the fluorophore positions are typically visualized as a fluorophore density map (Baddeley et al. 2010). The resulting super-resolution "image" strongly resembles a conventional fluorescence image, because it is likewise an estimate of the specimen fluorescence density. However the fluorophore density

* Correspondence: ejr36@cam.ac.uk
[1]Department of Chemical Engineering and Biotechnology, University of Cambridge, Pembroke Street, Cambridge CB2 3RA, UK
Full list of author information is available at the end of the article

map is typically an order of magnitude better resolved, because it is generated from fluorophore positions determined with high (~10-30 nm) precision.

To correctly interpret images of unknown or little-understood specimens – which is precisely the task for which super resolution imaging is most important – it is crucial to know the resolution of the system. When we refer to the localisation microscopy system, this includes the behaviour of the data fitting and visualisation algorithms that are applied to measurements of a particular specimen, as well as the optical instrument itself. Unfortunately, the resolution achieved by localisation microscopy can vary greatly between samples, because the accuracy of data fitting and visualisation depends on the quality of fluorescence images and background noise that are collected from the specimen. Therefore it is not sufficient to evaluate the instrument resolution using a calibration sample, and then state this as the resolution when studying other objects. Instead, it is important to estimate the particular resolution of each super-resolution image. This can be done empirically when measuring well-known samples such as actin filaments or sparse dye molecules on a coverslip – and this is often done for system calibration (Xu et al. 2012) – because the reconstructed image can be compared with the known structure. For unknown specimens it is necessary to make a blind assessment of how each localisation microscopy image is limited by its own source data. This paper discusses localisation microscopy as a Density Estimation technique (Parzen 1962, Silverman 1986), and thus identifies a resolution limit based on localisation precision – a quantity for which we can obtain a blind estimate from raw image data. Density Estimation theory also addresses the effect of localisation frequency: the best super-resolution image obtained with only one localisation per fluorophore tends to be less well resolved than one constructed with several detections, because successful visualisation of limited data requires broader smoothing. Other issues that can worsen the effective resolution include too-sparse labelling, fluorescent label width, and motion blur, and these effects can also be quantified.

Methods
Sample preparation
Pre-formed rabbit skeletal muscle actin filaments (Cytoskeleton Inc) were diluted in general actin buffer (Cytoskeleton Inc) to a concentration of 2 μM and incubated with a 1 μl in 200 μl dilution of 6.6 μM phalloidin-Alexa 647 solution (Invitrogen). These were allowed to adhere to LabTek glass chambers for 2 hours prior to being imaged. Before addition of actin filaments, LabTek glass chambers were treated with 2 M glycine for 30 min at 37°C and then coated with

0.05% poly L-lysine solution for 30 min at room temperature.

Hela cells (ATCC, CCL-2) were grown to 80% confluency in LabTek chambers in DMEM (Fisher, VX31966021) containing FBS (Fisher VX10500064). Cells were then starved in DMEM without FBS for 60 min followed by incubation with 1 μg/ml EGF-Alexa 647 (Fisher, VXE35351) in complete culture medium for 30 min at 37°C. Finally cells were fixed with 4% formaldehyde for 10 min and then washed with PBS.

Microscopy
Fluorescence images were taken on an Olympus IX71 inverted widefield microscope, with a 100x, 1.49 NA TIRF objective lens. To induce photoswitching, the LabTek chambers were filled with a "switching buffer" solution: 100 mM mercaptoethylamine (MEA) in phosphate buffered saline (PBS, pH 7.4), together with a glucose-enzyme oxygen scavenger (40 mg/ml glucose, 50 μg/ml glucose oxidase, 1 μg/ml catalase). Samples were illuminated with 642 nm laser light at 2 kW/cm². Stacks of 10^4 images were collected at 65 Hz, with 10 ms exposure times.

Image simulation and analysis
Simulated images were generated in MATLAB, after (Thompson et al. 2002). For the single fluorophore simulation, the expected number of photons arriving on each camera pixel was calculated by integrating a Gaussian PSF over the pixel, and allocating the corresponding proportion of N detected photons. To simulate quantum photon counting, this number was replaced with a random number generated from a Poisson distribution with the same mean. Independent, Gaussian random noise of standard deviation b was then added to each pixel. To simulate a structured object (lines of fluorophores at 5 nm spacing, crossing at 45°), a subset of the fluorophores was randomly "activated" for each camera frame, and the cumulative image of all the active fluorophores was generated as above. Gaussian noise was added at the end of this process.

Super-resolution images were reconstructed using a "Segmentation and Sparse Gaussian Fitting" algorithm, similar to (Wolter et al. 2010), implemented in MATLAB. In each camera frame, local maxima brighter than a threshold value were fit to a Gaussian PSF, by an iterative least squares method. For each spot, a localised position was obtained, as well as an estimate for the number of photons of fluorescence, N. The camera noise b was estimated on a framewise basis as the standard deviation of camera pixel brightness in regions without any bright spots. Average localisation precision was evaluated from N and b using Equation 1. Localised fluorophore density was visualised as a "histogram raster," with simulations

visualised in greyscale, and experimental results converted from grey to a colormap for clarity of viewing.

Results and discussion

Precision limit

The super-resolution "image" typically generated by a localisation microscope is a computational reconstruction of the localised fluorophore density. This reconstruction should be better resolved than a fluorescence image. A conventional fluorescence image suffers directly from the diffraction limit, which causes the specimen fluorescence to be blurred by convolution with the optical point spread function (PSF) of the microscope. By comparison, the reconstructed super-resolution image of fluorophore density is also imperfect – in this case the specimen fluorescence density is degraded by an analogous Localisation Error Function (LEF), which describes the uncertainty of determining fluorophore positions by localisation. Usually the LEF is much smaller than the PSF and the reconstructed images are much better resolved than standard fluorescence images. The difference arises because the PSF represents the spread of a single photon position, whereas the LEF represents the smaller spread in determining the centre position of N photons from the same point source.

Fluorescent molecules can be localised as precisely as 10 nm or better in practice, by measuring a few thousand photons, and resolutions of 20 nm can be achieved (Huang et al. 2008). However, the localisation precision Δx depends on properties that vary significantly between samples: primarily the number of photons collected from each fluorophore, and the background noise level. This dependence was derived by Thompson as Equation 1 (Thompson et al. 2002):
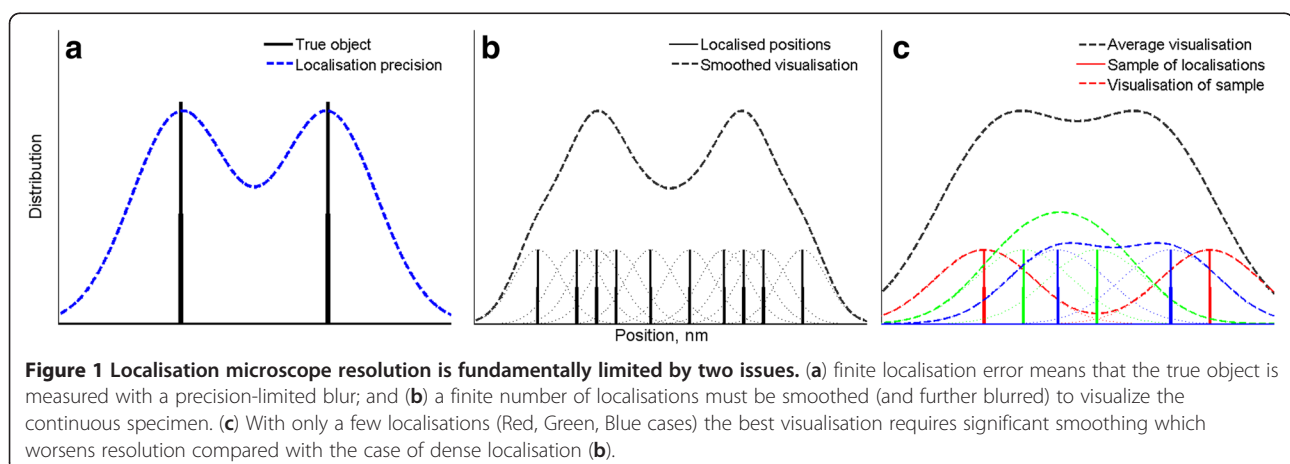
$$< (\Delta x)^2 > = \frac{s^2 + a^2/12}{N} + \frac{8\pi s^4 b^2}{a^2 N^2} \qquad (1)$$

Here, $<\Delta x^2>$ estimates the random variance of a localised fluorophore position, when the optical point spread function is approximately Gaussian with standard deviation s; a is the side-length of the detector pixels; N is the number of photons detected from the fluorophore; and b represents the number of photons of random camera noise in each pixel.

What image resolution is achievable with this localisation precision? Localisation microscope resolution is fundamentally limited by two steps, illustrated in Figure 1: (1) by imperfect determination of fluorophore positions; and (2) by visualising the specimen based on a finite number of fluorophore position samples. Localisation precision concerns the first step, and to quantify its effect on resolution we consider the model scenario of "frequent, unbiased localisation with Gaussian errors." Here, "frequent" measurement means that each fluorophore tends to be localised many times. "Unbiased" measurement means that localisations are detected from each individual fluorophore at the same average rate, and no spurious localisations are made. The positional errors are taken to be independently normally-distributed (assumed experimentally valid), so the probability density for a fluorophore with position $\mathbf{x_0}$ to be localised at position $\mathbf{x}$ is $p(\mathbf{x})$:

$$p(\mathbf{x}) = \left(\sqrt{2\pi}\sigma\right)^{-2} \exp\left(-|\mathbf{x} - \mathbf{x}_0|^2 / 2\sigma^2\right) \qquad (2)$$

Here, $p(\mathbf{x})$ is the LEF, which is defined to have a true variance of σ^2. The $<\Delta x^2>$ in Equation 1 provides an estimate of σ^2. In what follows we denote estimates of σ as $\hat{\sigma}$. The $p(\mathbf{x})$ stated here is two dimensional with circular symmetry. If necessary, $p(\mathbf{x})$ could be made elliptical or three dimensional by replacing the term in the exponential with a more general $\mathbf{x}'\mathbf{A}\mathbf{x}$, and adjusting the normalisation term.

Figure 1 Localisation microscope resolution is fundamentally limited by two issues. (**a**) finite localisation error means that the true object is measured with a precision-limited blur; and (**b**) a finite number of localisations must be smoothed (and further blurred) to visualize the continuous specimen. (**c**) With only a few localisations (Red, Green, Blue cases) the best visualisation requires significant smoothing which worsens resolution compared with the case of dense localisation (**b**).

Let a model specimen be defined by a static 2-dimensional fluorophore density, $\rho_0(\mathbf{x})$, comprised of Dirac delta functions representing M distinct fluorophores.

$$\rho_0(\mathbf{x}) = \sum_{i=1}^{M} \delta(\mathbf{x} - \mathbf{x}_i) \qquad (3)$$

In the model described above, the localisation microscope identifies fluorophore positions drawn from the probability distribution ρ_m.

$$\rho_m(\mathbf{x}) = \frac{1}{M}\rho_0 * p = c\sum_{i=1}^{M} \exp(-|\mathbf{x} - \mathbf{x}_i|^2/2\sigma^2) \qquad (4)$$

In this model, each fluorophore contributes the same probability weight to the measurement. The * denotes convolution, and c is a normalisation constant.

$$c = \left(2M\pi\sigma^2\right)^{-1} \qquad (5)$$

The localisation microscope samples a set of fluorophore localisations, which can be thought of as a collection of n delta functions, having the same form used to model the specimen in Equation 3 (but generally different positions due to finite localisation errors).

$$\rho_s(\mathbf{x}) = \sum_{j=1}^{n} \delta(\mathbf{x} - \mathbf{x}_j) \qquad (6)$$

For simple cases with only a few localisations, ρ_s can be directly visualised as a scatter plot. More typically, however, this set of fluorophore positions is used to generate a fluorophore density map of the specimen, by one of several possible methods (Baddeley et al. 2010). One effective visualisation is to construct a cumulative super-resolution image by plotting each localisation as a small blur (typically a Gaussian) of fluorescence intensity (Bates et al. 2007). This method is known in statistics as Kernel Density Estimation (Parzen 1962), but for simplicity and accessibility we may refer to its use in localisation microscopy as "Gaussian Visualisation." In this method, the point-like localisations are smoothed, by convolution with a suitable kernel function K, to estimate the fluorophore density of the specimen. Kernel Density Estimation obviates the hard-edge artefacts seen in histograms; however it requires the user to identify an optimal kernel width. Using a too-wide kernel for the visualisation can throw away resolution by oversmoothing regions of high sampling density. Adaptive-width methods of kernel density estimation can prevent this problem (Breiman et al. 1977), by plotting narrower kernels in densely sampled regions, and Baddeley's jittered histogram visualisation achieves the same effect.

The Gaussian Visualisation method is tractable for resolution analysis, because it involves a convolution analogous to a point spread function. In Equation 7 we therefore consider the super-resolution image described by ρ_v, which is obtained by smoothing the localisations ρ_s with a kernel K. In Gaussian Visualisation K is a 2-dimensional Gaussian, h_j scales its width to suit the local sampling density, and h_j^{-2} normalises the weight of each localisation in the reconstructed image.

$$\rho_v(\mathbf{x}) = \sum_{j=1}^{n} \frac{1}{h_j^2} K\left(\frac{\mathbf{x} - \mathbf{x}_j}{h_j}\right) \qquad (7)$$

In the model scenario of "frequent" localisations, $n \to \infty$ and ρ_s converges in distribution towards ρ_m, i.e. the localisations perfectly map out the measurement probability. The reconstructed super-resolution image, ρ_v, will then also converge to ρ_m, provided that we "draw" it using a suitably narrow kernel K – tending towards a delta function in the limit of infinite localisations. In this model, with "infinite" localisations, the localisation microscope has generated an image simply by convolving the true fluorophore density with the localisation error function $p(\mathbf{x})$.

Clearly the LEF blurs this ideal super-resolution image in the same manner as a PSF acts on a conventional image. Hence the image resolution can be determined analytically by studying the LEF with an optical resolution criterion. The Sparrow criterion is suitable because it can be evaluated for a Gaussian LEF, whereas the Rayleigh limit is only defined for Airy functions (Sparrow 1916). The Sparrow limit states that point sources are resolvable if they are sufficiently separated that their overlapping images (PSFs) produce a pattern with an intensity minimum lying between their positions. For a Gaussian PSF (or LEF) the Sparrow limit of resolution is satisfied for separations greater than 2σ. Crucially, we can estimate σ for each fluorophore localisation, by applying Equation 1 to the photon statistics (N and b) of the raw image data. This blind estimate ($\hat{\sigma}$) of σ indicates the resolution limit of any super-resolution image. Of course, the estimated localisation precisions typically encompass a spread of values; however since we are assuming independent, normally-distributed errors, we can obtain the estimate $\hat{\sigma}$ from the r.m.s. error (note (Holtzman 1950)).

$$\hat{\sigma} \approx \sqrt{\hat{\sigma^2}} = \sqrt{\sum_{j=1}^{n} \left\langle (\Delta \mathbf{x})^2 \right\rangle_j / n} \qquad (8)$$

If the fluorophore properties are homogenous (in terms of Signal-to-Noise, N and b), then we can take this average over the whole image; or if not then we could evaluate it on a regional basis. We can then state $2\hat{\sigma}$ as the Precision Limit of resolution for this particular super-resolution image. This resolution limit can be corroborated by simulations, as illustrated in Figure 2. It also provides a good indication of resolution that is achieved with real
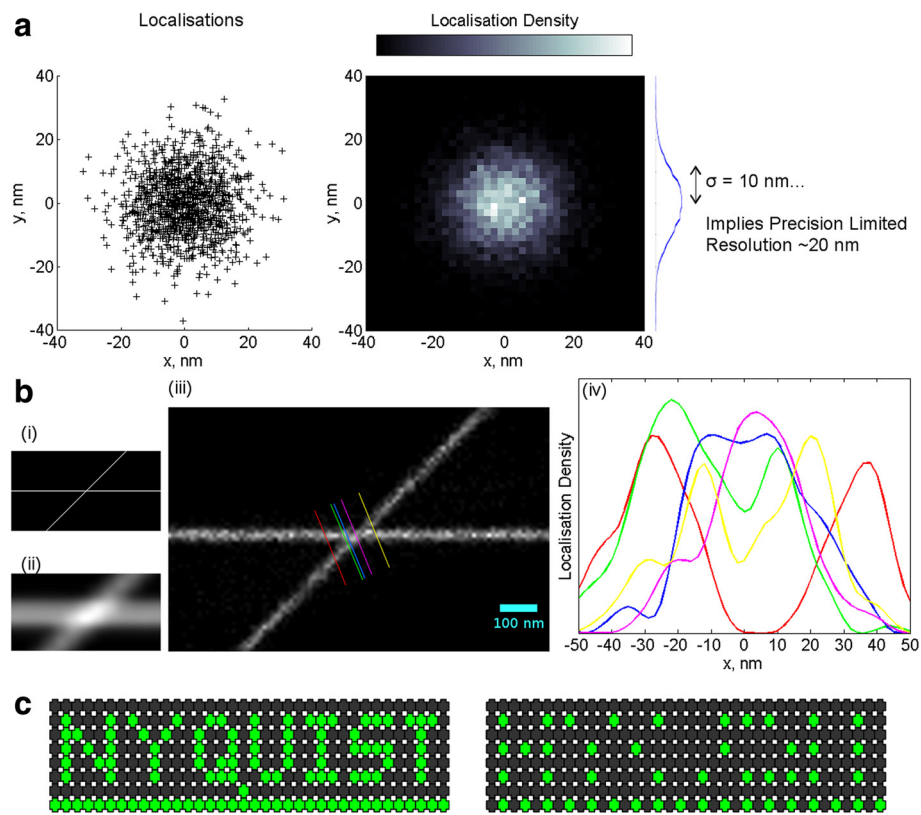
Figure 2 Localisation error, due to finite fluorescence Signal and camera Noise, means that multiple independent localisations of a fluorophore produce a blurred map of fluorescence density, analogous to a Point Spread Function. 10^4 images of a single fluorophore were simulated, each with $N = 10^3$ photons, and parameters ($b = 10$, and $a = s = 160$ nm) such that Equation 1 predicts a localisation error of 9.6 nm. After applying a super-resolution algorithm to these simulated images in MATLAB (using sparse Gaussian fitting similar to (Wolter et al. 2010)), the localised positions were found to have a standard deviation of 9.8±0.25 nm. The average Thompson precision estimated from the image data (without using the known parameters N and b) was similar, at 9.3 nm – with the difference due to slight overestimation of signal, N. This localisation error, normally-distributed, implies the Sparrow criterion of resolution is about 20 nm. (**b**) A simulated object (**i**) consists of 2 lines of fluorophores, at 5 nm spacing, crossing at 45°. Its diffraction limited image is shown in (**ii**), and in (**iii**) a super-resolution image of the localised fluorophore density was reconstructed from data simulated as for the single fluorophore images used in (**a**), but with random activation of fluorophore subsets. A mean of 80 localisations were obtained per fluorophore, which avoids significant sampling limitation. Cross sections through the super-resolved image show that a limiting Sparrow resolution of about 20 nm was achieved (blue line), at the Precision Limit expected from Equations 1 and 8 (**iv**). (**c**) Nyquist-limited resolution: if the localisation precision supports a resolution finer than the fluorescent label spacing, the microscope will (correctly) resolve a pointillist image. This means finer scale information about the underlying specimen is not found, and the size of features smaller than about twice the label spacing cannot be reliably measured.

experimental specimens (Figure 3), although real samples tend to be slightly less well resolved due to the finite number of localisations, which is discussed in the next section.

Some final remarks on precision-limited resolution are important. This blind estimate of the resolution limit, using Equations 1 and 8, is similar to the empirical resolution evaluated by localising a static point source several times, and reporting the FWHM ($\sim 2.4\ \hat{\sigma}$) of the spread of localised positions (Xu et al. 2012). Whereas the empirical method is a direct measure of resolution for calibration samples, the blind estimate is an indirect method that does not require a known specimen. Finally, the principle of precision-limited resolution is not restricted

to estimates of localisation error obtained by Equation 1. In Bayesian localisation of sparse fluorophores, for example, the Cramer-Rao bound can provide a valid estimate of Δx^2 for each fluorophore position (Kay 1993, Ober et al. 2004, Shaevitz 2009).

Sampling limit

Crucially, the Precision Limit is the smallest possible distance that can be resolved by localisation microscopy. This resolution is achieved only if the number of localisations is infinite. Sampling limitation worsens the resolution in practice, when the number of localisations is finite, and the extent of this resolution loss can be
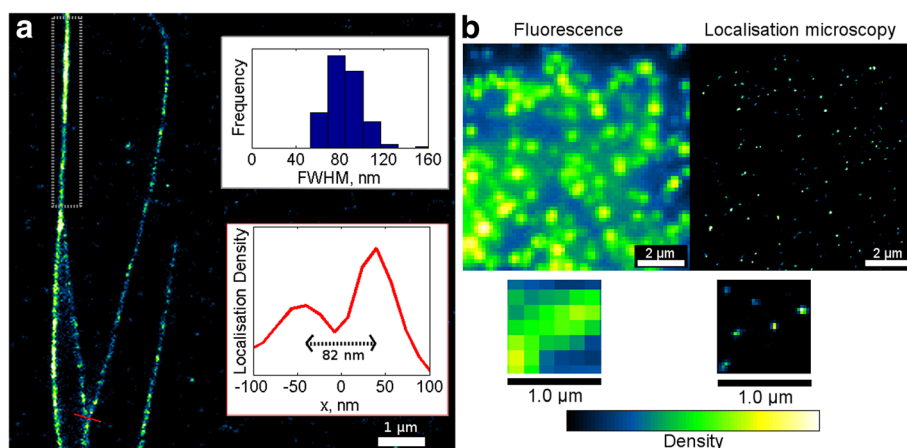
Figure 3 Super-resolution images of: (a) actin filaments stained with phalloidin-Alexa fluor 647 on glass. In (**a**) the Precision Limit of resolution is calculated as 53 nm and the limited sampling resolution is 80 nm, using Equations 1 and 8. The FWHM of the actin filament in the grey boxed area is plotted as an inset histogram to show the empirical resolution, and these FWHM values range from 56 nm to about 110 nm, with a mean of 79 nm. The empirical resolution is broadly consistent with the blind-assessed values that were calculated. The spread of FWHM arises from variable localisation precision, and from random permutations of the finite number of localisations that comprise each cross section. The localisation density of a cross section through two crossing filaments (red line) shows that features can be resolved at separations larger than the calculated resolution limit, as expected. (**b**) Epidermal growth factor (EGF) is conjugated to Alexa fluor 647 on HeLa cell surfaces, and bound to EGF receptors. Activation of EGF receptors results in dimerisation of receptors and clustering into pits and vesicles with diameters ranging from 50–150 nm. A Precision Limit of 60 nm was calculated from the raw data, which provides the user with confidence that the sub-diffraction sized clustering is a real structure.

studied by treating localisation microscopy as a "Density Estimation" task (Silverman 1986) within the model provided by Equations 1–7.

The perfect result of super-resolution would be to determine and visualise the exact fluorescence distribution of the specimen, ρ_0. However the finite localisation error means that the best visualisation achievable without deconvolution is ρ_m, analogous to the way a PSF blurs a conventional image. Given only a finite number of localisations, some smoothing must be applied to ρ_s to obtain the best visualisation, ρ_v. The Gaussian Visualisation method blurs the observations and further reduces the image resolution – so why is it necessary? The answer is given by the theory of Density Estimation, in which the objective is to minimise the difference between the visualisation, ρ_v, and the measured quantity ρ_m (Breiman et al. 1977). If a small sample of localisations is blurred too narrowly, then the visualisation will be spiky, so the difference between ρ_v and ρ_m may be drastic, and the super-resolution image could be misleadingly punctate. On the other hand, too coarse a smoothing process will worsen the resemblance of ρ_v to ρ_m, and also waste resolution. Determining the best width of a kernel smoothing filter is a central task in Density Estimation. The best kernel width will optimise some figure of merit, such as the expected mean square error between the true distribution (ρ_m) and the visualisation ρ_v. The optimum kernel can be identified if the underlying structure (ρ_m) is already known, but this is not usually

possible in practice. Some cases are well studied, however, and provide guidelines for the best kernel width. If ρ_m is known to be a single Gaussian, then an optimum smoothing kernel is itself a Gaussian of standard deviation h, (Silverman 1986).

$$h = \frac{0.9m}{n^{1/5}} \tag{9}$$

In this case m is the underlying standard deviation (estimated by $\hat{\sigma}$), and n is the number of localisations. The $n^{-1/5}$ term means this smoothing filter tends to a delta function as the number of localisations tends to infinity, which means the Precision Limit of resolution will be achieved in this case. For the case of a single localisation, the optimum smoothing filter is a Gaussian of width (standard deviation) $\sim \hat{\sigma}$. This visualisation is optimised for a specimen described by a single Gaussian ρ_m: it would provide a good visualisation when the object is a single fluorophore on a coverslip, or for well-separated fluorophores. Real microscopy specimens are likely to contain multiple structures with various length scales, and for this type of distribution an adaptive kernel width with superior properties was proposed by (Breiman et al. 1977).

$$h_j = \alpha d_{jk} \tag{10}$$

Here h_j is the kernel width which should be applied to the jth localisation, where d_{jk} is the kth nearest neighbour distance of the jth localisation, and α is a scaling constant.

Optimisation of α and k is discussed in (Breiman et al. 1977). The key message is that the optimum kernel width depends on both the local sampling density (distance between localisations) and the underlying structure.

Having evaluated the amount of smoothing required to produce the optimal Gaussian Visualisation of a specimen, the length scale of this visualisation can be used to estimate the best image resolution that can be achieved when the number of localisations is finite. An interesting case is the scenario in which each fluorophore is localised exactly once. This could happen as a consequence of measuring a short data series (e.g. due to rapid bleaching), or intentionally (e.g. in a molecule counting experiment designed to capture a single bright image of each fluorophore, before bleaching it). With only a single localisation per fluorophore, the sampling density will probably not tend towards the "infinite" value needed to achieve the Precision Limit of resolution. The "best" visualisation method indicated by Equation 9 is to smooth the localisations with a Gaussian of standard deviation $\hat{\sigma}$. Let us assume that a sample-dependent treatment with (Breiman et al. 1977), or another optimisation, does not significantly alter this value. In this case the average super-resolution image, $<\rho_v>$, is blurred by convolving the localised positions with a Gaussian visualisation kernel of standard deviation $\hat{\sigma}$. The localisation microscopy image is therefore degraded twice: first by localisation error, and second by the smoothing required for visualisation. Both are convolutions with a Gaussian of standard deviation $\sim \hat{\sigma}$, and this double convolution is equivalent to a single convolution with a Gaussian of standard deviation $\sqrt{2}\hat{\sigma}$ (Hirschman and Widder 1955).

This analysis leads to the remarkably simple result that, whereas the Precision Limit of resolution achieved with an infinite number of localisations is $2\hat{\sigma}$, an image produced with a single localisation per fluorophore has its expected resolution coarsened (assuming optimal visualisation) to about $2\sqrt{2}\hat{\sigma}$, or approximately $3\hat{\sigma}$. Importantly, this is an average or expected resolution: with a limited number of localisations there are many possible permutations of the super-resolution image, as shown in Figure 1c, but $2\sqrt{2}\hat{\sigma}$ is the average resolution. Furthermore, the optimal amount of smoothing required in Gaussian Visualisation is subjective and cannot be determined exactly without knowing the structure of the specimen beforehand. However, a Precision Limit of $2\hat{\sigma}$, and a limited sampling resolution of $3\hat{\sigma}$, can provide a useful blind estimate of the resolution achieved with a real specimen.

These resolution estimates are derived for localisation microscopy using a Gaussian Visualisation. How do alternative visualisation algorithms compare with this? Histogram visualisation (Wolter et al. 2010), for example, involves "binning" the localisations onto a grid of super-resolution pixels. The width of the histogram pixels effectively smoothes or blurs the localisation data, analogous to the smoothing involved in visualisation by convolving the localisations with a Gaussian. Based on this similarity, it seems plausible that a skilful histogram visualisation can achieve a resolution that is similar to the limit obtained by smoothing with a Gaussian, although the histogram is likely to be slightly worse due to hard-edge features. Certainly the histogram method asymptotes to the same Precision Limit, since in the case of infinite localisations a histogram can be plotted using infinitesimal bin widths and hence no loss of resolution due to plotting. The jittered histogram visualisation is based on a principle similar to an adaptive-width kernel, and should achieve the same resolution (Baddeley et al. 2010). A simple scatter plot, however, is a discontinuous visualisation and its resolution therefore cannot be directly compared with the continuous image visualisations considered above.

Nyquist (labelling) limit

Nyquist-limited resolution relates to the fluorescent labelling of a specimen. It is well known that discrete labelling constrains the useful resolution of any imaging method (Schroff et al. 2008). The essential idea is that features smaller than twice the typical separation of fluorescent labels cannot be reliably measured, due to aliasing, even if the precision-limited resolution of the microscope would be sufficient to do so, as shown in Figure 2c. Note that this limit arises from the distance between fluorophores on the sample. This is generally different from the distance between localisations, since a fluorophore might be localised several times; however if the fluorophore is static on the specimen then simply obtaining multiple localisations will not increase the underlying label density, and it is the spacing of the fluorophore labels that determines whether features are measurable. The fluorophore spacing can be estimated based on the known size of antibody labels, or by extrapolating from the density of localisations – in some cases this latter approach is simple, for example if there is precisely one localisation per labelled position (Shim et al. 2012).

Since the Nyquist limit is a concept from Fourier analysis, it strictly refers to subjects that are sampled at regularly spaced intervals. This is not true for localisation microscopy, where fluorophores are usually attached at random positions on their substrate. Further, specimens generally have very dense-labelling in some areas, and very sparse labelling in others. This means that the question of what resolution limit is imposed by label density is more complex than finding the average distance between fluorophores. There is a strong variation in the label density on different parts of a specimen, so fluorophore spacing and hence the Nyquist limit of

resolution must be evaluated for particular features. This is a sample-dependent issue, and the best way to evaluate the Nyquist limit of resolution should be considered on a case-by-case basis.

Nyquist-limited resolution can also arise in print or on-screen images. When super-resolution images are displayed as a square raster, the pixel width must be less than half of the resolution that needs to be illustrated. Therefore the value of $\hat{\sigma}$, from Equation 8, provides an indication of the best pixel size to use when plotting a super-resolution image.

Fluorophore width and corrected motion blur

So far the localisation microscope resolution has been evaluated with respect to the actual fluorescence of the sample. However, the purpose of super-resolution is usually to visualise some fluorescently-labelled structure, and this structure is in fact less precisely determined than the fluorescence due to the finite size of the fluorescent labels. If each fluorophore is localised with a precision σ, but the size of an antibody label causes the fluorescence to be displaced from the structure of interest by a distance d, then the effective localisation precision of the instrument is worsened. If the displacement of each localisation is in a random direction then the loss of precision can be modelled as a random walk. (A random displacement of the localised positions may occur for flexibly-attached fluorophores, or as an approximation in the case of rigidly-attached fluorophores – this approximation would be most valid for fluorophores that are more closely spaced than the localisation precision, σ, so that a localisation drawn from a particular area of the sample may come one of various fluorophores, each with static but random orientations, leading to a random offset in the localisation.) A similar loss of resolution can occur when drift-correction is applied to localisation microscopy data. In such cases, a fiducial marker is often used to track the localisation microscopy specimen, and any translational movement of the marker is subtracted from the localised fluorophore positions to mitigate motion blur. If the fiducial marker is itself localised with a finite random error, σ_{fid}, then its positional error worsens the effective precision of the corrected data. The instrument precision, fluorescent label size and uncertain drift correction can be combined as a random walk to estimate the effective localisation precision of the instrument, σ_{eff}.

$$\hat{\sigma}_{\mathbf{eff}}^2 = \hat{\sigma}^2 + d^2 + \hat{\sigma}_{\mathbf{fid}}^2 \qquad (11)$$

The range of 2-3 $\hat{\sigma}_{\mathbf{eff}}$ then provides a corrected estimate of microscope resolution for the sample in question. Ideally the fluorescent label size and any inexact

correction of motion blur should have a negligible effect, compared with fluorophore localisation precision.

Z-resolution in 3D localisation microscopy

The principle outlined so far is that the resolution of localisation microscopy depends on two factors, (a) localisation precision and (b) smoothing required for visualisation of finite data, and that a blind estimate of this resolution can be evaluated from the same source data that is processed to obtain a super-resolution image. This principle extends to three-dimensions; however the maths needed to evaluate the z-resolution varies between the different implementations of 3D localisation microscopy. Here we outline the general principles involved in assessing z-resolution. The task is simply to process the image data to estimate the standard error of each localised position in the z-direction, as well as in the lateral x and y directions. We also need to specify what it means to visualise a structure in three dimensions. In some methods of 3D localisation microscopy the visualisation may be a z-stack of super-resolution images, and in other methods the mean z-position of localisations may be indicated with a colour code (which is unsuitable for showing features at multiple z-positions). Neither of these methods is perfect. To evaluate the theoretical limit of z-resolution, it is possible to extend Equation 7 into three dimensions, so that the three-dimensional scatter plot of localisations (ρ_{s3}) is smoothed into a three dimensional scalar field of localisation density (ρ_{v3}) – probably with an ellipsoidal smoothing function. This representation can then be sliced or projected onto various planes for visualisation in 2D.

The method required to estimate the error of z-localisation depends on the technique used for 3D localisation. In the astigmatism-based method of 3D localisation, the vertical or horizontal ellipticity of a fluorophore image is used to evaluate its z-position. In this method, a Bayesian method of z-localisation is ideally suited to estimating the z-localisation error at the same time as the z-position (Shaevitz 2009). Alternatively the uncertainty in z-position could be evaluated by algebraic error propagation of the uncertainties in the x- and y-widths fitted to each PSF. In either case, once the uncertainties of the z-positions have been evaluated, the z-resolution follows from the method derived above for 2-dimensions. The Precision Limit of z-resolution is simply twice the estimated localisation error, and the limited sampling case is three times. Unfortunately, this resolution has to be evaluated on a regional basis, rather than being uniform over the image. This is because of an interesting drawback of astigmatism-based 3D localisation microscopy: the optical PSF widths, and hence the Precision Limits of resolution, are dependent on z-position and they are different in the x and y directions. The z-resolution

also varies with z-position, and is asymmetric. This makes it harder to express the resolution of a super-resolution image, even for a single z-slice. This drawback does not afflict the double-helix method of 3D localisation microscopy, since in that system the lateral and axial localisation precisions are translationally invariant over the depth of field (Badieirostami et al. 2010). In the double helix method, the uncertainty of z-positions could be estimated via a Fisher information approach, or by comparison with tabulated calibration data. Another 3D localisation microscopy technique is the biplane method: in this method, the ratio of image intensities measured from different focal planes within the sample is used to determine the z-position of a fluorophore (Juette et al. 2008). This technique also benefits from translationally invariant lateral resolution, and fairly homogenous axial resolution, however the uncertainty of z-position may have to be estimated via a lookup table of calibration data. The lateral localisation precisions could be estimated using the Thompson formula.

A final comment on z-resolution is that the image of a three-dimensional specimen may contain a much higher density of fluorophores than the image of a flat, two-dimensional specimen. In order to achieve sparse fluorescence images of individual fluorophores, subject to the constraint of limited photoswitching ratios, the spacing of fluorescent labels will need to be greater. This means that the Nyquist limit of resolution in 3D localisation microscopy may pose a greater problem than in 2D.

Mislocalisations and uncorrected motion blur
Localisation microscopy is predicated upon the ability to identify individual fluorophore positions, based on sparse fluorescence image data. Quality control methods are applied, either implicitly or explicitly, to exclude imprecise or spurious localisations, but some mislocalisations may persist which are not from single, photoswitching fluorophores. When the localised fluorescence density is visualized, these features may result in various artefacts. Some examples are: persistent fluorescent background features may be visualized as spots of concentrated fluorophore density; overlapping fluorescence signals from close, but separate, features may be visualized as a spurious blur of fluorophore density in between the features; and mobile fluorescent molecules that wander through the focal plane of the microscope may result in "salt and pepper noise" or "ant tracks." Uncorrected mechanical drifts of the specimen can generate "motion blur" in the reconstructed image – although this motion blur is in fact a correct visualization of a moving sample, it is unhelpful and worsens the resolution of the system. Lateral drift may be corrected using a method such as fiducial markers. Drift in the z-direction may also occur. As well as leading to a vertical streaking of the

visualisation – which resembles a worsened z-resolution – z-drift also causes features to move relative to the focal plane of the microscope. Features that lose focus have a broader optical PSF, and hence a worse localisation precision (after Equation 1), and therefore the Precision Limit of the (lateral) resolution is also worsened for specimens that experience z-drift. Fiducial markers may be able to correct the motion blur caused by z-drift, but can do nothing to prevent the loss of resolution that arises from broader PSFs and the consequent loss of (lateral) localisation precision.

Because mislocalisations arise from a failure of the assumption that underlies localisation microscopy – they occur when it is not possible to localise precise positions exclusively from individual fluorescent labels, and so Equations 2–7 do not properly describe the process – their impact on instrument resolution is difficult to quantify.

Conclusion
The resolution of localisation microscopy depends on the fluorescence properties of the specimen, specifically the number and precision of fluorophore localisations. Therefore it is important to assess the resolution achieved with unknown samples by analyzing their particular datasets, since it can differ significantly from the resolution evaluated for a calibration sample. Fluorophore positions determined by a localisation microscope have a finite error, and the spread of these errors is analogous to the optical point spread function of a traditional light microscope. Since the error of each fluorophore localisation can be estimated using Thompson's Equation 1, or as a Cramer-Rao bound, the average localisation precision (σ) of a fluorophore density reconstruction leads to a practical estimate of the resolution limit due to instrument precision of 2σ. This resolution limit is readily calculable from raw image data, and provides a useful guideline for the resolution achieved in a localisation microscopy image. Furthermore, in a case of only one localisation per fluorophore, the average resolution is worsened to about 3σ, because of the quantity of smoothing required to produce a realistic visualisation of the specimen using limited data.

As well as the precision and sampling limits of resolution, a separate constraint on the effective resolution of a localisation microscope can arise as Nyquist limitation due to fluorescent label density. Also, the effective localisation precision may be worsened by fluorophore width or motion blur correction. Spurious localisation data should be excluded by stringent quality control criteria, as it can otherwise generate artefacts that are hard to quantify.

Competing interests

The authors declare they have no competing interests.

Authors' contributions

EJR, DP, and CFK drafted the theory and manuscript. EJR produced the simulations and figures. EJR, AEK, MJE, and DM conceived of the study. Experimental development, imaging, and interpretation were performed by MJE and DM. All authors read and approved the final manuscript.

Author details

[1]Department of Chemical Engineering and Biotechnology, University of Cambridge, Pembroke Street, Cambridge CB2 3RA, UK. [2]Analytical Science Division, National Physical Laboratory, Hampton Road, Teddington TW11 0LW, UK.

References

Baddeley D, Cannell MB, Soeller C (2010) Visualisation of localization microscopy data, Microsc. Microanal 16:64–72

Badieirostami M, Lew MD, Thompson MA, Moerner WE (2010) Three-dimensional localization precision of the double-helix point spread function versus astigmatism and biplane. Appl Phys Lett 97:161103

Bates M, Huang B, Dempsey GT, Zhuang X (2007) Multicolor super-resolution imaging with photoswitchable fluorescent probes. Science 317:1749–1753

Betzig E, Patterson GH, Sougrat R, Lindwasser OW, Olenych S, Bonifacino JS, Davidson MW, Lippincott-Schwartz J, Hess HF (2006) Imaging intracellular fluorescent proteins at nanometer resolution. Science 313:1642–1645

Breiman L, Meisel W, Purcell E (1977) Variable kernel estimates of multivariate densities. Technometrics 19:135–144

Heilemann M, Linde S, Schuttpelz M, Kasper R, Seefeldt B, Mukherjee A, Tinnefeld P, Sauer M (2008) Angew. Chem Int Ed 47:6172–6176

Hell SW (2007) Far-field optical nanoscopy. Science 316:1153–1158

Hirschman II, Widder DV (1955) The Convolution Transform. Princeton University Press, Princeton University Press, New Jersey

Holtzman WH (1950) The unbiased estimate of the population variance and standard deviation. Am J Psychol 63:615–617

Huang B, Wang WQ, Bates M, Zhuang X (2008) Three-dimensional Super-resolution Imaging by Stochastic Optical Reconstruction Microscopy. Science 319:810–813

Juette MF, Gould TJ, Lessard MD, Mlodzianoski MJ, Nagpure BS, Bennett BT, Hess ST, Bewersdorf J (2008) Three-dimensional sub-100 nm resolution fluorescence microscopy of thick samples. Nat Methods 5:527–529

Kaminski GS, Linde S, Erdelyi M, Esbjorner EK, Klein T, Rees E, Bertoncini CW, Dobson CM, Sauer M, Kaminski CF (2011) In Situ Measurements of the Formation and Morphology of Intracellular β-Amyloid Fibrils by Super-Resolution Fluorescence Imaging. J Am Chem Soc 133:12902–12905

Kay SM (1993) Fundamentals of Statistical Signal Processing: Estimation Theory, Prentice-Hall

Ober RJ, Ram S, Ward ES (2004) Localisation Accuracy in Single-Molecule Microscopy. Biophys J 86:1185–1200

Parzen E (1962) On estimation of a probability density function and mode. Ann Math Stat 33:1065–1076

Schroff H, Galbriath CG, Galbriath JA, Betzig E (2008) Live-cell photoactivated localization microscopy of nanoscale adhesion dynamics. Nat Methods 5:417–423

Shaevitz JW (2009) Bayesian Estimation of the Axial Position in Astigmatism-Based Three-Dimensional Particle Tracking. International, Journal of Optics, 896208

Shim SH, Xia C, Zhong G, Babcock HP, Vaughan JC, Huang B, Wang X, Xu C, Bi GQ, Zhuang X (2012) Super-resolution fluorescence imaging of organelles in live cells with photoswitchable membrane probes. PNAS 109:13978–13983

Silverman BW (1986) Density estimation for statistics and data analysis, Chapman and Hall

Sparrow CM (1916) Astrophysical Journal 44:76–87

Thompson RE, Larson DR, Webb WW (2002) Precise nanometer localization analysis for individual fluorescent probes. Biophys J 82:2775–2783

Wolter S, Schuttpelz M, Tscherepanow M, Linde S, Heilemann M, Sauer M (2010) Real-time computation of subdiffraction-resolution fluorescence images, J. Microscopy 237:12–22

Xu K, Babcock HP, Zhuang X (2012) Dual-objective STORM reveals three-dimensional filament organisation in the actin cytoskeleton. Nat Methods 9:185–188