

Bernoulli generalized likelihood ratio test for signal detection from photon counting images

Mengya (Mia) Hu^{a,*}, He Sun^b, Anthony Harness^a,
and N. Jeremy Kasdin^c

^aPrinceton University, Department of Mechanical and Aerospace Engineering, Princeton,
New Jersey, United States

^bCalifornia Institute of Technology, Department of Computing and Mathematical Sciences,
Pasadena, California, United States

^cUniversity of San Francisco, College of Arts and Sciences, San Francisco,
California, United States

Abstract. Because exoplanets are extremely dim, an electron multiplying charge-coupled device operating in photon counting (PC) mode is necessary to reduce the detector noise level and enable their detection. Typically, PC images are added together as a co-added image before processing. We present a signal detection and estimation technique that works directly with individual PC images. The method is based on the generalized likelihood ratio test (GLRT) and uses a Bernoulli distribution between PC images. The Bernoulli distribution is derived from a stochastic model for the detector, which accurately represents its noise characteristics. We show that our technique outperforms a previously used GLRT method that relies on co-added images under a Gaussian noise assumption and two detection algorithms based on signal-to-noise ratio. Furthermore, our method provides the maximum likelihood estimate of exoplanet intensity and background intensity while doing detection. It can be applied online, so it is possible to stop observations once a specified threshold is reached, providing confidence for the existence (or absence) of planets. As a result, the observation time is efficiently used. In addition to the observation time, the analysis of detection performance introduced in the paper also gives quantitative guidance on the choice of imaging parameters, such as the threshold. Lastly, though our work focuses on the example of detecting point source, the framework is widely applicable. © *The Authors. Published by SPIE under a Creative Commons Attribution 4.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.* [DOI: [10.1117/1.JATIS.7.2.028006](https://doi.org/10.1117/1.JATIS.7.2.028006)]

Keywords: signal detection; photon counting mode; starshade; high contrast imaging; direct imaging; exoplanet detection.

Paper 20141 received Sep. 17, 2020; accepted for publication May 20, 2021; published online Jun. 17, 2021.

1 Introduction

Direct imaging of exoplanets is challenging; the flux ratio between an Earth-like exoplanet and its Sun-like host star is around 10^{-10} in reflected light at visible wavelengths.¹ A starshade or internal coronagraph can suppress the starlight and leave only the planet's light to be detected; however, the planets are extremely faint and detecting them is still a challenge. An Earth-like planet ranges from 28th to 30th magnitude or fainter. As a result, the signal can be smaller than the camera read noise. An electron multiplying charge-coupled device (EMCCD) can alleviate this problem by amplifying the signal in an electron-multiplication (EM) register, thus reducing the effective readout noise to less than 1 electron.² Unfortunately, at the same time, a new noise is introduced—the multiplicative noise from the amplification process. This can be overcome by operating in photon counting (PC) mode. PC mode reports a value of 1 or 0 in each pixel for each integration time by thresholding the value at the final stage. The value reported in the pixel is 1 if the number of electrons in a pixel is bigger than a chosen threshold and 0 otherwise. The binary value only indicates the existence of photons in the pixels during the integration time but does

*Address all correspondence to Mengya (Mia) Hu, mengyah@princeton.edu

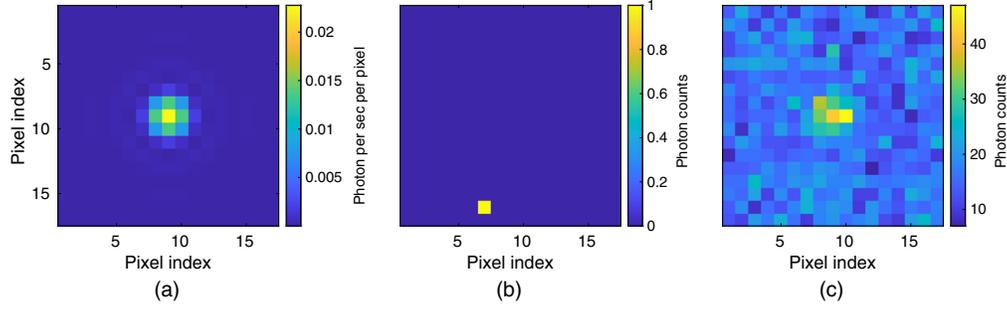


Fig. 1 (a) Simulated PSF for The Nancy Grace Roman Space Telescope without a starshade for a $1e-8$ Jansky (Jy) light source, 0.021 arc sec /pixel. The parameters and simulation process are described later in Secs. 2 and 3.3. (b) A PC image of (a) with 1-s integration time. (c) A co-added image from 2000 sequential PC images.

not reveal the exact number of those photons, so we need to choose the exposure time such that the expected photon count in any pixel is much less than 1.³ Examples of simulated PC images are shown in Fig. 1 (all the simulations mentioned in this work use a 1-s integration time) and details on the generation of simulated images can be found in Ref. 4.

Operation of an EMCCD in PC mode is fairly new and thus techniques are still developing. Available literature on the design and characteristics of EMCCD detectors can be found in Refs. 2 and 5–11. Previous work on image processing focuses mainly on image stacking and Bayesian estimation methods,^{12,13} which are applied to co-added PC images rather than designed for individual PC images.^{14,15} As those methods generally require a high signal-to-noise ratio (SNR), large numbers of PC images are needed, which take a long observation time. In our previous work, we presented an alternative methodology for co-added PC images¹⁶ which can efficiently detect even weak signals automatically. However, all these methods do not provide theoretical guidance on how to choose the integration times and the total number of PC images to combine into one co-added image.

In this work, we utilize a statistical model for the EMCCD to obtain a relationship between the detection probability and the photon rate. We use this distribution to formulate a Bernoulli distribution for the values of the same pixel on different PC images. Then, detection and estimation is performed. Consequently, detection and accurate intensity estimation can be achieved at low signal levels. We begin this paper with a description of the detector model, followed by the methodology of the Bernoulli generalized likelihood ratio test (BGLRT), and an application example. We finish with a comparison of the performance of our previous method, GLRT using co-added images,¹⁶ and the new BGLRT method using PC images directly. This work builds on the previous work introduced in Ref. 17, where the method is called the sequential generalized likelihood ratio test (GLRT). We change the method name to Bernoulli generalized likelihood ratio test to better reflect that the improvement of performance is from the usage of an accurate model based on the Bernoulli distribution for the detector.

2 Stochastic Model for EMCCDs in Photon Counting Mode

Hirsch et al.² present a detailed imaging process for an EMCCD and build a statistical model for each stage of the imaging process. Their statistical model is built as follows.² First, incident photons follow a Poisson process. Second, the EM register is represented by a gamma distribution. Third, readout noise is added using a Gaussian distribution. Fourth, a threshold is applied to give a binary result.

Hirsch et al. use some approximations to simplify and arrive at the equation for the probability distribution for the number of electrons in a pixel,

$$p(n_{ic}) = \begin{cases} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\lambda - \frac{(n_{ic} - B_{PC})^2}{2\sigma^2}\right) + \frac{2}{g} F_{\chi^2}\left(2\lambda; 4, \frac{2n_{ic}}{g}\right), & n_{ic} > 0 \\ \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\lambda - \frac{B_{PC}^2}{2\sigma^2}\right), & n_{ic} = 0, \end{cases} \quad (1)$$

where λ is the expected number of input electrons for the EM register,

$$\lambda = s \times q \times t + d \times t + \text{CIC}; \quad (2)$$

where s is the photon rate; q is the quantum efficiency; t is the integration time; d is the dark current; CIC is the clock-induced charge; n_{ic} is the number of counts after the EM register, including readout noise; $p(n_{ic})$ is the probability of counts n_{ic} ; σ is the standard deviation of the readout noise; g is the EM gain; B_{PC} is the bias in PC mode; $F_{\chi}(2\lambda; 4, \frac{2n_{ic}}{g})$ is the non-central χ^2 distribution for 2λ with 4 degrees of freedom and the non-centrality parameter $\frac{2n_{ic}}{g}$. Units for these parameters are given in Table 1.

An EMCCD decreases the read-out noise to sub-electron level by amplifying the signal before readout. However, the amplification process introduces an excess noise factor (ENF) that reaches $2^{\frac{1}{2}}$ at high gains.¹⁸ The effect on the SNR is the same as if the quantum efficiency of the EMCCD was halved.¹⁹ This can be overcome completely by operating in PC mode. PC mode applies a single threshold at the final stage and reports a binary result. The pixel registers a detected photon if its count level is higher than the threshold, and reports no photon otherwise. This solution alleviates the ENF without making any assumption on the signal's stability across multiple images.¹⁹

Combining models for all the stages described above produces the final model, which calculates the probability of getting a value of 1 in a detector pixel given a certain incident intensity. Thus, with a threshold of $B_{\text{PC}} + T \times \sigma$, set by the threshold parameter T , the probability of detecting a value of 1 is

$$f(s) = \int_{B_{\text{PC}} + T \times \sigma}^{\infty} p(n_{ic}) dn_{ic}, \quad (3)$$

where $f(s)$ is shortened for $f(s; q, t, \text{CIC}, d, g, B_{\text{PC}}, \sigma, T)$ (the detector parameters are not shown explicitly). The value in each pixel in each image only has two outcomes: either 1 or 0. The probability of getting a value of 1, $f(s)$, is decided by the ground truth of the flux, s . As the flux in each pixel is fixed for stable observations, the probability is fixed. That is to say, the measurement in each pixel in each image satisfies a Bernoulli distribution [the Bernoulli distribution is simply the probabilities of “heads” (p) and “tails” ($1 - p$) in a weighted coin flip]. In this paper, we assume that all the detector parameters, such as integration time, are fixed and the values are given in Table 1; so for each pixel, the imaging result follows a Bernoulli distribution whose probability of value 1 is only related to the flux value in the pixel. Figure 2 shows results for the detection probability given different flux levels and the calculated derivative of this probability. We use the detector parameters values given in Table 1, which are similar to the parameter values for the Nancy Grace Roman Space Telescope²⁰ and assume a starshade for starlight

Table 1 EMCCD detector parameters.

Parameter	Symbol	Value	Unit
Quantum efficiency	q	1	ph/e ⁻
Integration time	t	1	s
Clock-induced charge	CIC	0.01	e ⁻ pixel ⁻¹ frame ⁻¹
Dark current	d	2×10^{-4}	e ⁻ pixel ⁻¹ s ⁻¹
Electron-multiplying gain	g	2500	—
PC bias	B_{PC}	200	e ⁻ pixel ⁻¹ frame ⁻¹
Standard deviation of readout noise	σ	100	e ⁻ pixel ⁻¹ frame ⁻¹
Threshold parameter	T	5.5	—

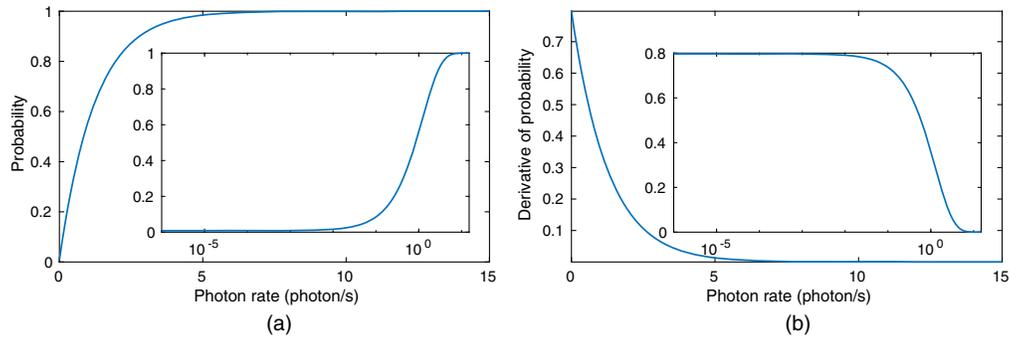


Fig. 2 Detection probability and its derivative. (a) Detection probability as a function of photon flux, $f(s)$, calculated by Eq. (3), where s is in units of photons/s. (b) Derivative of the detection probability. The inserts are the same plots shown with a logarithmic abscissa. We use the detector parameters values given in Table 1, which are similar to the parameter values for the Roman Telescope²⁰ and assume a starshade for starlight suppression.

suppression. The detection probability increases as the flux increases, where a very small flux or a very large flux tend to give deterministic 0 or 1 measurements.

The choice of imaging parameters (integration time, PC threshold, etc.) helps efficiently detect signals. The binary value only indicates the existence of photons in a pixel during the integration time but does not reveal the exact number of the photons. To utilize the photons efficiently, the exposure time is chosen so that the expected photon count in any pixel is less than 1. To decide a threshold for PC, we need to also pay attention to signal intensity and integration time, because they have similar influence on the result. Suppose that we have decided on the integration time and the expected range of signal intensity. To determine the imaging parameters, we can directly utilize the area under the receiver operating characteristic (ROC) curve (AUC) (details of ROC and AUC are in Sec. 3.4), which is a number reflecting the overall detection performance. For one threshold, we can calculate the average of the AUCs for different signal intensities in the expected intensity range. The threshold with the largest average AUC should be chosen.

3 Bernoulli Generalized Likelihood Ratio Test

The GLRT is a powerful tool for detection. The detection problem can be seen as deciding between the null hypothesis, H_0 , that there is no signal, and the alternative hypothesis, H_1 , that there is a signal. GLRT uses the ratio of the likelihood of observing the data under H_1 to the likelihood of observing the data under H_0 as the test statistic. If the ratio is large enough, we will decide H_1 . Compared to other likelihood ratio tests, the GLRT can include hypothesis tests which have unknown parameters in the two hypothesis. GLRT use maximum likelihood estimation (MLE) to estimate the unknowns and then calculates the maximum likelihood ratio. In other words, the MLEs of the unknown parameters such as signal intensity and background intensity, under both positive and negative hypotheses, are first required to be calculated. Then, a likelihood ratio test using the MLEs can be calculated and applied to decide whether a signal has been detected. Our previous work approximated the noise in co-added images by a Gaussian distribution.¹⁶ In this work, we directly use the detector's probability function derived in the previous section, which does not rely on co-added images with a large number of PC images. As a result, measurements in each pixel in each image follow a Bernoulli distribution.

3.1 Signal Estimation

In this section and Sec. 3.2, we examine an image area that has the size of PSF's central core (the main part of a signal). We explain how to apply the method on the whole image, which is normally larger than the PSF core, in Sec. 3.3. We begin with the simple model of the signal received by the detector,

$$s = \alpha \mathbf{x} + \beta \mathbf{1}, \quad (4)$$

where \mathbf{x} is a reference point spread function (PSF) of the telescope; α is the intensity of the planet relative to the source intensity of the reference PSF; β is the background light (such as dust and light leakage from starshade defects); and $\mathbf{1}$ is a column vector where each element equals one. We stack pixel values in the target area into a column vector for easier mathematical manipulation. That is to say, s and \mathbf{x} are column vectors.

In real life, due to the randomness of photons and the detector noise, we would not be able to observe s directly, but rather collect the noisy image output from the detector. Assume that N PC images have been collected, which is denoted as $\mathbf{y} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$. Given the PSF shape, $\mathbf{x} = (x_1, \dots, x_K)^T$, where the subscripts of x represent the pixel indices and K is the total number of pixels in a PSF, the conditional probability of an image $\mathbf{y}_n = (y_{n,1}, \dots, y_{n,K})^T$ is

$$p(\mathbf{y}_n | \alpha, \beta) = \prod_{k=1}^K [1 - f(\alpha x_k + \beta)]^{1-y_{n,k}} f(\alpha x_k + \beta)^{y_{n,k}}, \quad (5)$$

where the function $f(\cdot)$ is the probability function of detecting a one value in Eq. (3). The measurement in each pixel satisfies a Bernoulli distribution, and $[1 - f(\alpha x_k + \beta)]^{1-y_{n,k}} f(\alpha x_k + \beta)^{y_{n,k}}$ is the probability for the measurement in the k 'th pixel (for PC images, $y_{n,k}$ is either 1 or 0). The probability for the whole image is simply the product of the probability of every measurement. The interesting part is that the probabilities in different pixels are correlated because of the underlying signal model in Eq. (4). Figure 3 illustrates the meaning of the indices in the equation.

For this work, as a demonstration of the detection method, we use a PSF template as the underlying signal model. However, different signal models can be used and the later discussion of the detection method is still valid; the only difference would be detecting signal with different templates. This provides the Bernoulli GLRT with a broad applicability to various observation scenarios.

Furthermore, the conditional probability of N images, $\mathbf{y} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$, can be written as

$$\begin{aligned} L(\boldsymbol{\theta}) &= p(\mathbf{y} | \alpha, \beta) \\ &= \prod_{n=1}^N p(\mathbf{y}_n | \alpha, \beta) \\ &= \prod_{n=1}^N \prod_{k=1}^K [1 - f(\alpha x_k + \beta)]^{1-y_{n,k}} f(\alpha x_k + \beta)^{y_{n,k}} \\ &= \prod_{k=1}^K [1 - f(\alpha x_k + \beta)]^{N - \sum_{n=1}^N y_{n,k}} f(\alpha x_k + \beta)^{\sum_{n=1}^N y_{n,k}} \\ &= \prod_{k=1}^K [1 - f(\alpha x_k + \beta)]^{N_{0,k}} f(\alpha x_k + \beta)^{N_{1,k}}, \end{aligned} \quad (6)$$

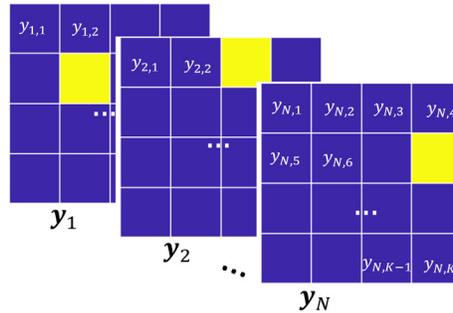


Fig. 3 Illustration for the variables in the model. K is the number of pixels in each PC image (with the size of PSF core). N is the number of PC images.

where $\boldsymbol{\theta}$ is the simpler vector representation for the two unknown parameters $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$; $N_{0,k} = N - \sum_{n=1}^N y_{n,k}$ and $N_{1,k} = \sum_{n=1}^N y_{n,k}$ represent the frequency of 0 and 1 occurring in N measurements of the k 'th pixel, respectively. The equality $N = N_{0,k} + N_{1,k}$ holds for every pixel. As the data \mathbf{y} is known and the parameters $\boldsymbol{\theta}$ are unknown, this probability function is a likelihood function for the unknown parameters [so it is denoted as $L(\boldsymbol{\theta})$ above].

Taking the natural logarithm of both sides of Eq. (6), the log-likelihood of the series of measurements is

$$l(\boldsymbol{\theta}) = \ln(L(\boldsymbol{\theta})) = \sum_{k=1}^K N_{0,k} \ln[1 - f(\alpha x_k + \beta)] + N_{1,k} \ln[f(\alpha x_k + \beta)]. \quad (7)$$

To estimate α and β , we can conduct an MLE based on Eq. (7) by taking the derivatives with respect to α and β and setting them equal to 0,

$$\begin{aligned} \frac{\partial l(\alpha, \beta)}{\partial \alpha} &= N \sum_{k=1}^K x_k f'(\alpha x_k + \beta) \frac{N_{1,k} - Nf(\alpha x_k + \beta)}{Nf(\alpha x_k + \beta)[1 - f(\alpha x_k + \beta)]} = 0, \\ \frac{\partial l(\alpha, \beta)}{\partial \beta} &= N \sum_{k=1}^K f'(\alpha x_k + \beta) \frac{N_{1,k} - Nf(\alpha x_k + \beta)}{Nf(\alpha x_k + \beta)[1 - f(\alpha x_k + \beta)]} = 0, \end{aligned} \quad (8)$$

where $f'(s) = \frac{df(s)}{ds}$. These equations can be solved for estimates of α and β using a gradient descent method. We denote the solution, i.e., the MLE, with a hat, such as $\hat{\alpha}$ and $\hat{\beta}$. As can be seen, the derivatives are weighted summations of the difference between the measurements and the distribution means, where $Nf(\cdot)$ and $Nf(\cdot)[1 - f(\cdot)]$ are, respectively, the mean and variance (each measurement satisfies a Bernoulli distribution). According to Fig. 2(b), the function $f'(s)$ is zero when $s \rightarrow +\infty$. That means that the method neglects the difference between different intensities that are too large in the estimation since the PC detector always gives 1 in these cases.

In the problem of parameter estimation, we obtain information about the unknown parameters from the observed data of the random variables from the probability distribution governed by the parameters. The Fisher information matrix is a way to quantify the amount of information that the observable random variables carry about the unknown parameters. In our case, the Fisher information matrix is

$$\begin{aligned} \mathbf{I}(\boldsymbol{\theta}) &= \text{Var}_{\boldsymbol{\theta}}\{\nabla l(\boldsymbol{\theta})\} \\ &= -E_{\boldsymbol{\theta}}\{\nabla^2 l(\boldsymbol{\theta})\}, \end{aligned} \quad (9)$$

where the notation Var means the variance, E means expectation, and

$$\nabla^2 l(\boldsymbol{\theta}) = \begin{bmatrix} N \sum_{k=1}^K x_k^2 g & N \sum_{k=1}^K x_k g \\ N \sum_{k=1}^K x_k g & N \sum_{k=1}^K g \end{bmatrix}, \quad (10)$$

and

$$\begin{aligned} g &= \frac{1}{Nf^2(1-f)^2} \{-N_{1,k}ff' + (N + N_{1,k})f^2f' - Nf^3f' - Nff'^2 \\ &\quad + Nf^2f'^2 + N_{1,k}ff'' - (N + N_{1,k})f^2f'' + Nf^3f''\}. \end{aligned} \quad (11)$$

For simpler notation, here, f is $f(\alpha x_k + \beta)$; f' is $f'(\alpha x_k + \beta)$; f'' is $f''(\alpha x_k + \beta)$; and $f''(s) = \frac{d^2f(s)}{ds^2}$.

With the Fisher information matrix, we can also derive the confidence intervals for the MLE,²¹

$$\hat{\alpha} \pm z \sqrt{(\mathbf{I}(\hat{\boldsymbol{\theta}})^{-1})_{11}}, \quad (12)$$

and

$$\hat{\beta} \pm z \sqrt{(\mathbf{I}(\hat{\theta})^{-1})_{22}}, \quad (13)$$

where z is the appropriate critical value (for example, 1.96 for 95% confidence), and the notation $(\mathbf{I}(\hat{\theta})^{-1})_{ii}$ means that we invert the Fisher information matrix first and then take the ii 'th component of the inverted matrix.

3.2 Signal Detection

The detection of an exoplanet signal can be modeled as a composite hypotheses testing problem. The null and alternative hypotheses are as follows:

$$H_0: \alpha = 0, \beta \geq 0 \quad \text{versus} \quad H_1: \alpha > 0, \beta \geq 0. \quad (14)$$

Here, since we have no prior information about the exoplanet intensity α and the background light β , the posterior probability of the observation under the two hypotheses cannot be calculated to decide which hypothesis is more likely. One solution for this challenge is to use the maximum likelihood. We can conduct a GLRT (or also called maximum-likelihood test), i.e., we determine whether there is an exoplanet signal based on the ratio,

$$\begin{aligned} R &= \frac{\max_{\alpha, \beta} p_1(\mathbf{y}|\alpha, \beta)}{\max_{\alpha, \beta} p_0(\mathbf{y}|\alpha, \beta)} \\ &= \frac{\max_{\alpha, \beta} \prod_{k=1}^K f(\alpha x_k + \beta)^{N_{1,k}} [1 - f(\alpha x_k + \beta)]^{N_{0,k}}}{\max_{\beta} \prod_{k=1}^K f(\beta)^{N_{1,k}} [1 - f(\beta)]^{N_{0,k}}} \\ &= \frac{\prod_{k=1}^K f(\hat{\alpha}_1 x_k + \hat{\beta}_1)^{N_{1,k}} [1 - f(\hat{\alpha}_1 x_k + \hat{\beta}_1)]^{N_{0,k}}}{\prod_{k=1}^K f(\hat{\beta}_0)^{N_{1,k}} [1 - f(\hat{\beta}_0)]^{N_{0,k}}}, \end{aligned} \quad (15)$$

where $p_0(\cdot)$ and $p_1(\cdot)$ are the probability under hypotheses H_0 and H_1 , respectively, and the two maximal probabilities and the corresponding parameters (α and β) are calculated using the estimation algorithm in Sec. 3.1. R is the generalized likelihood ratio. For easier calculation, we usually take the log of R , that is, the log-likelihood ratio. In a real space mission, we can conduct detection while sequentially collecting images. After receiving every new image, we update the test ratio in Eq. (15); when $R \geq \pi_U$ or $R \leq \pi_L$, we conclude H_1 or H_0 is accepted and stop taking new images, otherwise we take a new image and repeat the detection procedure, where π_U and π_L are thresholds chosen beforehand.

The thresholds can be determined according to users' desired true positive and false alarm rates. According to Wilks' theorem, the probability distribution of $2R$ under H_0 , i.e., twice the ratio, is approximately a chi-squared distribution with 1 degree of freedom.²² There is no simple closed-form theoretical solution for the true positive rate and false alarm rate given a threshold for this model, so we numerically calculate them via Monte Carlo simulation. Given the planet intensity that users want to detect and the number of PC images that they plan to take, they can numerically compute the true positive and false alarm rate for each threshold using this detection algorithm. They can choose a desirable true positive and false alarm rate value pair and thus its corresponding threshold.

We simulate a signal (signal only) with intensity $1e-8$ Jy (for reference, Venus is $2.99e-8$ Jy and Earth is $4.85e-9$ Jy if viewed from 10 pc at $0.552 \mu\text{m}$) and run multiple trials to get the statistics of the method's performance, shown in Fig. 4. The Bernoulli GLRT method correctly estimates the intensity of the signal quickly. The value of the log-likelihood ratio is high and increases quickly with the increasing number of observations. This means that it is possible to confidently detect the existence of a signal with just a few images and the more observations the larger the gain in confidence for the detection.

The following is a note on using Bernoulli (rather than binomial) in the method name. Each count in each pixel follows a Bernoulli distribution, while the sum of photon counts in each pixel follows a binomial distribution. Equation (6) describes the probability for a specific

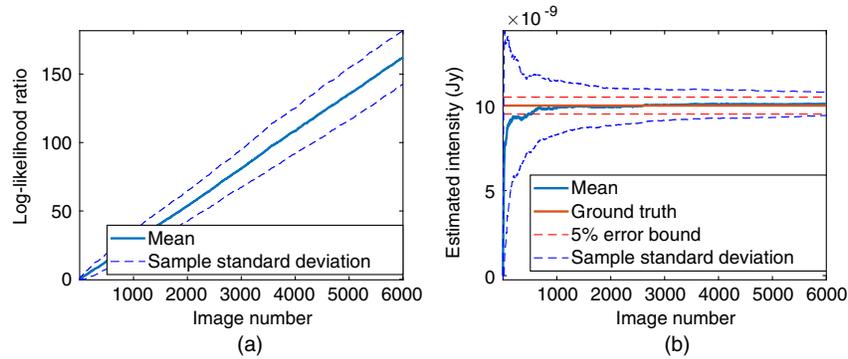


Fig. 4 Statistics of the maximum likelihood planet intensity estimation from 100 trials. The blue solid line is the average of all trials with an increasing number of observations and the blue dashed lines are the sample standard deviation band. (a) The log-likelihood ratio with an increasing number of observations. The more images that are available, the bigger the log-likelihood ratio is. Thus, the greater the confidence of the existence of a true signal. (b) The estimated intensity with increasing number of images. The red solid line is the true intensity. The red dashed lines are the 5% error band. The blue dashed line is the standard deviation band. The method results in a good estimate.

count sequence rather than their sum. Furthermore, the binomial distribution and Bernoulli distribution only differ by a binomial coefficient, which will not change the MLE theoretically. In our likelihood ratio in Eq. (15), the coefficient will also be cancelled out from the numerator and denominator, so Bernoulli and binomial will have the same detection result theoretically. In our algorithm, we include the PC images sequentially rather than work on a co-added image. (The MLE calculated using the first k PC images would be the initial value for the first $k + 1$ PC images and the process repeats until all the N images are included.) Thus, “Bernoulli distribution” is the more accurate name for our algorithm.

Theoretically, a cold start on a co-added image using BGLRT (i.e., we directly calculate the MLE for the co-added image from all the N PC images with a chosen set of initial values) should have the same result as calculating sequentially. However, calculating sequentially converges to better results in practice.

3.3 Multi-Signal Detection in an Image

The hypothesis testing process described in the previous section assumes the planet PSF is in the center of the set of K test pixels (the size of PSF’s central core), which will be called an image window from here on. In reality, an image much bigger than the size of a PSF core will be examined for an unknown number of planets in an unspecified location. Thus, the detection procedure is repeated for each set of K pixels across the whole image. This generates a log-likelihood ratio map for the image.

In this section, we demonstrate the performance of the Bernoulli GLRT detection method on simulated starshade images. The imaging system in all simulations in this paper is the Nancy Grace Roman Space Telescope with a starshade. The pixel size of the detector is $0.021 \text{ arc sec} \times 0.021 \text{ arc sec}$. The optical model to calculate the starshade diffraction uses Fresnel diffraction theory²³ to propagate light past the starshade. A detailed description of the simulation process can be found in Ref. 4.

We apply the Bernoulli GLRT method to simulated starshade images to detect possible exoplanet signals. As shown in Fig. 5(a), we observe our solar system from 10 pc away with a Starshade-Roman Telescope system, where two observable planets are in view, Venus and Earth. The Sun is present but is not recognizable because the starshade successfully suppressed its light. For this demonstration, we assume no local or extrasolar zodiacal dust. Here, we use a fixed integration time (1 s) for each single exposure and wavelength at $\lambda = 0.552 \mu\text{m}$ with a $0.128\text{-}\mu\text{m}$ bandwidth. More details about the simulation are available in Ref. 4. After including the PC mode detector described in Sec. 2, PC images are produced, an example of which is given

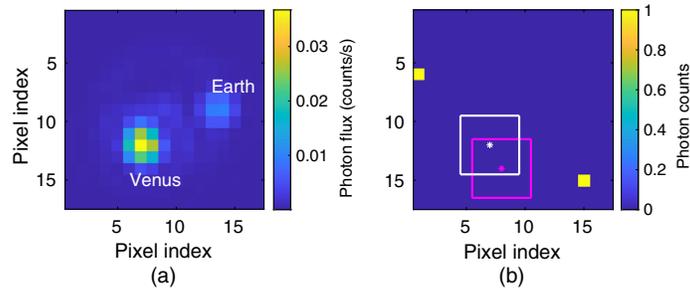


Fig. 5 (a) Simulated noiseless image for the solar system from 10 pc away with a Starshade-Roman Telescope system. This is the ground truth for the detection problem. Only the Sun, Earth, and Venus are included. The brighter signal is Venus and the dim one is Earth. The Sun's light is successfully suppressed by the starshade and thus can hardly be recognized. (b) One example of a PC image with the detector model described in Sec. 2. Examples of an image window are also shown. The white box is of the size of PSF core. It is the area used to detect whether there is a planet at the position of the white asterisk, i.e., pixel (7, 12). The magenta box is for the magenta asterisk, i.e., pixel (8, 14).

in Fig. 5(b). As the image size is much bigger than the size of a PSF core, the detection procedure is repeated for each image window with the size of a PSF core across the whole image. Examples of an image window are also shown in Fig. 5(b). The white box is of the size of PSF. It is the area used to calculate the likelihood ratio and then decide whether there is a planet at the position of the white asterisk, i.e., pixel (7, 12). The magenta box is for the magenta asterisk, i.e., pixel (8, 14). Repeating the process, we will get the likelihood ratio at each pixel and thus generates a log-likelihood ratio map for the whole image.

Though the method is applied to each PC image separately rather than on a co-added image, there is insufficient space to show all PC images one-by-one here. Thus, to give a sense of what the data look like, the co-added images are shown in Figs. 6(a)–6(c). Their corresponding log-likelihood ratio maps given by Bernoulli GLRT are shown in Figs. 6(d)–6(f). Venus is located at pixel (7, 12) and Earth at (14, 9). However, as we see in Figs. 6(d)–6(f), not only the pixel (7, 12), and the pixel (14, 9), where the planets are located, have high log-likelihood ratio values, but the pixels around them also have high log-likelihood ratio values compared to the background. This

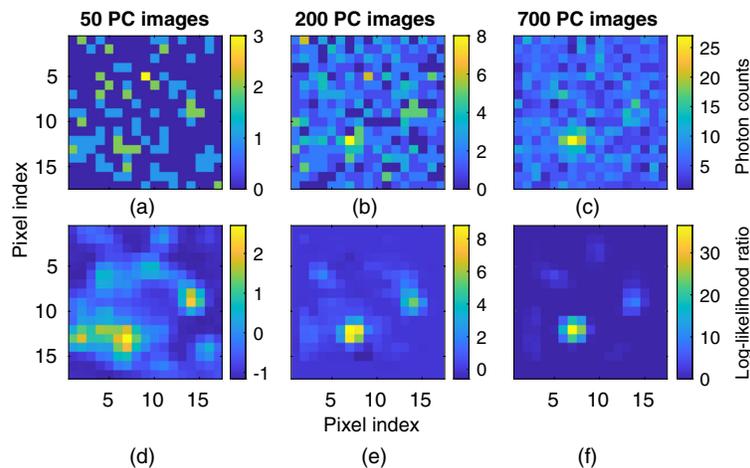


Fig. 6 Results of the new Bernoulli-based GLRT. (a) Co-added image with 50 sequential PC images for Fig. 5(a). (b) Co-added image with 200 sequential PC images for Fig. 5(a). (c) Co-added image with 700 sequential PC images for Fig. 5(a). (d) Log-likelihood ratio of each pixel using the 50 sequential PC images in (a). (e) Log-likelihood ratio of each pixel using the 200 sequential PC images in (c). (f) Log-likelihood ratio of each pixel using the 700 sequential PC images in (e). The change of log-likelihood ratio of Venus, Earth, and a background pixel with increasing number of observations is shown in Fig. 7.

makes sense as an image window centered on these pixels close to the signal position includes part of the signal. For example, for pixel (8, 14), the magenta asterisk, its image window, the magenta box in Fig. 5(b), covers most part of Venus. Our methods only consider two cases. The first is that a signal is at the center of the window. The second one is that no signal exists in the window (only constant background). Therefore, the bright partial signal will have a high log-likelihood ratio. However, the image window only includes part of the signal, so the light distribution does not fit the PSF template perfectly. That is to say, it does not follow hypothesis 1, so the log-likelihood ratio is lower than that in the real signal pixel. To decide the exact position of the signal, after thresholding the log-likelihood map, we will just choose the center of the detected area's circumscribed circle as the position of the planet. An example is shown in Fig. 7. This planet position estimation should be more robust, especially when only a partial signal is in the image. Our method detects the existence of the signals quickly and separates them from the background successfully. The more observations, the larger the gain in confidence for the detection.

The changes of log-likelihood ratio at Venus, Earth, and a background pixel at (5, 5) after each observation are shown in Fig. 8. As the figures demonstrate, our method can detect the existence of the signals quickly and separate them from the background successfully. The corresponding false alarm rate for the log-likelihood value of Venus and Earth for the three cases are in Table 2. Details on finding the false alarm rate are given in Sec. 3.4. The false alarm rate is how likely a background pixel will have a log-likelihood equal or bigger than the chosen threshold value and thus get wrongly considered as a detection. The smaller the false alarm rate is, the more confidence we have for the detection. As shown by the example in this section, the more observations, the more confidence we gain for the detection.

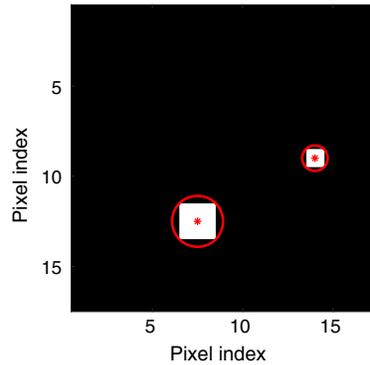


Fig. 7 Example of position estimation. Binary detection image after thresholding the log-likelihood ratio map in Fig. 6(e) with a threshold of 5. The red circles are the minimal bounding circles of the detected area. The centers of the circles can be used to estimate the planet positions.

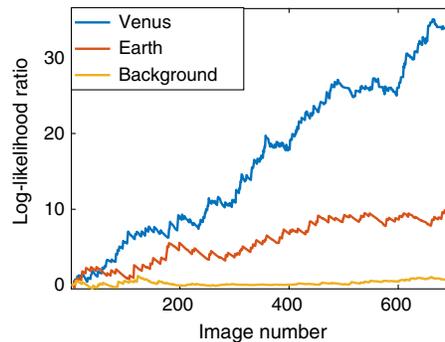


Fig. 8 Log-likelihood ratio of Venus, Earth, and a background pixel at (5, 5) with increasing number of observations for Fig. 5(a) using Bernoulli-based GLRT.

Table 2 False alarm rate (FA) for the cases in Fig. 6 using BGLRT and Gaussian GLRT (GGLRT).

Case	Venus BGLRT	Venus GGLRT	Earth BGLRT	Earth GGLRT
Fig. 6(a)	0.326	0.655	0.076	0.110
Fig. 6(b)	0.000	0.002	2.000e-04	0.004
Fig. 6(c)	0.000	0.000	0.000	0.000

When two signals overlap, the light distribution in the search area centered on the signals will be distorted. This violates the hypothesis that the image area contains a PSF-shaped signal and constant background. Therefore, the estimation and detection performance may be influenced. If there is not too much overlap, the intensity estimate for both signals will be higher than the true values and the position estimate will be biased toward each other. If the signals are close enough, the algorithm may consider them as one signal. Future work can look into solving this problem by expanding the image area and modeling overlapping signals. Moreover, astronomers can take another set of images after a while, because the signals' relative positions are likely to change and will be separate. In this work, we will not further consider the problem.

In the case of non-uniform background, the algorithm's performance depends on how well we know the non-uniform background. The current algorithm finds the signal looking like the PSF template against the uniform background in an area of the size as PSF core. We traverse the whole image by checking the PSF-core sized area one by one. Thus, the assumption of the algorithm is that the background is locally constant but can be non-uniform in the whole image. If the non-uniform background changes slowly spatially, a constant background may still be a good approximation and the detection will not be influenced much. If we have zero knowledge about the distribution of the bright background close to the planets, the results are affected by the distorted signal. However, if we have some knowledge about the non-uniform background, we can accommodate it in the model. For example, it is a reasonable approximation to assume the face-on exozodiacal light being axisymmetric. We have developed an iterative GLRT to do detection and estimation,²⁴ which is essentially an EM algorithm. We iteratively estimate either the planets' signal or the exozodiacal dust first, and then use the estimation as a known factor to estimate the other until both estimates converge.²⁴

3.4 Comparison with Other Methods

In Refs. 16 and 24, we introduced a GLRT method that was applied to co-added images generated from the combination of many frames. That method assumed the values in each pixel followed a Gaussian distribution, so the summation of frames was necessary in order to build enough signal in each pixel for the Gaussian assumption to be valid. Alternatively, the Bernoulli distribution in the new BGLRT method works for single PC images, which makes it easier to use online as it can process even a few PC images.

The underlying relationship between the probability in the Bernoulli distribution and flux intensity can be derived theoretically based on the detector model or directly measured in experiment. No approximation is used and thus the method is accurate and efficient at extracting information. In Figs. 9(a)–9(c), we show the T map using the Gaussian GLRT for the co-added images of Fig. 6.¹⁶ The T value is a proxy for the log-likelihood ratio that is easier to compute. The T value in a pixel is given by

$$T(\mathbf{x}_k) = (N - 2)(L_G(\mathbf{x}_k)^{\frac{2}{N}} - 1), \quad (16)$$

where $L_G(\mathbf{x}_k)$ is the log-likelihood ratio of the co-added image \mathbf{x}_k under a Gaussian assumption. The T value consists of basic operations (matrix multiplication, summation, and division) of the image and parameters, so it is easier to calculate than L_G .^{16,24} The T value is used to compared with a threshold to get a detection result in the Gaussian GLRT. Figure 9 shows that the T map is noisier than the log-likelihood map from the Bernoulli GLRT. The corresponding false alarm

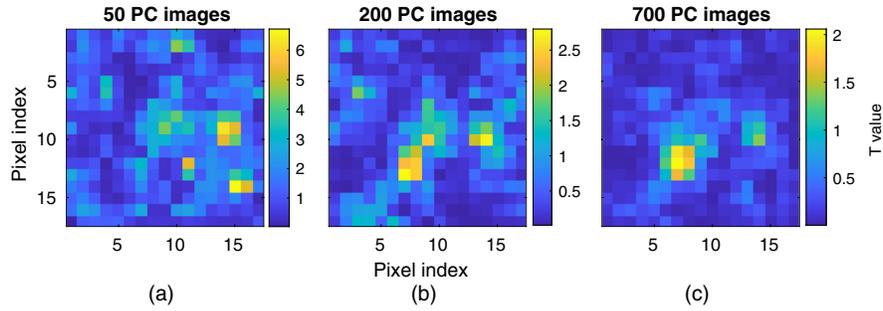


Fig. 9 Results of a Gaussian-based GLRT for images in Fig. 6. (a) T value of each pixel for the co-added image with 50 sequential PC images for Fig. 6(a). (b) T value of each pixel for the co-added image with 200 sequential PC images for Fig. 6(b). (c) T value of each pixel for the co-added image with 700 sequential PC images for Fig. 6(c).

rates for the T value of Venus and Earth for the three cases are given in Table 2. We generally have higher confidence using Bernoulli GLRT.

The above comparison indicates that the Bernoulli model for PC images contributes to the performance improvement. To further analyze the contribution from GLRT, we also compared the Bernoulli GLRT with the performance of detection based on the SNR from our Bernoulli model. With the MLE of signal intensity and estimated standard deviation of the MLE derived from the Fisher information matrix in Sec. 3.1, we define their ratio as the SNR from our model. We will refer to this SNR definition as Bernoulli SNR (BSNR). For each pixel, we utilize the image window centered at it and calculate the estimated signal intensity and the standard deviation as in Sec. 3.1. Then, we calculate the ratio. After repeating the process for all the pixels in the image, we obtain an SNR map. The SNR maps for the three example co-added images in Fig. 6 are shown in Fig. 10 to help visually compare the performance with the BGLRT method.

We also compared the Bernoulli GLRT with the performance of the detection method based on the SNR map implemented in pyKLIP.²⁵ For each pixel, the algorithm masks its surrounding pixels within the signal area in question (we chose the size of signal area as the size of PSF core) and then computes the standard deviation using the rest of pixels in concentric annuli. The ratio of the pixel value and this standard deviation is the SNR value at this pixel. The width of the annuli used is the diameter of the PSF core in this paper. Repeating the process for all the pixels in the image generates an SNR map. The SNR maps for the three example co-added images in Fig. 6 are shown in Fig. 11 to help visually compare the performance with the GLRT methods.

To further demonstrate the different properties of the methods, we compare their ROC curves for the detection of Venus and Earth. It is difficult to analytically derive a closed-form relationship between the false alarm rate and true positive rate and the threshold chosen in the new Bernoulli GLRT, unlike our previous GLRT,¹⁶ which assumes Gaussian noise. Thus, we use

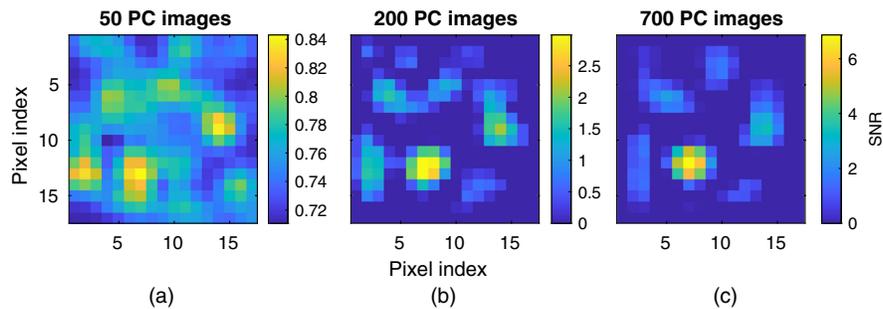


Fig. 10 The SNR map based on the Bernoulli model for images in Fig. 6. (a) The SNR map of each pixel for the co-added image with 50 sequential PC images for Fig. 6(a). (b) The SNR map of each pixel for the co-added image with 200 sequential PC images for Fig. 6(b). (c) The SNR map of each pixel for the co-added image with 700 sequential PC images for Fig. 6(c).

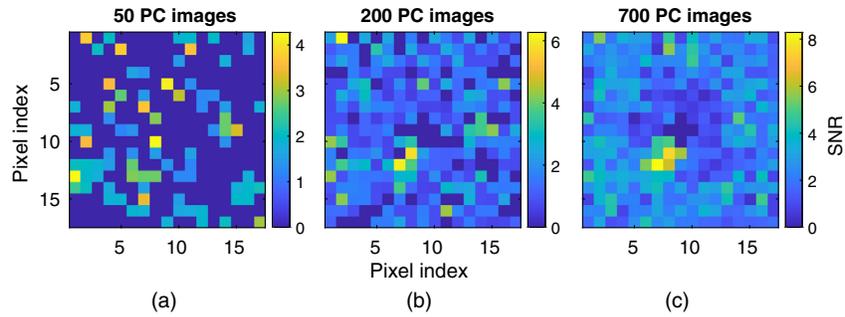


Fig. 11 The SNR map from the pyKLIP package²⁵ for images in Fig. 6. (a) SNR of each pixel for the co-added image with 50 sequential PC images as shown in Fig. 6(a). (b) SNR of each pixel for the co-added image with 200 sequential PC images as shown in Fig. 6(b). (c) SNR of each pixel for the co-added image with 700 sequential PC images as shown in Fig. 6(c).

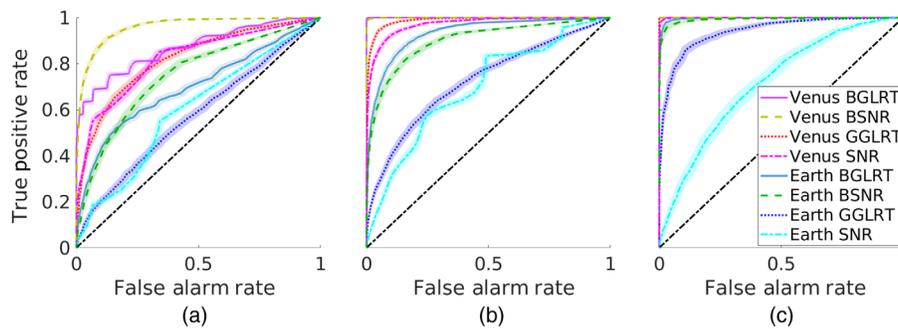


Fig. 12 ROC curves with confidence interval for Earth and Venus detection with the four different methods. “BGLRT” means the result used the Bernoulli GLRT described in this paper, which processes each PC image sequentially using the GLRT based on a Bernoulli distribution. “BSNR” means the detection method based on our BSNR, which utilizes the final MLEs from our Bernoulli model and estimated standard deviation from Fisher information matrix. “GGLRT” in the legends means the result used the previous GLRT,¹⁶ which processes co-added images and assumes Gaussian noise. “SNR” means the detection method is based on SNR implemented in pyKLIP, which is applied to co-added images. The shaded region behind each ROC curve is its 95% confidence interval. (a) ROC curves using 50 PC images calculated from 5000 trials. (b) ROC curves using 200 PC images calculated from 5000 trials. (c) ROC curves using 700 PC images calculated from 2000 trials.

Monte Carlo simulations, using multiple thresholds for each simulation, to calculate the ROC curves, shown in Fig. 12. The simulation is the same as described in Sec. 3.3, which is the starshade images with the Sun, Venus, and Earth. We compare the performance of the methods using different numbers of total images, which is denoted as n_{pc} . We apply the Bernoulli GLRT on the set of images and obtain the log-likelihood ratio map; we also apply the Gaussian GLRT to get the false alarm rate map. We also apply BSNR and SNR from pyKLIP on the images. We apply a set of different thresholds with the resulting detection or missed detection of Earth, Venus, and a background pixel. We run a large number of trials, which is denoted as n_{trials} , and record the ratio of detection of Earth and Venus as the true positive rate for Earth and Venus, and record the ratio of detection of the background pixel as the false positive rate. For each ROC curve, we also calculate the confidence interval shown as the shaded areas in Fig. 12. The confidence interval is standard Monte Carlo statistics, which is detailed in the Appendix. As can be seen in Fig. 12, the confidence interval is tight.

For Fig. 12(a), we ran 5000 trials with 50 PC images generated for Venus and Earth in each trial. The Bernoulli GLRT and BSNR were applied to the 50 PC images sequentially; the Gaussian GLRT and SNR from pyKLIP were applied to the co-added image of the 50 PC images. Results for different decision thresholds were recorded and combined into the ROC curve. For Fig. 12(b), we ran 5000 trials and 200 PC images for each trial. For Fig. 12(c),

Table 3 Comparison of AUC for BGLRT, BSNR, GGLRT, and SNR method from pyKLIP.²⁵

	Venus BGLRT	Venus BSNR	Venus GGLRT	Venus SNR	Earth BGLRT	Earth BSNR	Earth GGLRT	Earth SNR
50PC	0.8697	0.9594	0.8205	0.8224	0.6991	0.7488	0.5753	0.5991
PC	0.9999	0.9996	0.9878	0.9714	0.9392	0.9002	0.7245	0.6953
PC	1	1	1	0.9993	0.9970	0.9920	0.9401	0.7159

we ran 2000 trials and 700 PC images for each trial. The performance for Venus is better than Earth when we use the same method and the same number of PC images. One reason is that Venus is brighter than Earth. Moreover, Venus is further away from the star at the center, so the influence from the residual starlight is smaller compared to that for Earth. The performance for both Venus and Earth is better with more images. It is easy to understand that the method performs better with more information, i.e., more images. Generally speaking, the Bernoulli GLRT outperforms the other methods.

For easier comparison, we also list the AUC for all the curves in Table 3. AUC is an aggregate measure of performance across all possible thresholds. It can be interpreted as the probability that the model ranks a random positive example higher than a random negative example. AUC is 1 if the model's decisions are all correct and 0 if all wrong. More comparisons of GLRT with Gaussian assumption and the SNR method can be found in our other work.²⁴ Overall, Bernoulli GLRT outperforms GLRT with Gaussian assumption and the SNR method.

3.5 Early Stopping for Observation When No Planet Exists

As shown in the previous examples, the Bernoulli GLRT successfully detects a planet even for short integration times. Another important problem is to know when to stop when no detection is made so time wasted on a system with no planets is minimized. As no closed-form analytical solution for Bernoulli GLRT can be derived among false alarm rate and intensity and number of PC images, which will define the total time needed, it is impossible to solve for the time needed to reach a specific false alarm rate given the planet intensity via inverting a function. Instead, we will rely on numerical calculation via Monte Carlo simulation.

First, three parameters are specified: the minimum planet intensity to be detected, the maximum false alarm rate that can be accepted, and the minimum true positive rate that is acceptable. Then, for a different number of images, the ROC is calculated via Monte Carlo simulation. Finally, the minimum number of PC images that can reach the requirements are chosen. For example, if we assume the dimmest planet has the same intensity as Earth, the maximum acceptable false alarm rate is 0.16 and the minimum acceptable true positive rate is 0.85. The acceptable false alarm rate and true positive rate pairs are in the shaded green region in Fig. 13.

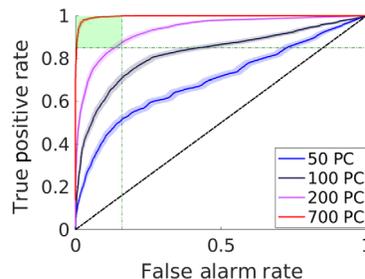


Fig. 13 ROC with confidence interval for Earth using Bernoulli GLRT with different number of PC images. The horizontal green dotted line is the acceptable true positive rate (TP) and the vertical one is the acceptable false alarm rate (FA) region. Thus, the acceptable true TP and acceptable FA region is the shaded green area. The ROC with 200 PC is the first ROC reaching the acceptable region with the fewest possible images.

We calculate the ROCs with different numbers of PC images and find that 200 PC images' ROC is the first one to reach the green region, as shown in Fig. 13. Thus, after taking 200 PC images and if there is still no signal detected, the telescope system could move on to observing another target star with confidence there is no planet.

4 Conclusions

In this paper, we presented a model and signal processing approach based directly on single PC images. Previous work on PC image processing for the most part only uses co-added PC images rather than operating on individual PC images. Such methods do not provide theoretical guidance for choosing the total number of PC images to combine into one co-added image as in our previous work on GLRT under Gaussian assumption.²⁴ Our method, the BGLRT, here directly works with each PC image, so there is no need to choose the hyper parameter, the number of PC images to combine into one co-added image. Furthermore, previous work on PC image processing assumes Gaussian noise for co-added images, but the method here directly uses the detector model, which more accurately represents its noise characteristics. In this paper, we showed that such a method outperforms our previous GLRT method using co-added images under a Gaussian noise assumption, and the SNR method commonly used in high-contrast imaging. We also compared the BGLRT with the BSNR, which also uses the estimation results from our Bernoulli model. The BGLRT and the BSNR have similar performance (BGLRT is a little better in most cases) and are better than other methods, which indicates the performance gain in detection is mostly the result of the improved model for the imaging process. Furthermore, our method provides the maximum likelihood estimate of exoplanet intensity and background intensity.

This approach maximizes the utilization of information presented in each observation. We applied this method to simulated starshade images and successfully detected Venus and Earth. Directly processing the PC images online helps allocate the observation time efficiently. We can compare the log-likelihood ratio with thresholds chosen beforehand after each observation and stop accordingly. As we can decide the existence or lack of planet efficiently in this way, we can move to other planetary systems if there is no planet or make sure that enough information is gathered if there is a planet. In addition to the observation time, the analysis of detection performance introduced could also give quantitative guidance on the choice of imaging parameters, such as the threshold.

The Bernoulli model for PC images and the resulting GLRT for hypothesis testing are a general model and detection method that can be significantly expanded for different applications. This work focuses on the simple case to demonstrate the concept, but the framework is more general. Through Eq. (5), the likelihoods in different pixels are related by the underlying template for signals we want to detect, i.e., a PSF template plus constant background as in Eq. (4). Thus, the resulting GLRT decides whether such template exists in the observed area. However, this is just one particular application example of the model and the method. By replacing x in Eq. (4), different signal templates can be chosen to appropriately match the observation scenario; the remaining steps in the derivation and the application of the Bernoulli GLRT are the same.

5 Appendix: Confidence Interval for ROC Curves

The confidence interval of a proportion \hat{p} is given by²⁶

$$\left(\hat{p} - z_{1-\eta/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n_{\text{trials}}}}, \hat{p} + z_{1-\eta/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n_{\text{trials}}}} \right). \quad (17)$$

This confidence interval is based on Fisher information. Thus, for a point $(\hat{p}_{\text{false}}, \hat{p}_{\text{positive}})$ on the ROC curve, we take the confidence interval using the two points,

$$\left(\hat{p}_{\text{false}} - z_{1-\eta/2} \sqrt{\frac{\hat{p}_{\text{false}}(1-\hat{p}_{\text{false}})}{n_{\text{trials}}}}, \hat{p}_{\text{positive}} + z_{1-\eta/2} \sqrt{\frac{\hat{p}_{\text{positive}}(1-\hat{p}_{\text{positive}})}{n_{\text{trials}}}} \right) \quad (18)$$

and

$$\left(\hat{p}_{\text{false}} + z_{1-\eta/2} \sqrt{\frac{\hat{p}_{\text{false}}(1-\hat{p}_{\text{false}})}{n_{\text{trials}}}}, \hat{p}_{\text{positive}} - z_{1-\eta/2} \sqrt{\frac{\hat{p}_{\text{positive}}(1-\hat{p}_{\text{positive}})}{n_{\text{trials}}}} \right), \quad (19)$$

where $z_{1-\eta/2}$ is the upper critical value for the standard normal distribution with $\eta/2$ area to the right of it.

Acknowledgments

This work was supported by Caltech-JPL NASA under Grant No. NNN12AA01C. The authors would like to thank the anonymous reviewers for the insightful comments, especially for suggesting the comparison with methods based on SNR defined with our Bernoulli model results. The authors have no relevant financial interests in the paper and no other potential conflicts of interest to disclose.

References

1. W. A. Traub and B. R. Oppenheimer, *Direct Imaging of Exoplanets*, pp. 111–156, University of Arizona, Tucson (2010).
2. M. Hirsch et al., “A stochastic model for electron multiplication charge-coupled device—‘from theory to practice’,” *PLoS One* **8**(1), e53671 (2013).
3. T. R. Marsh, *High-Speed Optical Spectroscopy*, pp. 75–94, Springer, Dordrecht (2008).
4. M. Hu et al., “Simulation of realistic images for starshade missions,” *Proc. SPIE* **10400**, 104001S (2017).
5. G. A. de Vree et al., “Photon-counting gamma camera based on an electron-multiplying CCD,” *IEEE Trans. Nucl. Sci.* **52**(3), 580–588 (2005).
6. M. C. DeSantis et al., “Precision analysis for standard deviation measurements of immobile single fluorescent molecule images,” *Opt. Express* **18**(7), 6563–6576 (2010).
7. C. Mackay et al., “High-speed, photon-counting CCD cameras for astronomy,” *Proc. SPIE* **7742**, 774202 (2010).
8. O. Daigle et al., “Characterization results of EMCCDs for extreme low-light imaging,” *Proc. SPIE* **8453**, 845303 (2012).
9. A. N. Wilkins et al., “Characterization of a photon counting EMCCD for space-based high contrast imaging spectroscopy of extrasolar planets,” *Proc. SPIE* **9154**, 91540C (2014).
10. L. K. Harding et al., “Technology advancement of the CCD201-20 EMCCD for the WFIRST coronagraph instrument: sensor characterization and radiation damage,” *J. Astron. Telesc. Instrum. Syst.* **2**(1), 011007 (2015).
11. T. Plakhotnik, A. Chennu, and A. V. Zvyagin, “Statistics of single-electron signals in electron-multiplying charge-coupled devices,” *IEEE Trans. Electron Devices* **53**(4), 618–622 (2006).
12. E. Lantz et al., “Multi-imaging and Bayesian estimation for photon counting with EMCCDs,” *Mon. Not. R. Astron. Soc.* **386**, 2262–2270 (2008).
13. K. B. W. Harpse, M. I. Andersen, and P. Kjgaard, “Bayesian photon counting with electron-multiplying charge coupled devices (EMCCDS),” *Astron. Astrophys.* **537**, A50 (2012).
14. C. A. G. Gonzalez et al., “VIP: vortex image processing package for high-contrast direct imaging,” *Astron. J.* **154**(1), 7 (2017).
15. R. Soummer, L. Pueyo, and J. Larkin, “Detection and characterization of exoplanets and disks using projections on Karhunen–Loève eigenimages,” *Astrophys. J. Lett.* **755**(2), L28 (2012).
16. M. Hu, A. Harness, and N. J. Kasdin, “Image processing methods for exoplanets detection and characterization in starshade observations,” *Proc. SPIE* **10698**, 106985K (2018).
17. M. Hu, H. Sun, and N. J. Kasdin, “Sequential generalized likelihood ratio test for planet detection with photon-counting mode,” *Proc. SPIE* **11117**, 111171K (2019).

18. M. Stanford and B. Hadwen, “The noise performance of electron multiplying charge-coupled devices,” *IEEE Trans. Electron Devices* **50**(5), 1227–1232 (2003).
19. O. Daigle et al., “Extreme faint flux imaging with an EMCCD,” *Publ. Astron. Soc. Pac.* **121**(882), 866 (2009).
20. M. J. Rizzo et al., “Simulating the WFIRST coronagraph integral field spectrograph,” *Proc. SPIE* **10400**, 104000B (2017).
21. T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, pp. 111–156, Springer, New York (2009).
22. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses,” *Ann. Math. Stat.* **9**(1), 60–62 (1938).
23. D. Sirbu, “Occulter-based high-contrast exoplanet imaging: design, scaling, and performance verification,” PhD Thesis, Princeton University, Princeton, New Jersey (2014).
24. M. Hu et al., “Exoplanet detection in starshade images,” *J. Astron. Telesc. Instrum. Syst.* **7**(2), 021214 (2021).
25. J. Wang et al., “pyKLIP: PSF subtraction for exoplanets and disks,” Astrophysics Source Code Library, record ascl:1506.001 (2015).
26. J. P. M. De Sá, *Applied Statistics Using SPSS, STATISTICA, MATLAB and R*, pp. 92–93, Springer, New York (2007).

Mengya (Mia) Hu is a PhD candidate in mechanical and aerospace engineering at Princeton University. Her research focuses on the image simulation and signal detection of space telescope systems with starshades. She graduated from the Department of Thermal Science and Energy Engineering, University of Science and Technology of China in 2015 and was awarded the highest honor of the university, the “Guo Mo-Ruo Scholarship.”

He Sun is a postdoctoral researcher of computing and mathematical sciences at California Institute of Technology. He received his PhD from Princeton University in 2019 and his BS degree from Peking University in 2014. His research focuses on adaptive optics and computational imaging, especially their applications in astrophysical and biomedical sciences, such as exoplanet and black hole imaging.

Anthony Harness is an associate research scholar in the Mechanical and Aerospace Engineering Department at Princeton University. He received his PhD in astrophysics from the University of Colorado Boulder in 2016. He currently leads the experiments at Princeton validating starshade optical technologies.

N. Jeremy Kasdin is the assistant dean for engineering at the University of San Francisco and the Eugene Higgins Professor of Mechanical and Aerospace Engineering, emeritus, at Princeton University. He received his PhD from Stanford University in 1991. After being the chief systems engineer for NASA’s Gravity Probe B spacecraft, he joined the Princeton Faculty in 1999 where he researched high-contrast imaging technology for exoplanet imaging. From 2014 to 2016, he was vice dean of the School of Engineering and Applied Science at Princeton. He is currently the adjutant scientist for the coronagraph instrument on NASA’s Wide Field Infrared Survey Telescope.