The Astrophysical Journal, 813:66 (15pp), 2015 November 1 doi:10.1088/0004-637X/813/1/66
© 2015. The American Astronomical Society. All rights reserved.

DETECTING UNSPECIFIED STRUCTURE IN LOW-COUNT IMAGES

Nathan M. Stein[1], David A. van Dyk[2], Vinay L. Kashyap[3], and Aneta Siemiginowska[3]
[1] Department of Statistics, The Wharton School, University of Pennsylvania, 400 Jon M. Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104-6340, USA; natstein@wharton.upenn.edu
[2] Statistics Section, Imperial College London Huxley Building, South Kensington Campus, London SW7 2AZ, UK; dvandyk@imperial.ac.uk
[3] Smithsonian Astrophysical Observatory, 60 Garden Street, Cambridge, MA 02138, USA; vkashyap@cfa.harvard.edu, asiemiginowska@cfa.harvard.edu
Received 2015 January 13; accepted 2015 August 31; published 2015 October 29

ABSTRACT

Unexpected structure in images of astronomical sources often presents itself upon visual inspection of the image, but such apparent structure may either correspond to true features in the source or be due to noise in the data. This paper presents a method for testing whether inferred structure in an image with Poisson noise represents a significant departure from a baseline (null) model of the image. To infer image structure, we conduct a Bayesian analysis of a full model that uses a multiscale component to allow flexible departures from the posited null model. As a test statistic, we use a tail probability of the posterior distribution under the full model. This choice of test statistic allows us to estimate a computationally efficient upper bound on a p-value that enables us to draw strong conclusions even when there are limited computational resources that can be devoted to simulations under the null model. We demonstrate the statistical performance of our method on simulated images. Applying our method to an X-ray image of the quasar 0730+257, we find significant evidence against the null model of a single point source and uniform background, lending support to the claim of an X-ray jet.

Key words: galaxies: jets – methods: data analysis – methods: statistical – quasars: individual (0730+257) – techniques: image processing – X-rays: general

1. INTRODUCTION

Detecting scientifically meaningful structure in digital images is a ubiquitous and notoriously difficult problem. Typically, image analysis algorithms in high-energy astronomy are optimized for the detection and characterization of point sources. However, this strategy fails when confronted with complex extended structures at many scales. Optical observations often reveal rich and irregularly structured emission associated with a variety of objects in the universe, such as galaxies, nebulae, clusters of stars, or clusters of galaxies. The X-ray emission of these objects is often as rich as the optical, but the Poisson nature of the observed images makes the emission hard to discern. The X-ray images are often sparse and may require binning to expose the emission features, but binning lowers the resolution and potentially leads to loss of the smaller scale structures. Detecting irregular X-ray emission is thus challenging, and there has been no principled method to date to assess the statistical significance of arbitrary irregular features in X-ray images.

Source detection algorithms, such as `celldetect` (Calderwood et al. 2001) and `wavdetect` (Freeman et al. 2002) in Chandra Interactive Analysis of Observations (CIAO), work quite well for detecting point sources, but not for unspecified irregular emission. The CIAO `vtpdetect` algorithm (Ebeling & Wiedenmann 1993) can identify extended regions by looking at the distribution of tesselation areas and imposing a threshold cut, but does not otherwise determine the significance of the detected sources. Moreover, `vtpdetect` can spuriously combine the diffuse emission with embedded point sources, resulting in the confusion of the emission components. Other techniques used by astronomers include direct two-dimensional fitting of image features with pre-defined models, and qualitative analysis of residuals from such fits. Many studies also rely on maximum entropy-based image deconvolution techniques (e.g., Richardson 1972; Lucy 1974), but these typically do not yield unique fits and do not provide associated uncertainties. A Bayesian method that constructs a representation of an image using a basis set of regions (`pixons`; Pina & Puetter 1992) has also been tried on astronomical images, but again, without a means to evaluate the significance of identified regions. More generally, powerful computational tools such as Markov chain Monte Carlo (MCMC) enable researchers to fit more and more sophisticated models that capture the complexities of astronomical sources and instruments. On its own, however, MCMC is better suited for fitting a model than for choosing between models or for detection problems (see, however, Weinberg 2012). Thus new tools are needed for quantifying statistical significance or goodness of fit in the context of complicated MCMC-fitted models.

This paper addresses the problem of detecting image structure not adequately accounted for by a baseline model that represents features known to be present in the image; from a statistical perspective the baseline model serves as the null hypothesis. This formulation is useful in a wide range of applications. Here and in a forthcoming companion paper (K. McKeough et al. 2015, in preparation), we consider the problem of detecting X-ray jets emitted from quasars. The baseline model includes only the quasar and a flat background, with no additional emission representing the jet. Because it is difficult to specify parametric models that adequately capture the range of possible appearances of X-ray jets, we use a multiscale model that allows flexible, nonparametric departures from the baseline. Another possible application is detecting dynamic behavior, such as the time evolution of a supernova remnant. In these cases, the baseline model could be constructed using an earlier image, and the goal would be to test whether a later image represented a significant departure from the earlier one. Such applications extend beyond astronomy. Detecting changes in images over time is important in fields ranging from medical imaging to surveillance; see

Radke et al. (2005) for a review. Finally, we might be interested in detecting fine structure blurred by a point spread function (PSF), such as when analyzing filament structure in coronal loops in images of the Sun (McKeough et al. 2014). In this case, the baseline model could include readily apparent low-frequency structure, and the goal would be to detect mid-frequency departures from this model, after removing high-frequency noise.

Much previous work has attempted to quantify the statistical uncertainty of inferred features in noisy images. In functional magnetic resonance imaging, for example, Friston et al. (1995) proposed "statistical parametric maps," pixel-wise significance tests with subsequent adjustments for multiplicity based on Gaussian random fields. In astronomy, Ripley & Sutherland (1990) used spatial stochastic processes to directly model structures in spiral galaxies, and Esch et al. (2004) obtained uncertainty estimates for reconstructed X-ray images using a Bayesian model known as EMC2 that included multiple levels for instrumental effects, background contamination, and a multiscale hierarchical representation of the source. Sutton & Wandelt (2006) and Sutter et al. (2014) used Bayesian models to perform image reconstruction with uncertainty estimates using radio interferometry data. Bayesian methods have also been employed to quantify the uncertainty in the large-scale structure of the universe (Jasche & Wandelt 2013) and in secondary anisotropies of the cosmic microwave background radiation (Bull et al. 2014). Friedenberg & Genovese (2013) proposed a multiple testing procedure for detecting sources in astronomical images. Other approaches can be found in the computer vision literature; see for instance Godtliebsen et al. (2004), Holmström & Pasanen (2012), and Thon et al. (2012).

Rather than estimating the uncertainty in inferred image features, we focus on the more fundamental problem of feature detection. Specifically, we adopt a hypothesis testing framework to address the statistical question of whether there is sufficient evidence to conclude that the structure observed in an image is unlikely to have been produced by chance under the baseline (null) model. This framework ensures that we can control the probability of a false positive result, i.e., declaring that there is significant additional structure in the image beyond the baseline model, when in fact there is none. Our test statistic is a tail probability of a Bayesian posterior distribution under a full statistical model that includes both the baseline model and a multiscale component for structure not included in the baseline. This distinguishes our method from existing goodness-of-fit tests for inhomogeneous (baseline) Poisson processes (e.g., Guan 2008). We do not frame our approach in terms of Bayesian model selection (e.g., using Bayes factors) because the flexible full model for additional emission beyond the baseline is intentionally weakly specified, making it especially difficult to reliably apply Bayes factors due to their sensitivity to prior distributions (e.g., van Dyk 2012, pp. 141–146).

Our framework provides a reference distribution with which we can quantify how inferences given the observed data differ from inferences given data generated under the null model, when all analyses are performed under the full model. Computationally, this is accomplished by simulating multiple replicate images under the null model, fitting the full model to each, and computing the test statistics for each. This gives us a reference distribution for the test statistic that we can compare with the test statistic computed from the observed image in order to determine the statistical significance of apparent structure in the image. A primary novelty of our method is its use of an upper bound on the p-value that enables us to obtain statistical significance with a limited number of replicate images. Because each replicate image is fit under the fully Bayesian model, limiting their number is important for controlling the computational demands of the method.

We use a Bayesian model to infer image structure because it provides a principled way to account for a PSF, varying probability of photon detection, and varying amounts of smoothing at different image resolutions. This model builds on Esch et al.'s EMC2 in that it uses the same multiscale representation but for a different purpose. Whereas Esch et al. (2004) used this multiscale model to fully represent the source, we include a baseline model for the source and use the multiscale model of EMC2 to flexibly capture added structure beyond the baseline model. Combining this extension with the formal statistical testing proposed here leads to a new statistical software package called LIRA[4] (Low-count Image Reconstruction and Analysis). Like EMC2, LIRA deploys MCMC for fully Bayesian model fitting of the overall multilevel statistical model.

This article is organized into five sections. We begin in Section 1.1 with a simulated example that illustrates the scientific problem we aim to solve and serves as a running example for the remainder of the paper. In Section 2 we formulate the baseline and full models that we compare using a formal hypothesis testing framework in Section 3. A set of simulation studies and an analysis of the X-ray jet associated with the 0730+257 quasar illustrate our proposed methods in Section 4. We conclude with discussion in Section 5 and technical details in the Appendix.

1.1. A Simulated Example

As a concrete example of the scientific problems this paper addresses, we consider jet detection in an X-ray image of a quasar. This example applies also to detecting any secondary faint source in the vicinity of a bright point source, such as multiple sources, halos, non-uniform faint emission, and so on. Quasar jets extend out to distances on the order of 100 kiloparsecs from a supermassive black hole, and can trace the history of the black hole's activity and power (Urry & Padovani 1995; Harris & Krawczynski 2006). Many such jets have been observed for the first time in X-rays with the *Chandra* X-ray Observatory, where some quasar images show a bright point source and a much fainter jet feature. The jet surface brightness is non-uniform, so brighter knots and fainter emission are often seen along the jet.

Figure 1(a) is a simulated ground-truth image of a quasar, modeled as a single bright point source with a jet composed of two fainter point sources. Figure 1(b) is a simulated observation designed to mimic the degradation of Figure 1(a) due to a detector's PSF and Poisson noise from a limited exposure time. Both images are 64×64 pixels. The quasar was simulated as a two-dimensional spherical Gaussian with standard deviation 0.5 pixel and 500 expected counts. The jet was composed of two elliptical Gaussian components each with ellipticity 0.5, standard deviation 0.5 pixel, and 20 expected counts. The simulated background was uniform across the image with 200 total expected counts (approximately 0.05 counts per pixel).

[4] LIRA is a package for the R statistical programming language (r-project.org) and is available at github.com/astrostat/LIRA.

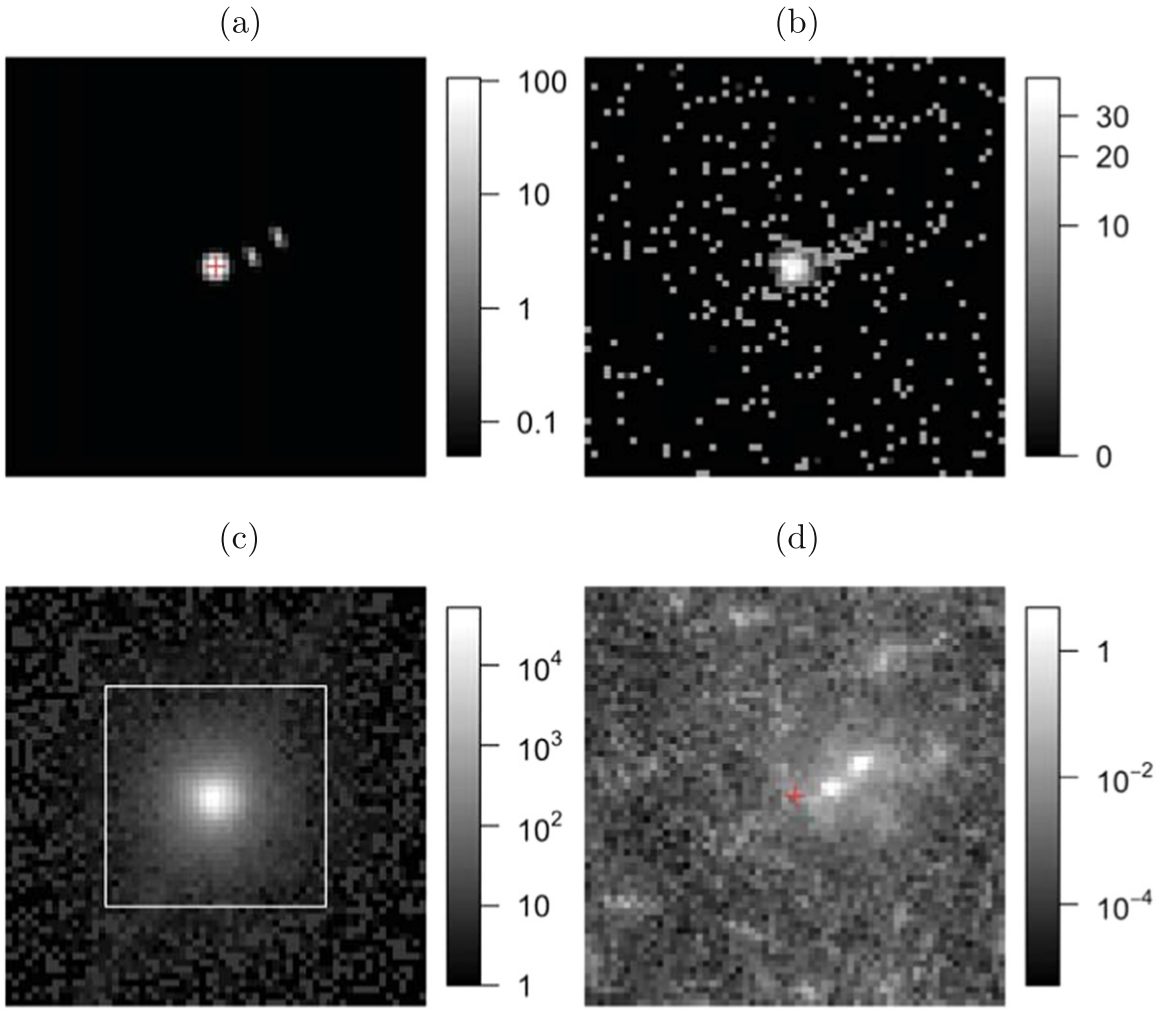


Figure 1. (a) True underlying intensity for a simulated quasar with a jet consisting of two additional Gaussian sources. (b) Simulated observation for the quasar and jet in (a), after convolving (a) with the PSF in (c) and introducing Poisson noise. (c) Un-normalized PSF for the analysis of simulated and observed quasars. The full PSF was used to simulate data, but only the portion inside the white square was used in analyses under the full model. (d) Pixel-wise posterior means of the added component $\tau_1 \mathbf{\Lambda}_1$ in the full model, given the simulated data in (a). The red + in (d) identifies the location of the simulated quasar in (a).

We aim to compare a baseline model that posits that Figure 1(b) was generated from a background-contaminated single point source (i.e., just a quasar, with no jet) with a full model that allows for unspecified structure in the image beyond the single point source. Further, we aim to quantify the statistical evidence in the image for choosing between the two models. Our method specifically avoids parametric modeling of the extended structure (here the jet) because in practice we often do not wish to specify the precise nature of possible departures from the simple model. More generally, we want to flexibly detect and quantify the evidence for departures from simple baseline models of images that are observed under the imperfect conditions that often arise in high-energy astrophysics, such as photon scatter, background contamination, and non-negligible Poisson noise.

2. MODELS AND HYPOTHESES

2.1. The Statistical Model

We consider an image composed of n photon counts arranged into a grid of pixels; we denote the counts in the n pixels by $\boldsymbol{y}_{\text{obs}} = (y_1, \ldots, y_n)$. If the two-dimensional image written in matrix form has l rows and m columns, then in our vectorized notation it has dimension $n = lm$. We model the image as a superposition of two Poisson processes. The first is intended to represent known or presumed aspects of the image, which could include anything from background noise to complicated structures of interest. For example, if we aim to quantify the evidence for a jet extending from a known point source, as in Section 1.1 and Figure 1(b), then the first Poisson process would consist of the point source and background contamination. We refer to this first Poisson process as the *baseline component*. The second Poisson process is intended to account for image features unexplained by the first process and is called the *added component*. In the example of testing for a jet, the added component would model the hypothesized jet.

Because of blurring and varying instrument response across the detector, the distribution of the counts observed in detector pixel i is

$$y_i \overset{\text{indep}}{\sim} \text{Poisson}\left(\sum_{j=1}^{n} P_{ij} A_j \left(\mu_{0j} + \mu_{1j}\right)\right), \tag{1}$$

where $\boldsymbol{P}$ is the $n \times n$ PSF, with (i, j) element P_{ij} denoting the probability that a photon in location j is observed in pixel i; $\boldsymbol{A} = (A_1, \ldots, A_n)$ is the detector efficiency, with A_j equal to the

probability that a photon in pixel j is detected; and $\boldsymbol{\mu}_0 = (\mu_{01}, \ldots, \mu_{0n})$ and $\boldsymbol{\mu}_1 = (\mu_{11}, \ldots, \mu_{1n})$ are, respectively, the intensities of the baseline and added components. The representation of the PSF in Equation (1) is quite general in that each column of $\boldsymbol{P}$ is the vectorized PSF for a particular source pixel. Thus, this representation allows the PSF to vary across the source. Throughout this paper, we assume that $\boldsymbol{P}$ and $\boldsymbol{A}$ are known, that $\boldsymbol{\mu}_1$ is an unknown parameter of interest, and that $\boldsymbol{\mu}_0$ includes some known structure but is partially unknown. When creating the simulated observation in Figure 1(b), we set $\boldsymbol{A} = (1, \ldots, 1)$ and used the PSF in Figure 1(c). Details on this PSF appear in Section 4.1.

We parameterize component i ($i = 0$ or 1) as $\boldsymbol{\mu}_i = \tau_i \boldsymbol{\Lambda}_i$, where $\tau_i = \sum_{j=1}^{n} \mu_{ij}$ is the expected photon count in component i, and $\boldsymbol{\Lambda}_i = (\Lambda_{i1}, \ldots, \Lambda_{in}) = \boldsymbol{\mu}_i / \tau_i$ is the proportion of τ_i that is expected in each pixel. The baseline component is often parameterized in terms of a lower dimensional parameter vector $\boldsymbol{\nu}$, in which cases we write $\boldsymbol{\Lambda}_0 = \boldsymbol{\Lambda}_0(\boldsymbol{\nu})$. The parameter $\boldsymbol{\nu}$ may include unknown aspects of posited image structure, such as the location of a point source. In practice, some parameters in $\boldsymbol{\nu}$ may be well-constrained by the data, and fixing these parameters at their estimates causes no problems, but we also consider the situation in which there is substantial uncertainty in at least some components of $\boldsymbol{\nu}$; see Section 3.5. If the baseline component is fully specified except for its total intensity τ_0, then $\boldsymbol{\nu}$ will be empty.

In fitting the model in Equation (1), τ_1 and $\boldsymbol{\Lambda}_1$ describe the added component and are of direct scientific interest. The parameter τ_0 is intertwined with these parameters in that the image's total count constrains its total intensity, $\tau_0 + \tau_1$, and thus the fitted τ_0 will decrease as the fitted τ_1 increases. We denote the unknown parameters $\boldsymbol{\theta} = (\boldsymbol{\theta}_0, \boldsymbol{\theta}_1)$, where $\boldsymbol{\theta}_0 = (\tau_0, \boldsymbol{\nu})$ are the parameters of the baseline component and $\boldsymbol{\theta}_1 = (\tau_1, \boldsymbol{\Lambda}_1)$ are the parameters of the added component. Typically, $\boldsymbol{\nu}$ is a nuisance parameter, at least in the context of searching for added structure beyond the baseline component.

2.2. Bayesian Inference

We adopt a Bayesian framework to fit the image parameters, $\boldsymbol{\theta}$, given the observed photon counts, $\boldsymbol{y}_{\rm obs}$. In particular, we quantify our state of knowledge before having seen the data using a *prior distribution* and that after having seen the data using a *posterior distribution*. Bayes' Theorem allows us to transform the prior distribution into the posterior distribution by conditioning on the observed counts. In particular, the theorem states that the posterior distribution of $\boldsymbol{\theta}$ given $\boldsymbol{y}_{\rm obs}$ is

$$\pi\left(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs}\right) = \frac{\mathcal{L}\left(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}\right)\pi\left(\boldsymbol{\theta}\right)}{\pi\left(\boldsymbol{y}_{\rm obs}\right)}, \tag{2}$$

where $\pi(\boldsymbol{\theta})$ is the joint prior distribution of $\boldsymbol{\theta}$, $\mathcal{L}(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta})$ is the likelihood function of $\boldsymbol{y}_{\rm obs}$ given $\boldsymbol{\theta}$, and $\pi(\boldsymbol{y}_{\rm obs}) = \int \mathcal{L}(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}$ is the normalizing constant that ensures that $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$ integrates to one.

While prior distributions can be used to incorporate external information about the likely values of model parameters, they can also be used to enforce relationships among parameters. We use the added component, for example, to represent structure in an image that does not appear in the baseline component. If we did not impose any constraint on $\boldsymbol{\Lambda}_1$, random fluctuation from the baseline component and unstructured background would be indistinguishable from genuine image structures that are missing in the baseline. As detailed in Section 2.3, we use the prior distribution of $\boldsymbol{\Lambda}_1$ to specify a multiscale smooth structure that characterizes the added components that can be identified by our procedure.

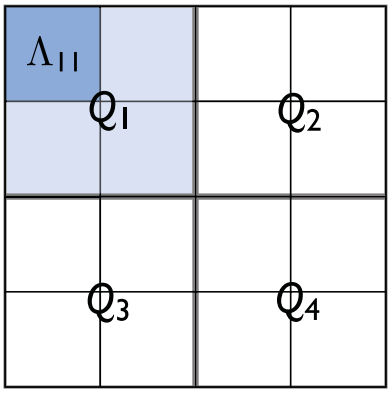


Figure 2. A schematic representation of the multiscale decomposition for an image consisting of a 4 × 4 grid of pixels. For instance, $\Lambda_{11} = \phi_{11}\phi_{211}$, where ϕ_{11} is the proportion expected in the first quadrant of the expected total count across the entire image, i.e., $\phi_{11} = \sum_{j \in Q_1} \Lambda_{1j}$, and ϕ_{211} is the proportion of the expected counts in the first quadrant expected in its first pixel, i.e., $\phi_{211} = \Lambda_{11} / \sum_{j \in Q_1} \Lambda_{1j}$.

Assuming the prior distributions for $(\tau_0, \boldsymbol{\theta}_1)$ and $\boldsymbol{\nu}$ are independent, we can write the posterior distribution as

$$\pi\left(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs}\right) \propto \mathcal{L}\left(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}\right)\pi\left(\tau_0, \boldsymbol{\theta}_1\right)\pi\left(\boldsymbol{\nu}\right), \tag{3}$$

where we have omitted the denominator of Equation (2) because it is a constant determined by the numerator. Since the likelihood function is specified by Equation (1), we need only set $\pi(\tau_0, \boldsymbol{\theta}_1)$ and $\pi(\boldsymbol{\nu})$. Insofar as the baseline model is well specified, the uncertainty in $\boldsymbol{\nu}$ and hence the sensitivity of the final result to the choice of $\pi(\boldsymbol{\nu})$ are both limited. For example, when testing for an X-ray jet in an image of a quasar, $\boldsymbol{\nu}$ could consist of the location and amplitude of the quasar point source and the intensity of a constant background. These parameters are well constrained by the data and thus relatively insensitive to the choice of prior $\pi(\boldsymbol{\nu})$. In practice, we typically use the default settings in CIAO's Sherpa software (Freeman et al. 2001) and fit $\boldsymbol{\nu}$ via maximum likelihood (equivalent to the posterior mode under a uniform $\pi(\boldsymbol{\nu})$); see Section 2.4. We do not further discuss the choice of $\pi(\boldsymbol{\nu})$ in this paper, but the choice of $\pi(\tau_0, \boldsymbol{\theta}_1)$ is central to the general problem and is the topic of Section 2.3.

2.3. The Prior Distributions

The relative intensity of the added component, $\boldsymbol{\Lambda}_1$, is unknown and must be estimated from the data. Following Esch et al. (2004) and Connors & van Dyk (2007, pp. 101–117), we place a multiscale smoothing prior distribution on $\boldsymbol{\Lambda}_1$; see also Nowak & Kolaczyk (2000). This prior distribution provides a flexible class of models while ensuring stability in the fit. We illustrate the structure of the prior distribution by considering an image composed of a simple 4 × 4 grid of pixels, as in

Figure 2. First, we reparameterize $\boldsymbol{\Lambda}_1$ using the decomposition

$$\Lambda_{1i} = \left(\sum_{j \in Q_{k(i)}} \Lambda_{1j} \right) \left(\frac{\Lambda_{1i}}{\sum_{j \in Q_{k(i)}} \Lambda_{1j}} \right), \quad (4)$$

where $k(i)$ indicates which quadrant of the image contains pixel i, and $Q_{k(i)}$ is the collection of pixels in that quadrant. This formulation allows us to hierarchically specify the multiscale smoothing prior distribution. At the first level of the hierarchy we model the proportion of the expected total count in each of the four quadrants of the added component. Let $\boldsymbol{\phi}_1 = (\phi_{11}, \ldots, \phi_{14})$ represent these proportions, i.e.,

$$\phi_{1k} = \sum_{j \in Q_k} \Lambda_{1j}, \quad k = 1, \ldots, 4. \quad (5)$$

The first subscript of ϕ represents the level of the hierarchy, here a one, and the second represents the quadrant number; see Figure 2.

We formulate the prior distribution to encourage fitted values of the expected quadrant proportions $\boldsymbol{\phi}_1$ that are similar to each other, thus encouraging smoothing in the added component. Mathematically, this is accomplished using a Dirichlet prior distribution,[5]

$$\boldsymbol{\phi}_1 \sim \text{Dirichlet}\left\{ (\psi_1, \psi_1, \psi_1, \psi_1) \right\}.$$

Under this distribution, the larger ψ_1 is, the smoother the reconstruction is at this level of the hierarchy/resolution in the image; for this reason ψ_1 is called a *smoothing parameter*. Similarly, at the second level of the hierarchy/resolution, we model the expected pixel counts within each quadrant as a proportion of the total expected quadrant count, i.e., we model

$$\phi_{2k\ell} = \frac{\Lambda_{1i}}{\phi_{1k}} = \frac{\Lambda_{1i}}{\sum_{j \in Q_k} \Lambda_{1j}}, \quad i = \ell\text{th pixel in } k\text{th quadrant}; \quad (6)$$

see Figure 2. Again, we use a Dirichlet distribution: $\boldsymbol{\phi}_{2k} \sim \text{Dirichlet}\{(\psi_2, \psi_2, \psi_2, \psi_2)\}$, $k = 1, \ldots, 4$, where $\boldsymbol{\phi}_{2k} = (\phi_{2k1}, \ldots, \phi_{2k4})$ with subscripts representing the level of resolution, the quadrant, and the pixel within quadrant. We may use a different smoothing parameter in this level of the hierarchy than in the first level (i.e., ψ_1 may differ from ψ_2) to allow for different structures at different image resolutions.

For larger images, we can continue this hierarchy using different smoothing parameters for the Dirichlet distribution at different levels of resolution. In this way, we might expect little smoothing at the lowest level of resolution and more smoothing at higher levels. A small value of ψ_1 minimizes smoothing across the four quadrants of the image, while larger values of ψ_k for $k > 1$ encourage more smoothing at level k of the hierarchy. Esch et al. (2004) suggests using cycle spinning to prevent visual artifacts that arise from a fixed multiscale decomposition. Cycle spinning consists of randomly translating the origin of the multiscale grid while iteratively updating parameter estimates; for details, see Esch et al. (2004).

[5] In our representation, a (four-dimensional) symmetric Dirichlet distribution with parameter ψ has probability density function $\text{pdf}(x_1, x_2, x_3, x_4) = \frac{\Gamma(4\psi)}{\Gamma(\psi)^4} \prod_{i=1}^4 x_i^{\psi-1}$; the mean of x_i is 1/4; the standard deviation of x_i is $\sqrt{3/\{16(4\psi+1)\}}$; and the correlation of x_i and x_j is $-1/3$.

Esch et al. (2004) apply this hierarchical prior distribution to derive fitted Bayesian X-ray images in the absence of a baseline model. They include a hyperprior distribution to fit the smoothing parameters $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_D)$, where D is the number of scales in the multiscale decomposition. This strategy alleviates the need to specify the values of the smoothing parameters. We follow their recommendation and use the hyperprior distribution $\pi(\boldsymbol{\psi}) \propto \prod_{i=1}^D \exp(-1000\psi_i^3)$, which encourages small values of ψ_i, and so imposes less smoothing, but is not so heavily concentrated near zero as to cause numerical problems. Using this specification of the added component, we confine attention to images that are cropped to $2^D \times 2^D$ pixels for some integer D.

Because of their different roles in our model, we place different prior distributions on τ_1 and τ_0. First, τ_1 specifies the total expected count from the added component, and its prior distribution must be flexible enough to allow for values near zero if the baseline model is adequate and for large values if the baseline model is not adequate. We accomplish this using a Gamma distribution[6] with mean and standard deviation equal to 20. This distribution exhibits significant skewness, with substantial probability near zero and appreciable probability extending to large values.

Second, τ_0 specifies the total expected count from the baseline component. In practice, the observed image typically provides plenty of information to constrain τ_0, since we usually observe at least 100 counts across the entire image and the baseline component is a reasonable description of at least some major image features. Thus, we use a relatively diffuse prior distribution, specifically, the improper distribution, $\pi(\tau_0) \propto \tau_0^{0.001-1}$ (it can be shown that under very mild conditions, the posterior distribution will be proper when $\pi(\tau_0) \propto \tau^{\epsilon-1}$ for any $\epsilon > 0$).

2.4. The Null and Alternative Hypotheses

We are interested in comparing two models for the image. The first corresponds to the hypothesis that the baseline component fully represents the image and no added structure is needed. The second hypothesis stipulates that the baseline component is insufficient and there is significant structure in the image that can be represented by the added component. We refer to these two hypotheses as the *null hypothesis* and the *alternative hypothesis*, respectively. (Up until now we have referred to these two models as the baseline and full models, respectively. From here on we will employ the more formal terminology, i.e., null and alternative hypothesis/model.) Statistically our goal is to quantify the evidence in the image for deciding between these hypotheses. This choice can be formalized using the notation of Section 2.1: the alternative hypothesis is specified in Equation (1) and the null hypothesis arises as the special case where $\tau_1 = 0$. Thus, our hypotheses are

$$H_0\text{: } \tau_1 = 0 \quad (7)$$

$$H_A\text{: } \tau_1 \sim \pi(\tau_1), \quad (8)$$

where $\pi(\tau_1)$ is the prior distribution for τ_1 (under the alternative hypothesis).

[6] A Gamma distribution with shape parameter a and rate parameter b has probability density function $\text{pdf}(x) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}$, mean a/b, and standard deviation $\sqrt{a}/b$.

In the example of Section 1.1 and Figure 1(b), under the null hypothesis, the image is assumed to have an underlying intensity that consists of the baseline component of a flat background and a single point source representing a quasar with no jet. An image of the baseline component (not shown) would look like Figure 1(a) without the two fainter point sources. Under the alternative hypothesis, the assumed underlying intensity is a weighted sum of the quasar-only baseline component and the added multiscale component that allows for additional structure beyond the quasar point source.

The likelihood function and the posterior distribution under the alternative hypothesis, i.e., H_A in Equation (8), are described in Equations (2)–(3). Similarly, we let

$$\mathcal{L}_0(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}_0) = \mathcal{L}(\boldsymbol{y}_{\rm obs}|\tau_1 = 0, \boldsymbol{\theta}_0)$$

and

$$\begin{aligned}\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs}) &\propto \mathcal{L}_0(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}_0)\pi(\boldsymbol{\theta}_0)\\ &= \mathcal{L}_0(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}_0)\pi(\tau_0)\pi(\boldsymbol{\nu})\end{aligned}$$

denote the likelihood function and posterior distribution, respectively, under the null hypothesis, i.e., H_0 in Equation (7). (Note that if $\tau_1 = 0$, then the likelihood function does not depend on $\boldsymbol{\Lambda}_1$ and we do not attempt to estimate it.)

When fitting the alternative model, sometimes we fix the nuisance parameters of the baseline component, $\boldsymbol{\nu} = \hat{\boldsymbol{\nu}}$, perhaps estimating them in a preliminary analysis. In this case, we work with the conditional posterior distribution of $(\tau_0, \boldsymbol{\theta}_1)$ given $\hat{\boldsymbol{\nu}}$ rather than the full posterior distribution $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$. We denote this conditional posterior distribution $\pi(\tau_0, \boldsymbol{\theta}_1|\boldsymbol{y}_{\rm obs}, \hat{\boldsymbol{\nu}})$. When there is little posterior uncertainty in $\boldsymbol{\nu}$, this "plug-in posterior distribution" approximates the marginal posterior distribution, $\pi(\tau_0, \boldsymbol{\theta}_1|\boldsymbol{y}_{\rm obs}) = \int \pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})d\boldsymbol{\nu}$. This is the approach taken in K. McKeough et al. (2015, in preparation), where the location and amplitude of each quasar as well the intensity of a uniform background are fit using Sherpa (Freeman et al. 2001) in separate preliminary analyses; these fitted values are then used to set the relative intensities, $\boldsymbol{\Lambda}_0(\hat{\boldsymbol{\nu}})$. Finally, we fit the alternative model in LIRA, conditioning on $\boldsymbol{\Lambda}_0(\hat{\boldsymbol{\nu}})$, but leaving the scale factor τ_0 as a free parameter.

3. TESTING FOR STRUCTURE

3.1. Statistical Hypothesis Testing

Although we employ Bayesian methods for model fitting, we consider the classical *hypothesis testing* paradigm for model selection. The test is conducted using a test statistic, denote by $T(\boldsymbol{y}_{\rm obs})$, which is chosen so that larger values of $T(\boldsymbol{y}_{\rm obs})$ are indicative of an added component in the image. In particular, larger values are less likely to have been obtained as a random fluctuation under the null hypothesis. Thus, if $T(\boldsymbol{y}_{\rm obs})$ is large enough we decide there is sufficient evidence to conclude that the null hypothesis is inappropriate and there is added structure in the image beyond the baseline component. In this framework, we must determine a threshold for $T(\boldsymbol{y}_{\rm obs})$ such that values of $T(\boldsymbol{y}_{\rm obs})$ greater than the threshold are sufficient evidence to declare detection of structure beyond the baseline component. In the hypothesis testing framework, this is done by limiting the probability of a false detection. Thus, the detection threshold, $T^\star$, is the smallest value such that

$$\Pr\left(T(\boldsymbol{y}_0) \geqslant T^\star|\tau_1 = 0, \boldsymbol{\theta}_0\right) \leqslant \alpha, \tag{9}$$

where $\boldsymbol{y}_0$ is a random replicate image generated under the null hypothesis and α is the maximum allowed probability of a false detection.

Conversely, we can compute the probability under the null hypothesis of observing an image as extreme or more extreme than the observed image, as quantified by the test statistic, i.e.,

$$p = \Pr\left(T(\boldsymbol{y}_0) \geqslant T(\boldsymbol{y}_{\rm obs})|\tau_1 = 0, \boldsymbol{\theta}_0\right). \tag{10}$$

This is called a *p-value* and small values, e.g., less than 0.05 or 0.01, are taken as evidence that the image was not generated under the null hypothesis and thus are generally interpreted as evidence in favor of structure in the image beyond the baseline component. Although very popular in practice, *p*-values are criticized on theoretical grounds from both frequentist and Bayesian perspectives (e.g., Berger & Delampady 1987; Zhao 2014).

3.2. The Test Statistic

Before we can compute the *p*-value in Equation (10), we need to choose a test statistic $T(\boldsymbol{y}_{\rm obs})$. For a test statistic to be useful, it should provide discrimination between the null and alternative hypotheses. To motivate our choice of test statistic, consider the parameter

$$\xi = \tau_1/(\tau_1 + \tau_0), \tag{11}$$

the proportion of the total image intensity that is due to the added component. If the baseline component fits the data poorly, we expect more of the observed counts to be attributed to the added component, corresponding to large values ξ. On the other hand, if the data are generated under the null hypothesis with $\tau_1 = 0$, we expect more of the observed counts to be attributed to the baseline component, corresponding to ξ near zero. (Formally, under the null hypothesis, $\xi = 0$. Nonetheless, its fitted value under the alternative hypothesis will typically be small but positive even if the data are generated under the null hypothesis.) Thus, ξ is a good candidate for discriminating between the null and alternative hypotheses.

Unfortunately, ξ is a parameter, not a statistic; that is, it cannot be computed directly as a function of the data $\boldsymbol{y}_{\rm obs}$. However, the posterior distribution of ξ under the alternative hypothesis, conditional on the data, can be computed from the data. This motivates us to use a feature of the posterior distribution of ξ as a test statistic. In particular, our test statistic is a posterior tail probability of ξ. Given a threshold c, we let

$$T_c(\boldsymbol{y}_{\rm obs}) = \Pr\left(\xi \geqslant c|\boldsymbol{y}_{\rm obs}\right), \tag{12}$$

where the probability is taken with respect to $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$, the posterior distribution under the alternative hypothesis. To some readers, it may seem more natural to use the fitted value of ξ as a test statistic, but as we discuss in Section 3.4, there are advantages to using the tail probability, $T_c(\boldsymbol{y}_{\rm obs})$. This choice allows us to treat c as a tuning parameter and thereby to select a more powerful test statistic than the fitted value of ξ. Although $T_c(\boldsymbol{y}_{\rm obs})$ involves computations under the posterior distribution, it is a true statistic in that (for a fixed prior distribution and a given c) it is a function only of the data, and not of any unknown parameters. Computing $T_c(\boldsymbol{y}_{\rm obs})$ with respect to the

plug-in posterior distribution $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs}, \hat{\boldsymbol{\nu}})$ also leads to a valid test statistic in the same sense.[7]

When the alternative hypothesis is true, we expect large values of ξ to explain the data better than small values, and we therefore expect a high posterior probability that ξ exceeds an appropriate value of c; that is, we expect a large value of $T_c(\boldsymbol{y}_{\rm obs})$. Conversely, when the null hypothesis is true, $T_c(\boldsymbol{y}_{\rm obs})$ is typically small because small values of ξ tend to better explain the data. Thus, $T_c(\boldsymbol{y}_{\rm obs})$ behaves differently under the null and alternative hypotheses, making it a reasonable candidate for a test statistic to discriminate between the two hypotheses.

It can be useful to substitute other choices of $\xi = h(\boldsymbol{\theta})$ in Equation (12) to define different test statistics $T_c(\boldsymbol{y}_{\rm obs})$. For instance, often we are interested detecting departures from the baseline model in a specific, known region of the image; see K. McKeough et al. (2015, in preparation) for numerous examples. In this case, it is possible to design a test statistic targeted at the given region. Let R be a collection of pixel indices defining a region of interest on the image. To test whether there is a significant departure from the null hypothesis in the region R, we can use as a test statistic a posterior tail probability of the parameter

$$\xi_R = \frac{\sum_{j\in R}\tau_1\Lambda_{1j}}{\sum_{j\in R}\left(\tau_1\Lambda_{1j} + \tau_0\Lambda_{0j}\right)}, \tag{13}$$

the fraction of the total intensity in R attributed to the added component. In particular, our test statistic is

$$T_{R,c}(\boldsymbol{y}_{\rm obs}) = \Pr\left(\xi_R \geqslant c|\boldsymbol{y}_{\rm obs}\right). \tag{14}$$

Because it only considers pixels in R, $T_{R,c}(\boldsymbol{y}_{\rm obs})$ only has power to detect departures from the null hypothesis that manifest in the region of interest. However, by ignoring regions of the image where little or no departure from the null hypothesis is expected, $T_{R,c}(\boldsymbol{y}_{\rm obs})$ may be more powerful than the image-wide $T_c(\boldsymbol{y}_{\rm obs})$ for detecting departures concentrated in R.

3.3. A Fully Specified Null Hypothesis

Because the probabilities in Equations (9)–(10) depend on the unknown parameters $\boldsymbol{\theta}_0$, the probability of a false detection and hence $T^\star$ and the p-value cannot be computed. This is a nuisance and why parameters that are unknown under the null hypothesis, like τ_0 and $\boldsymbol{\nu}$, are called *nuisance parameters*. We set aside this difficulty for the moment to focus on statistical issues that arise in the absence of nuisance parameters, but return to it in Section 3.5.

In particular, we start by assuming that $\boldsymbol{\theta}_0$ is known, the null hypothesis has no unknown parameters, and $\mathcal{L}_0(\boldsymbol{y}_{\rm obs}|\boldsymbol{\theta}_0) = \mathcal{L}_0(\boldsymbol{y}_{\rm obs})$. In practice, $\boldsymbol{\theta}_0$ is never known exactly, but in some situations we may be able to estimate it with high enough precision that we can treat it as fixed and known. However, great care must be taken when fixing $\boldsymbol{\theta}_0$. If it is fixed at an inappropriate value, evidence against the null hypothesis may not indicate that the null model is inappropriate so much as that the fixed values of $\boldsymbol{\theta}_0$ are inappropriate. For example, τ_0 should not be fixed at an arbitrary value. If $\tau_0\boldsymbol{\Lambda}_0(\boldsymbol{\nu})$ is fixed, then the expected total count under the null hypothesis is also fixed. If this expected total count differs significantly from the observed total count, the null hypothesis can be rejected on this basis alone. Generally speaking, if τ_0 is to be fixed, it should be set equal to a reasonable estimate of the total expected count after adjusting for detector inefficiencies, for example to

$$\hat{\tau}_0 = \sum_{i=1}^{n} y_i \Big/ \sum_{i=1}^{n}\sum_{j=1}^{n} P_{ij}A_j\Lambda_{0j}(\hat{\boldsymbol{\nu}}), \tag{15}$$

and this should only be done if the total count is large. For example, when testing for a quasar jet, we fit the location and amplitude of the point source and the intensity of a constant background using Sherpa (Freeman et al. 2001). In Section 4 and K. McKeough et al. (2015, in preparation), we use the fitted values of these parameters to fix $\tau_0\boldsymbol{\Lambda}_0(\boldsymbol{\nu}) = \hat{\tau}_0\boldsymbol{\Lambda}_0(\hat{\boldsymbol{\nu}})$ in the null model. When fitting the alternative model, however, we recommend never fixing τ_0, because this would leave no flexibility to reduce the emission attributed to the baseline component and increase the emission attributed to the added multiscale component. Instead, as described at the end of Section 2.4, under the alternative model, we fix $\boldsymbol{\nu} = \hat{\boldsymbol{\nu}}$ according to the fitted values from Sherpa, but allow τ_0 to be estimated.

With ξ defined in Equation (11), larger values of $T_c(\boldsymbol{y})$ are more unusual under the null hypothesis. The p-value in Equation (10) simplifies to

$$p = \Pr\left(T_c(\boldsymbol{y}_0) \geqslant T_c(\boldsymbol{y}_{\rm obs})|\tau_1 = 0\right), \tag{16}$$

where $\boldsymbol{y}_0$ is a randomly generated image under the null hypothesis and $\boldsymbol{y}_{\rm obs}$ is the fixed observed image. Because there are no nuisance parameters, this p-value can in principle be computed exactly.

3.4. Computing the Statistical Significance

The primary advantage of using a tail probability of $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$ rather than a fitted value of ξ or the likelihood ratio test as the test statistic is computational. We describe this advantage here.

Although we cannot directly evaluate the tail probability $T_c(\boldsymbol{y}_{\rm obs}) = \Pr(\xi \geqslant c|\boldsymbol{y}_{\rm obs})$ under the alternative model, we can estimate it numerically via MCMC. Even if we could compute $T_c(\boldsymbol{y}_{\rm obs})$ directly, we would need to compute the probability in Equation (10) to evaluate the p-value, and this too is most easily obtained through Monte Carlo simulation.

MCMC involves obtaining L correlated draws $\xi^{(1)}_{\rm obs}, \ldots, \xi^{(L)}_{\rm obs}$ from the posterior distribution under the alternative model, $\pi(\xi|\boldsymbol{y}_{\rm obs})$. This can be accomplished using the LIRA package, which relies on the Gibbs sampling algorithm described in Esch et al. (2004). Specifically, LIRA delivers a correlated sample $\boldsymbol{\theta}^{(1)}_{\rm obs}, \ldots, \boldsymbol{\theta}^{(L)}_{\rm obs}$ from the full posterior distribution $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$, and we then compute each $\xi^{(\ell)}_{\rm obs}$ as a function of each $\boldsymbol{\theta}^{(\ell)}_{\rm obs}$, say $\xi^{(\ell)}_{\rm obs} = h(\boldsymbol{\theta}^{(\ell)}_{\rm obs})$. With the MCMC sample in hand we estimate $T_c(\boldsymbol{y}_{\rm obs})$ as

$$\widehat{T}_c(\boldsymbol{y}_{\rm obs}) = \frac{1}{L}\sum_{\ell=1}^{L} 1\left\{\xi^{(\ell)}_{\rm obs} \geqslant c\right\}, \tag{17}$$

where $1\{\cdot\}$ is the indicator function that equals one if its argument is true and is zero otherwise.

[7] Indeed, the resulting test statistic is valid whether or not $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs}, \hat{\boldsymbol{\nu}})$ is a good approximation of $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs})$, because the posterior distribution $\pi(\boldsymbol{\theta}|\boldsymbol{y}_{\rm obs}, \hat{\boldsymbol{\nu}})$ can in principle be computed as a function of the data. See Section 4.2.2 for discussion of the trade-offs involved when using computational approximations of test statistics.

A straightforward way to estimate the p-value is to simulate M independent replicate images under the null hypothesis and then fit the alternative model and compute the test statistic for each. This can be accomplished via the following method.

Direct P-value Method:

For $j = 1, \ldots, M$,

1. Simulate $\boldsymbol{y}_0^{(j)} \sim \mathcal{L}_0(\boldsymbol{y}_0)$;
2. Fit the alternative model to $\boldsymbol{y}_0^{(j)}$ by running LIRA to obtain L correlated draws $\xi^{(j,1)}, \ldots, \xi^{(j,L)}$ from $\pi(\xi|\boldsymbol{y}_0^{(j)})$; and
3. Compute the estimated test statistic $\widehat{T}_c(\boldsymbol{y}_0^{(j)})$ using Equation (17) with $\xi_{\text{obs}}^{(\ell)}$ replaced with $\xi^{(j,\ell)}$.

Finally, estimate the p-value with the Monte Carlo p-value,

$$\hat{p} = \frac{1 + \sum_{j=1}^{M} 1\left\{\widehat{T}_c\left(\boldsymbol{y}_0^{(j)}\right) \geqslant \widehat{T}_c(\boldsymbol{y}_{\text{obs}})\right\}}{1 + M}, \tag{18}$$

recommended in Davison & Hinkley (1997). Equation (18) adds one to the numerator and denominator of the naive Monte Carlo p-value,

$$\hat{p}_{\text{naive}} = \frac{\sum_{j=1}^{M} 1\left\{\widehat{T}_c\left(\boldsymbol{y}_0^{(j)}\right) \geqslant \widehat{T}_c(\boldsymbol{y}_{\text{obs}})\right\}}{M}. \tag{19}$$

We use $\hat{p}$ instead of $\hat{p}_{\text{naive}}$ to guarantee that a testing procedure that rejects the null hypothesis when the p-value is less than or equal to a pre-specified α has false positive rate no greater than α. The naive Monte Carlo p-value $\hat{p}_{\text{naive}}$ does not control the false positive rate in this manner.[8] For further discussion and a numerical demonstration, see Section 4.3.

Fitting our Bayesian imaging model via MCMC using LIRA is computationally expensive. Unfortunately, the Direct P-value Method requires us to run LIRA M times and M must be very large to have any chance of achieving a high level of statistical significance (i.e., a low value of $\hat{p}$) because $\hat{p} \geqslant 1/(1 + M)$. This requires devoting M times the computational resources to analyzing simulated replicate images as analyzing the observed image. This is very computationally expensive and often unacceptable in applied work. In practice, often the reconstructed image under the alternative model is of primary interest and the statistical p-value is intended as an additional check to prevent over-interpreting apparent structure in the noise as discovery of real structure in the signal. Although the p-value is often of secondary interest, preventing over-interpretation of images is of course critically important. Nonetheless, devoting 1000 times (or more) the computing time to computing a p-value is often infeasible.

To reduce the computational requirements of this significance test, we propose estimating not the p-value but an upper bound on it. In the Appendix we show that

$$p \leqslant \frac{\gamma}{T_c(\boldsymbol{y}_{\text{obs}})} = u, \tag{20}$$

[8] Under the null hypothesis, $\Pr(\hat{p}_{\text{naive}} = i/M) = 1/(M+1)$, for $i = 0, 1, \ldots, M$. Thus, the true false positive rate using the naive Monte Carlo p-value is $\Pr(\hat{p}_{\text{naive}} \leqslant \alpha) = (\lfloor M\alpha \rfloor + 1)/(M+1)$, which is greater than α for some choices of α (such as $\alpha = j/M$, if $j \in \{0, 1, \ldots, M-1\}$).

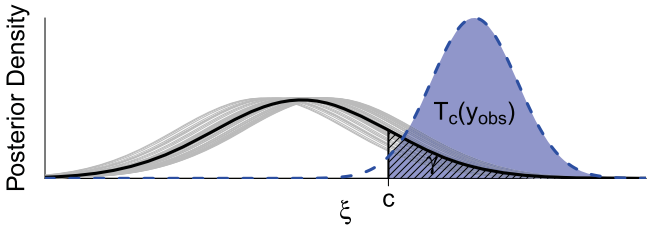


Figure 3. A schematic illustration of the quantities used to compute the upper bound u in Equation (20). The gray solid lines represent posterior densities $\pi(\xi|\boldsymbol{y}_0^{(j)})$ given a sample of images $\boldsymbol{y}_0^{(j)}$ simulated under the null; the black solid line is $g(\xi)$, computed as the average of the posterior densities shown in gray; and the blue dashed line is the posterior density $\pi(\xi|\boldsymbol{y}_{\text{obs}})$ given an observed image $\boldsymbol{y}_{\text{obs}}$. To compute $\hat{u}$, we fix γ (the hatched area under $g(\xi)$ to the right of c); compute c as the $(1-\gamma)$ quantile of $g(\xi)$, i.e., $\int_c^1 g(\xi)d\xi = \gamma$, from an MCMC sample from $g(\xi)$; and use this value of c to compute $T_c(\boldsymbol{y}_{\text{obs}})$ (the area under $\pi(\xi|\boldsymbol{y}_{\text{obs}})$ to the right of c, shaded in blue) from an MCMC sample from $\pi(\xi|\boldsymbol{y}_{\text{obs}})$.

where $\gamma = \Pr(\xi \geqslant c)$ under the distribution

$$g(\xi) = E\left\{\pi(\xi|\boldsymbol{y}_0)\right\} = \sum_{\boldsymbol{y}_0} \pi(\xi|\boldsymbol{y}_0)\,\mathcal{L}_0(\boldsymbol{y}_0), \tag{21}$$

the expectation under the null hypothesis of the posterior distribution under the alternative hypothesis. To compute the upper bound, u, on the p-value for a given fixed value of γ, the denominator of Equation (20) must be estimated, and doing so involves two sources of uncertainty: (i) estimating the quantile c of $g(\xi)$ and (ii) estimating the tail probability under the alternative hypothesis given $\boldsymbol{y}_{\text{obs}}$; see Figure 3. The upper bound can be computed using the following method.

Upper Bound Method:

For $j = 1, \ldots, M$,

1. Simulate $\boldsymbol{y}_0^{(j)} \sim \mathcal{L}_0(\boldsymbol{y})$;
2. Fit the alternative model to $\boldsymbol{y}_0^{(j)}$ by running LIRA to obtain L correlated draws $\xi^{(j,1)}, \ldots, \xi^{(j,L)}$ from $\pi(\xi|\boldsymbol{y}_0^{(j)})$;

Then, set $\hat{c}$ equal to the estimated $(1-\gamma)$ quantile of $g(\xi)$, using the LM posterior draws $\xi^{(j,\ell)}$, $j = 1, \ldots, M$; $\ell = 1, \ldots, L$. For instance, $\hat{c}$ may be set equal to the $(LM\gamma)$ th largest value among all of the $\xi^{(j,\ell)}$. Finally, compute the estimated test statistic $\widehat{T}_{\hat{c}}(\boldsymbol{y}_{\text{obs}})$ using Equation (17) with c replaced with the estimate $\hat{c}$. Our estimate of the upper bound is

$$\hat{u} = \gamma / \widehat{T}_{\hat{c}}(\boldsymbol{y}_{\text{obs}}). \tag{22}$$

Equation (22) reveals one of the advantages of using a posterior tail probability as a test statistic. If we were to replace Equation (11) with another choice of (an always non-negative) ξ, we could use its fitted value as a test statistic and derive an upper bound on the appropriate p-value as in the Appendix, but this upper bound could be quite large and there would be no remedy. However, using the tail probability $T_c(\boldsymbol{y})$ as a test statistic allows us flexibility in the choice of the tuning parameter c. We use this flexibility to fix the numerator in Equation (22) to a reasonable small value (for more on the choice of γ, see Section 4.2), enabling the possibility of obtaining a small upper bound. Whereas $\hat{p} \geqslant 1/(1+M)$, it is possible to obtain a $\hat{u}$ much less than $1/(1+M)$ if γ is chosen appropriately. This is the primary advantage of the Upper Bound Method. It allows us to establish statistical significance

with fewer simulated replicate images and thus can be appreciably faster in practice.

3.5. Null Hypothesis with Unknown Parameters

In this section, we consider settings in which the parameters of the null model have non-negligible uncertainties. When such uncertainty exists, fixing the null model by substituting estimates for these unknown parameters can lead to problems. As mentioned in Section 3.3, we might reject the null hypothesis because the parameters have been fixed at inappropriate values, not because the null model itself is incorrect.

When there is non-negligible uncertainty in the parameters of the null model, $\mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)$ depends on $\boldsymbol{\theta}_0$. We assume that under the null hypothesis, we can obtain the posterior distribution $\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})$ of these nuisance parameters under a prior $\pi_0(\boldsymbol{\theta}_0)$. With this approach, the p-value in Equation (16) can be calculated for each fixed value of $\boldsymbol{\theta}_0$. We denote this

$$p(\boldsymbol{\theta}_0) = \Pr\big(T_c(\boldsymbol{y}_0) \geqslant T_c(\boldsymbol{y}_{\rm obs})|\boldsymbol{\theta}_0\big), \tag{23}$$

where the probability is taken over the sampling distribution $\mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)$ for a fixed value of the parameter $\boldsymbol{\theta}_0$. A Bayesian posterior predictive p-value (ppp-value) averages $p(\boldsymbol{\theta}_0)$ over the posterior for $\boldsymbol{\theta}_0$ under the null hypothesis and is given by

$$\begin{aligned}\text{ppp-value} &= E\left\{p(\boldsymbol{\theta}_0)|\boldsymbol{y}_{\rm obs}\right\} \\ &= \int p(\boldsymbol{\theta}_0)\,\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})\,d\boldsymbol{\theta}_0.\end{aligned} \tag{24}$$

See Rubin (1984), Meng (1994), and Gelman et al. (1996), among others, for discussions of the properties of ppp-values.[9]

We can compute an upper bound on the ppp-value using an expression similar to Equation (20)

$$\text{ppp-value} \leqslant \frac{\gamma}{T_c(\boldsymbol{y}_{\rm obs})} = u_{\rm ppp}, \tag{25}$$

where now $\gamma = \Pr(\xi \geqslant c|\boldsymbol{y}_{\rm obs})$ under the distribution[10]

$$\begin{aligned}g(\xi|\boldsymbol{y}_{\rm obs}) &= E\left\{\pi(\xi|\boldsymbol{y}_0)|\boldsymbol{y}_{\rm obs}\right\} \\ &= \sum_{\boldsymbol{y}_0}\pi(\xi|\boldsymbol{y}_0)\int \mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)\,\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})\,d\boldsymbol{\theta}_0.\end{aligned} \tag{26}$$

To implement the Direct P-value and Upper Bound Methods in this setting, we only need to modify each method's Step 1, replacing it with

1. Simulate $\boldsymbol{\theta}_0^{(j)} \sim \pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})$, followed by $\boldsymbol{y}_0^{(j)} \sim \mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0^{(j)})$.

This requires, of course, that we can obtain draws from $\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})$, the posterior distribution of the unknown parameters under the null hypothesis. In practice, we can approximate this posterior distribution using the estimated uncertainties from fitting the baseline component in Sherpa (Freeman et al. 2001). The implementation of Step 2, while notationally the same, is more complicated when there are unknown parameters under the null hypotheses because these parameters must be fit.

4. NUMERICAL RESULTS

In this section, we extend the example in Section 1.1 and investigate the performance of our proposed method on images, both simulated and real, of quasars with possible jets.

4.1. Three Simulated Images

We begin by analyzing the simulated image of Section 1.1 and compare the results to analyses of two other simulated images, one with a weaker jet and one with a stronger jet. Recall that in the image of Section 1.1 the jet was composed of two additional Gaussian components, each with 20 expected counts. We call this the "medium jet." The simulated images for the weak and strong jets were constructed identically to the medium jet image, except that each additional Gaussian component had 10 expected counts in the weak jet and 35 in the strong jet.

Each ground truth image was convolved with the same PSF to obtain the Poisson intensity in each pixel, with which we sampled to generate the simulated observations. The PSF (Figure 1(c)) was generated with SAOTrace[11] to reflect the observation conditions of the quasar analyzed in Section 4.4, which was observed by the *Chandra* X-ray Observatory.

To analyze each simulated image, we created a baseline component $\tau_0\boldsymbol{\Lambda}_0$ for each jet strength that consisted of the simulated quasar with the correct location and intensity (analogous to a real data analysis in which these parameters are well constrained by the data and can be treated as known) and a uniform background. The expected count due to background tabulated in $\tau_0\boldsymbol{\Lambda}_0$ was set to the sum of the actual background and the jet components, so that the total expected count in the null model was equal to the total in the true image for each jet strength. This prevents rejection of the null hypothesis purely on the basis of the total count. In this simulation the null hypothesis was fully specified with no unknown parameters. The posterior means of $\tau_1\Lambda_{1j}$ for the medium jet are shown in Figure 1(d), in which the two sources not included in the baseline component are clearly visible.

We also simulated $M = 50$ images under the null hypothesis for each jet strength. For each simulated and replicate image, we fit the alternative hypothesis via MCMC using the LIRA package, obtaining 1800 draws from the posterior distribution after discarding the initial 200 steps as burn-in. The posterior distributions of ξ for the medium jet are shown in the middle panel of Figure 4. As expected, $\pi(\xi|\boldsymbol{y}_{\rm obs})$ tends to be to the right of the $\pi(\xi|\boldsymbol{y}_0^{(j)})$, because a higher fraction of observed counts are attributed to the added component $\tau_1\boldsymbol{\Lambda}_1$ with $\boldsymbol{y}_{\rm obs}$ than with most of the $\boldsymbol{y}_0^{(j)}$.

Figure 5 displays the estimated p-value upper bounds, $\hat{u}$, for the simulated images with weak, medium, and strong jets, as a function of the upper tail probability, γ, of $g(\xi)$. For the weak jet, $\hat{u}$ is consistent with the null hypothesis. However, for the medium and strong jets, if γ is in an appropriate interval, $\hat{u}$ appears to reveal significant evidence of inadequacy of the null

[9] Meng (1994) investigates the frequency properties of ppp-values under the prior predictive distribution $\int \mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)\,\pi_0(\boldsymbol{\theta}_0)\,d\boldsymbol{\theta}_0$. Under this distribution, the ppp-value is more concentrated around 0.5 than a uniform distribution, and a test that rejects the null hypothesis when the ppp-value $\leqslant \alpha$ will have false positive rate no greater than 2α. Similarly, Sinharay & Stern (2003) and Bayarri & Castellanos (2007) discuss the conservativeness of ppp-values in hierarchical models.

[10] The distribution $\int \mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)\,\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})\,d\boldsymbol{\theta}_0$ is the *posterior predictive* distribution of $\boldsymbol{y}_0$ given $\boldsymbol{y}_{\rm obs}$ under the null hypothesis, so $g(\xi|\boldsymbol{y}_{\rm obs})$ is the posterior predictive expectation under the null hypothesis of the posterior distribution of ξ under the alternative hypothesis.

[11] http://cxcoptics.cfa.harvard.edu/SAOTrace/Index.html

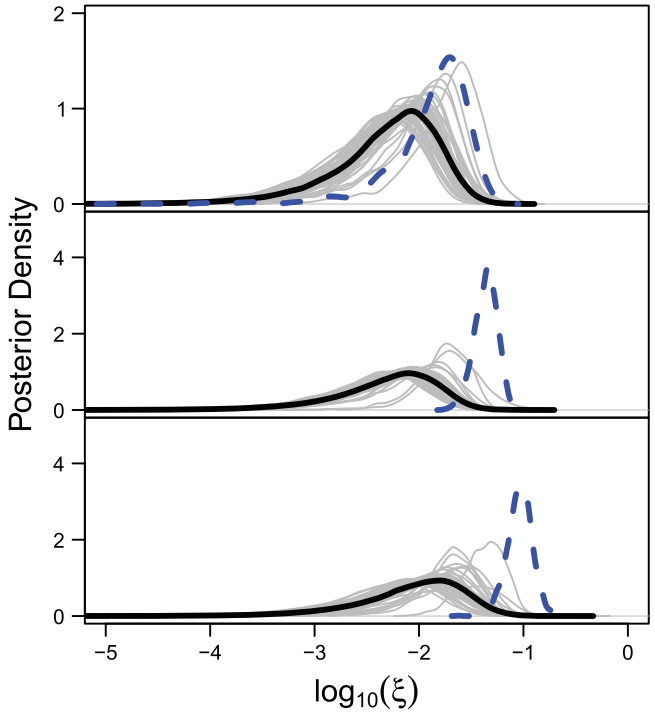


Figure 4. Estimated posterior distributions for the simulated quasar with a weak jet (top panel), simulated quasar with a medium-strength jet (middle panel), and observed quasar discussed in Section 4.4 (bottom panel). The dashed blue lines are $\pi(\xi|\boldsymbol{y}_{\rm obs})$, the posterior under the alternative hypothesis given the observed data; the light gray solid lines are $\pi(\xi|\boldsymbol{y}_0^{(j)})$, the posterior under the alternative hypothesis given images $\boldsymbol{y}_0^{(j)}$ simulated under the null hypothesis; and the heavy black solid lines are $g(\xi)$, the expectation under the null hypothesis of the posterior distribution under the alternative hypothesis.

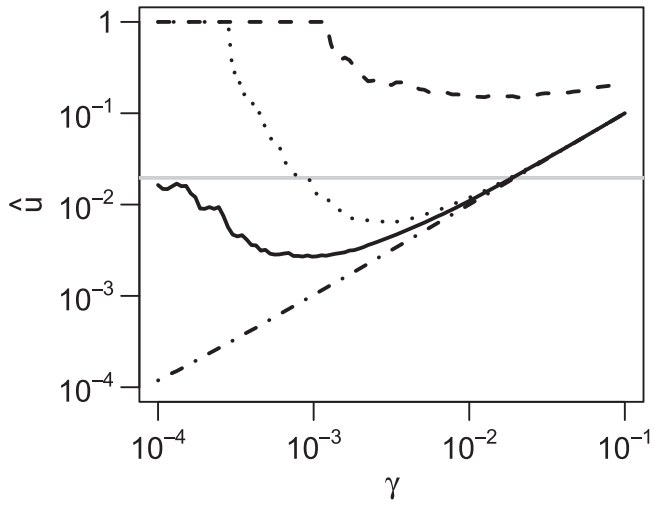


Figure 5. Estimated upper bounds $\hat{u}$ for varying values of the upper tail probability, γ, of $g(\xi)$, for the simulated jets of Section 4.1 (weak: dashed, medium: dotted, strong: dashed–dotted) and the observed quasar (solid black line) of Section 4.4. The solid gray line is at $1/51$, the minimum achievable $\hat{p}$ as defined in Equation (18) when $M = 50$.

hypothesis. We discuss bootstrap estimation of the uncertainty in $\hat{u}$ in Section 4.4.

Using the Direct P-value Method, the best (i.e., lowest) p-value that can be achieved with $M = 50$ is $\hat{p} = 1/51$. If we were to use $\hat{p}$ as defined in Equation (18), this minimum value would be achieved for the medium simulated jet when $c < 0.06$ and for the strong jet when $c < 0.14$. For the weak jet, the smallest $\hat{p}$ that would be achieved is $2/51$, when $c < .04$. The Upper Bound Method can lead to higher significance than the Direct P-value Method for the medium and strong jets, but not for the weak jet. The top panel of Figure 4 helps illustrate this, as $\pi(\xi|\boldsymbol{y}_{\rm obs})$ for the weak jet is centered slightly to the right of most of the $\pi(\xi|\boldsymbol{y}_0^{(j)})$, but the right tail of $\pi(\xi|\boldsymbol{y}_{\rm obs})$ does not extend much past the right tail of $g(\xi)$, preventing the upper bound from being very low.

4.2. Practical Implementation Issues

4.2.1. Choosing γ

As seen in Figure 5, the Upper Bound Method generates a family of upper bounds $\hat{u}(\gamma)$ corresponding to different choices of $\gamma \in [0, 1]$. It would not be valid to simply report the minimum of these upper bounds, $\min_{0\leqslant\gamma\leqslant 1}\hat{u}(\gamma)$, as the statistical significance. This would be analogous to computing multiple p-values and only reporting the most significant (i.e., smallest) one. If multiple tests are performed, each with a bounded probability of returning a false detection, the probability of *at least one* false detection among the multiple tests increases with the number of tests. This phenomenon is often referred to as the *look elsewhere effect* (e.g., Gross & Vitells 2010; van Dyk 2014). Likewise, rejecting the null hypothesis if the minimum of multiple p-values is below some threshold α does not guarantee that the probability of false detection is less than α.

Thus, to implement the Upper Bound Method in practice, we need a procedure to choose γ in order to eliminate these multiple comparisons problems and allow us to report a single upper bound. Suppose we could compute $T_c(\boldsymbol{y})$ exactly, given c and $\boldsymbol{y}$. This seems a reasonable simplifying assumption if we run MCMC until the effective sample sizes of each posterior sample are sufficiently large and do not use an extreme value of c. Under this assumption, the Monte Carlo error in $\hat{u}$ comes exclusively from the uncertainty in $\hat{c}$ due to only having M draws $\boldsymbol{y}_0^{(1)}, \ldots, \boldsymbol{y}_0^{(M)}$ from $\mathcal{L}_0(\boldsymbol{y}_0)$. By Equation (22), $\hat{u} \geqslant \gamma$, so we should choose γ as small as possible in order to increase the chance of achieving a small $\hat{u}$ and hence a high statistical significance. However, there is a trade-off: the Monte Carlo error in estimating the quantile c grows as γ becomes smaller, leading to larger Monte Carlo error in estimating u. One approach to selecting γ is to estimate the Monte Carlo error in $\hat{u}$ for a range of values of γ and choose the smallest γ for which this error value is acceptably small. Choosing γ based on the estimated *uncertainty* in $\hat{u}$ alleviates some of the multiple comparisons concerns that would arise if we chose γ based on the estimated upper bounds themselves. The Monte Carlo error can be estimated via bootstrap by resampling with replacement from $\{\boldsymbol{y}_0^{(1)}, \ldots, \boldsymbol{y}_0^{(M)}\}$. We demonstrate this in Section 4.4; see Figure 7. In our analyses, we found that values of γ in the range of 0.005–0.01 appeared reasonable when $M = 50$.

4.2.2. Implementation of a Suite of MCMC Samplers

The reason that we aim to reduce the number, M, of replicate images that must be analyzed is not just to save computer time, but also because each replicate must be fit using MCMC, which can be temperamental in practice. Indeed we can never precisely compute $T_c(\boldsymbol{y})$, but only an estimate of the tail probability in Equation (12) with $\widehat{T}_c(\boldsymbol{y})$. Monte Carlo error affects the variance of this estimate and lack of MCMC convergence causes bias. Even a biased estimator, however,

can be a valid test statistic, although perhaps a less powerful one. (A similar view of computational errors when using Monte Carlo to evaluate the sampling distribution of a likelihood ratio test was espoused by Protassov et al. 2002.) The key is that to guarantee a valid statistical test, precisely the same procedure must be implemented for each replicate image as for the observed images. Thus, the same methods for choosing starting values and diagnosing MCMC convergence must be performed for each image. The danger lies in the temptation to carefully implement and monitor MCMC for the observed image, but not for the M replicate images. Such differences in implementation may result in a systematic bias that may be mistaken as a statistically significant difference, simply because the test statistic function $\widehat{T}_c(\cdot)$ applied to replicate images $\boldsymbol{y}_0^{(j)}$ is different than the function applied to the observed image $\boldsymbol{y}_{\rm obs}$.

4.3. False Positive Rate and Statistical Power

To investigate the statistical properties of our proposed method, we applied our procedure to a total of 6000 simulated observations. Specifically, we simulated 1000 observed images under each of three null hypotheses (i.e., with no jet, but with differing background rates) and 1000 observed images under each of the three specific alternative hypotheses used in the simulations of Section 4.1 (i.e., with a weak, medium, or strong jet). The null hypotheses were not identical for the different jet strengths in that the background intensities under the null hypotheses were higher when the alternative contained a stronger simulated jet. All simulated observed images were generated as in Section 4.1. We considered a testing procedure that first chooses a tail probability γ of $g(\xi)$ and a nominal significance level α, and rejects the null hypothesis if the estimated upper bound $\hat{u} \leqslant \alpha$, where $\hat{u}$ is defined in Equation (22). We evaluated this procedure for a range of values of γ.

To approximate $g(\xi)$, we sampled $M = 50$ replicate images under each null hypothesis. A full simulation would require fitting the null model to each of the 6000 simulated observations, generating 50 replicate images from the fitted null model of each, and performing a full Bayesian analysis of all 300,000 resulting images. To reduce the computational demands of this simulation study, we performed an approximate simulation. We generated 500 additional replicate images from each of the three fixed null hypotheses (one for each background rate) and analyzed each replicate image using MCMC. For each simulated observed image, we resampled 50 replicate images without replacement from this collection of 500 replicate images, and used the corresponding 50 MCMC runs to construct $g(\xi)$.

We compared this approach to two Direct P-value Methods. The first uses the Davison & Hinkley (1997) Monte Carlo p-value of Equation (18) and the second uses the naive Monte Carlo p-value of Equation (19). Both Direct P-value Methods use the same test statistic as the Upper Bound Method and thus require an estimate of c. To ensure fair comparisons, for both direct methods and for every value of γ, we set $c = \hat{c}$, the estimated $(1-\gamma)$ quantile of $g(\xi)$, and reject the null hypothesis if the estimated p-value is less than or equal to α.

[12] The estimated *false positive rate* is the fraction of images simulated under the null hypothesis for which the null hypothesis was incorrectly rejected and thus structure was falsely detected. The estimated *statistical power* is the fraction of images simulated under the alternative hypothesis for which the null hypothesis was rejected and thus true structure was detected.

Table 1 presents the estimated false positive rate and statistical power[12] for the three procedures for a range of choices of γ and α. Because $\hat{u} \geqslant \gamma$, it is impossible to reject the null hypothesis using the upper bound approach if $\gamma > \alpha$, so we do not include such choices in our comparisons.

If we knew the true upper bound u in Equation (20), then rejecting the null hypothesis only when $u \leqslant \alpha$ would lead to a conservative procedure. That is, the actual false positive rate would be less than or equal to the nominal significance level α. Using an estimate $\hat{u}$ rather than the exact value introduces the possibility that the procedure is no longer conservative and that the false positive rate is no longer controlled at the nominal level. However, from the results for the Upper Bound Method (UB) in Table 1, we see that in these simulations the estimated upper bound procedure is conservative: the false positive rate is never greater than α under any of the settings considered. The Direct P-value Method using $\hat{p}$ (DP1) is also conservative. The Direct P-value Method based on $\hat{p}_{\rm naive}$ (DP2) is not conservative, and the actual false positive rate exceeds the nominal level in all cases.

Because of the conservativeness of the Upper Bound Method, it suffers from low power when the simulated jet is weak. In these cases, the Direct P-value Methods have a much higher chance of correctly rejecting the null hypothesis. However, $\hat{p} \geqslant 1/(1+M)$, so rejection is only possible using $\hat{p}$ if $\alpha \geqslant 1/(1+M)$; the power drops to zero when α is smaller than $1/(1+M)$, which equals $1/51$ in our simulation. Using $\hat{p}_{\rm naive}$ allows us to reject at much lower nominal levels α and thus achieve reasonable power even when α is very small, but without properly controlling the actual false positive rate. When the jet is strong enough, the Upper Bound Method dominates the Direct P-value Methods in that it is able to achieve high power of detection at high confidence levels (low α), while conservatively controlling the false positive rate.

4.4. Data Analysis

The X-ray jet associated with the 0730+257 quasar (redshift $z = 2.868$) was observed by *Chandra* (ACIS-S detector) on 2009 December 12 (ObsID 10307) for about 20 ks. We reprocessed the *Chandra* data in CIAO (Fruscione et al. 2006) using the calibration database CALDB 4.5.7,[13] binned the original event files, selecting only the events in the energy range of 0.5–7 keV, and created a 64×64 pixel image with pixel size of 0.246 arcsec centered on the quasar; see Figure 6(a). The PSF, shown in Figure 1(c), was binned to the same scale. The baseline component used in fitting the alternative model included a Gaussian model of the quasar with standard deviation of 0.5 and 225 expected counts and a uniform background with 44 expected counts. The baseline component and simulated null images were created using Sherpa (Freeman et al. 2001).

Figure 6(b) shows the posterior means of $\tau_1 \Lambda_{1j}$. There appears to be additional structure beyond the baseline component. The posterior distributions, $\pi(\xi|\boldsymbol{y}_{\rm obs})$ and $\pi(\xi|\boldsymbol{y}_0^{(j)})$, are shown in the bottom panel of Figure 4; $\pi(\xi|\boldsymbol{y}_{\rm obs})$ appears slightly more extreme relative to $g(\xi)$ than does the corresponding posterior distribution for the simulated medium jet in the middle panel of Figure 4. From Figure 5, it appears that the strength of evidence for additional emission

[13] http://cxc.cfa.harvard.edu/caldb/

Table 1
False Positive Rate and Power Estimated from 1000 Images Simulated under the Null and 1000 Images Simulated under the Alternative for Each Jet Strength

Jet Strength	γ (%)	α (%)	False Positive Rate (%)			Power (%)		
			UB	DP1	DP2	UB	DP1	DP2
20	1.0	2.0	0.1	2.0	3.9	29.8	**74.0**	81.7
	0.5	2.0	0.4	2.0	3.9	41.4	**70.9**	80.7
		1.0	0.0	0*	2.0	**18.1**	0*	72.0
	0.1	2.0	0.8	2.0	3.9	48.8	**67.8**	78.6
		1.0	0.5	0*	2.0	**33.2**	0*	66.8
		0.5	0.0	0*	2.0	**19.3**	0*	66.8
40	1.0	2.0	0.1	2.0	3.9	**99.7**	**100.0**	100.0
	0.5	2.0	0.3	2.0	3.9	**99.7**	**100.0**	100.0
		1.0	0.1	0*	2.0	**97.6**	0*	100.0
	0.1	2.0	0.8	2.0	3.9	**99.6**	**99.8**	100.0
		1.0	0.2	0*	2.0	**98.4**	0*	99.8
		0.5	0.1	0*	2.0	**96.2**	0*	100.0
70	1.0	2.0	0.1	2.0	3.9	**100.0**	**100.0**	100.0
	0.5	2.0	0.4	2.0	3.9	**100.0**	**100.0**	100.0
		1.0	0.0	0*	2.0	**100.0**	0*	100.0
	0.1	2.0	0.6	2.0	3.9	**100.0**	**100.0**	100.0
		1.0	0.3	0*	2.0	**100.0**	0*	100.0
		0.5	0.2	0*	2.0	**99.9**	0*	100.0

Note. Jet strengths are total expected counts in the simulated jet. The null hypothesis was rejected if $\hat{u} \leqslant \alpha$ for the upper bound method (UB); if $\hat{p} \leqslant \alpha$ for the first direct P-value method (DP1); and if $\hat{p}_{\text{naive}} \leqslant \alpha$ for the second direct P-value method (DP2). False positive rates for DP1 and DP2 were calculated analytically. Because the null distribution was constructed from 50 images simulated under the null, for DP1 $\hat{p} \geqslant 1/51$. Thus, if $\alpha < 1/51$, it is impossible for the DP1 approach to reject the null hypothesis; these cases are identified with asterisks. Boldface indicates the methods that lead to the highest power (within Monte Carlo uncertainty) while ensuring that the false positive rate is less than the nominal significance level α.

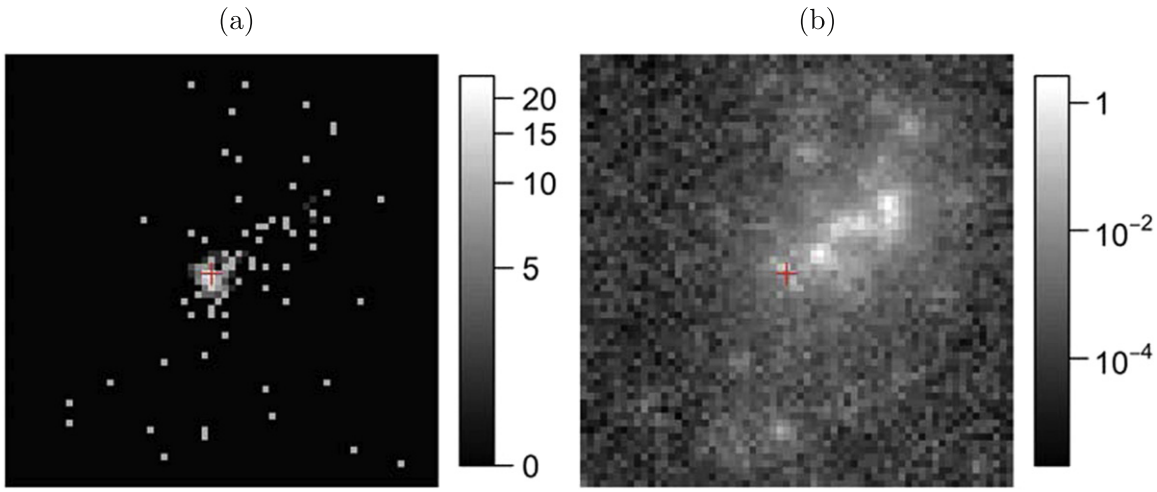


Figure 6. (a) X-ray observation of the 0730+257 quasar and its possible jet. The image is centered on the location of the quasar, and its width and height are both 15.7 arcsec. (b) Pixel-wise posterior means of the added component in the alternative model. The red + in (a) and (b) identifies the location of the quasar obtained from fitting the null model in Sherpa.

beyond $\mathbf{\Lambda}_0$ is between that in the simulations with the medium jet and with the strong jet.

Here, we perform a global test of adequacy of the null hypothesis, based on ξ given in Equation (11), computed for the entire image. In K. McKeough et al. (2015, in preparation), we perform region-specific tests for this same observed image, based on ξ_R given in Equation (13), for several regions R chosen with guidance from radio observations.

We use bootstrap resampling to investigate the uncertainty in the estimated upper bound $\hat{u}$. In particular, we resampled the $\{\mathbf{y}_0^{(j)}, j = 1, \ldots M\}$ and their corresponding posterior samples, $\{\xi^{(j,1)}, \ldots, \xi^{(j,L)}\}$, to obtain 1000 bootstrap replications of our sample of size LM from $g(\xi)$. This bootstrap procedure treats the posterior samples, $\{\xi^{(j,1)}, \ldots, \xi^{(j,L)}\}$, given each null dataset $\mathbf{y}_0^{(j)}$, as fixed and estimates the uncertainties in $\hat{u}$ due to only having $M = 50$ null datasets. Figure 7(a) shows the margins of error (MoE) for approximate 95% bootstrap confidence intervals, $\exp\{\ln(\hat{u}) \pm 2\, s\}$, where s is the bootstrap standard deviation of $\ln(\hat{u})$. As expected, the uncertainty in $\hat{u}$ grows as γ decreases. This is also reflected in Figure 7(b), which shows that the variability in $\hat{u}$ renders it unreliable for very small values of γ.

Based on the estimated MoE in Figure 7(a), we choose to set $\gamma = 0.005$, since it is a small value that nonetheless allows a

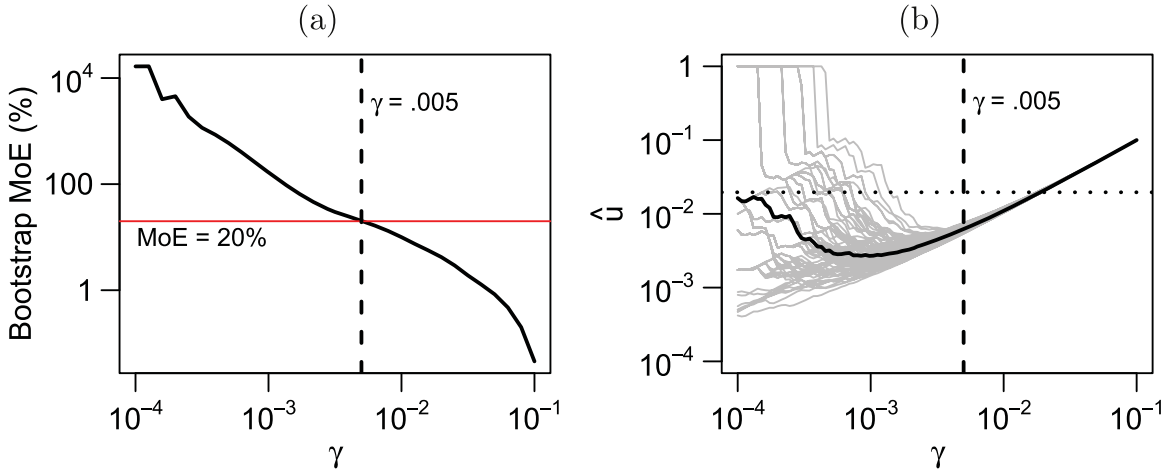


Figure 7. (a) Bootstrap margins of error (MoE) of the estimated upper bound $\hat{u}$ for different values of γ (solid black line). (b) The estimated upper bound $\hat{u}$ (solid line) and for 100 bootstrap replications (gray solid lines), obtained by sampling with replacement from the original null replicate images. The dotted horizontal line is at the minimum obtainable direct Monte Carlo p-value with $M = 50$, $\hat{p} = 1/51$. The dashed vertical line in (a) and (b) at $\gamma = 0.005$ corresponds to a bootstrap MoE of 20% (red horizontal line in (a)).

relatively precise estimate of $\hat{u}$, with MoE = 20% and a 95% bootstrap confidence interval of $(\hat{u}/1.2, 1.2\hat{u})$. This leads to an estimated threshold of $\hat{c} = 0.073$ and an estimated upper bound on the p-value of $\hat{u} = 0.0062$. For comparison, using $\gamma = 0.005$, the estimated upper bounds for the weak, medium, and strong jet simulations of Section 4.1 were 0.1837, 0.0076, and 0.00501, respectively. Thus, this analysis suggests significant inadequacy in the null hypothesis for the quasar, lending plausibility to the claim of a jet.

5. DISCUSSION

We have presented a method for computing the statistical significance of departures from a null model of an image. The test statistic is based on the posterior distribution under a Bayesian model that accounts for a PSF, detector inefficiencies, and Poisson noise, making it appropriate for low-count images in high-energy astrophysics. The Bayesian model allows for flexible departures from the null model via an added multiscale component. Because we use a posterior tail probability as a test statistic, we can compute an upper bound on a p-value that enables achieving high significance levels even when we have limited resources to devote to computations under the null hypothesis. We apply this method to an observed image of the 0730+257 quasar and find significant evidence of additional structure beyond the quasar and (flat) background, supporting the claim of an X-ray jet.

The simulations in Section 4.3 illustrate the trade-off between statistical efficiency and computational efficiency that our proposed Upper Bound Method navigates. The Upper Bound Method sacrifices some statistical efficiency, as seen in the reduced power relative to the Direct P-value Methods when there is a weak jet (the first, second, and fourth rows of Table 1). In return, the Upper Bound Method gains computational efficiency by enabling us to draw stronger conclusions when there are constraints on the number M of simulated null images that we can afford to analyze. In particular, the Upper Bound Method enables testing at significance levels α smaller than $1/(1 + M)$ (which the Direct P-value Method based on $\hat{p}$ cannot do), while ensuring that the false positive rate is no larger than the nominal significance level α (which the Direct P-value Method based on $\hat{p}_{\text{naive}}$ cannot do). Put another way, for a fixed computational time, the Upper Bound Method allows valid testing at smaller significance levels than does the Direct P-value Method.

Because of this sacrifice in statistical power, the Upper Bound Method may be too conservative to recommend if the goal is detection of very weak signals. Of course, if the signal is actually weak, the p-value under any test will most likely not be very small and can be computed with a smaller number of null simulations, so a direct Monte Carlo $\hat{p}$ may not be too computationally expensive. The real advantage of the upper bound approach is that high significance levels can be achieved for moderate or strong signals without extreme demands for simulation under the null, but with control of false positive rates.

We emphasize that it is only advantageous to use the Upper Bound Method instead of the Direct P-value Method when, due to computational constraints, we can only afford a modest M. The direct Monte Carlo estimates $\hat{p}$ and $\hat{p}_{\text{naive}}$ converge to the correct p-value as $M \to \infty$, while the estimated upper bound $\hat{u}$ converges to a conservative bound. In a constrained setting, however, the behavior as $M \to \infty$ is less relevant than the behavior for small M. When M is small, $\hat{p}$ may be more conservative than $\hat{u}$ (e.g., the entries in Table 1 in which the false positive rate and power of the Direct P-value Method based on $\hat{p}$ are exactly zero).

The advantages of the Upper Bound Method relative to the Direct P-value Method depend on the reliability of the estimate $\hat{c}$. The Monte Carlo estimate $\hat{p}$ is based on M draws of the test statistic $\widehat{T}_c(\boldsymbol{y}_0^{(j)})$ given a fixed c. The L draws from the posterior distribution conditional on $\boldsymbol{y}_0^{(j)}$ are only used to compute $\widehat{T}_c(\boldsymbol{y}_0^{(j)})$. In contrast, the upper bound $\hat{u}$ requires an estimate, $\hat{c}$, of a quantile of $g(\xi)$, and the LM posterior draws given the M simulated replicate images can be viewed as a cluster sample from $g(\xi)$. Whereas a large value of M is always needed to obtain a small $\hat{p}$, it may be possible to obtain an accurate estimate of a small upper bound with a relatively small M. (Technically, this requires the posterior variance of ξ for any given $\boldsymbol{y}_0$ to dominate the variability of the posterior expectation of ξ as a function of $\boldsymbol{y}_0$; see Figure 8.)

There are many avenues for future work. We are especially interested in exploring extensions of this method to

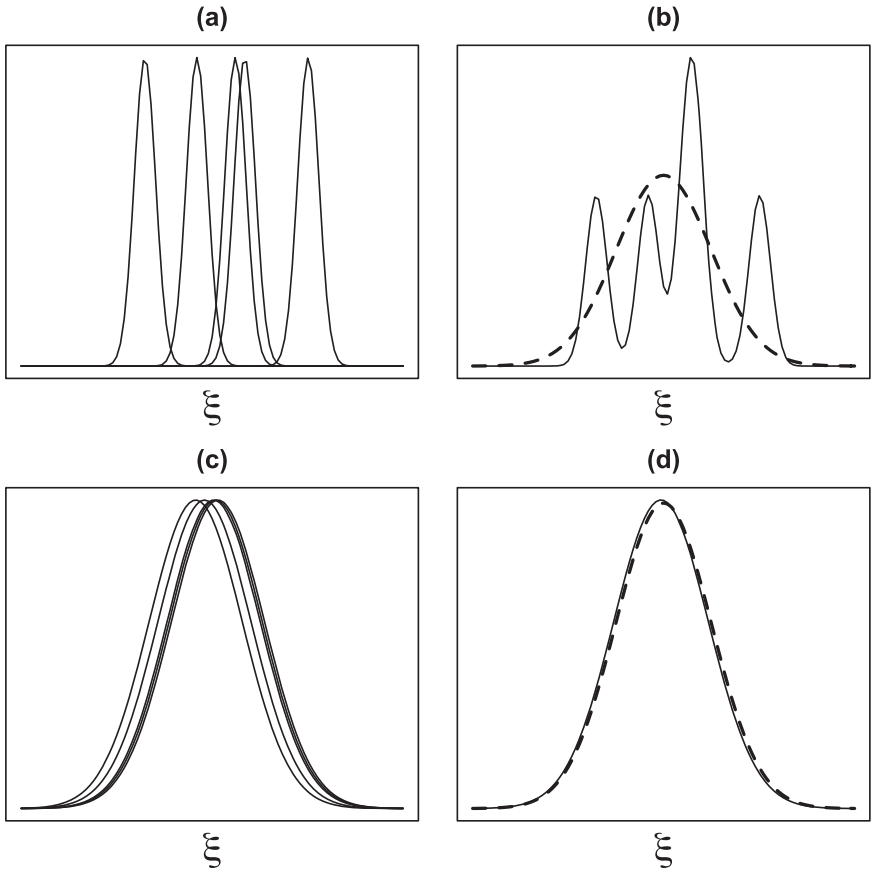


Figure 8. Illustration of the possibility of obtaining an accurate estimate of a small upper bound $\hat{u}$ with small M. Panels (a) and (c) display posterior densities $\pi(\xi|\mathbf{y}_0^{(j)})$ for $j = 1, \ldots, 5 = M$, where $\mathbf{y}_0^{(j)}$ are replicate images. Panels (b) and (d) show the estimated $\hat{g}(\xi) = M^{-1}\sum_{j=1}^{M}\pi(\xi|\mathbf{y}_0^{(j)})$ (solid lines) using the samples from Panels (a) and (c), respectively, along with the true $g(\xi)$ (dashed lines). In (a) and (b), there is more variation between posterior distributions of ξ for different replicate images $\mathbf{y}_0^{(j)}$ than within the posterior of ξ for any given $\mathbf{y}_0^{(j)}$, and the resulting estimate of $g(\xi)$ is not close to the truth with such a small M; in this case, we cannot accurately estimate quantiles of $g(\xi)$. In (c) and (d), the reverse is true, and even with only $M = 5$ replicate images, we can accurately estimate quantiles of $g(\xi)$.

automatically identify localized regions of significant departures from the null. This is important in astronomy because of the prevalence of low-count images and the great temptation to (over) interpret features in smoothed images (whether smoothed by eye or by an algorithm) as newly discovered objects. We are also interested in effects of model misspecification, in particular of the PSF, which may have non-negligible uncertainty. Finally, we seek to extend this type of analysis to cases where some property defining the source (e.g., the prevalence of substructure in a solar coronal loop) is tested, not just its intensity.

This work was conducted under the auspices of the CHASC International Astrostatistics Center. CHASC is supported by NSF grants DMS 1208791, DMS 1209232, DMS 1513492, DMS 1513484, DMS 1513546, and SI's Competitive Grants Fund 40488100HH0043. In addition, D.v.D. acknowledges support from a Wolfson Research Merit Award provided by the British Royal Society and from a Marie-Curie Career Integration Grant provided by the European Commission, and V.K. and A.S. from a NASA contract to the Chandra X-ray Center NAS8-03060. We used data obtained from the Chandra Data Archive and software provided by the Chandra X-ray Center (CXC) in the application packages CIAO and Sherpa. Finally, we thank CHASC members for many helpful discussions, especially Xiao-Li Meng and Kathryn McKeough.

APPENDIX
TECHNICAL DETAILS: COMPUTING u

In this Appendix, we derive the p-value upper bounds of Equations (20) and (25). We begin with the case in which the null hypothesis contains no unknown parameters. Because $T_c(\mathbf{y})$ is non-negative, Markov's inequality[14] yields

$$p \leqslant \frac{E\{T_c(\mathbf{y}_0)\}}{T_c(\mathbf{y}_{\rm obs})} = u, \tag{27}$$

where the expectation $E\{T_c(\mathbf{y}_0)\} = \sum_{\mathbf{y}_0} T_c(\mathbf{y}_0)\mathcal{L}_0(\mathbf{y}_0)$. Using the definition of T_c, we can rewrite this expectation as

$$\begin{aligned} E\{T_c(\mathbf{y}_0)\} &= \sum_{\mathbf{y}_0}\Pr(\xi \geqslant c|\mathbf{y}_0)\,\mathcal{L}_0(\mathbf{y}_0) \\ &= \sum_{\mathbf{y}_0}\left\{\int_c^1 \pi(\xi|\mathbf{y}_0)\,d\xi\right\}\mathcal{L}_0(\mathbf{y}_0) \\ &= \int_c^1\left\{\sum_{\mathbf{y}_0}\pi(\xi|\mathbf{y}_0)\,\mathcal{L}_0(\mathbf{y}_0)\right\}d\xi \\ &= \Pr(\xi \geqslant c), \end{aligned} \tag{28}$$

[14] Markov's inequality states that if X is a non-negative random variable and $a > 0$, then $\Pr(X \geqslant a) \leqslant E(X)/a$.

where the probability in Equation (28) is taken with respect to $g(\xi)$ as given in Equation (21).

If the null hypothesis contains unknown parameters $\boldsymbol{\theta}_0$, then we obtain Equations (25) and (26) as follows. For each fixed $\boldsymbol{\theta}_0$, Markov's inequality yields

$$p(\boldsymbol{\theta}_0) \leqslant \frac{E\left\{T_c(\boldsymbol{y}_0)|\boldsymbol{\theta}_0\right\}}{T_c(\boldsymbol{y}_{\rm obs})},$$

where the expectation in the numerator is taken with respect to $\mathcal{L}_0(\boldsymbol{y}_0|\boldsymbol{\theta}_0)$. Averaging over $\pi_0(\boldsymbol{\theta}_0|\boldsymbol{y}_{\rm obs})$, the ppp-value is bounded by

$$\text{ppp-value} \leqslant \frac{E\left\{T_c(\boldsymbol{y}_0)|\boldsymbol{y}_{\rm obs}\right\}}{T_c(\boldsymbol{y}_{\rm obs})} = u_{\rm ppp}.$$

The expectation

$$E\left\{T_c(\boldsymbol{y}_0)|\boldsymbol{y}_{\rm obs}\right\} = \Pr\left(\xi \geqslant c|\boldsymbol{y}_{\rm obs}\right), \tag{29}$$

where the probability is taken with respect to the distribution in Equation (26).

REFERENCES

Bayarri, M. J., & Castellanos, M. E. 2007, StaSc, 22, 322
Berger, J. O., & Delampady, M. 1987, StaSc, 2, 317
Bull, P., Wehus, I. K., Eriksen, H. K., et al. 2014, arxiv:1410.2544
Calderwood, T., Dobrzycki, A., Jessop, H., & Harris, D. E. 2001, in Astronomical Data Analysis Software and Systems X 238, The Sliding-cell Detection Program for Chandra X-ray Data, ed. F. R. Harnden, F. A. Primini & H. E. Payne (San Francisco, CA: ASP), 443
Connors, A., & van Dyk, D. A. 2007, in Statistical Challenges in Modern Astronomy IV, Vol. CS371, ed. E. Feigelson & G. Babu (San Francisco, CA: ASP)
Davison, A. C., & Hinkley, D. V. 1997, Bootstrap Methods and their Application (Cambridge: Cambridge Univ. Press)
Ebeling, H., & Wiedenmann, G. 1993, PhRvE, 47, 704
Esch, D. N., Connors, A., Karovska, M., & van Dyk, D. A. 2004, ApJ, 610, 1213
Freeman, P., Doe, S., & Siemiginowska, A. 2001, Proc. SPIE, 4477, 76
Freeman, P. E., Kashyap, V., Rosner, R., & Lamb, D. Q. 2002, ApJS, 138, 185
Friedenberg, D. A., & Genovese, C. R. 2013, J. Am. Stat. Assoc., 108, 456
Friston, K. J., Holmes, A. P., Worsley, K. J., et al. 1995, Human Brain Mapping, 2, 189
Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, Proc. SPIE, 6270, 62701V
Gelman, A., Meng, X.-L., & Stern, H. 1996, Stat. Sin., 6, 733
Godtliebsen, F., Marron, J. S., & Chaudhuri, P. 2004, Image Vision Comput., 22, 1093
Gross, E., & Vitells, O. 2010, EPJC, 70, 525
Guan, Y. 2008, Biometrika, 95, 831
Harris, D. E., & Krawczynski, H. 2006, ARA&A, 44, 463
Holmström, L., & Pasanen, L. 2012, Technometrics, 54, 16
Jasche, J., & Wandelt, B. D. 2013, ApJ, 779, 15
Lucy, L. B. 1974, AJ, 79, 745
McKeough, K., Kashyap, V., & McKillop, S. 2014, in 2014 AGU Fall Meeting (San Francisco, CA: AGU), abstract SH13C-4127
Meng, X.-L. 1994, AnSta, 22, 1142
Nowak, R. D., & Kolaczyk, E. D. 2000, ITIT, 46, 1811
Pina, R. K., & Puetter, R. C. 1992, PASP, 104, 1096
Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V., & Siemiginowska, A. 2002, ApJ, 571, 545
Radke, R., Andra, S., Al-Kofahi, O., & Roysam, B. 2005, ITIP, 14, 294
Richardson, W. H. 1972, JOSA, 62, 55
Ripley, B. D., & Sutherland, A. I. 1990, RSPTA, 332, 477
Rubin, D. B. 1984, AnSta, 12, 1151
Sinharay, S., & Stern, H. S. 2003, J. Stat. Plann. Inference, 111, 209
Sutter, P. M., Wandelt, B. D., McEwen, J. D., et al. 2014, MNRAS, 438, 768
Sutton, E. C., & Wandelt, B. D. 2006, ApJS, 162, 401
Thon, K., Rue, H., Skrøvseth, S. O., & Godtliebsen, F. 2012, Comput. Stat. Data Anal., 56, 49
Urry, C. M., & Padovani, P. 1995, PASP, 107, 803
van Dyk, D. A. 2012, in Statistical Challenges in Modern Astronomy V, ed. E. Feigelson & G. Babu (Berlin: Springer)
van Dyk, D. A. 2014, Annu. Rev. Stat. Appl., 1, 41
Weinberg, M. D. 2012, BayAn, 7, 737
Zhao, S. 2014, PhD thesis, Univ. California