

Fourier Ring Correlation with a single image

Tip ten Brink

Fourier Ring Correlation with a single image

by

Tip ten Brink

Student number: 4927192
Project duration: September 13, 2021 – May 23, 2022
Thesis committee: Prof. dr. B. Rieger, AS, primary supervisor
Dr. P. M. Visser, EEMCS, supervisor
Dr. ir. F. H. van der Meulen, EEMCS
Prof. dr. S. Stallinga, AS

Abstract

As super-resolution methods make it possible to capture images at a resolution beyond the diffraction limit, they have no straightforward measure for optical resolution. Consequently, signal-to-noise ratio-based methods to determine resolution such as Fourier Ring Correlation (FRC) have seen increased use. We look at a new method, called 1FRC, as it requires only a single image instead of two (2FRC, the original method). It splits each pixel value into two values according to a binomial distribution, producing two images usable in a standard FRC routine. We consider for which noise modalities and conditions the results are equivalent to using 2FRC. If the image noise is only Poisson-distributed we derive mathematically that 1FRC and 2FRC are equivalent. Using simulations on Siemens star test images we find that the mean squared error between 1FRC and 2FRC curves is small for images containing only Poisson noise and a combination of Poisson and low variance Gaussian noise. However, when the mean variance ratio μ/σ^2 (image mean divided by variance of a pixel) is less than one, the 1FRC curve no longer goes to zero at high frequencies, but instead fluctuates at an elevated level. We see that this elevated 1FRC curve level is exactly $1 - \mu/\sigma^2$. For $\mu/\sigma^2 < 0.973$ the difference in resolution exceeds one standard deviation of the 1FRC resolution. From this point we the Kolmogorov-Smirnov test confidently states that the 1FRC pixel sum distributions are not equal to 2FRC distributions. We also test 1FRC on an experimental dataset made with varying STED intensity. The computed resolution curves fit well to the modified Abbe equation for STED. However, the 1FRC resolutions are up to 30% better than resolutions obtained using decorrelation analysis.

Contents

List of Symbols	iii
1 Introduction	1
2 FRC and imaging	3
2.1 FRC	3
2.2 Image detection.	5
2.2.1 Probability distribution of pixel counts	6
3 1FRC	7
3.1 Binomial split	7
3.2 Binomial split of Poisson distribution	7
3.3 Poisson-Gauss distribution.	8
3.3.1 Analytical approximation	9
3.3.2 Analysis of pure Gaussian binomial split	9
3.3.3 Moments of compound binomial-Poisson-Gauss distributions	10
4 Evaluation of 1FRC	12
4.1 Siemens star	12
4.2 Mean variance ratio	14
4.2.1 Theoretical computation from moments	15
4.2.2 Comparison with signal-to-noise ratio	15
4.3 Comparing with 2FRC	16
4.4 Pixel sum method to compare image distributions	16
4.4.1 Decomposition using Linnik's theorem	16
4.4.2 Pixel sum statistical test	17
5 Simulations	18
5.1 Comparing 1FRC and 2FRC simulation results.	18
5.1.1 Images with simulated Poisson noise	18
5.1.2 Images with simulated Poisson-Gauss noise	20
5.2 Pixel sum method to investigate distributions of split images	23
5.3 Relation between 1FRC discrepancy and moments	26
5.3.1 Compound distribution moment simulations	26
5.3.2 Discrepancy relation	27
5.3.3 Maximum Gaussian variance	30
5.3.4 Heuristic modification to ensure correct resolution	31
6 Experimental data	33
6.1 STED dataset.	33
6.2 1FRC on STED images	33
7 Conclusion	38
A Appendix A	40
A.1 Implementation	40
A.1.1 Binomial split	40
A.1.2 1FRC	40

List of Symbols

d	Resolution
A	First image resulting from binomial split
B	Second image resulting from binomial split
N	Random variable representing a pixel in a measurement, with subscript $_{pg}$ indicating a Poisson-Gauss, $_{g}$ a pure Gaussian and $_{p}$ a pure Poisson probability distribution
λ	Expectation parameter of a Poisson distribution
I	Input image before binomial split
p	Probability of success in binomial distribution
a	Expectation parameter of a Gaussian distribution
g	Standard deviation parameter of a Gaussian distribution
S	Random variable representing a pixel following a binomial split, with subscript $_{pg}$ indicating a compound binomial-Poisson-Gauss, $_{g}$ a compound binomial-Gaussian and $_{p}$ a compound binomial-Poisson probability distribution
I_{avg}	Average pixel intensity of an entire image
μ	Average/mean value
σ^2	Variance of an arbitrary pixel
μ/σ^2	Mean variance ratio
Σ	Sum of all pixels in an image
\hat{F}_m	Empirical cumulative distribution function with number of samples m
$\langle 1\text{FRC} \rangle_{\text{high } f}$	Average plateau level where a 1FRC curve tends to at high spatial frequencies, specifically the latter half of frequencies

Introduction

Measuring optical resolution is an important routine within the field of microscopy. Without an accurate way to measure resolution, it becomes impossible to objectively compare different optical setups. For the longest time, microscopes have been held back by the physical limit of diffraction. When light enters an imaging system through an aperture, this aperture acts as a secondary source of light, causing diffraction. As a result, light is not exactly focused into a point, but instead forms a diffraction pattern, which puts a limit on resolution. In 1873, Ernst Abbe found that this limit could be calculated as follows:

$$d = \frac{\lambda_{\text{wave}}}{2\text{NA}}, \quad (1.1)$$

where d is the resolution, λ_{wave} is the wavelength of the light used and NA is the numerical aperture, which characterizes the range of angles that can enter the camera, also taking into account the medium in which the lens works.

Other effects, including aberration and construction imperfections, can negatively influence resolution. However, most modern lenses, especially at smaller apertures, are diffraction-limited. Consequently, any method to improve resolution would have to go beyond the diffraction limit. In the past three decades, a variety of methods have done exactly that. These methods are collectively called super-resolution microscopy or optical nanoscopy.

One of the first super-resolution methods was stimulated emission depletion (STED) microscopy.¹ In this method, a fluorophore is excited at the focal point. Around this focal point, a STED beam (of a longer wavelength than the primary excitation) induces stimulated emission, depleting the excited state and selectively deactivating the fluorescence in the primary excitation wavelength. As the area around the target is now no longer fluorescent, the target itself becomes more resolvable. This improves resolution. The stimulated-emission-depletion scales with the intensity of the beam. As a result, the resolution depends on this parameter. Other super-resolution methods are also greatly dependent on multiple parameters. Consequently, it is precisely super-resolution microscopy that requires an effective process to determine resolution.

One of the techniques that has been used to determine resolution based on acquired images is Fourier Ring Correlation (FRC). This technique was originally introduced by Saxton and Baumeister in 1982², as well as by Van Heel et al.^{3,4,5} It has since been applied in a wide variety of fields, primarily in cryo-EM, but more recently also in super-resolution microscopy.⁶ However, determining the resolution using FRC requires two independent images that differ only by noise content. These images are acquired by different means, such as splitting the odd and even frames or by randomly assigning short time series to one of two groups. However, in practice, it is not always possible to do this. Furthermore, in many cases images have already been made without any attention to applying methods to make multiple images that are usable for FRC.

We discuss a new method, due to B. Rieger and S. Stallinga (department of Imaging Physics, TU Delft), which we call 1FRC. This method, under certain conditions, allows the computation of the resolution using a single image, hence the name (we will use the name 2FRC to indicate the original technique using two separately captured images).

The core of this 1FRC method involves randomly splitting all pixels of the single image, element-wise, into two new images. This random split occurs by sampling a binomial distribution for each pixel, which is why we call it a binomial split. The distribution uses a value of one-half for the probability parameter and the pixel value (representing the photon count) as the number of trials parameter. The two images are then used in a standard FRC routine to compute the resolution. For the resulting resolution to be meaningful, it is required that 1FRC is equivalent to performing two separate measurements and

performing a standard FRC directly. This is only the case if the two images resulting from the binomial split have the same probability distribution as if two independent measurements were taken.

The problem we consider in this report is under which conditions 1FRC gives similar results to 2FRC. We specifically look at how the 1FRC performs under different noise modalities. We study this by finding when the binomial split preserves the probability distribution and by directly comparing 1FRC and 2FRC curves and their resulting resolutions. We also apply 1FRC to experimental data.

If the original image contains only Poisson noise, we can mathematically derive that the split images will also consist of only Poisson noise. These new images will look like independent noise realizations. For other probability distributions, this is much more difficult. As it is not always known exactly what noise is contained in an arbitrary microscopy image, it is important to study when it is possible to apply the 1FRC method.

We will describe the process of the original Fourier Ring Correlation technique and the new 1FRC technique. We will also describe the theory behind the binomial split for Poisson noise. The application of 1FRC to a combination of Poisson and Gaussian noise, a common noise modality arising from, for instance, a combination of shot noise and electronic readout noise, will also be studied. For this noise, we are unable to derive analytically whether the binomial split works. Instead, we characterize its statistical moments. These moments are important for the computation of the mean variance ratio after a binomial split, which is the image mean divided by the variance of a single pixel.

To evaluate the performance of 1FRC, we use Siemens star images in our simulations, as these allow us to confirm if the FRC resolutions are in the right order of magnitude. We quantify the difference between 1FRC and 2FRC using the mean squared error for different average image intensities and different noise levels. We also use the two-sample Kolmogorov-Smirnov test on the sums of all pixels in an image to see if the distributions of the original image and the split image are the same, specifically for the case of Poisson-Gauss noise, for which we have no analytical result. We discuss a discrepancy, which presents itself as an elevated FRC curve plateau at high frequencies. This discrepancy occurs when the mean variance ratio is unequal to zero. We then show a linear relationship between this elevated plateau and the mean variance ratio. We quantify the value of the mean variance ratio for which the difference in resolution results falls outside one standard deviation of the 1FRC resolution. Finally, we apply 1FRC to experimental STED data gathered by the Heilemann Group (Johann Wolfgang Goethe-University). This data includes images of the same samples captured at varying STED intensity. We evaluate the performance of 1FRC on this data in two ways. We fit the 1FRC resolutions to the modified Abbe equation for STED, which depends on the STED intensity. We also compare the resolutions to those obtained using an alternative technique called decorrelation analysis.⁷

FRC and imaging

Fourier Ring Correlation (FRC) is a technique used to determine the resolution of images. FRC allows determining the resolution without detailed knowledge of the optical setup. Instead, it requires two images that differ only in their noise content, which for each image must be independently realized. We will seek to expand on a new technique, called 1FRC, which requires only a single image. This new technique also performs an ordinary FRC computation, which we will explain in detail in the following chapter.

2.1. FRC

We are concerned with two-dimensional images made with optical microscopes. For this reason, we constrain ourselves to the two-dimensional FRC. The more general, three-dimensional case (known also as Fourier Shell Correlation) will not be discussed. In two dimensions, we consider two functions representing images A and B : $f_A(\mathbf{x})$, $f_B(\mathbf{x})$, where \mathbf{x} is the two-dimensional location vector. Each image contains noise that is independent of the noise of the other image. The Fourier Ring Correlation is then calculated as follows:

$$\text{FRC}(q) = \frac{\sum_{\mathbf{q} \in \text{circle}} \widehat{f}_A(\mathbf{q}) \widehat{f}_B(\mathbf{q})^*}{\sqrt{\sum_{\mathbf{q} \in \text{circle}} |\widehat{f}_A(\mathbf{q})|^2} \sqrt{\sum_{\mathbf{q} \in \text{circle}} |\widehat{f}_B(\mathbf{q})|^2}}. \quad (2.1)$$

Here q is the radial component of the location vector \mathbf{q} in Fourier space. $\widehat{f}(\mathbf{q}) = \mathcal{F}\{f(\mathbf{x})\}$, for images A and B , where \mathcal{F} indicates the Fourier transform operator. We use a discrete Fourier transform, defined as:

$$F_k = \sum_{n=0}^{N-1} x_n \cdot e^{-\frac{i2\pi}{N}kn}, \quad (2.2)$$

where $\{x_k\}$ is a sequence of complex numbers transformed to a new sequence of complex numbers $\{F_k\}$. No normalization is necessary, as the FRC denominator is already a normalization. To compute the two-dimensional Fourier transform, two sequential Fourier transforms as in equation (2.2) are performed, where $\{x_k\}$ are the rows and then the columns of a image. As we are in Fourier space after the transformation, q is the spatial frequency (in m^{-1}), where $q = 0$ corresponds to zero spatial frequency. We can see an example of a 2D Fourier transform in figure 2.1. For an image of $N \times N$ pixels, the FRC is computed for a radius up to $N/2$, as beyond this the radial circle no longer fits inside the square image and results in no useful value. For the numerator, the complex conjugate of $\widehat{f}_B(\mathbf{q})$ times $\widehat{f}_A(\mathbf{q})$ is computed for each pixel on the circle at radius q , which are then summed. The denominator normalizes the result.




Figure 2.1: **Fourier transformation of a 2D image.** The Fourier transformation of a square, 2D image is shown. Contrast has been significantly increased for better visibility. We use an image captured using STED as input, specifically the Series C sample at maximum STED intensity, which we will show in figure 6.2d. We see higher intensity at the center of the image, which corresponds to the lower spatial frequencies, which contain most of the image content. A Fourier "ring" over which we sum in equation (2.1) is shown in pink, with the corresponding spatial frequency q indicated by the blue dashed line.

The Fourier Ring Correlation (FRC) is a cross-correlation measure, individually evaluated at each spatial frequency, which we will show now. Cross-correlation is defined as follows:

$$(f \star g)(\tau) := \int_{-\infty}^{\infty} f(x)^* g(x + \tau) dx, \quad (2.3)$$

where the superscript '*' indicates the complex conjugate. The Fourier transform is useful because of two important properties relating cross-correlation and convolution. Here we define convolution as:

$$(f * g)(x) := \int_{-\infty}^{\infty} f(\tau) g(x - \tau) d\tau. \quad (2.4)$$

The following first identity shows the relationship between cross-correlation and convolution:

$$f(x) \star g(x) = f(-x)^* * g(x). \quad (2.5)$$

The convolution theorem states that convolution equals multiplication in the Fourier domain, giving us the following property:

$$\mathcal{F}\{f(x) \star g(x)\} = \mathcal{F}\{f(-x)^* * g(x)\} = \mathcal{F}\{f(-x)^*\} \cdot \mathcal{F}\{g(x)\} = \mathcal{F}\{f(x)\}^* \cdot \mathcal{F}\{g(x)\}. \quad (2.6)$$

The final expression equals the expression inside the sum of the FRC. Consequently, we can see the FRC at a particular frequency as equivalent to calculating a normalized cross-correlation summed over that frequency.

Cross-correlation is a measure of similarity. Most noise sources for images, such as shot noise, electronic noise and thermal noise, are broadband, meaning they are widely distributed over the frequency spectrum. However, the meaningful signal tends to be concentrated at lower frequencies, as it

contains patterns and structures whose spatial frequency is usually in the lower region. Furthermore, it is rare for it to vary so significantly to leave behind lots of traces in the higher frequencies.

Next, we note that correlation will be high for the frequencies that are dominated by the meaningful signal, as they should be similar since image A and B should contain the same object. Very differently, the noise is random and dissimilar for each image, meaning frequency components dominated by noise will show little to no correlation. Combining these two facts, the FRC will start out near one at $q = 0$, as the (usually zero-mean) noise will barely affect the overall constant level of the image, meaning the values are most likely equal. As noise starts to dominate and the images become more dissimilar at increasing frequency levels, the FRC drops to zero.

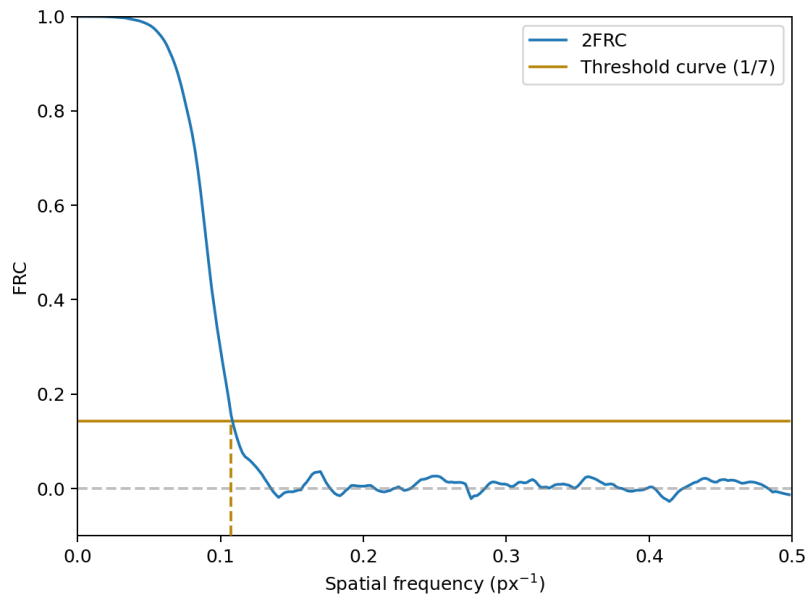


Figure 2.2: **Determining FRC resolution.** We show a smoothed FRC curve. The correlation starts at 1, as the two independently measured images should have the same mean and zero spatial frequency corresponds to the image mean. As the spatial frequency increases, noise starts to dominate the signal and the correlation drops steeply, eventually fluctuating around zero as there is no signal left at these frequencies. A $1/7$ threshold is employed to determine the resolution. The FRC curve intersects this threshold at a spatial frequency of 0.107 px^{-1} , i.e. the resolution is 9.31 pixels.

At some spatial frequency, the FRC becomes "low", indicating the images are now dissimilar and noise dominates. If this only occurs at a high frequency, this means the information content of the images dominates the noise for lower frequencies and could be resolvable. If this occurs at a much lower frequency, this means noise dominates much sooner and the image is resolvable for fewer frequencies. These frequencies can be converted to a measure of distance. The point at which the FRC becomes "low" is the resolution of the image, since it cannot be resolved for higher frequencies (lower spatial distances) as noise dominates the signal for these frequencies. To make "low" more precise and make the FRC comparable to other resolution measures, a standard threshold value is used. The resolution is calculated by determining at which spatial frequency the intersection of the FRC curve with the threshold curve is located. Commonly used thresholds curves include the $1/2$ -bit information curve, σ -based curves (usually 3σ) as well as the constant value of $1/7$.⁸ There is no "best" curve, as FRC is fundamentally a measure of information content and is not completely equivalent to resolution. Based on empirical research, the value of $1/7$ (0.143) seems to correspond closely to visible resolution and will be used in further results.⁶

2.2. Image detection

Before the new method is introduced, we seek to gain more insight in the general process of capturing images. This allows to better understand the applicable constraints and use them to define our new method.

Photons hitting the camera trigger electrons inside the camera, which are then counted. The number

of photons that are assumed to correspond to each measured count is known as the gain. The images we use are assumed to already be gain-corrected, i.e. the number of electronic counts has already been converted to the number of photons. This is important, as a gain can change the pixels' probability distribution. Generally, images are not gain-corrected as a multiplication factor does not matter much for analysis. However, for the 1FRC, assumptions are made about the probability distribution, which no longer hold if the gain is not corrected.

2.2.1. Probability distribution of pixel counts

In general it is a good assumption that the photons appear independently from each other at a constant rate. In this situation, shot noise is an important source of noise. Shot noise is a Poisson process, meaning the number of photons counted can be modelled by a Poisson distribution.

We now make this more precise. Our image is a uniform, square grid of pixel values. Assume a pixel value N follows the Poisson distribution with parameter λ . The Poisson distribution models independent events with a fixed rate, with its output being the number of events in a given time. The parameter λ is the expected number of events in that time (the rate). In the rest of our report, λ will refer only to the expectation parameter of a Poisson distribution.

Consider an observed value of n for our random variable N ($n \in \mathbb{N}$). The probability of this occurring is (as known from the Poisson probability mass function):

$$\mathbb{P}(N = n) = \frac{\lambda^n e^{-\lambda}}{n!}. \quad (2.7)$$

This probability distribution forms the basis for the new method, as it is well-behaved after subjecting it to a binomial split, which we will introduce in the next chapter.

Now that we have introduced Fourier Ring Correlation (FRC), we will discuss our new method, 1FRC. We begin by introducing the binomial split that makes it work and then apply it to Poisson and Poisson-Gauss noise.

3.1. Binomial split

In this report, we use a binomial split to acquire two new images, which we call A and B , from a single image I , for later use in FRC. This is done by splitting individual pixel values according to the binomial distribution.

Consider a pixel (in some image I) modeled by a random variable N , which follows some probability distribution. When performing a binomial split, we sample from the binomial distribution with N as the number of trials and $p \in (0, 1)$ as the probability of success. The result of this sampling is a new random variable X , defined as:

$$X \sim \text{Binomial}(N, p), \quad (3.1)$$

where the binomial distribution is defined as:

$$\mathbb{P}(X = x) = \binom{n}{x} p^x (1 - p)^{n-x}, \quad (3.2)$$

with parameters $(n \in \mathbb{N}, p)$ and $x \in \mathbb{N}$, since the binomial distribution is discrete and the number of successes is nonnegative. The probability distribution of X is known as a compound probability distribution, as a parameter of the distribution is itself a random variable distributed according to some probability distribution. The result of this sampling is the new pixel value in the first new image, A . To obtain the pixel value in image B , Y , the difference between the original image I and the pixel value in image A is calculated (since $X \leq N$, we have that Y is nonnegative):

$$B = I - A, \quad (3.3)$$

so for a single pixel, $Y = N - X$.

The standard computation of the FRC involves making two measurements of the same object, with independently realized noise. For equivalency between the original method (which we will call 2FRC) and the new method, 1FRC, the split images A and B must follow the probability distribution they would draw from if they were measured independently. Therefore, when M is a pixel in the independent measurement following some probability distribution P and X and Y are arbitrary pixels of A , B , respectively, we must have:

$$M \sim P \Rightarrow (X \sim P \text{ and } Y \sim P). \quad (3.4)$$

In other words, for equivalency with 2FRC, images A and B should be realizations of the same probability distribution as an independently realized half-set of the original image.

3.2. Binomial split of Poisson distribution

We seek the probability distribution of the split pixel value X in image A after performing a binomial split on the Poisson-distributed pixel value N . We want X to also be a Poisson random variable. We calculate the joint $\mathbb{P}(N = n \text{ and } X = x)$ using conditional probabilities as follows:

$$\mathbb{P}(N = n \text{ and } X = x) = \mathbb{P}(X = x | N = n) \cdot \mathbb{P}(N = n) = \binom{n}{x} p^x (1-p)^{n-x} \cdot \frac{\lambda^n e^{-\lambda}}{n!}. \quad (3.5)$$

Here, $\mathbb{P}(X = x | N = n)$ is the conditional probability mass function of the binomial distribution, where we know N has been realized as n .

To now calculate the marginal distribution $\mathbb{P}(X = x)$, the joint probability is summed for all possible values of N , which range from x to infinity. This is because $x \leq N$, since x is sampled from a binomial distribution with N as the number of trials and there cannot be more successes than trials. Consequently, $N \geq x$.

$$\begin{aligned} \mathbb{P}(X = x) &= \sum_n \mathbb{P}(X = x | N = n) \cdot \mathbb{P}(N = n) \\ &= \sum_{n=x}^{\infty} \binom{n}{x} p^x (1-p)^{n-x} \cdot \frac{\lambda^n e^{-\lambda}}{n!} \\ &= \sum_{n=x}^{\infty} \frac{1}{x!(n-x)!} \left(\frac{p}{1-p}\right)^x (1-p)^n e^{-\lambda} \lambda^n \\ &= \frac{1}{x!} \left(\frac{p}{1-p}\right)^x e^{-\lambda} \sum_{n=x}^{\infty} \frac{(1-p)^n \lambda^n}{(n-x)!}. \end{aligned} \quad (3.6)$$

By performing the substitution $k = n - x$, we rewrite the final sum, resulting in:

$$\mathbb{P}(X = x) = \frac{1}{x!} \left(\frac{p}{1-p}\right)^x e^{-\lambda} (1-p)^x \lambda^x \sum_{k=0}^{\infty} \frac{((1-p)\lambda)^k}{k!}. \quad (3.7)$$

The infinite sum can be recognized as the power series definition of the exponential function with variable $(1-p)\lambda$. This allows us to rewrite the probability as follows:

$$\begin{aligned} \mathbb{P}(X = x) &= \frac{e^{-\lambda}}{x!} \left(\frac{p}{1-p}\right)^x (1-p)^x \lambda^x e^{(1-p)\lambda} \\ &= \frac{(p\lambda)^x e^{-p\lambda}}{x!}. \end{aligned} \quad (3.8)$$

This is equal to the probability mass function of a Poisson distribution with parameter $p \cdot \lambda$. For $p = 0.5$, we have a Poisson distribution with half the mean, which is equivalent to capturing half the original image I independently. This means we satisfy our requirement in equation (3.4). Consequently, we expect 1FRC to work well on images with purely Poisson noise.

3.3. Poisson-Gauss distribution

In addition to the inherent Poisson noise of the photon generation process, the readout process often introduces additional noise, which is generally modelled by Gaussian noise. Gaussian noise follows the normal distribution, which is shown with mean a and variance g^2 in the following equation:

$$f(n) = \frac{1}{\sqrt{2\pi g^2}} e^{-\frac{1}{2} \frac{(n-a)^2}{g^2}}. \quad (3.9)$$

For the rest of this report, a will only refer to the mean of the Gaussian noise and g^2 to the Gaussian variance.

In this case, the observed photon counts follow a compound distribution characterized by adding two random variables that are Gaussian and Poisson. If we take N_{pg} to be the resulting random variable, we can define it as:

$$N_{pg} = P + G, \quad (3.10)$$

if P is a Poisson random variable and G is a Gaussian random variable. The resulting probability density function is a convolution of the Poisson and Gaussian probability functions. From hereon we call this the Poisson-Gauss distribution. Assuming a gain of unity and zero mean Gaussian noise, we give the probability density as follows⁹:

$$p_{N_{pg}}(n) = A \frac{1}{\sqrt{2\pi g^2}} \sum_{q=0}^{\infty} \frac{1}{q!} e^{-\lambda} \lambda^q e^{-\frac{1}{2} \frac{(n-q)^2}{g^2}}. \quad (3.11)$$

Here A is a normalization constant, λ is the expected number of photons in that pixel and g^2 is the variance of the Gaussian readout noise.

To determine whether the binomials split is valid also for the Poisson-Gauss distribution, we want to compute the probability distribution of the compound random variable S_{pg} , defined as:

$$S_{pg} \sim \text{Binomial}(N_{pg}, p), \quad (3.12)$$

which is similar to equation (3.1), but N replaced by N_{pg} . For this purpose, we seek to evaluate the marginal probability $\mathbb{P}(S = s)$ from the joint probability as with the pure Poisson distribution (the process in equations (3.5)-(3.8)).

However, equation (3.11) is a difficult expression not in a closed form (due to the infinite sum), making it difficult to work with. We were unable to work out the marginal probability into a workable expression that allows it to be compared to the original probability distribution, which should be of the same type for the binomial split to be valid. This means we must resort to approximations and simulations to study the behavior after the split empirically.

3.3.1. Analytical approximation

For small values of g^2 compared to λ (as discussed by Huang et al.⁹), the expression in equation 3.11 can be approximated by a more manageable analytic function:

$$p_N(n) = \frac{e^{-(\lambda+g^2)} (\lambda + g^2)^n}{\Gamma(n + 1)} \quad (3.13)$$

If we assume a gain of unity and zero mean Gaussian noise (as we did before) and only look at integer offsets and variance, we ensure that n is integer, allowing us to write equation (3.13) as follows, since $\Gamma(z + 1) = z!$ for integer z :

$$\mathbb{P}(N = n) = \frac{e^{-(\lambda+g^2)} (\lambda + g^2)^n}{n!} \quad (3.14)$$

This equation is an ordinary Poisson distribution with parameter $\lambda + g^2$. Since it is a Poisson distribution, the result derived in section 3.2 also holds. This means that, in cases where equation (3.13) is a good approximation, we can assume the binomial split is mathematically valid also for Poisson-Gauss noise. However, g^2 is not always small compared to λ and in that case the approximation does not help us. We shall not quantify exactly for which values of g^2 this is the case and instead look directly at the performance of the 1FRC for different values g^2 compared to λ .

3.3.2. Analysis of pure Gaussian binomial split

Since we could not analytically derive the distribution function of S_{pg} , defined as in equation (3.12), we seek to simplify the situation by considering noise that is only Gaussian. The new compound random variable S_g is defined as:

$$S_g \sim \text{Binomial}(N_g, p), \quad (3.15)$$

where N_g is a Gaussian random variable with a probability distribution function as defined in equation (3.9). However, this requires modifying the normal distribution as it involves mixing a discrete distribution (binomial) and a continuous distribution (normal), as can be seen in the incorrect equations below:

$$\begin{aligned}\mathbb{P}(S = s \text{ and } N = n) &= \mathbb{P}(S = s | N = n) \cdot \mathbb{P}(N = n) = \binom{n}{s} p^s (1-p)^{n-s} \cdot \frac{1}{g\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{n-\lambda}{g}\right)^2} \\ \mathbb{P}(S = s) &= \sum_n \mathbb{P}(S = s \text{ and } N = n) dn = \sum_s \binom{n}{s} p^s (1-p)^{n-s} \cdot \frac{1}{g\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{n-\lambda}{g}\right)^2} dn.\end{aligned}\quad (3.16)$$

In the final part of equation (3.16), the normal distribution component becomes a different distribution, as only integer values of n would be evaluated.

In practice, the data is rounded before a binomial split is performed. An example of a probability mass function that represents a discrete normal distribution is¹⁰:

$$\Phi((n+1-\lambda)/g) - \Phi((n-\lambda)/g). \quad (3.17)$$

Here, Φ is the cumulative distribution function of a standard normal, which has no analytical form. For increasing values of λ , we can expect the effect of rounding to become negligible. Furthermore, for high n ($np(1-p) \geq 10$) it is well-known that a binomial distribution (n, p) is approximated by a normal distribution with mean np and variance $np(1-p)$. Subsequently, we instead try to compute the probability density function by using two normal distributions instead, which are both continuous:

$$\begin{aligned}p_{S,N}(s, n) &= p_{S|N}(s|n)p_N(n) = \frac{1}{\sqrt{2\pi np(1-p)}} e^{-\frac{1}{2}\frac{(s-np)^2}{np(1-p)}} \cdot \frac{1}{g\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{n-\lambda}{g}\right)^2} \\ p_S(s) &= \int_0^\infty np_{S|N}(s|n)p_N(n)dn = \int_0^\infty \frac{1}{2\pi g\sqrt{np(1-p)}} e^{-\frac{1}{2}\left(\frac{(s-np)^2}{np(1-p)} + \left(\frac{n-\lambda}{g}\right)^2\right)} dn.\end{aligned}\quad (3.18)$$

For $n < 0$ we assume that λ will be sufficiently high that the probability is near zero (for example, for $\lambda = 100$ and $g^2=500$, this probability is less than 0.01%. For lower levels of g^2 , the probability becomes even smaller). From this integral, we can make no more progress. Furthermore, if we seek to numerically evaluate it, this gives us no hints as to why Gaussian noise leads to problems for 1FRC. This would require further analyzing the resulting probability distribution (which we would only know numerically) when subjected to Fourier Ring Correlation, which itself is non-trivial to compute even on known distribution functions.

3.3.3. Moments of compound binomial-Poisson-Gauss distributions

While we could not analytically derive the distribution functions of S_{pg} (compound binomial and Poisson-Gauss) and S_g (compound binomial and pure Gaussian), we can instead characterize the distributions by looking at their moments. From the law of total variance we know, if we have two random variables N and S :

$$\text{Var}(S) = \mathbb{E}[\text{Var}(S|N)] + \text{Var}(\mathbb{E}[S|N]). \quad (3.19)$$

Furthermore, the expected value for a random variable S , distributed according to a compound probability distribution with parameter N a random variable, can be computed in a straightforward way:

$$\mathbb{E}[S] = \mathbb{E}_N[\mathbb{E}_S[S|N]]. \quad (3.20)$$

Pure Gauss

We first look at the case where S_g is a binomial random variable (n, p) and N_g is a purely Gaussian random variable (a, g^2) that determines the binomial parameter n . This makes S_g 's distribution a compound probability distribution. However, the support of N_g includes any real value, while the support of S_g consists of only the non-negative integers. However, with a high a and corresponding g^2 , negative values can be mostly eliminated. Furthermore, we are only interested in large values of a , as the range of outcomes is large and even after rounding values to the nearest integer the curve remains quite smooth and well-normalized. Under these conditions, the normal distribution and binomial distribution

have nearly identical support, justifying the use of the law of total variance. We now calculate the variance of our random variable S_g :

$$\begin{aligned}\text{Var}(S_g) &= \mathbb{E}[N_g p(1-p)] + \text{Var}(N_g p) \\ &= ap(1-p) + g^2 p^2 \\ &= \frac{1}{4}a + \frac{1}{4}g^2 \quad (p = 0.5).\end{aligned}\tag{3.21}$$

We also compute the expectation, which is $\mathbb{E}[S] = \mathbb{E}[Np] = p\lambda$, equal to $a/2$ for $p = 0.5$. We see that for the range $g^2 < a/4$, the variance is less than the expectation.

Poisson-Gauss

We can also derive the variance of the Poisson noise with added Gaussian noise. In this case the parameter n of our random variable S_{pg} is determined by a different random variable, N_{pg} . This variable N_{pg} is the sum of a Poisson random variable P (parameter λ) and Gaussian random variable G (mean $a = 0$, variance g^2). We compute:

$$\begin{aligned}\text{Var}(S_{pg}) &= \mathbb{E}[N_{pg} p(1-p)] + \text{Var}(N_{pg} p) \\ &= p(1-p)\mathbb{E}[P+G] + p^2\text{Var}(P+G) \\ &= p(1-p)(\mathbb{E}[P] + \mathbb{E}[G]) + p^2(\text{Var}(P) + \text{Var}(G)) \\ &= p(1-p)\lambda + p^2(\lambda + g^2) \\ &= \frac{1}{2}\lambda + \frac{1}{4}g^2 \quad (p = 0.5),\end{aligned}\tag{3.22}$$

where we assume that P and G are independent, in which case the variance of the sum is the sum of the variances. For $p = 0.5$, the value is $\lambda/2$ more than only Gaussian noise (if we associate λ with a). More importantly, since the expected value is $\mathbb{E}[S] = \mathbb{E}[Np] = p\mathbb{E}[P] + \mathbb{E}[G] = p\lambda$, equal again to $\lambda/2$, we see that the variance is always greater than or equal to (the latter only if there is no Gaussian noise, i.e. $g^2 = 0$) the expected value.

Comparison to pure Poisson

For Poisson noise, we have already mathematically derived the exact distribution after performing a binomial split. We can still compute the moments as above for comparison:

$$\begin{aligned}\text{Var}(S_p) &= \mathbb{E}[N_p p(1-p)] + \text{Var}(N_p p) \\ &= p(1-p)\mathbb{E}[N_p] + p^2\text{Var}(N_p) \\ &= p(1-p)\lambda + p^2(\lambda) \\ &= \frac{1}{2}\lambda \quad (p = 0.5).\end{aligned}\tag{3.23}$$

Here, the expectation is simply $\mathbb{E}[S] = \mathbb{E}[Np] = p\lambda$, equal to $\lambda/2$ for $p = 0.5$. We see that the expectation and variance are equal. We can also observe that for $g^2 = 0$, the moments of Poisson-Gauss reduce to the moments of pure Poisson noise. Furthermore, if $\lambda \gg g^2$, we have:

$$\frac{1}{2}\lambda + \frac{1}{4}g^2 \approx \frac{1}{2}\lambda.\tag{3.24}$$

In that case, we see that the moments again reduce to Poisson and that the variance is approximately equal to the expectation.

Evaluation of 1FRC

While we can mathematically justify the use of 1FRC on Poisson-distributed images, we must also evaluate it on simulated and real-world images. Furthermore, due to the complexity of the Poisson-Gauss distribution, no mathematical derivation could be made that shows 1FRC works also on images with Poisson-Gauss noise. The analytical approximation for Poisson-Gauss is Poisson and we know 1FRC works for Poisson, but we must still quantitatively show under which conditions.

We first introduce the primary sample image for our simulations, after which we discuss the mean variance ratio, an important quantity related to the coefficient of variance. We then discuss how we compare 1FRC and 2FRC curves using the mean squared error (MSE). Finally, we explain a method used to determine the distribution of a sample image by analyzing the distribution of the sum of its pixel values.

4.1. Siemens star

For most simulations, we will use a Siemens star as test image, shown in figure 4.1. Its primary use is that the resolution of the image can be discerned visually by inspecting the center of the Siemens star. At infinite resolution, the arms only meet exactly at the center. However, at some point away from the center, the arms cannot be distinguished as the distance between them has become too small. At this radius, which we call r , the actual distance between the arms can be calculated, which offers a value close to the image's resolution.




Figure 4.1: **Siemens star**. A Siemens star is used for most simulation images as it is a simple object that also provides a convenient way to estimate the resolution visually. This can be done by looking at the distance between two arms at the radius from the center where the arms can just be resolved. This Siemens star image is 512x512 pixels and has ten arms. The distance between the arms at the edge of the image is 80.1 pixels, decreasing linearly towards the center.

The distance between n_{arms} (white) arms at the star's outer radius can be calculated by constructing a regular polygon with $2n_{\text{arms}}$ sides as the star has arms. Then, the length of its side is calculated by

considering a $2n_{\text{arms}}$ -polygon, which consists of $4n_{\text{arms}}$ right triangles with inner angle $2\pi/4n_{\text{arms}}$. We can then derive the arm side length L for a star with n_{arms} arms and radius R to be:

$$L = 2 \sin\left(\frac{\pi}{2n_{\text{arms}}}\right) R. \quad (4.1)$$

Since the distance between the arms decreases linearly to zero from L as you near the center, an estimate for the resolution d of the image is:

$$d = \frac{r}{R} L = 2r \sin\left(\frac{\pi}{2n_{\text{arms}}}\right), \quad (4.2)$$

where r is the distance at which the arms can no longer be resolved (where they start to blur together). To ensure a clear resolution cut-off and to prevent sharp edges throwing off the Fourier transform and achieve greater similarity to a real image, a Gaussian filter is applied to the simulation images. The Gaussian filter smooths the image by removing high spatial frequency components. This is shown in figure 4.2.




Figure 4.2: Base Siemens star with Gaussian filter ($\sigma = 3$ pixels). Average pixel intensity $I_{\text{avg}} = 130$.

Furthermore, the background is set to a nonzero level. If this is not done, the Poisson noise would also be zero there, meaning the Poisson noise will have little impact on the resolution of the Siemens star arms. An example of an image with Poisson noise is shown in 4.3a.




Figure 4.3: **Determining resolution of a Siemens star image.** (a)-(c) show a 512x512 pixel Siemens star. (a) shows the full image without markings. The image has an average pixel intensity I_{avg} of 130. It is Gaussian filtered as in image 4.2 with a strength of $\sigma = 3$ pixels. Poisson noise was added afterward. The purple line in (b) shows the maximum distance between the stars L , while the red line shows the radius r at the point where the arms can no longer be distinguished. The blue line shows the the vertical location where the distance between the arms indicates the resolution. The resolution d can then be determined according to equation (4.2). Figure (c) shows the same blue and red lines as (b), but with the area outside a circle around the red line blacked out. To estimate the resolution, the largest circle where the arms cannot be resolved should be chosen. We cannot resolve the arms in (c).

4.2. Mean variance ratio

To evaluate different strengths of Gaussian noise added to simulation images, we use the following quantity:

$$\frac{\mu}{\sigma^2}, \quad (4.3)$$

where μ is the mean of the signal and σ^2 is the variance of an arbitrary pixel (so not including inherent variance from the structure of the image, but including e.g. the variance caused by performing a binomial split). We will use these symbols exclusively for these quantities in the rest of this report. We continue to use g^2 for the variance of just the Gaussian noise. The mean variance ratio μ/σ^2 is related to the signal-to-noise ratio and coefficient of variance. We will use this quantity as it strongly influences the performance of the 1FRC, independent of factors such as image size or the height of the input pixel values.

4.2.1. Theoretical computation from moments

We derived the moments of the binomial distribution compounded with the Gaussian, Poisson and Poisson-Gauss distributions in section 3.3.3. Subsequently, we can compute the theoretical values for the mean variance ratio, all for a binomial split with $p = 0.5$.

For pure Gaussian noise:

$$\begin{aligned}\frac{\mu}{\sigma^2} &= \frac{\frac{1}{2}a}{\frac{1}{4}a + \frac{1}{4}g^2} \\ &= \frac{2a}{a + g^2},\end{aligned}\quad (4.4)$$

for pure Poisson noise:

$$\frac{\mu}{\sigma^2} = \frac{\frac{1}{2}\lambda}{\frac{1}{2}\lambda} = 1, \quad (4.5)$$

and finally for Poisson-Gauss noise:

$$\begin{aligned}\frac{\mu}{\sigma^2} &= \frac{\frac{1}{2}\lambda}{\frac{1}{2}\lambda + \frac{1}{4}g^2} \\ &= \frac{2\lambda}{2\lambda + g^2}.\end{aligned}\quad (4.6)$$

We will use these equations extensively to compute the mean variance ratio, which allows us to predict when problems with 1FRC occur in the case of high variance Gaussian noise.

4.2.2. Comparison with signal-to-noise ratio

A metric more widely used than the mean variance ratio is the signal-to-noise ratio (SNR), which we define as follows, where we assume I to correspond to μ :

$$\text{SNR (in dB)} = 20 \log_{10} \left(\frac{I}{\sigma} \right). \quad (4.7)$$

As μ is often constant in the experiment, with the Gaussian variance set so a desired μ/σ^2 is achieved, we prefer to compute the SNR with μ/σ^2 and μ . If we set $\mu/\sigma^2 = v$, we can compute:

$$\begin{aligned}\frac{\mu}{\sigma^2} &= v \\ \sigma^2 &= \frac{\mu}{v} \\ \sigma &= \sqrt{\frac{\mu}{v}},\end{aligned}\quad (4.8)$$

Since we can calculate:

$$\begin{aligned}\frac{\mu}{\sigma} &= \frac{\mu}{\sqrt{\frac{\mu}{v}}} \\ &= \sqrt{\mu \cdot v},\end{aligned}\quad (4.9)$$

the mean variance ratio $v = \mu/\sigma^2$ as in equation (4.3) can be converted to SNR as follows:

$$\text{SNR} = 20 \log_{10} (\sqrt{\mu \cdot v}). \quad (4.10)$$

We will use the SNR to make the values we use for the mean variance ratio more insightful, as the SNR is more widely used.

4.3. Comparing with 2FRC

Another method to evaluate the performance of the 1FRC is comparing the curve directly to a 2FRC curve, as our goal is achieving the same results. We compare both the resolution estimates as well as the difference between the two curves.

We quantify the difference between the curves using the mean squared error (MSE):

$$\text{MSE} = \frac{1}{N_{\text{curve}}} \sum_{i=1}^{N_{\text{curve}}} |y_{1\text{FRC},i} - y_{2\text{FRC},i}|^2. \quad (4.11)$$

Here, N_{curve} is the number of points in the FRC curves, and $y_{1\text{FRC},i}$ and $y_{2\text{FRC},i}$ are the y -values of the 1FRC and 2FRC curves, respectively. While the noisiness of the curves near the end does contribute to the MSE, we empirically found the difference between the curves at the start to be greater than the noise at the end, justifying also computing the MSE over the entire curve.

4.4. Pixel sum method to compare image distributions

In section 3.2 we derived that if a pixel is Poisson distributed and it is subjected to a binomial split, it remains Poisson-distributed. This is a requirement for 1FRC to work, as we seek results equivalent to the 2FRC, where the two independent images will also be Poisson-distributed. However, we are unable to mathematically derive the same fact for Poisson-Gauss noise. Independent measurements for 2FRC will both have Poisson-Gauss distributions, so we must also have that a Poisson-Gauss image remains Poisson-Gauss-distributed after a binomial split. Again, we aim to satisfy equation (3.4).

In a mathematical derivation, we can consider an arbitrary pixel. However, in our simulation we aim to show that an entire picture remains Poisson-Gauss-distributed. If we can use all pixels in the image at once on a realistic test image, we get more information than from simulating a single pixel. We do this by performing a statistical test against the sum of all pixel values in an image, the pixel sum.

We again define an original image I , with a binomial split of this image giving two new images A and B . We will aim to show that the distributions of the pixel sums of A and I are identical (after compensating for a different mean). For this, we will make use of Raikov's theorem¹¹ and Linnik's theorem.¹² We will state Linnik's theorem, from which it is known Raikov's theorem immediately follows.

4.4.1. Decomposition using Linnik's theorem

First, let G be a normal random variable with parameters a, g^2 according to equation (3.9). Next, let P be a Poisson random variable with parameter λ according to equation (2.7). Suppose G is independent of P and let us construct the random variable

$$Z = G + P. \quad (4.12)$$

Suppose we have another decomposition of Z into the sum of two independent quantities Z_1 and Z_2 :

$$Z = Z_1 + Z_2. \quad (4.13)$$

Linnik's Theorem. A convolution of Gaussian and Poissonian laws can only be decomposed into similar convolutions. More precisely: Let $g^2 > 0$ and $\lambda > 0$. For the existence of a decomposition of the form (4.13) it is necessary (and, trivially, sufficient) that there should exist the decompositions

$$Z_j = Z_{j1} + Z_{j2}, \quad (4.14)$$

where the random variables Z_{jk} are all independent and moreover the Z_{j1} are normal and the Z_{j2} are Poissonian random variables. Here Z_j is either Z_1 or Z_2 as in equation (4.13).

Raikov's theorem is analagous to Linnik's theorem and follows directly from it. However, for Raikov's theorem we have that Z is Poisson-distributed and Z_1 and Z_2 are both also Poisson random variables.

We can use Linnik's theorem by considering the sum of all pixel values in an image (each a random variable P_i). Let Σ be this sum, defined as:

$$\Sigma = P_1 + P_2 + \dots + P_{N \times N} \quad (4.15)$$

If we show that the sum Σ follows a Poisson-Gauss (or Poisson) distribution, that must mean that according to Linnik's (or analogously according to Raikov's) theorem any decomposition into independent random variables must consist of Poisson-Gauss (or Poisson) random variables. So we decompose Σ into two random variables:

$$\Sigma = \Sigma_1 + \Sigma_2. \quad (4.16)$$

We repeat this process again for Σ_1 and Σ_2 , and then keep applying it for each resulting decomposition, eventually arriving at a decomposition of $N \times N$ random variables, exactly as in equation (4.15). As we have repeatedly applied Linnik's theorem, we know each pixel P_i must also be distributed as a Poisson-Gauss random variable. Analogously, we can use Raikov's theorem if the sum follows a Poisson distribution. In that case, each individual pixel follows the Poisson distribution.

We now have a single value, the pixel sum Σ , which represents the entire image, that we can perform a statistical test on to determine if it is indeed the Poisson-Gauss (or Poisson) distribution.

4.4.2. Pixel sum statistical test

The Kolmogorov-Smirnov (K-S) test is used to evaluate the total pixel sum Σ (as in equation (4.15)). To perform a K-S test, the empirical (cumulative) distribution function (ECDF) must be computed. The ECDF converges to the underlying cumulative distribution function. The ECDF \hat{F}_m is defined as follows for m independent and identically distributed observations X_i :

$$\hat{F}_m(x) = \frac{\# \text{ of observations } \leq x}{m} = \frac{1}{m} \sum_{i=1}^m 1_{[-\infty, x]}(X_i), \quad (4.17)$$

where $1_{[-\infty, x]}(X_i)$ is the indicator function, equal to 1 if $X_i \leq x$ and equal to 0 otherwise. Based on this, the one-sample K-S test uses the following test statistic:

$$T_m = \sup_x |\hat{F}_m(x) - F(x)|, \quad (4.18)$$

where T_m is the test statistic and $F(x)$ is the known distribution function. However, using a known distribution function is not practical, as the exact parameters of the Poisson-Gauss/Poisson distribution that the pixel sum would be compared to are unknown. Furthermore, the rounding and clipping applied prior to the binomial split modifies the distributions slightly, meaning that a pure Poisson-Gauss/Poisson distribution is not an appropriate comparison, as the K-S test is very sensitive. Therefore, the two-sample K-S test is performed, which uses the following test statistic, with $\hat{F}_{1,k}(x)$ the ECDF of the first sample and $\hat{F}_{2,k}(x)$ the ECDF of the second sample:

$$T_{m,k} = \sup_x |\hat{F}_{1,m}(x) - \hat{F}_{2,k}(x)|. \quad (4.19)$$

To apply the two-sample K-S test to the pixel sum, we take one sample from binomially split images and another from two images with independent noise. We can then compute the test statistic based on the ECDF of both samples. Furthermore, this allows us to compute a p -value. For low values, we can reject with confidence the fact they have the same distribution, which would undermine our expectation that the binomial split preserves the distribution.

Unfortunately, no test exists that allows the opposite, i.e. rejecting that the two distributions are different. This is a consequence of the fact that there can always be tiny differences between distributions that no test statistic could pick up. As a result, the best we can do is to provide evidence that the distributions before and after the binomial split are not too dissimilar to be the same distribution.

A possible objection to this method is that as a result of the central limit theorem, these sums (which are of the order of tens of thousands to millions of pixels) might approximate a simple normal distribution. To verify if this is the case, the sums are also to be compared to realizations of normally distributed pixels. However, even if they are indeed nearly normally distributed, this is not a problem as this will also be the case for the distribution of the pixels in the non-split image.

Simulations

In this chapter, we evaluate the performance of 1FRC for test images with simulated Poisson and Poisson-Gauss noise. We first compare 1FRC curves to 2FRC curves using the Fréchet distance and also look at the difference between their resolution estimates. We then use the pixel sum method to investigate the probability distribution of the images after the binomial split to justify the use of 1FRC. This is done both for images with simulated Poisson and Poisson-Gauss noise. Finally, we investigate the problems that occur with the 1FRC when Gaussian noise with high variance is added to images.

5.1. Comparing 1FRC and 2FRC simulation results

5.1.1. Images with simulated Poisson noise

The performance of 1FRC and 2FRC is compared using the images introduced in section 4.1. We use the 512×512 pixels Siemens star of figure 4.1. A Gaussian filter is then applied as in figure 4.2. Next, Poisson noise is added to the picture, where each pixel value is the parameter for a draw from a Poisson distribution, with the new pixel value set to the result from that draw. An example of added Poisson noise can be seen in figure 5.1, for two values of the average pixel intensity I_{avg} (100 and 1000).




Figure 5.1: **Two examples of the 512x512 Siemens star simulation images.** For (a) we have that average pixel intensity $I_{avg} = 100$ and for (b) $I_{avg} = 1000$. Both have a Gaussian filter of $\sigma = 3$ pixels applied before Gaussian noise is applied.

We can estimate the resolution using the Siemens star as seen in figure 4.3. This is a subjective method and not highly accurate, but serves as a good baseline and provides an order of magnitude. In figure 5.2 we compare the 1FRC and 2FRC for different values of the average pixel intensity I_{avg} , where the background pixels are set to 55% of the white foreground star pixels. The value of 55% is constant throughout this report and was chosen empirically. This value allows a meaningful estimation of the resolution using the visual method of section 4.1, as at that level the Gaussian filter and Poisson noise have a visible effect on the resolution.

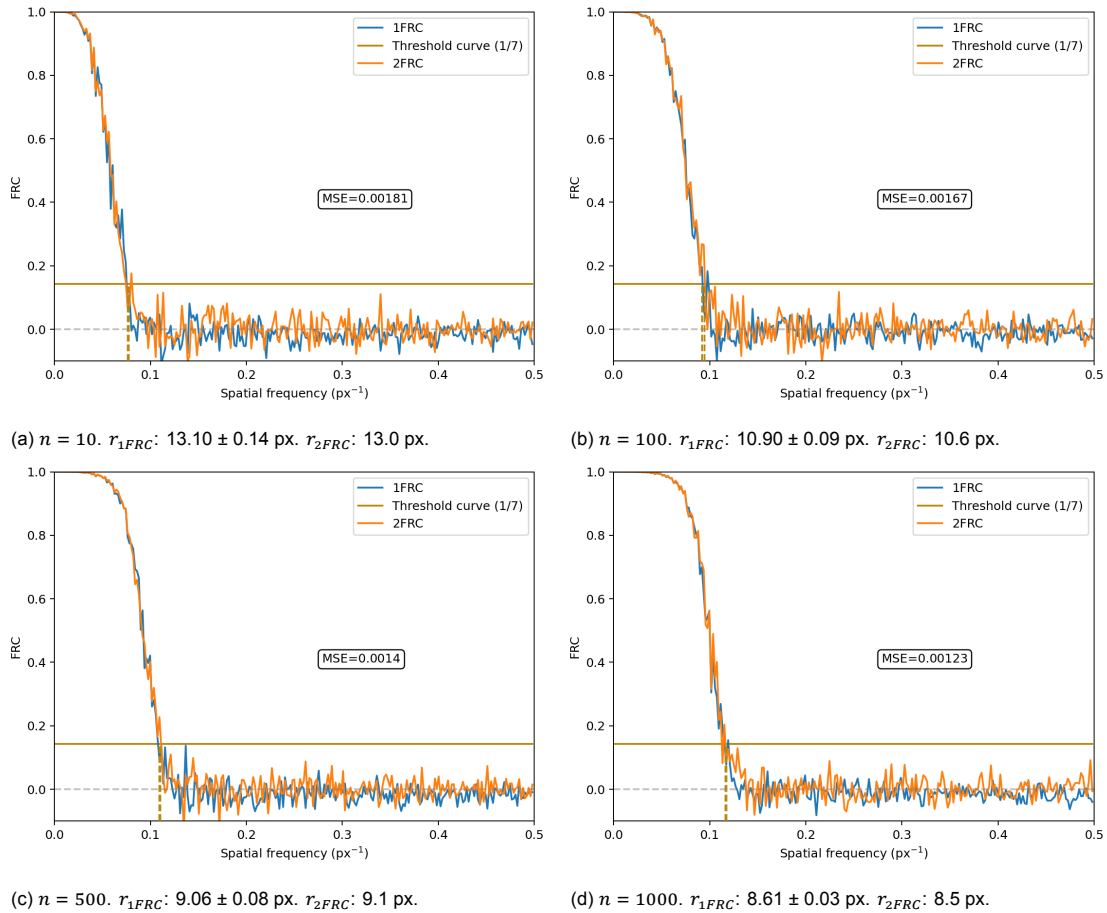


Figure 5.2: **Comparison of 1FRC and 2FRC for simulated Poisson noise.** Figures (a)-(d) contain FRC curves computed using a Siemens star image as a base with simulated Poisson noise added. Before noise is added, the image is adjusted so the average pixel count equals a certain intensity I_{avg} . The background is a constant 55% of the intensity of the foreground. Furthermore, a Gaussian filter with $\sigma = 3$ pixels is also applied before the noise is added. The 1FRC resolution r_{1FRC} is computed by averaging five 1FRC curves per image (by splitting five different times). The uncertainty is the standard deviation of the resolution results of the individual curves. The MSE between the curves is less than 0.002 for all four values of I_{avg} . The Poisson noise is a limiting factor, as reducing the count I_{avg} worsens the resolution. The 1FRC and 2FRC methods agree well for this simulation. On average, the 1FRC resolution r_{1FRC} differed 1.3% from the 2FRC resolution r_{2FRC} for these four figures.

Figure 5.2 shows the basic comparison of 1FRC and 2FRC for simulated Poisson noise added to a Siemens star image. The simulation is performed for four average pixel values I_{avg} (10, 100, 500 and 1000). The background is 55% of the pixel value of the foreground. This value is chosen to allow for significant variance in the background to increase the influence of the Poisson noise on the resolution.

A Gaussian filter with $\sigma = 3$ pixels is used to ensure a cutoff of spatial frequencies, which allows for a well-defined resolution for a simulation image. The specific value ensures enough blur that the resolution is not so high that determining the resolution visually using the Siemens star becomes too difficult. Furthermore, the value also does not cause too much blur so as to completely overshadow the effect of the Poisson noise on the resolution.

The MSE between the curves never exceeds a value of 0.002. The 1FRC resolutions differ 1.3% from the 2FRC resolutions on average. Three times the 1FRC gives worse resolutions, once it is indistinguishable.




Figure 5.3: **1FRC vs. 2FRC resolutions for different I_{avg} .** Different values of the average pixel value I_{avg} are used to generate different versions of the Siemens star (again, with the background pixels set to 55% of the foreground). For each value of I_{avg} , the average of ten 2FRC and 1FRC resolution computations is taken. Furthermore, the 1FRC curve is the average of three 1FRC curves for each of those computations.

We can see in figure 5.3 that the resolutions computed by 1FRC and 2FRC lie very close together across a wide range of signal levels. Ten FRC curves were computed for each value of I_{avg} . On average, 1FRC resolution gives 0.4% lower values than the 2FRC.

For each image, the resolution was also estimated using the Siemens star method by looking at the radius where the arms could no longer be resolved. For all pictures, this resulted in resolutions of the same order of magnitude as the resolutions given by the 1FRC. This is expected, as the 1FRC corresponds well with the 2FRC and the performance of 2FRC has already been proven.

5.1.2. Images with simulated Poisson-Gauss noise

Similar to images generated with Poisson noise, we can also compare FRC curves made of images with simulated Poisson-Gauss noise. However, while the choice of parameter for Poisson noise is clear, as it relates directly to the count, this is not the case for the normal distribution. The Gaussian noise is mostly caused by readout noise, which can vary significantly depending on the setup. For this reason, different variance levels are simulated, namely for $\mu/\sigma^2 = 0.99, 0.98, 0.97, 0.95, 0.9, 0.5$ (1 is equivalent to pure Poisson noise), where this quantity is to be understood as in equation (4.3), i.e. the final image mean and split image pixel variance. Two examples of the effect of Gaussian noise on an image can be seen in figure 5.4. The mean of the Gaussian noise is always taken to be zero, as we also did in section 3.3. If we convert these values to a more standard measure, such as the signal-to-noise ratio (SNR) as defined in equation (4.10), we must also take the average pixel intensity I_{avg} (which is equal to μ) into account. For $I_{avg} = 500$, we have that the SNR values are 26.95 dB, 26.90 dB, 26.86 dB, 26.77 dB, 26.53 dB, 23.98 dB for the above values of μ/σ^2 , respectively.




Figure 5.4: Two examples of the 512x512 Siemens star simulation images with average pixel intensity $I_{\text{avg}} = 100$. For (a) we have that $\mu/\sigma^2 = 0.99$ and for (b) $\mu/\sigma^2 = 0.5$. Both have a Gaussian filter of $\sigma = 3$ pixels applied before Gaussian noise is applied.

Figure 5.5 shows six different comparisons between 1FRC and 2FRC for images with Poisson-Gauss noise. For the first time, we see evidence that the binomial split is not always equivalent to the 2FRC. We see that, instead of dropping to zero, the 1FRC curve fluctuates around a nonzero plateau at higher frequencies. The higher the variance, the higher the final plateau. In further sections, we will seek to further quantify and investigate this difference between the 1FRC and 2FRC curves. However, it is clear that for low, practical values of the Gaussian noise (a)-(c), the 1FRC and 2FRC curves are nearly indistinguishable and have similar resolution values.

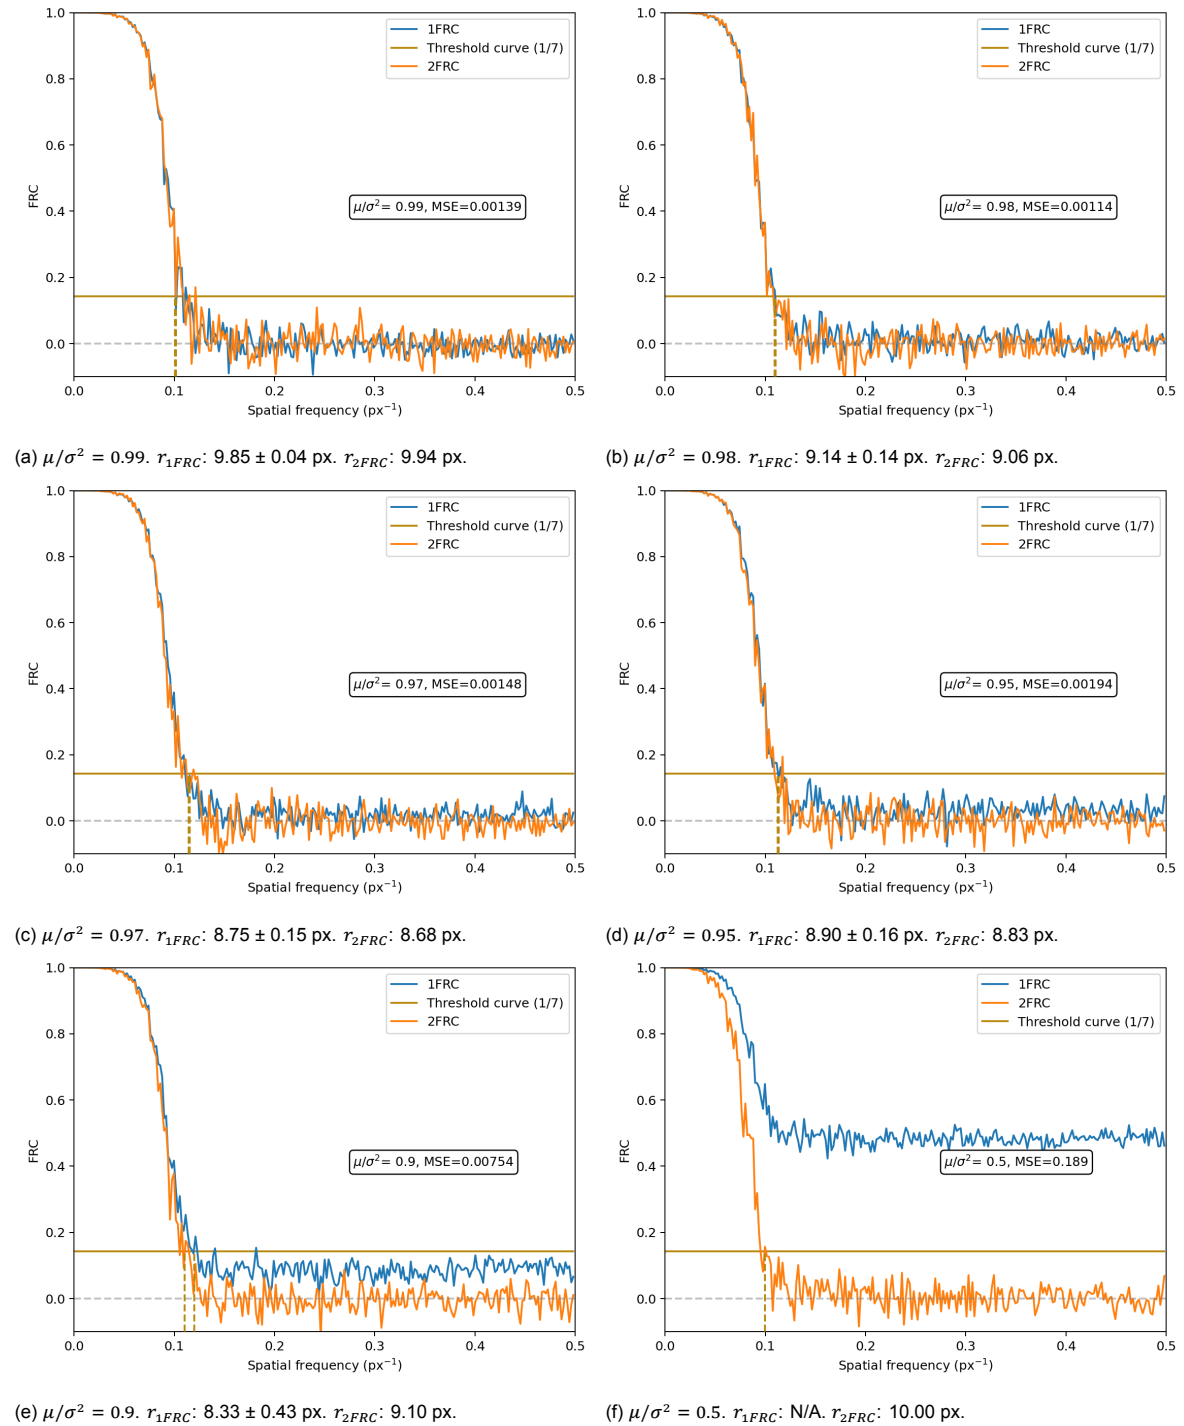


Figure 5: **Comparison of 1FRC and 2FRC for simulated Poisson-Gauss noise.** (a)-(e) contain FRC curves computed using a Siemens star image as a base with simulated Poisson and Gaussian noise added, for different values of $\mu/\sigma^2 = 0.99, 0.98, 0.97, 0.95, 0.9, 0.8, 0.5$. These values are chosen to highlight the region where we are interested in, namely the case of low Gaussian variance where the results can still be accurate. The average pixel value is set to $I_{avg} = 500$. As before, the background pixels are set to 55% of the foreground pixels. Furthermore, a Gaussian filter with $\sigma = 3$ pixels is also applied before the noise is added to the image. The 1FRC resolution r_{1FRC} is computed by averaging five 1FRC curves per image (by splitting five different times). The uncertainty is the standard deviation of the resolution results of the individual curves. For curves (a)-(c), not much difference can be observed. However, for (d) the MSE starts rising faster and (e) shows a significant difference between the resolutions. For (f) it even becomes impossible to compute the 1FRC resolution as there is no crossing with the threshold.

5.2. Pixel sum method to investigate distributions of split images

From observing the 1FRC and 2FRC curves and comparing their resolutions and the MSE between them, we have established that images with Poisson noise agree well, while problems occur for images with Poisson-Gauss noise once the variance becomes high. We will seek to further study these findings by comparing the probability distribution of images after a binomial split to the probability distribution they should have if they were independently captured (as in 2FRC).

Due to Linnik's theorem applied to the case of Poisson-Gauss noise (and Raikov's in the case of pure Poisson noise), we know that if the sum of all pixel values in an image follows a Poisson-Gauss distribution, the individual pixels must do so as well. We will use the empirical distribution function (ECDF) of pixel sums of many images for the Kolmogorov-Smirnov (K-S) test to determine whether the binomially split pixel sums have the same distribution as the reference pixel sums. The series of reference pixel sums is computed from images H_j (where j is 1, 2, ..., k). No binomial split will be performed as they represent the 2FRC case. The images all have a fixed average pixel value of $I_{\text{avg}}=250$. The images H_j differ only in their noise content, which is an independent realization of Poisson-Gauss (or Poisson) noise. Subsequently, we compute the pixel sum Σ , giving us the series $\Sigma_{H,k}$. From all these pixel sums, we compute the reference ECDF $\hat{F}_{H,k}$.

The series of 1FRC pixel sums is computed from images A_i (where i is 1, 2, ..., m). These images are identical to the original images H_j , except that they have a fixed average pixel value of $I_{\text{avg}}=500$. This is double the value for H_j and compensates for the fact that the binomial split exactly halves the average pixel value (as we use $p = 0.5$). Like the reference series, Poisson-Gauss (or Poisson) noise is added to the image. Unlike for the reference series, we now perform a binomial split on each image and then compute the pixel sums. This gives us the series $\Sigma_{A,m}$, from which we compute the 1FRC ECDF $\hat{F}_{A,m}$.

We do not necessarily need to have that $m = k$, as this is not necessary to compute the K-S test. Furthermore, as the computation of $\hat{F}_{A,m}$ is much more computationally intensive due to having to perform a binomial split for each image, it might be useful to have $k > m$ to improve the accuracy of the K-S test. For the images, we use the Siemens star with Gaussian filter (figure 4.2). However, as this routine is computationally intensive, a smaller version is used, of only 64×64 pixels.

From the ECDFs $\hat{F}_{A,m}$ and $\hat{F}_{H,k}$ we compute the Kolmogorov-Smirnov (K-S) test statistic and compute the corresponding p -value. If the p -value of the K-S test performed on these two sample series is below 0.05, the null hypothesis that these samples come from the same distribution is rejected (and we conclude the distributions are different). For higher values, the test sees little distinction and the null hypothesis is not rejected. However, this is not sufficient statistical proof that they are from the same distribution, only that they are not so different as to rule out them being identical. It does, however, give confidence that they could come from the same distribution.

Poisson results

We compute the ECDFs $\hat{F}_{A,m}$ and $\hat{F}_{H,k}$ for $m, k=10,000$, which are shown in figure 5.6. While we have already shown, both mathematically and from simulations, that 1FRC and 2FRC correspond closely for pure Poisson noise, we will first demonstrate the pixel sum method on pure Poisson noise. Consequently, we only add Poisson noise to the images, with the pixel value as the parameter. We also compute a third ECDF $\hat{F}_{G,l}$ with $l=10,000$, which is similar to $\hat{F}_{H,k}$ but instead of Poisson noise we use normally distributed noise with mean and variance equal to the pixel value, approximating the Poisson distribution.

The distributions of the pixel sums of the split and original images correspond closely in the figure. The sums with Poisson-like normal distributed noise added are also similar. This is a result of the fact that for high values, the normal distribution is a close approximation of the Poisson distribution. This does not undermine the validity of this method, as the non-split sums are also very close to the normal distribution sums.

Small random fluctuations can have large impacts on the p -value of the K-S test performed on the ECDFs. This is a consequence of the fact that the K-S test statistic looks only at the maximum difference between the ECDFs. For this reason, we compute the ECDFs on a large number of samples (10,000). Unfortunately, increasing the amount of samples increases the computation time significantly. Due to the fluctuating p -value, the entire experiment is run ten times to quantify the uncertainty. From this, the average p -value is 0.61 ± 0.17 .




Figure 5.6: **Pixel sum empirical distribution function of Poisson noise images.** (a) shows the empirical distribution functions $\hat{F}_{A,m}$ (split images), $\hat{F}_{H,k}$ (halved independent images) and $\hat{F}_{G,l}$ (halved independent images with normal approximation of Poisson), where $m, k, l=10,000$. Furthermore the p -values are computed for the Kolmogorov-Smirnov test between the pixel sums of the images A and H as well as for A and G , with values 0.769 and 0.5, respectively. (b) shows the average over 5 times of the absolute difference between $\hat{F}_{A,m}$ and $\hat{F}_{H,k}$ for 10,000 samples. The K-S test statistic is computed from the maximum of this graph, which in this case is near 0.01 and concentrated in the center.

The null hypothesis is that the distributions are equal. As such, it cannot be rejected for $p=0.61$. This further corroborates the validity of the 1FRC method for Poisson noise. However, the uncertainty is still high and a K-S test cannot prove that distributions are indeed equivalent. Furthermore, if they were indeed exactly equal, we would expect a p -value closer to 1.

Poisson-Gauss results

In figure 5.7, we compute the ECDFs for the split images A_i and halved images H_j with Poisson-Gauss noise instead of pure Poisson noise, again with 10,000 samples per ECDF. The graphs of the ECDFs are shown for three different variance levels of the added Gaussian noise. Here, the split clearly affects the distribution for higher levels of variance. In practice, finding such high levels of variance is rare, but it still shows that for Poisson-Gauss noise, the binomial split only approximately preserves the distribution.

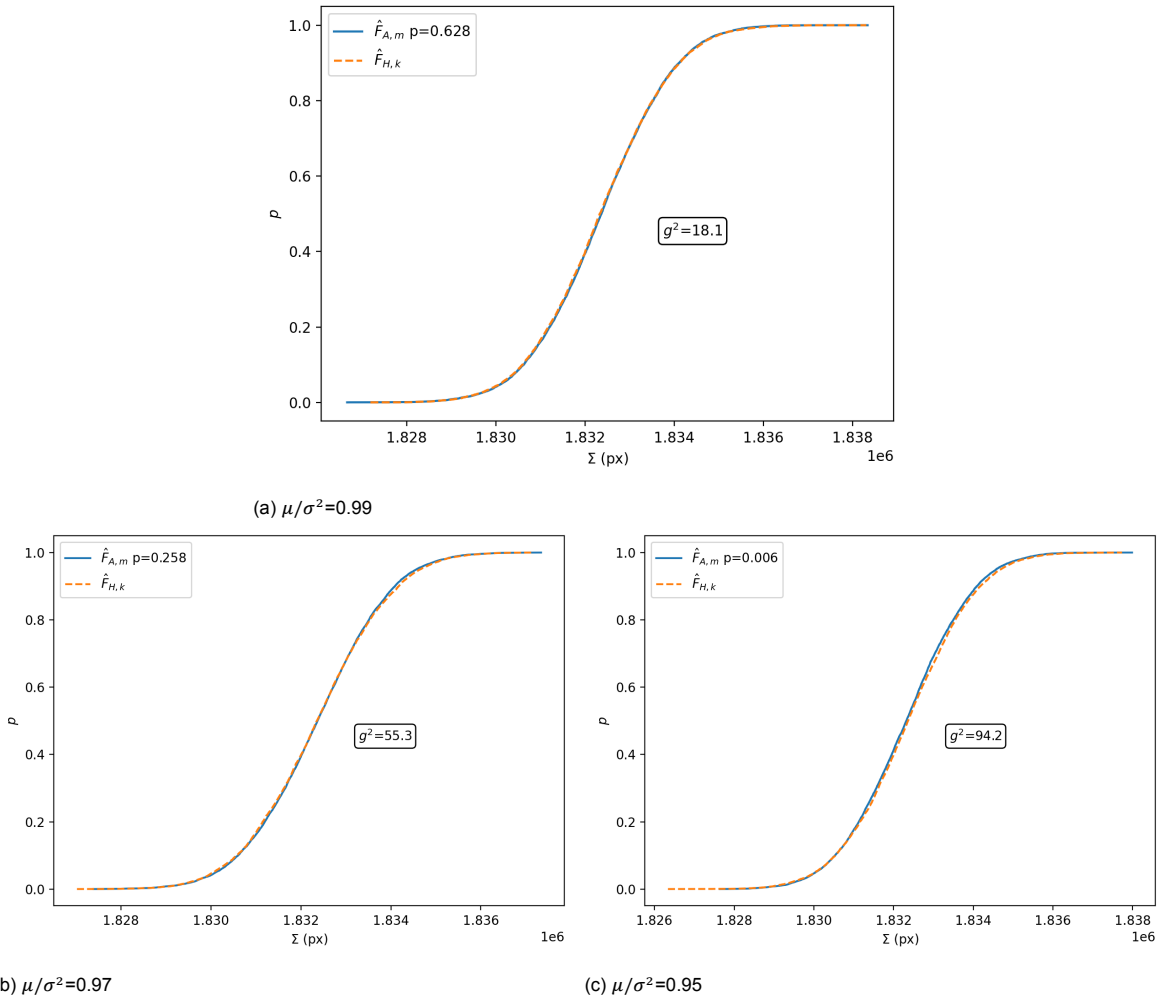


Figure 5.7: **Pixel sum empirical distribution function for Poisson-Gauss noise images.** (a)-(c) show the empirical distribution functions $\hat{F}_{A,m}$ (split images), $\hat{F}_{H,k}$ (halved independent images) and corresponding p -values for three different values of μ/σ^2 , the split image A 's mean divided by the split single pixel variance. The corresponding variance of the Gaussian noise g^2 is shown in the diagram, which is also applied to the non-split image H . As the variance level increases, the ECDF curves correspond less closely and the p -value is below 0.05 for $\mu/\sigma^2=0.95$, which corresponds to a Gaussian variance of $g^2=94.2$ (compare this to the image mean of 894). Both ECDF curves are computed from 10,000 samples each ($m, k=10,000$).

Figure 5.8 shows a more extreme example of high variance. In this picture, the final variance is around two times the image mean. At this point, the difference curve between the two ECDFs also takes a more stable form, as can be seen in subfigure 5.8b.




Figure 5.8: **Pixel sum empirical distribution function for Poisson-Gauss noise images with very high Gaussian variance.** (a) shows the empirical distribution functions $\hat{F}_{A,m}$ (split images), $\hat{F}_{H,k}$ (halved independent images). Both ECDF curves are computed from 10,000 samples each ($m, k=10,000$). (b) shows the average over 5 times (differences of 5 runs are averaged) of the absolute difference between $\hat{F}_{A,m}$ and $\hat{F}_{H,k}$, each computed from 10,000 samples. For this high level of variance, it can be seen that the ECDF curves diverge strongly near the edges, indicating the binomial split distorts the variance.

5.3. Relation between 1FRC discrepancy and moments

We have seen that if the Gaussian variance reaches a high level, the 1FRC no longer drops to zero, but instead levels off at some value greater than zero. This value is dependent on the variance and mean of the final distribution. In section 3.3.3 we derived the following general relations for the variance of the compound probability distributions of S_g (pure Gaussian binomial split) and S_{pg} (Poisson-Gauss binomial split):

$$\begin{aligned} \text{Var}(S_g) &= \frac{1}{4}a + \frac{1}{4}g^2 \quad (p = 0.5) \\ \text{Var}(S_{pg}) &= \frac{1}{2}\lambda + \frac{1}{4}g^2 \quad (p = 0.5). \end{aligned} \quad (5.1)$$

We will first empirically verify these findings.

5.3.1. Compound distribution moment simulations

In section 3.3.3, we derived the moments of compound probability distributions of pure Gaussian and Poisson-Gauss random variables subjected to a binomial split. Those results assume an ideal normal distribution. However, we showed earlier it is not ideal as it must be rounded and clipped to account for values lower than zero. To test if the calculation of moments is applicable also to this slightly modified distribution, we perform two simulations, one for the compound binomial-pure Gaussian random variable S_g , the other for the compound binomial-Poisson-Gaussian random variable S_{pg} . In the first, 100 million values are drawn from a normal distribution ($a = 1000$, $g^2 = 10$) and in the second, 100 million values from the sum of a draw from a Poisson distribution (parameter $\lambda = 1000$) and normal distribution (zero mean, variance $g^2 = 100$). These are then rounded to the nearest integer and negative values set to zero. For each resulting value, a draw from a binomial distributions is made, with the rounded, clipped normal sample as the n parameter.

Random variable	Sample mean	Sample variance
N_g	1000.0000 ± 0.0002	10.084 ± 0.002
N_{pg}	999.999 ± 0.003	1100.1 ± 0.01

Table 5.1: **Rounded normal and Poisson-Gauss distribution.** Mean and variance of 100 million samples from a normal distribution N_g (mean $a=1000$, variance $g^2=10$) and sum of a normal distribution (zero mean, variance $g^2=100$) and Poisson distribution ($\lambda=1000$), rounded to the nearest integer and with negative values clipped to zero. The experiment was conducted 10 times to compute the standard deviation of the variance and mean.

We now look at the mean and variance of the underlying distributions, the rounded and clipped normal and Poisson-Gauss distributions. The results can be seen in table 5.1. The means are barely influenced by the rounding and clipping, falling easily within one standard deviation. However, the variance is slightly higher than expected for both distributions, but the discrepancy is very small.

Table 5.2 shows the simulation results for the moments of the compound probability distributions. Again, the values are very close, with only the sample variance slightly outside a single standard deviation, but not too far out to disregard the moments.

Random variable	Var(S)	Sample variance	$\mathbb{E}[S]$	Sample mean
S_g	252.5	252.59 \pm 0.04	500	500.000 \pm 0.002
S_{pg}	525	525.07 \pm 0.04	500	499.999 \pm 0.002

Table 5.2: **Binomial from normal and Poisson-Gauss.** Computed variance and expected value of S_g and S_{pg} , binomial random variables with parameters ($n, p = 0.5$) where n is a realization of random variable N_g for S_g , which is normally distributed (mean $\alpha=1000$, variance $g^2=10$), but rounded to the nearest integer and with negative values clipped to zero. For S_{pg} , n is a realization of random variable N_{pg} , which is the sum of a normal distribution (zero mean, variance $g^2=100$) and Poisson distribution (parameter $\lambda=1000$), again rounded to the nearest integer and with negative values clipped to zero. The experiment was conducted 10 times to compute the standard deviation of the variance and mean.

Now that we know we can accurately estimate the moments of the compound probability distributions, even after rounding and clipping, we will look at the relation between expectation, variance and the 1FRC discrepancy.

5.3.2. Discrepancy relation

Empirically, we found that for both the pure Gaussian noise and Poisson-Gauss noise, the 1FRC begins to give different results from the 2FRC once the pixel variance after the split σ^2 exceeds the mean after the split μ , as defined in section 4.2. This different result manifests itself as an elevated plateau for the FRC values at high spatial frequencies, which we denote as $\langle 1FRC \rangle_{\text{high } f}$. Generally, it fluctuates around zero for higher spatial frequencies, with the average near zero. However, for higher relative levels of variances, this average is nonzero.

We plot the discrepancy relation in figure 5.9. The plot does not extend beyond 1 as for Poisson-Gauss noise, we know from equation (4.6) that μ/σ^2 cannot be higher than 1. In the plot, we first determine which values of μ/σ^2 we want to show and then compute the associated value of g^2 from equation (4.6) for Poisson-Gauss noise and equation (4.4) for pure Gaussian noise. The $\langle 1FRC \rangle_{\text{high } f}$ is then calculated by taking the average of the 1FRC values in the second half of the spatial frequency range. Figure 5.9 shows how the erroneous increase in the 1FRC vs the 2FRC (of which the latter stays around zero for any level of variance) decreases as μ/σ^2 decreases. This increase is identical for pure Gauss and Poisson-Gauss at identical levels of μ/σ^2 . To illustrate that this effect is an inherent product of the 1FRC, we also show this for a constant image input. Even for such a picture, the elevated plateau can be observed. In fact, we can see the increase can be predicted using the computed value of μ/σ^2 . We can write the observed relation as follows:

$$\langle 1FRC \rangle_{\text{high } f} = 1 - \frac{\mu}{\sigma^2} \quad (5.2)$$

Here, μ/σ^2 is computed using equations (4.4) (pure Gaussian noise) and (4.6) (Poisson-Gauss noise). This requires knowledge of the variance parameter g^2 of the Gaussian noise.

In figure 5.10, we compute $\hat{\mu}/\hat{\sigma}^2$, where $\hat{\mu}$ is the sampled mean of the split image and $\hat{\sigma}^2$ is the sample variance of a constant value pixel in the split image. This is different from the previous figure, as we now compute $\hat{\mu}/\hat{\sigma}^2$ from observed pixel values, instead of pre-computing it using known image and noise parameters. To compute $\hat{\sigma}^2$, we calculate the variance of a sample of pixels that have the same value in the input image, as the noise and binomial split work independently on each pixel. This is not possible if the noise is already applied, as in that case it is unknown which pixels have the same value in the input image. We see that for pixel samples with a different value, the curve is different, although the beginning and starting points remain the same. We are unable to determine if this is a result of some sampling artifact or an actual effect relevant to 1FRC. We do see that if we look at the variance of pixels that are close to the overall image mean, we again have the relationship of equation (5.2).

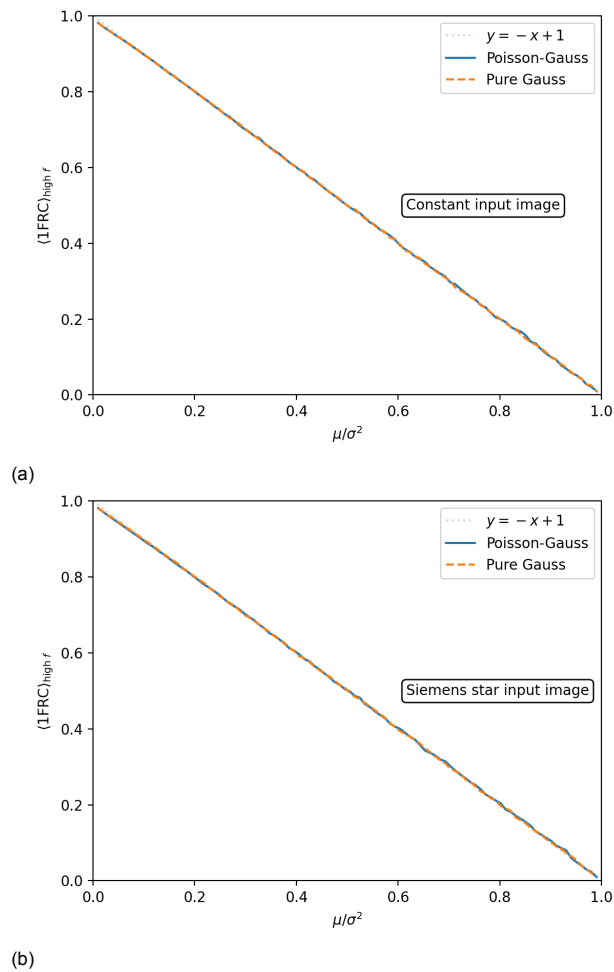


Figure 5.9: **Incorrect 1FRC from variance, computed mean variance ratio.** Average of the second half of the 1FRC curve plotted against μ/σ^2 . This value is computed using equation (4.6) for Poisson-Gauss noise and equation (4.4) for pure Gaussian noise, using a known value for g^2 as the Gaussian noise variance parameter. The parameters a and λ correspond to the input images' mean. For (a) the input image is 512 by 512 pixels of constant value 100. For (b) the input image is a 512 by 512 Siemens star with Gaussian filter applied ($\sigma = 3$ pixels), with $I_{\text{avg}} = 100$ and the background pixels set to 55% of the foreground. For the Poisson-Gauss noise, first Poisson noise is added to the constant image. Next, zero mean Gaussian noise is added with a variance set for the desired final value of μ/σ^2 . For pure Gaussian noise, only the Gaussian noise is added. μ/σ^2 does not go higher than 1 as this is impossible for Poisson-Gauss noise. The relation is very close to a simple straight line of the form $y = -x + 1$, with the maximum divergence from the straight line at 0.0078 for pure Gaussian noise and 0.0090 for Poisson-Gauss noise in the case of (a). For (b), the maximum divergence is equal to 0.0080 for pure Gaussian noise and 0.011 for Poisson-Gauss noise.

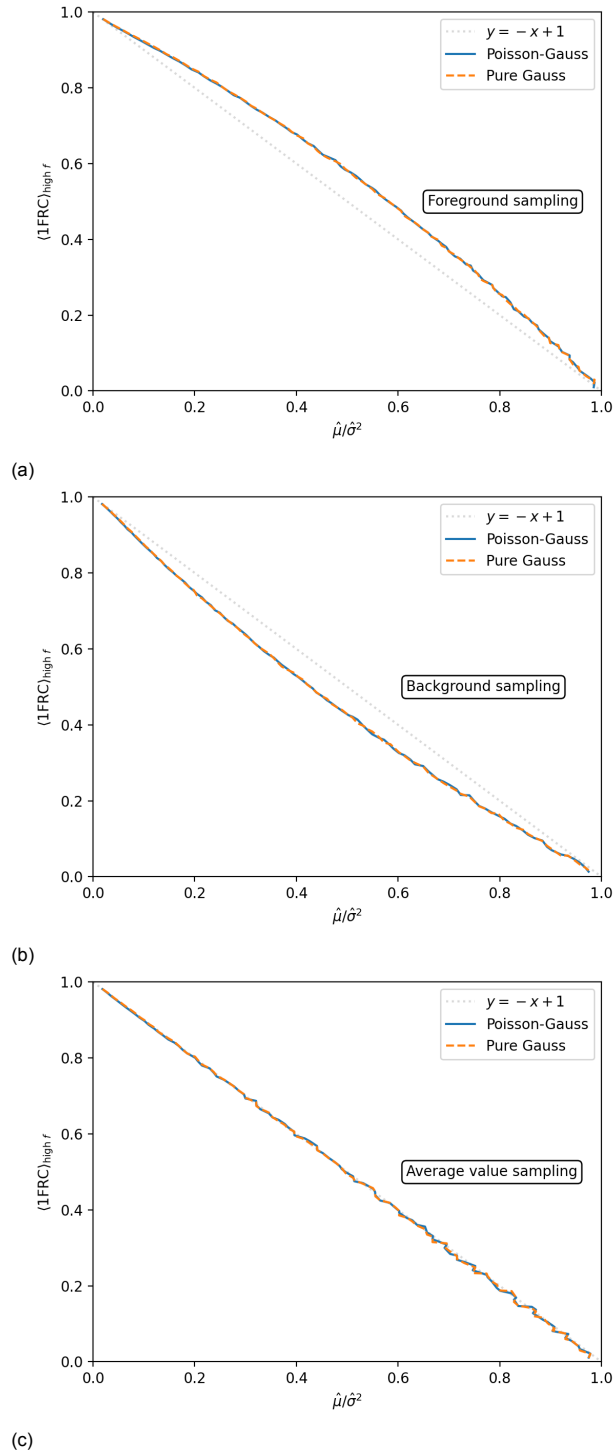


Figure 5.10: **Incorrect 1FRC from variance, sample mean variance ratio.** Average of the second half of the 1FRC curve plotted against $\hat{\mu}/\hat{\sigma}^2$, where $\hat{\mu}$ is the sampled mean of the split image and $\hat{\sigma}^2$ is the sample variance of a constant value pixel in the split image. The input image is a 512 by 512 Siemens star, as for figure 5.9b. Taking the variance over the entire picture would include variance inherent in the image, namely the difference between the star foreground and background. To remedy this, the variance is taken only over 50,000 different pixels of equal value (all in the star foreground for (a), all in the star background for (b) and for pixels with an absolute difference of less than two from the image mean for (c)). Furthermore, this is done for 5 different noise realizations and subsequent splits for each Gaussian variance g^2 input value. Since only around 1,200 pixels have a value near the image mean, for (c) we perform 50 different noise realizations and splits instead of 5. While the starting points are on the straight line, a curve can be observed for (a) and (b), with the maximum divergence from the line for (a) equal to 0.083 for pure Gaussian noise and 0.085 for Poisson-Gauss noise, a factor 10 larger than in the case the computed μ/σ^2 is used. For (c), we see that it is much closer to the line, with a maximum divergence of 0.016 for pure Gaussian noise and 0.017 for Poisson-Gauss noise.

In the previous figures we have plotted μ/σ^2 between 0 and 1 as these are the only values allowed by equation (4.6). However, equation (4.4) allows values between 0 and 2. While pure Gaussian noise is not the primary type of noise we are interested in, it is useful to see whether the relation holds more generally. We plot the results in figure 5.11. We observe that the relation indeed holds more generally for pure Gaussian noise. In fact, we can state that if this relation holds for all types of noise, 1FRC only works in the narrow region where μ/σ^2 is roughly one, making this an important condition for its use.




Figure 5.11: **Incorrect 1FRC from variance, pure Gaussian noise.** Average of the second half of the 1FRC curve plotted against μ/σ^2 , with the same input image and parameters as figure 5.9b. However, here we plot μ/σ^2 between 0 and 2. We see that the relation of equation (5.2) continues to hold, with 1FRC indicating full anticorrelation for $\mu/\sigma^2 = 2$. At this μ/σ^2 , we see from equation (4.4) that $g^2 = 0$. This means that in this case no noise is added at all. Since we are simply splitting an image with no noise, we are in effect only halving the image mean. Since we compute the second image by subtracting the first from the original image, described by equation (3.3), with no randomness full anticorrelation can be expected.

We now know that for μ/σ^2 higher than the FRC threshold value (e.g. 1/7), the 1FRC becomes useless, as the 1FRC will no longer intersect the threshold. In earlier sections, we saw that even for values close to 1, the discrepancy leads to inaccurate resolution results. More research is needed to determine the underlying process that creates this relation, which only applies to 1FRC, not standard 2FRC.

5.3.3. Maximum Gaussian variance

We have now seen in many ways that for increasing variance levels of Gaussian noise, the 1FRC plateau at higher frequencies does not drop to zero as the 2FRC always does. For zero variance Gaussian noise, we can see that the situation reduces to pure Poisson, in which case 1FRC works well. Furthermore, the approximation in section 3.3.1 suggests the same. We seek to quantify for what level of Gaussian variance this is the case. We do this by comparing the standard deviation of 1FRC SD_{1FRC} over different noise realizations (and subsequent binomial splits) to the absolute difference in resolution Δd between 1FRC and 2FRC. We plot this in figure 5.12. In the figure, we observe that the difference crosses the 1FRC standard deviation at $\mu/\sigma^2 = 0.973$. If we compare this to figure 5.5c, where $\mu/\sigma^2 = 0.97$, 0.973 seems like a reasonable value from which we can say that for lower values, the 1FRC gives incorrect results.

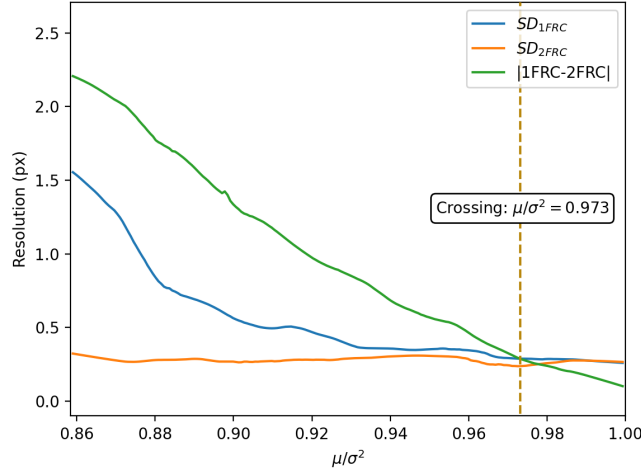


Figure 5.12: **Standard deviation 1FRC and 2FRC vs. difference 1FRC and 2FRC.** We measured standard deviation of 1FRC and 2FRC resolution (using the 1/7 threshold) by computing the sample standard deviation across 20 Poisson-Gauss noise realizations (and 20 subsequent binomial splits for the 1FRC) for each value of the mean variance ratio μ/σ^2 , which we compute in steps of 0.0007. The curves are smoothed using LOESS (with the fraction of points considered equal to 0.3).¹³ We only plot for $\mu/\sigma^2 > 0.86$ as for lower values $(1FRC)_{high f}$ lies above the threshold value. For $\mu/\sigma^2 > 0.973$, the absolute difference between 1FRC and 2FRC exceeds the 1FRC standard deviation. From this point on, the difference cannot be explained by simple uncertainty and we characterize 1FRC as unreliable compared to the 2FRC. The input image is a 512 by 512 Siemens star with Gaussian filter applied ($\sigma = 3$), with $I_{avg} = 200$. For this image mean, we know from equation (4.6) that at $\mu/\sigma^2 = 0.973$, we have $g^2 = 11.1$. Equivalently, we have SNR = 22.9 dB.

It is not immediately clear if the intersection between Δd and SD_{1FRC} occurs at constant g^2 , SNR or $\mu/\sigma^2 = 0.973$ (constant for different values of I_{avg}). There is high uncertainty in the μ/σ^2 value found from the intersection in the plot, as heavy smoothing is required to remove the fluctuations in the curves. Furthermore, this is a computationally expensive process, as many realizations must be computed for many values of μ/σ^2 . From empirical investigation, we found that for $I_{avg} = 70$ we have that $\mu/\sigma^2 = 0.979$ ($g^2 = 3.00$, SNR = 18.3 dB) and for $I_{avg} = 1000$, $\mu/\sigma^2 = 0.970$ ($g^2 = 61.9$, SNR = 29.9 dB). For these different image means, g^2 and the SNR vary far more than μ/σ^2 does. These facts make it plausible that $\mu/\sigma^2 = 0.973$ is a more fundamental limit, although simulations on far more noise realizations and for finer values of μ/σ^2 are necessary to make this statement more confident.

Another way to quantify for which value of μ/σ^2 the 1FRC results become too inaccurate is by using the pixel sum statistical test. If the p -value drops below 0.05, we can confidently reject that the two distributions are equal. Unfortunately, the p -value of the Kolmogorov-Smirnov test we employ is very sensitive, leading to significant uncertainties and different values in each simulation. If we look at figure 5.7, we can see that the p -value drops below 0.05 between $\mu/\sigma^2 = 0.97$ and $\mu/\sigma^2 = 0.95$. This is in rough accordance with our 0.973 result from the intersection we found in figure 5.12. It is important to note that the widely used statistical significance level of 0.05 is quite arbitrary.

5.3.4. Heuristic modification to ensure correct resolution

Without modification, the resolution cannot be computed for high variance levels as the 1FRC curve has no intersections with the threshold. However, a simple transformation of the curve can ensure that it does. Consider a 1FRC curve $f(x)$ with initial value $f(0)$ (this should be 1). If we define \bar{f} to be the average level of the end of the curve (in this example we use the second half, but by investigating the slope of the curve a more accurate point could be found), we can calculate a transformed curve:

$$f^*(x) = \frac{f(x) - \bar{f}}{f(0) - \bar{f}} \quad (5.3)$$

This moves the final average down while transforming the curve so it still starts at 1. Figure 5.13 shows the result. While it does allow for the computation of a resolution, the curve remains slightly shifted,

which was not the case for low variance 1FRC curves without this transformation. More investigation and mathematical justification is necessary to evaluate the performance of this *ad hoc* transformation.




Figure 5.13: **Transformed 1FRC with high variance.** A 1FRC curve with high variance Gaussian noise is indicated with blue. We use the same input image and parameters as in figure 5.5f. This figure also includes the transformed curve $f^*(x)$, modified according to equation (5.3). The resolution for the 2FRC curve is 10.0 pixels, while for the modified 1FRC curve it is 9.7 pixels.

Experimental data

We have shown the performance of 1FRC for simulated data, but not yet for experimental data. Here, we will use 1FRC to determine the resolution of a series of images captured using stimulated emission depletion (STED) microscopy and try to show a well-known relation between STED intensity and resolution. STED is a common technique used for super-resolution microscopy, which is the primary area of interest for applying 1FRC. STED was one of the first super-resolution methods to be introduced.¹

6.1. STED dataset

The dataset was compiled by the Heilemann Group at the Johann Wolfgang Goethe-University (Institute of Physical and Theoretical Chemistry) in Frankfurt am Main. Three different biological samples were imaged, all using Nile red membrane stain. The samples were imaged at different STED laser settings intensity settings, all using 561 nm excitation and a 775 nm STED laser.


One of the images of the first series, series *A*, is shown in figure 6.2. Six pictures were taken for this series, with STED intensity settings of 0%, 13%, 25%, 50% and 100% (relative to the maximum intensity). The object has a size of $22.55 \times 22.55 \mu\text{m}$ and is imaged at 1024×1024 pixels. The numerical aperture was 1.2.

The second series (series *B*) and third series (series *C*) both consist of 13 images of 20.41×20.41 and $20.06 \times 20.06 \mu\text{m}$, respectively. Both were imaged at 1024×1024 pixels with a numerical aperture of 1.4. For these series, STED intensity settings included 0%, 25%, 50% and 100%.

This dataset was not acquired in a way that allowed standard FRC; there is only a single image available per measurement. For this reason, the 1FRC is a good fit. The results are also compared with another recent technique utilizing decorrelation analysis.⁷


6.2. 1FRC on STED images

To showcase the utility of the 1FRC method, it is applied to real world data. For this, a data set of STED images is used. This data set consists of three series, each corresponding to a different biological sample. Figure 6.1 shows the 1FRC for the first of those, consisting of six images taken at increasing intensity levels of the STED microscope.



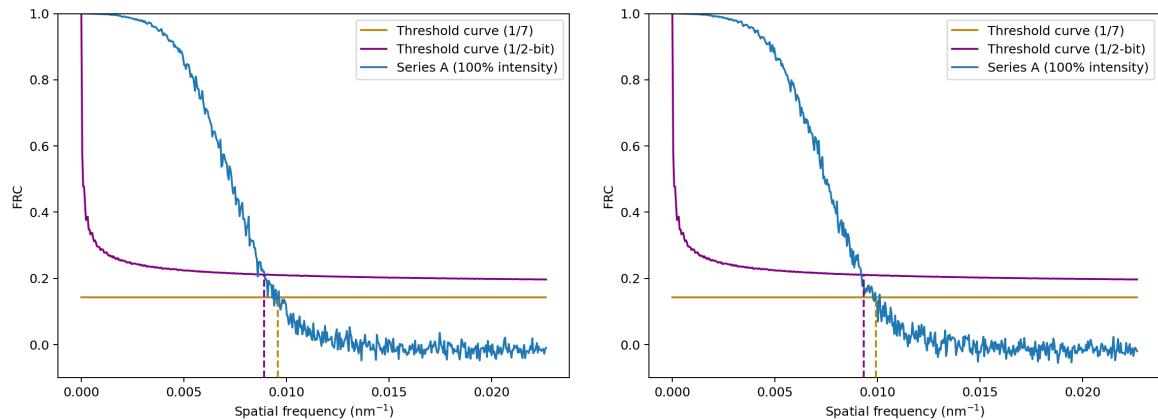
(a) STED intensity: 0%. $d_{1/7}$: 206 ± 6 nm. $d_{1/2\text{-bit}}$: 226 nm. Diffraction limit: 234 nm.

(b) STED intensity: 13%. $d_{1/7}$: 158 ± 1 nm. $d_{1/2\text{-bit}}$: 169 nm.



(c) STED intensity: 25%. $d_{1/7}$: 139 ± 2 nm. $d_{1/2\text{-bit}}$: 148 nm.

(d) STED intensity: 50%. $d_{1/7}$: 124 ± 4 nm. $d_{1/2\text{-bit}}$: 129 nm.



(e) STED intensity: 100%. $d_{1/7}$: 103 ± 1 nm. $d_{1/2\text{-bit}}$: 113 nm.

(f) STED intensity: 100%. $d_{1/7}$: 102 ± 1 nm. $d_{1/2\text{-bit}}$: 108 nm.

Figure 6.1: **Series A 1FRC curves.** Six 1FRC curves of real-world STED data (Series A) are shown here, as well as the corresponding computed resolutions for two standard threshold curves (1/7=0.143 and 1/2-bit information). Each curve was computed from an image taken using different STED parameters, with the STED intensity as the parameter of interest. The intensity values are 0%, 13%, 25%, 50% and twice 100%, respectively. It can be seen from the curves that the resolution improves as the intensity increases.




Figure 6.2: **STED images.** Low-resolution versions of the images in the STED dataset. (a) and (b) show Series *A*, the first at 0% STED intensity (meaning it is equivalent to a standard confocal microscope image and is diffraction limited), the latter is taken at maximum STED intensity. $d_{(a)} = 223 \pm 2.6$ nm and $d_{(b)} = 108 \pm 1.2$ nm. (c) and (d) show Series *B* and *C* at maximum STED intensity, respectively. $d_{(c)} = 89.7 \pm 0.78$ nm and $d_{(d)} = 108 \pm 1.8$ nm. We use the 1/2-bit threshold for these resolutions.

We expect from the theory of STED microscopy that the resolution can be computed using a modified version of the Abbe equation (1.1):

$$d = \frac{\lambda_{\text{exc}}}{2NA\sqrt{1 + \zeta}}, \quad (6.1)$$

where λ_{exc} is the excitation wavelength and ζ is the STED saturation factor, equal to $I/I_{\text{saturation}}$.¹⁴ After some empirical investigation, we found that the ζ scales with the square root of the aforementioned STED intensity. As we know the excitation has a wavelength of 561 nm, we use the following fit function:

$$d = \frac{561 \text{ nm}}{2NA\sqrt{1 + \alpha I_{\text{STED}} \%^{1/2}}}, \quad (6.2)$$

where α is our fit parameter and $I_{\text{STED \%}}$ is the STED intensity setting of the experiment.

Table 6.1 shows the 1FRC resolutions next to the associated STED intensity. Resolution clearly improves as intensity increases. The table also contains a comparison with the resolution computed using decorrelation analysis. This alternative technique gives more conservative results for this situation.

It is hard to judge which is more accurate, as there exists no objective resolution nor theoretical maximum resolution that allows a choice between the two. However, for the 0% STED intensity figure, the microscope should behave as an ordinary confocal microscope, which should be bound by the diffraction limit. Based on the NA of 1.2 of the setup and the roughly 561 nm wavelength, the diffraction limit can be estimated to be around 233 nm.

The fact that the 1FRC result lies below this means that for this specific situation, it most likely overestimates the resolution. However, there should be very few other limiting factors, so a resolution of 308 nm seems also too conservative.

STED intensity	Saturation factor ζ (fit)	1FRC resolution (nm)	Decorr. resolution (nm)
0%	N/A	223 ± 2.6	308
13%	1.1 ± 0.20	169 ± 1.8	203
25%	1.6 ± 0.41	148 ± 1.1	167
50%	2.3 ± 0.82	129 ± 1.7	152
100%	3 ± 1.6	112 ± 1.1	120
100%	3 ± 1.6	108 ± 1.2	127

Table 6.1: Comparison of STED intensity, 1FRC resolution (using the 1/2-bit threshold) and resolution computed with decorrelation analysis⁷ for the Series A STED images (in ascending order). The saturation factor as estimated from a fit to equation (6.2) is also included. No value is reported for zero STED intensity as these were not included in the fit. As expected, the resolution increases as STED intensity increases. There is a clear difference between decorrelation analysis and 1FRC, with the difference larger for worse resolutions / lower intensities.

A more comprehensive analysis of the computed resolutions and the corresponding STED intensities is shown in figure 6.3, which also contains fits for all three measurement series. These fits are not ideal, as there is only a single fit parameter and it is unknown if the square root of the STED intensity scales exactly with the saturation factor ζ (a fact we found empirically). However, there is still a good agreement. These curves provide evidence that there is indeed a dropoff in improved resolution as the STED intensity is increased, although they are not accurate enough to challenge the exact formula, as it is just as likely that there are issues with the 1FRC method or specific images.

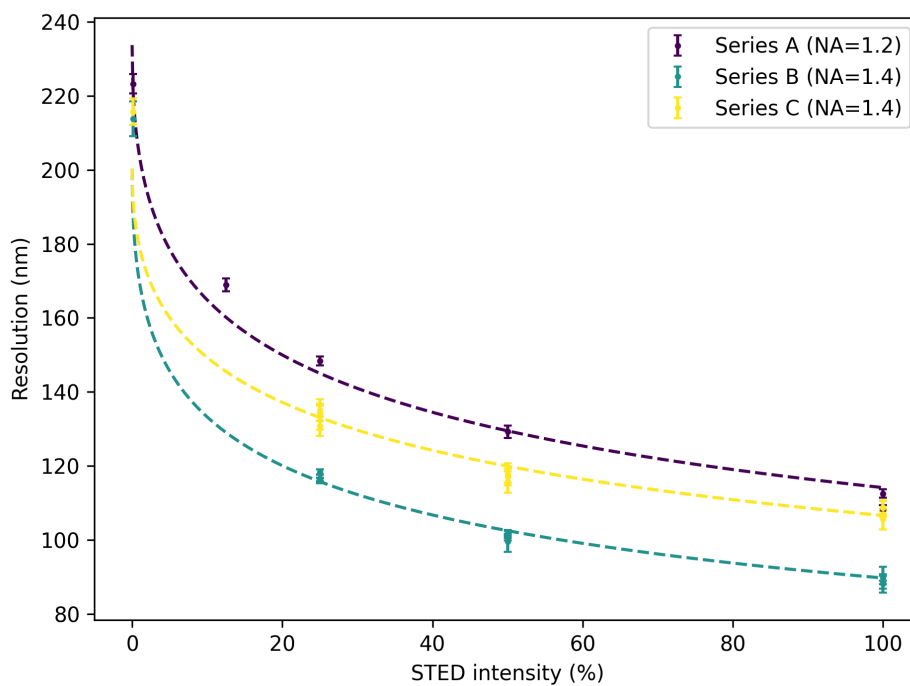


Figure 6.3: **STED resolutions and fits.** The resolutions are computed using 1FRC of Series A, B and C, which are each a series of STED images at different STED intensities. The STED wavelength is 775 nm, while the excitation wavelength is 561 nm. These curves are fits to the theoretical STED resolution equation, specifically to equation (6.2).

Conclusion

Fourier Ring Correlation with two images generated from a single measurement, which we call 1FRC, is studied for different noise modalities, specifically pure Poisson noise and a sum of Poisson and Gaussian noise (Poisson-Gauss noise). To investigate problems that occur for Poisson-Gauss noise, we also withok at pure Gaussian noise. For Poisson noise, we derive exactly that a binomial split preserves the Poisson distribution, validating the use of 1FRC. Furthermore, we show with simulations that 1FRC gives results that are nearly identical to using two independent measurements (2FRC). By using a Siemens star as a test image, we verify that Fourier Ring Correlation gives the resolution in the right order of magnitude.

If we add Poisson noise combined with Gaussian noise with high variance, we see that the 1FRC in its current form is not a suitable alternative to using two independent measurements and using the standard 2FRC routine. If we increase the variance of the noise, the FRC curve no longer drops to zero, but instead levels off at a nonzero value, even for higher frequencies (with high we mean the latter half of the spatial frequencies). If it stays above the threshold, the resolution can no longer be computed at all. Using a Kolomgoronov-Smirnov test on the sum of all pixels in an image, we found that for these high levels of Gaussian noise, we can reject the hypothesis that the distribution after a binomial split is equal to the distribution of a reference image (equivalent to an independent measurement as in 2FRC). This is required for the use of 1FRC to make sense. However, as the uncertainty in the p -value results of the test are high, we cannot exactly quantify what counts as "high" variance using this method alone.

As we are unable to analytically derive the probability distribution of Poisson-Gauss noise (after a binomial split) to find the cause of the discrepancy in the 1FRC curve plateau at high frequencies, we instead looked at its statistical moments. We observe that the mean variance ratio is a good predictor for the discrepancy. The mean variance ratio μ/σ^2 is the image mean divided by the variance of one pixel (over many noise and split realizations). We can directly compute μ/σ^2 from the Gaussian variance parameter and the average pixel intensity of the input image. The elevated curve level where the 1FRC tends to for high spatial frequencies is linearly dependent on μ/σ^2 , with the curve plateau at zero for $\mu/\sigma^2 = 1$ (equivalent to zero Gaussian noise, i.e. pure Poisson noise) and 1 near $\mu/\sigma^2 = 0$ (Gaussian noise completely dominates the signal). By comparing the difference between the resolutions computed by 1FRC and 2FRC Δd to the standard deviation of 1FRC resolution (over many noise and split realizations), we found that for $\mu/\sigma^2 < 0.973$, this difference Δd is greater than one standard deviation of the 1FRC resolution. This value was computed for an image with an average pixel intensity of 200, but remained consistent for other intensities, while the values where this happens for other measures such as the absolute Gaussian variance and SNR varied significantly. Simulations on more average intensities and more values of $\mu/\sigma^2 < 0.973$ are necessary to confirm this value works for all those situations.

We conclude that for levels of Gaussian noise where we have $0.973 < \mu/\sigma^2 \leq 1$, the 1FRC gives nearly identical results to 2FRC. This value is corroborated by the pixel sum statistical tests, as they transitioned to a p -value below 0.05 in the range $0.95 < \mu/\sigma^2 < 0.97$. It is important to note that the p -value for these tests showed a p -value of 0.61 ± 0.17 even in the case of pure Poisson noise, which indicates significant uncertainty in the p -value. We also found that for pure Gaussian noise, the linear relation between μ/σ^2 and the 1FRC curve plateau at high frequencies continues even for $\mu/\sigma^2 > 1$. The 1FRC curve goes to -1 (full anticorrelation) for the extreme case of zero Gaussian variance ($\mu/\sigma^2 = 2$). This indicates that the 1FRC most likely only works when the mean variance ratio is close to 1. Further research is required to determine why the binomial split method generates correlation at high frequencies for $\mu/\sigma^2 < 1$.

On experimental data, gathered using STED microscopy (one particular super-resolution method),

the 1FRC curve behaved well, giving more optimistic results than results obtained using decorrelation analysis, which can also be applied to a single image. The difference between 1FRC and decorrelation analysis lies between 15-30%, with 1FRC always giving lower resolutions. For images captured using zero STED intensity (equivalent to confocal microscopy), the results were better than the theoretical diffraction limit of the setup (up to 5% better). This indicates 1FRC results might be a slightly optimistic measure, although they remain in the right order of magnitude. The resolution values for different STED intensities agree well with a fit to the theoretical curve. These good results indicate that for this particular dataset, there was most likely little Gaussian noise. This gives good confidence that similar STED datasets can also be analyzed with 1FRC, making it a useful tool regardless of its problems with high variance Gaussian noise.



Appendix A

A.1. Implementation

All analysis was done using Python. For most image transformations (e.g. Fourier transform), the DIPlib package was employed¹⁵, while NumPy was used as a utility tool.

Furthermore, the Python package FRC¹⁶, built specifically for this project, was used for the implementation of the "binomial split" for 1FRC. This package uses an extension module written in the Rust programming language, called RustFRC¹⁷ and also created for this project. Rust has similar performance to C++, but avoids its common memory pitfalls and generally makes it much easier to write multithreaded code.

Furthermore, a third package called AnalyzeFRC¹⁸ was made to help with the analysis of the curves, with significant attention paid to configuration, so it can also be used by other researchers.

Locally weighted smoothing (LOESS) was performed using the LOESS Python package. Their implementation is described in Cappellari et al (2013).¹³

A.1.1. Binomial split

A binomial split, necessary for the calculation of the 1FRC, is a computationally expensive task as it requires a separate computation for each array element (a pixel/voxel). An array (2D or 3D) is mapped in-place, with each element value serving as parameter n for a sample from a binomial ($n, p = 0.5$) distribution. The result of the sampling is the new array element value. This mapping takes place in parallel, speeding up the computation significantly.

The result is an array with, on average, half the intensity of the original image. To generate the second array, the generated array is subtracted from the original image, resulting in two arrays of similar intensity.

A.1.2. 1FRC

An image is loaded as an array-like structure. It is important the image is square (or cubic), as this is required for the 2D discrete Fourier transform. The image is subjected to a binomial split, resulting in two images with half the intensity of the original image.

Tukey windows (cosine-tapered windows) are applied to these two images, with parameter $\alpha = 1/8$. These windows are necessary to prevent DFT-artifacts occurring on the image edges. The windowed images are then Fourier transformed (DFT with the origin in the center). The FRC is then calculated using the Fourier transformations for each, resulting in a single curve.

To find the appropriate spatial frequencies, the x-axis values are divided by the pixel dimension of the image. These then have the unit $1/d$ where d is the size of the original pixel.

To calculate the resolution, the intersection for the FRC curve with the fixed threshold of $1/7$ is calculated. As no exact function exists for the FRC curve, the intersection is calculated by a heuristic that looks for the center of the x-values of FRC curve values that are close to the threshold curve.

Bibliography

- [1] S. W. Hell and J. Wichmann. “Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy”. In: *Optics Letters* 19.11 (June 1994), pp. 780–782. DOI: 10.1364/OL.19.000780.
- [2] W. O. Saxton and W. Baumeister. “The correlation averaging of a regularly arranged bacterial cell envelope protein”. In: *Journal of Microscopy* 127.2 (1982), pp. 127–138. DOI: 10.1111/j.1365-2818.1982.tb00405.x.
- [3] M. Van Heel et al. In: *Structure and function of invertebrate respiratory proteins*. Ed. by E. J. Wood EMBO Workshop (1982: Leeds). Life chemistry reports, 0278-6281. Supplement 1. 1983, pp. 69–73.
- [4] M. Van Heel and M. Stöffler-Meilicke. “Characteristic views of E. coli and B. stearothermophilus 30S ribosomal subunits in the electron microscope.” In: *The EMBO Journal* 4.9 (1985), pp. 2389–2395. DOI: 10.1002/j.1460-2075.1985.tb03944.x.
- [5] M. Van Heel. “Similarity measures between images”. In: *Ultramicroscopy* 21.1 (1987), pp. 95–100. DOI: 10.1016/0304-3991(87)90010-6.
- [6] R.P.J. Nieuwenhuizen et al. “Image Resolution in Optical Nanoscopy”. In: *Nature Methods* 10.6 (2013), pp. 557–562. DOI: 10.1038/nmeth.2448.
- [7] A. Descloux, K. S. Größmayer, and A. Radenovic. “Parameter-free image resolution estimation based on decorrelation analysis”. In: *Nature Methods* 16.9 (2019), pp. 918–924. DOI: 10.1038/s41592-019-0515-7.
- [8] M. Van Heel and M. Schatz. “Fourier shell correlation threshold criteria”. In: *Journal of Structural Biology* 151.3 (Sept. 2005), pp. 250–262. ISSN: 1047-8477. DOI: 10.1016/j.jusb.2005.05.009.
- [9] F. Huang et al. “Video-rate nanoscopy using sCMOS camera-specific single-molecule localization algorithms”. In: *Nature Methods* 10.7 (2013), pp. 653–658. DOI: 10.1038/nmeth.2488.
- [10] D. Roy. “The Discrete Normal Distribution”. In: *Communications in Statistics - Theory and Methods* 32.10 (2003). 1871, pp. 1871–1883. ISSN: 0361-0926. DOI: 10.1081/STA-120023256.
- [11] D. Raikov. “On the decomposition of Poisson laws”. In: *Dokl. Akad. Nauk SSSR*. Vol. 14. 1. 1937, pp. 9–12.
- [12] Y. V. Linnik. “On the Decomposition of the Convolution of Gaussian and Poissonian Laws”. In: *Theory of Probability & Its Applications* 2.1 (1957), pp. 31–57. DOI: 10.1137/1102002.
- [13] M. Cappellari et al. “The ATLAS^{3D} project - XX. Mass-size and mass- σ distributions of early-type galaxies: bulge fraction drives kinematics, mass-to-light ratio, molecular gas fraction and stellar initial mass function”. In: *MNRAS* 432 (2013), pp. 1862–1893. DOI: 10.1093/mnras/stt644. eprint: 1208.3523.
- [14] S. W. Hell. “Toward fluorescence nanoscopy.” In: *Nature Biotechnology* 21.11 (2003). 1347, pp. 1347–55. ISSN: 1087-0156. DOI: 10.1038/nbt895.
- [15] C. Luengo et al. *DIPlib (PyDIP)*. Version 3.2.0. 2022. URL: <https://github.com/DIPlib/diplib3>.
- [16] T. ten Brink. *FRC*. Version 0.1.0. 2022. DOI: 10.5281/zenodo.6461444.
- [17] T. ten Brink. *RustFRC*. Version 1.1.3. 2022. DOI: 10.5281/zenodo.6461451.
- [18] T. ten Brink. *AnalyzeFRC*. Version 0.2.3. 2022. DOI: 10.5281/zenodo.6457673.