The Astronomical Journal, 169:329 (22pp), 2025 June https://doi.org/10.3847/1538-3881/adc931
© 2025. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

Auto-BUQ: Uncertainty Quantification for the Boundaries of Segmented Events

Jue Wang[1], Vinay L. Kashyap[2], Thomas C. M. Lee[1], David A. van Dyk[3], and Andreas Zezas[4]
[1] Department of Statistics, University of California, Davis, One Shields Avenue, Davis, CA 95616, USA; jujwang@ucdavis.edu
[2] Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA
[3] Statistics Section, Department of Mathematics, Imperial College London, 180 Queen's Gate, London, SW7 2AZ, UK
[4] Physics Department, University of Crete, P. O. Box 2208, GF-710 03, Heraklion, Crete, Greece
Received 2024 July 29; revised 2025 March 12; accepted 2025 March 27; published 2025 May 22

Abstract

We present a new method to estimate the boundary of extended sources in high-energy photon lists and to quantify the uncertainty in the boundary. This method extends the graphed seeded region growing method developed by M. Fan et al. Here, we describe how an unambiguous boundary of a centrally concentrated astronomical source may be defined by first spatially segmenting the photon list, then forcibly merging the segments until only two segments—an extended source and its background—remain, and finally constructing a boundary as the connected outer edges of the Voronoi tessellation of the photons included in the source segment. The resulting boundary is then modeled using Fourier descriptors to generate a smooth curve, and this curve is bootstrapped to generate uncertainties. We apply the method to photon event lists obtained during the observations of galaxies NGC 2300 and Arp 299. We demonstrate how the derived extent and enclosed flux of NGC 2300 obtained with Chandra and XMM-Newton are comparable. We also show how complex internal structure, as in the case of Arp 299, may be subsumed to construct a compact boundary of the object.

Unified Astronomy Thesaurus concepts: X-ray astronomy (1810); X-ray sources (1822); Spatial point processes (1915); Astrostatistics (1882); Astrostatistics strategies (1885); Astrostatistics techniques (1886); Voronoi tessellation (1952); Astronomy data modeling (1859); Extended radiation sources (504)

1. Introduction

> Who in the rainbow can draw the line where the violet tint ends and the orange tint begins? Distinctly, we see the difference of the colors, but where exactly does the one first blendingly enter into the other?
>
> Billy Budd, Sailor, Chapter 21 Herman Melville

A common challenge in observational astrophysics is the identification of the boundary of an irregular region of extended emission. We consider the boundary to be the detectable outline of the surface brightness distribution of a source as recorded in an observation. Planetary nebulae, star clusters, star-forming regions, jets, galaxies, and other cosmic sources exhibit irregular shapes that cannot be easily quantified, and are not amenable to descriptions with simple parametric formulae. The situation is exacerbated by background contamination, instrumental effects, and statistical noise. In practice, the shape or extent of such objects is often characterized by their approximate size or scale. (For instance, the "radius" of a globular cluster is defined as the average radial distance at which the cluster's surface brightness falls to half its maximum value; see, e.g., S. Chatterjee et al. 2013.) The boundaries of astronomical sources are thus often defined in a manner that is "data driven." It is therefore important to specify the boundary in a replicable manner and provide uncertainty bounds on it. The definition of the boundary of extended sources has been the subject of several efforts in optical astronomy, especially in the case of objects showing complex surface brightness distributions that cannot be described adequately with an analytic model. A commonly used method is the `SExtractor` (E. Bertin & S. Arnouts 1996), which is based on comparing the intensity of image pixels with a background map. However, these methods do not provide the uncertainty on the source boundary, which is important for estimating uncertainties in the measured photometry resulting from the selection of the aperture and also for comparing the source extent between different observations (e.g., across wavelengths or as a function of time for evolving sources).

The advent of high-resolution imaging in high-energy astrophysics poses particular challenges for data analyses. X-ray and γ-ray data exhibit complex features that are intricate to discern and characterize; the extraction of feature boundaries from event lists is challenging, especially in scenarios characterized by sparsely distributed photon counts (e.g., D. van Dyk et al. 2006). To address this challenge, event-list segmentation algorithms have been developed (e.g., S. D. Byers & A. E. Raftery 2002; J. D. Scargle 2002).

Recently, in M. Fan et al. (2023), we introduced an innovative nonparametric event-list segmentation algorithm designed to partition a field of view into distinct emission components. The method constructs the segmentation using seeded region growing (R. Adams & L. Bischof 1994) on a graph (`SRGonG`, Seeded Region Growing on a Graph). Starting with a Voronoi tessellation built on the location of the recorded photons, `SRGonG` sequentially grows an initial set of seeds along the dual graph of the Voronoi tessellation until every photon is assigned to exactly one seed, thus forming a segmentation of the full event list. In practice, a large number

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

of seeds are used to over-segment the list. This allows the combination of a greedy merging algorithm and the Bayesian information criterion (BIC) to selectively merge segments. In this way, we arrive at the penalized profile maximum likelihood estimate of the number of segments, the segmentation itself, and the segment-wise constant intensity function across the field of view. SRGonG is tailored to identify both point sources and diffuse spatial structures, even with sparse event-list data.

The new methods developed in this paper fall into two categories. First, we introduce modifications to SRGonG to enhance its performance and stability, particularly in the low signal-to-noise ratio (S/N) regime. Second, we develop a new set of algorithms and statistical methods that together can be used to quantify the uncertainties associated with the segment boundaries found in SRGonG. We refer to this new set of methods as Auto-BUQ (i.e., Boundary of Segmented Events with Uncertainty Quantification).

Although the simulation studies conducted in M. Fan et al. (2023) demonstrate the ability of SRGonG to separate event-list photons that are associated with an extended source from photons from a diffuse background, unsurprisingly, the quality of the segmentation tends to degrade when the S/N is low. This degradation is observed both in the general cases (as measured by the median of the adjusted Rand index, or ARI; see Section 2.2.1) and the worst-case scenario (as measured by the lower bound of the ARI). To address this, we modify the greedy algorithm used by SRGonG to merge segments within the over-segmented event list, known as the greedy merger. Specifically, we develop the randomized merger and the beam merger, which improve the performance of SRGonG under both a general case and the worst case, respectively. To enable the user to flexibly balance performance and computation time, we also devise a hybrid merging process that optimally switches between the greedy merger and randomized/beam merger.

Quantifying the uncertainty of a source boundary computed with SRGonG is a challenging task, owing both to the data-driven nature of SRGonG and the type of the statistical noise (an inhomogeneous Poisson process). Indeed, we have not identified any existing uncertainty quantification method in the astrophysical literature for extended object boundary estimation when the input data consist of high-energy photon lists.

The quantification of uncertainty in medical image segmentation has garnered significant attention, especially at the pixel/voxel level (e.g., H. B. Hamu Goldberg et al. 2018). These methods, however, typically rely on probabilistic atlases constructed from training sets of aligned images with expert annotations. Deep learning methods for probabilistic image segmentation have also been developed to encode prediction confidence and quantify ambiguity in binary segmentation problems (e.g., Y. Gal & Z. Ghahramani 2015; S. A. A. Kohl et al. 2018; Y. Ding et al. 2020). These methods model the pixel-wise class probability for a given image and visualize the uncertainty impact with probability maps or heat maps.

The methods developed in this paper are fundamentally different. SRGonG is not trained on a sample with expert annotations, and there is no ground-truth training set available. Thus, we must develop novel methods to quantify uncertainty in the SRGonG segmentation. Our methods focus on the case where a single extended source resides in the interior of the field of view. In practice, the source brightness and the intensity of the surrounding field vary continuously, leading SRGonG to produce a segmentation composed of more than two regions. Thus, the first step of our methods is to produce a single boundary estimate of the extended source. We propose Expanded SRGonG (ESRGonG), an extension of SRGonG that (i) separates a single boundary from within an SRGonG result and (ii) extracts the boundary from the overall segmentation. We also propose an Automated Boundary Uncertainty Quantification (Auto-BUQ) method to quantify the uncertainty in the ESRGonG-generated boundary.

To quantify ESRGonG's uncertainty, we first parameterize the boundary using Fourier descriptors (W. Burger & M. J. Burge 2016) and choose the complexity of the parameterization (i.e., the number of parameters) with BIC. With the parameterization in hand, we are able to develop inferential tools for computing both parameter estimates (endowed with geometric interpretations) and global confidence regions that effectively capture the uncertainty in the shape of the extended source. A particular advantage of using Fourier descriptors is that we can further characterize the source in terms of its location, area, and eccentricity, and quantify the uncertainty of these characterizations. For example, the first term of the Fourier-descriptor spectrum represents the centroid of the segment, and hence its location. Additional terms represent the semiaxes of an ellipse approximating the area of the segment.

Our uncertainty quantification relies on resampling techniques (B. Efron 1979). The nature of the photon list, however, prohibits the direct application of the standard bootstrap. Choosing events with replacement from the event list invariably produces multiple identical events, and thus any such resampled data set is qualitatively different from the observed data. Thus, we develop a novel nonparametric spatial point bootstrap method that is applicable to event lists.

Our goal is to develop a method to compute a well-defined boundary of an irregular, extended astronomical object from high-energy photon lists. This enables, for example, the derivation of consistent sizes and mass estimates of galaxies for observations made across several telescopes and across several epochs.

In Section 2, we provide the necessary statistical background related to SRGonG and the modifications to SRGonG that we propose here. Section 3 reviews the Fourier descriptor parameterization, and Section 4 introduces the main novel methodology of the paper, Auto-BUQ: a parametric estimate of the boundary of a spatially extended source and an asymptotically minimal global confidence region for this boundary. We validate our proposed methods with a simulation study using synthetic data in Section 5, and with experiments involving real data of galaxies obtained with the Chandra X-Ray Observatory and XMM-Newton in Section 6. Finally, we summarize our work and discuss future extensions in Section 7.

2. Review and Improvements for SRGonG

2.1. SRGonG: Seeded Region Growing on Graphs

Instruments like the Chandra X-Ray Observatory record the energy, sky coordinates, and arrival time of X-ray photons. In this paper, we confine our attention to the locations recorded by the sky coordinates, i.e., we ignore the instrumental point-spread function (PSF). We summarize our notation in Table 1, and describe them in detail below. We model the recorded sky coordinates of a

Table 1
Glossary of Notation

Variables Describing the Data	
$\mathcal{F}$	Field of view, a bounded domain in $\mathbb{R}^2$ that contains all of the photons being considered
n	Number of photons in $\mathcal{F}$
$\boldsymbol{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$	Observed locations of the n photons, denoting their spatial coordinates
Variables describing the segmentation	
$\hat{\zeta}$	Estimate of a generic parameter ζ
$\mathcal{V}_i$	Voronoi cell defined by photon i
K	Number of segments within $\mathcal{F}$
$\boldsymbol{S} = \{\mathcal{S}_k, k = 1,\ldots,K\}$	Partitions of $\mathcal{F}$; $\mathcal{S}_k$ is the domain for the k^{th} segmented region
$\lambda(\boldsymbol{x})$	Poisson intensity at location $\boldsymbol{x}$
$\boldsymbol{\lambda} = \{\lambda_1, \ldots, \lambda_K\}$	Collection of the Poisson intensities over each of the segments
$\boldsymbol{\Lambda}$	Total integrated intensity over $\mathcal{F}$
$\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$	The Voronoi estimator of the Poisson intensity across Voronoi cell i, with $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i} = 1/\text{Area}(\mathcal{V}_i)$
Area($\mathcal{V}_i$)	Area of Voronoi cell $\mathcal{V}_i$
Num($\mathcal{S}_k$)	Number of photons in $\mathcal{S}_k$
Area($\mathcal{S}_k$)	Area of $\mathcal{S}_k$
$\mathcal{L}(K, \boldsymbol{S}, \boldsymbol{\lambda} \mid \boldsymbol{X})$	Log-likelihood of the model
$\mathcal{L}_{\text{profile}}(K, \boldsymbol{S} \mid \boldsymbol{X})$	Profile log-likelihood obtained by replacing $\boldsymbol{\lambda}$ in $\mathcal{L}(K, \boldsymbol{S}, \boldsymbol{\lambda} \mid \boldsymbol{X})$ with its estimate $\hat{\boldsymbol{\lambda}}$
Variables describing the SRGonG method	
m_{R}	Number of strata used during Voronoi-area stratified sampling to specify seeds
m_{vorthr}	Minimum number of photons required for a Voronoi-area derived seed to be accepted
M_S	Number of seeds
m_{seg}	User-defined number of parameters required to model each individual nonbackground Segment in SRGonG algorithm
M_B	Size of merging buffer in beam merger
M_C	Size of merging candidate list in randomized merger
M_R	Number of repetitions of merging procedure in randomized merger
M_s	Number of random/beam merging steps in hybrid merging process
Variables describing the boundary	
$s(t; \boldsymbol{\theta}, J)$	Parameterized function of extended object boundary.
$\mathcal{S}_{\hat{s}(J)}$	Interior of $\hat{s}(J)$, the estimate of the boundary with J pairs of FDs
$\boldsymbol{\theta}$	FD parameters in $s(t; \boldsymbol{\theta}, J)$
J	Number of FDs in $s(t; \boldsymbol{\theta}, J)$
C	A generic segment boundary
C_1	Boundary curve of inner segment obtained from ESRGonG
L	Number of uniformly sampled points on s
$x_l, y_l, l = 0\ldots L-1$	Coordinates of the l^{th} uniformly sampled points on s
$\boldsymbol{z} = \{z_l, l = 0\ldots L-1\}$	Set of complex signals with $z_l = x_l + iy_l$
$\boldsymbol{G} = \{G_j, j = 0\ldots L-1\}$	Fourier spectrum derived by applying DFT to $\boldsymbol{z}$
Re(G_j), Im(G_j)	Real and imaginary part of G_j
Variables describing the simulations	
μ_B	Expectation of simulated background counts over $\mathcal{F}$
μ_S	Expectation of total counts in the simulated stepped source
μ_{SC}, μ_{SH}	Expectation of total counts in the central circle and halo of stepped source
c	Simulated source contrast $\equiv \frac{\text{Halo area} + \text{Circle area}}{\text{Circle area}} \frac{\mu_{SC}}{\mu_S}$
N_B	Number of bootstrap replications
N_M	Number of Monte Carlo replications

collection of the n photons over a field of view, $\mathcal{F}$,

$$\boldsymbol{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\},$$

as an inhomogeneous Poisson process. This means that the number of photons, $n \sim \text{Pois}(\boldsymbol{\Lambda})$, and that given n, the sky coordinates are distributed $\boldsymbol{x}_i \overset{\text{iid}}{\sim} \lambda(\boldsymbol{x})/\boldsymbol{\Lambda}$, where $\lambda(\boldsymbol{x})$ is an integrable intensity function over $\mathcal{F}$, and $\boldsymbol{\Lambda} = \int_{\mathcal{F}} \lambda(\boldsymbol{x}) d\boldsymbol{x}$. We further assume that $\mathcal{F}$ can be partitioned into K segments, denoted $\boldsymbol{S} = (\mathcal{S}_1, \ldots, \mathcal{S}_K)$, where $\lambda(\boldsymbol{x})$ is constant with intensity λ_k on each $\mathcal{S}_k$; we denote $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_K)$. Thus, the intensity function can be expressed as

$$\lambda(\boldsymbol{x}) = \sum_{k=1}^{K} \lambda_k \mathbb{1}_{\mathcal{S}_k}(\boldsymbol{x}), \tag{1}$$

where $\mathbb{1}_{\mathcal{S}_k}(\boldsymbol{x})$ is an indicator function that is equal to 1 if $\boldsymbol{x} \in \mathcal{S}_k$ and is equal to 0 otherwise.

The SRGonG algorithm (M. Fan et al. 2023) is a novel event-list segmentation algorithm for obtaining penalized maximum likelihood estimates of $\lambda(\boldsymbol{x})$, $\boldsymbol{S}$, and K under an BIC-type penalty. The algorithm starts with the Voronoi diagram of the sky coordinates and uses a seeded region growing (SRG) step to combine Voronoi cells into an over-segmented[5] data set. SRG requires that each cell have a well-defined set of neighbors, and M. Fan et al. (2023) proposed using the dual graph of the Voronoi cell, known as the Delaunay triangulation[6] for this purpose.

2.1.1. Seed Specification

SRGonG requires a set of initial seeds, each with coordinates in $\mathcal{F}$. M. Fan et al. (2023; see their Section 3.2) suggested the seeds include a regularly spaced grid of photons from $\boldsymbol{X}$ and perhaps a small group of photons neighboring each grid point (where neighbors are determined by the Delaunay triangulation). This is the seeding strategy we use with the simple geometry adopted for simulations (see Section 5 below). If the set of seeds is too sparse to capture particular structures, it can be augmented with seeds from areas of $\mathcal{F}$ with high photon density (local maxima), since such areas are likely to be associated with sources. M. Fan et al. (2023) showed that seeds must be overspecified (i.e., starting from a state of over-segmentation), but not excessively so, as greedy mergers can lead to a reduction in computational efficiency and segment recoverability; specifying seeds using a context dependent heuristic provides a reasonable starting point for the merging process. They suggest, for example, grouping the Voronoi cells into m_{R} strata with similar Voronoi cell areas, collecting the events within each stratum into isolated groups (clustering them strictly by separation, i.e., a group is defined by requiring that the "within" variance is smaller than the "between" variance), and selecting as seeds a representative event from each group with $> m_{\text{vorthr}}$ events. This strategy ensures that the set of seeds includes a selection of the smallest cells (likely point sources), the largest cells (likely background), and everything in between. This is the strategy we adopt to analyze real data (see Section 6 below).

[5] By over-segmented, we mean a segmentation that contains far more segments than are physically meaningful.

[6] The Voronoi diagram and Delaunay triangulation are dual structures. The Voronoi cell for a photon is the spatial region that is closer to this photon than to any other photon in the data set. The Delaunay triangulation is constructed from the Voronoi diagram by connecting the centers of adjacent Voronoi cells (e.g., M. Fan et al. 2023; see also Section 2.2).

2.1.2. SRG Step

The penalized likelihood function favors segments that include Voronoi cells with similar estimated intensities. Following M. Fan et al. (2023), let $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ denote the estimated intensity for the cell corresponding to the photon $\boldsymbol{x}_i$, i.e., $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ is the reciprocal of the area of this cell. Similarly, the estimated intensity of segment $\mathcal{S}$ is $\mathrm{Num}(\mathcal{S})/\mathrm{Area}(\mathcal{S})$, where $\mathrm{Num}(\mathcal{S})$ is the number of photons with sky coordinates in $\mathcal{S}$, and $\mathrm{Area}(\mathcal{S})$ is the area of $\mathcal{S}$. To obtain the over-segmentation, the SRG step starts with an overly large set of M_S seeds and, at each iteration, aggregates one unclustered photon, $\boldsymbol{x}_i$, into one segment, $\mathcal{S}_k$, where the (neighboring) photon and segment (i and k) are selected to minimize the change in penalized log-likelihood, i.e., by minimizing the difference in their log estimated intensities,

$$\delta(i, \mathcal{S}_k) = |\log \widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i} - \log\{\mathrm{Num}(\mathcal{S}_k)/\mathrm{Area}(\mathcal{S}_k)\}|. \quad (2)$$

This procedure continues until all photons in $\mathcal{F}$ are assigned to exactly one segment.

Since we aim to have at least one seed for each potential segment, but do not know the true number and location of the segments, we recommend using a large number of seeds, which leads to an over-segmentation of the data.

Greedy Merger Step. `SRGonG` uses a greedy algorithm to sequentially merge pairs of segments, reducing the estimate of K until the penalized profile likelihood[7] is maximized. The penalized profile log-likelihood value (BIC) is defined as

$$\mathrm{BIC}(K, \boldsymbol{S}) = -2\mathcal{L}_{\mathrm{profile}}(K, \boldsymbol{S}|\boldsymbol{X}, n) + Km_{\mathrm{seg}}\log n, \quad (3)$$

with

$$\mathcal{L}_{\mathrm{profile}}(K, \boldsymbol{S}|\boldsymbol{X}, n) = \sum_{k=1}^{K} \mathrm{Num}(\mathcal{S}_k)\log(\mathrm{Num}(\mathcal{S}_k)/\mathrm{Area}(\mathcal{S}_k)) - n - \log n!, \quad (4)$$

where m_{seg} is a user-defined number of parameters required to model each individual nonbackground segment. In a step with K segments, if two neighboring segments, $\mathcal{S}_i$ and $\mathcal{S}_j$, are merged, the change in penalized profile log-likelihood is

$$\Delta\mathrm{BIC}_{K,i,j} = 2\,\mathrm{Num}(\mathcal{S}_i)\log\frac{\mathrm{Num}(\mathcal{S}_{i\cup j})\mathrm{Area}(\mathcal{S}_i)}{\mathrm{Area}(\mathcal{S}_{i\cup j})\mathrm{Num}(\mathcal{S}_i)} + 2\mathrm{Num}(\mathcal{S}_j)\log\frac{\mathrm{Num}(\mathcal{S}_{i\cup j})\mathrm{Area}(\mathcal{S}_j)}{\mathrm{Area}(\mathcal{S}_{i\cup j})\mathrm{Num}(\mathcal{S}_j)} + m_{\mathrm{seg}}\log n. \quad (5)$$

Thus, the pair of segments to be merged, $\mathcal{S}_i$ and $\mathcal{S}_j$, are selected to minimize the $\Delta\mathrm{BIC}_{K,i,j}$ resulting from their merger. By stepping through this procedure, we obtain a sequence of segmentations, with $k = 1, \ldots, M_S$ segments. The segmentation within this sequence with the largest penalized profile likelihood is returned as the output of the greedy algorithm.

[7] As noted, M. Fan et al. (2023) used a BIC-type penalty function that favors fewer segments; see Equation (3). They also profiled over $\boldsymbol{\lambda}$, replacing each λ_k with its maximum likelihood estimate given K and $\boldsymbol{S}$.

2.2. Improving the Greedy Merger Step in `SRGonG`

Unfortunately, the greedy merger step described in Section 2.1 is not robust to low S/N.

It is well known that stepwise greedy model selection algorithms, e.g., forward stepwise regression (B. Efron et al. 2004) and the matching pursuit (MP) algorithm (S. G. Mallat & Z. Zhang 1993), in general, provide suboptimal solutions. From extensive simulations, we find this to be particularly problematic when the sample size, n, or S/N are low or when there is a large number of fitted segments. On the other hand, an exhaustive search (over all possible segmentations) leads to the optimal solution but would be extremely time-consuming when there is a large set of possible segmentations. We, therefore, propose two alternative Merger Steps to improve the quality of the segmentation, one based on a randomized search and one based on a beam search. Both search strategies achieve a reasonable trade-off between computation time and segmentation quality.

Randomized Merger Step. A randomized search introduces randomization to each merging step by allowing suboptimal steps and therefore expanding the search space. In each merge step, with say K associated segments, we label each segment $1 \leqslant i < j \leqslant K$ and compute $\Delta\mathrm{BIC}_{K,i,j}$ for each neighboring pair, (i, j). Keeping track of the M_C candidate pairs (i, j) with the smallest $\Delta\mathrm{BIC}_{K,i,j}$, with equal probability, we randomly select one pair to merge. In this way, we obtain a sequence of nested segmentations, $\{\hat{S}(K), K = 1,\ldots,M_S\}$, each with a BIC value; we set $\widehat{K}$ and $\hat{\boldsymbol{S}}(\widehat{K})$ to the values of the nested level with the smallest BIC as computed with Equation (3). This procedure is repeated M_R times, resulting in M_R segmentations $\hat{\boldsymbol{S}}^{(m_r)}(\widehat{K}^{(m_r)})$, for $m_r = 1, \ldots, M_R$. Finally, we select the fitted segmentation to be the $\hat{\boldsymbol{S}}^{(m_r)}$ and $\widehat{K}^{(m_r)}$ that minimizes the BIC across all M_R segmentations.

Beam Merger Step. A beam search is another possible solution and provides more flexibility than the greedy merge step. It is a heuristic search algorithm that explores the space of possible segmentations by expanding the most promising nodes in a limited set; it has been widely used in fields like speech recognition (C. Tillmann & H. Ney 2003). We use one buffer to store M_B distinct possible segmentation for each $1 \leqslant K \leqslant M_S$. Initially, for $K = M_S$, the buffer contains one segmentation. For a fixed $K < M_S$, we compute the BIC of all possible mergers of all of the segmentations, $\hat{\boldsymbol{S}}^{(m_b)}(K + 1)$, $m_b = 1, \ldots, M_B$ in the Level $K + 1$ buffer, where the mergers we consider are segmentations of size K formed by merging two neighboring segments in one of the segmentations of size $K + 1$ in the Level $K + 1$ buffer. The distinct mergers with the smallest M_B BIC values are stored in the Level K buffer. This is repeated until $K = 1$ and the segmentation with the smallest BIC value across all buffers is returned. When $M_B = 1$, the beam search is the same as the greedy search. As M_B increases, the result converges to that of an exhaustive search.

Hybrid Merging Process. The computational times for the randomized merger and the beam merger steps scale with M_S times M_R and M_B, respectively. In practice, a large value of M_S impacts algorithmic efficiency. We propose a hybrid merging process that combines the greedy merger step and either the randomized or beam merger step and offers more flexibility between computation cost and model performance. Specifically, we execute the greedy merger step until it reaches $K = M_s$, where M_s is a predetermined hyperparameter.

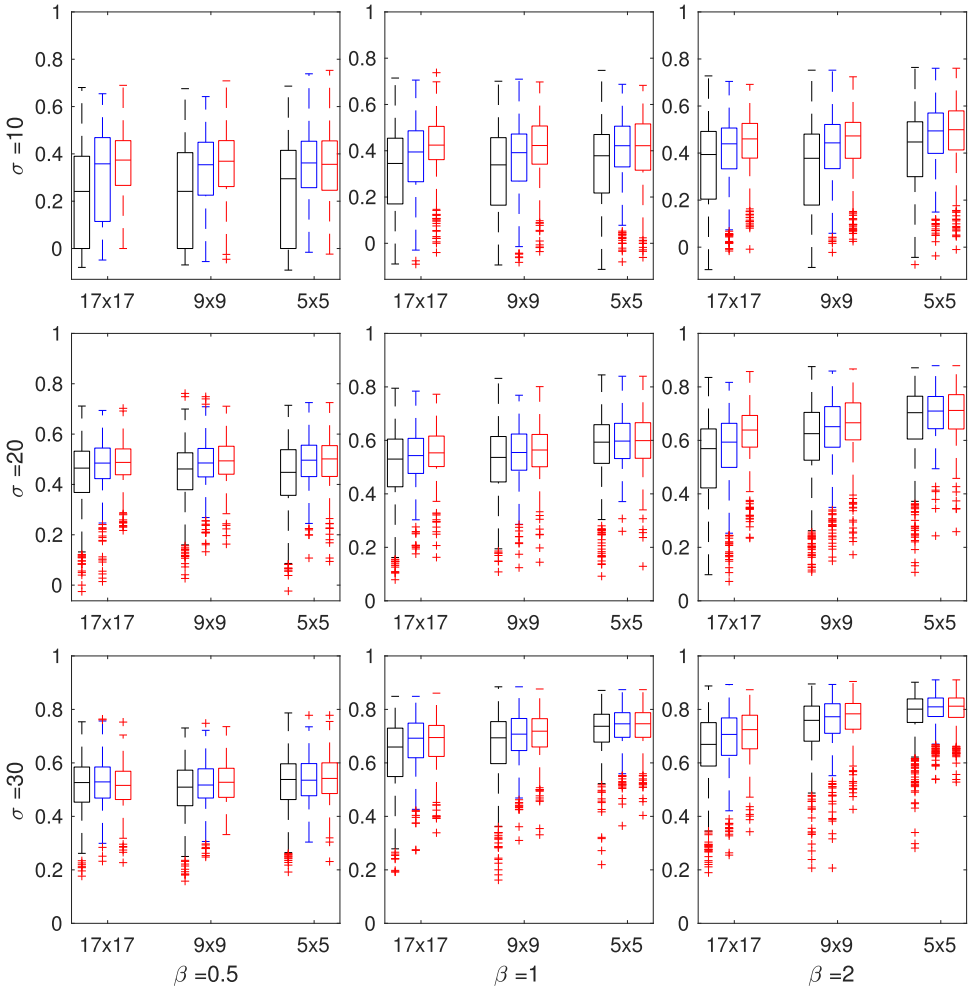


Figure 1. Simulation study comparing the greedy (black), beam (blue), and randomized (red) merger steps for a circular disk source with four embedded point sources. Box plots (with rectangles indicating the 25%, 50%, and 75% quartiles, whiskers extending to 1.5× the interquartile range, and samples beyond which are marked as outliers with '+') of the adjusted Rand index (ARI) obtained for each of 500 replications of each simulation setting are shown. Each panel corresponds to a specific set of (σ, β) denoting source strength and contrast, and each group within the panel corresponds to the number of seeds chosen to be on a grid. The ARI is systematically better (either with a narrower distribution or with systematically higher values) for randomized and beam merging strategies compared to the greedy merge. To minimize computation cost, we adopt a "hybrid" method that starts with greedy and ends with randomized or beam merge. The hybrid method's performance and computational cost are expected to be between that of the greedy and the randomized/beam merge methods, depending on when the switch is made and the performance of the greedy merge and the randomized/beam merges presented here serve as lower and upper bounds for the performance of hybrid merge.

Subsequently, we transition to either the randomized or beam merger step. The value of M_s should be set to be significantly larger than a reasonable guess as to the true underlying K; we typically set $M_s = 150$. In Appendix A, we introduce an alternative ad hoc approach for selecting M_s.

2.2.1. Simulations

We adopt the first simulation experiment as in M. Fan et al. (2023) to compare the greedy, random, and beam merger steps under several scenarios. In all simulation settings, four pointlike sources of area 0.08 are embedded within a circular extended source of area 0.2. Each pointlike source has $\beta\sigma$ counts, the diffuse extended source has $10\beta\sigma$ uniformly distributed counts, and there are 1000β background counts spread uniformly across $\mathcal{F}$.

Our experiments are organized into a $3 \times 3 \times 3$ design, where the first factor, β, takes on values 0.5, 1, and 2, representing small, medium, and large sample size; the second factor, σ, takes on values 10, 20, and 30, representing small, medium, and large S/N; the third factor is the specification of grid size and seed size for the SRG Step.[8] We generate 500 replicate data sets for each of 9 (β, σ) pairs of simulation settings and apply the three settings of the seed specification to each data set. We set the beamwidth to $M_B = 100$, the randomized candidate size to $M_C = 3$, and the iteration number to $M_R = 100$.

As in M. Fan et al. (2023), performance is evaluated via the ARI.[9] Figure 1 summarizes the median ARI of the 500

[8] We consider three combinations of grid size and seed size: (i) a 5×5 grid with and a seed size of 5, (ii) a 9×9 grid with a seed size of 3, and (iii) a 17×17 grid with and a seed size of 1. We use a neighborhood size of 50 for the local-maxima seed specification in addition to the regular grid.

[9] The Rand index (RI; W. M. Rand 1971) quantifies how well two segmentations agree (e.g., the fitted and ground-truth segmentations). Specifically, any pair of photons is classified into either (a) the same fitted segment and the same ground-truth segment, (b) different fitted segments and different ground-truth segments, or (c) neither (a) or (b). The RI equals the number of photon pairs in class (a) or (b) divided by the total number of photon pairs. Thus, a perfect match between a fitted segmentation the ground-truth results in $RI = 1$. The adjusted Rand index (ARI; L. Hubert & P. Arabie 1985) corrects the RI by accounting for the chance overlap of segments and takes values in the range $-1 < \text{ARI} < +1$ with larger values corresponding to better alignment between two segmentations.

replicates for each of the nine simulation settings and three seed specifications, with error bars representing the range of the first and third quartiles of the ARI. When the S/N is low ($\sigma = 10$), both the beam and randomized merger steps generally achieve a better ARI than the greedy merger Step. The ARI interquartile range for both the beam and the randomized merger step is smaller than that obtained with the greedy merger step. This illustrates the two alternative merger steps are more robust to noise and more stable than the greedy merger step when the S/N is low. Although the improvement is not significant with a 5×5 grid for the other two S/N settings, we observe both the beam and the randomized merger step still achieve higher lower bounds than that obtained with the greedy Merger step. We also observe an increasingly larger boost in ARI with a larger number of seeds, and the improvement of randomized or beam merger over the greedy merger is most significant for 17×17 grid.

As discussed in Section 2.1, we recommend using a large number of initial seeds, and it is precisely the case that using the randomized or the beam merger step yields the largest improvement over the greedy merger step.

Theoretically, the computing times of the randomized and beam merger steps equal the computing time of the greedy merger step times M_R and M_B, respectively. In practice, however, the randomized merger step can proceed in parallel, and is faster than the beam merger step when M_R and M_B are equal. As an illustration, the average computing times of the greedy, randomized, and beam merger steps with a 5×5 grid of seeds, $\sigma = 30$, and $\beta = 2$ are 0.117, 0.667, and 2.622 s, respectively (using a 2021 Macbook Pro with M1 Pro Chip and 16 GB memory).

3. Fourier Descriptors

3.1. Review of Coordinate-based Fourier Descriptors

Fourier descriptors (FDs; W. Burger & M. J. Burge 2013) parameterization is a highly flexible method to model closed curves in two-dimensional space (J. Mayer et al. 2010). A generic segment boundary, C, can be parameterized via FDs, i.e., C can be expressed as function s: $\mathbb{R} \to \mathbb{R}^2$,

$$\begin{aligned} s(t;\boldsymbol{\theta}, L) &= \begin{bmatrix} s_x(t;\boldsymbol{\theta}, L) \\ s_y(t;\boldsymbol{\theta}, L) \end{bmatrix} \\ &= V_L(t)^T\boldsymbol{\theta}, \qquad \text{for } t \in [0, 1], \end{aligned} \tag{6}$$

where $\boldsymbol{\theta}$ are the parameters of the FD, and $V_L(t)^T\boldsymbol{\theta}$ is a matrix representation of the parameterization; see Appendix C for details of the construction of $V_L(t)$. To derive $s(t; \boldsymbol{\theta}, L)$, we apply the discrete Fourier transformation (DFT) to C; specifically, we select L evenly spaced points along C, resulting in a two-dimensional sequence of coordinates $\{(x_l, y_l), l = 0 \ldots L-1\}$, where L can be arbitrarily large. The Fourier spectrum is the complex-valued vector, $\boldsymbol{G} = (G_0, \cdots, G_{L-1})$, of length L derived by applying the DFT to the set of complex signals $\boldsymbol{z} = (z_0, \cdots, z_{L-1})$, where $z_l = x_l + i \cdot y_l$, i.e.,

$$\begin{aligned} G_j &= \frac{1}{L}\sum_{l=0}^{L-1} z_l \cdot \exp\left\{-i2\pi j\frac{l}{L}\right\}, \\ &\text{for } j = 0, \cdots, L-1. \end{aligned} \tag{7}$$

Each G_j denotes an FD with frequency j. Because $\boldsymbol{G}$ is periodic, we have $G_j = G_{j \bmod L}$ for $j < 0$. The FD parameterization of C is a smooth interpolating curve, $s(t; \boldsymbol{\theta}, L)$ given by[10]

$$\begin{aligned} s_x(t;\boldsymbol{\theta}, L) = &\sum_{j=-\lfloor (L-1)/2 \rfloor}^{\lfloor L/2 \rfloor} [\mathrm{Re}(G_j) \cdot \cos(2\pi j \cdot t) \\ &- \mathrm{Im}(G_j) \cdot \sin(2\pi j \cdot t)], \\ s_y(t;\boldsymbol{\theta}, L) = &\sum_{j=-(\lfloor L-1)/2 \rfloor}^{\lfloor L/2 \rfloor} [\mathrm{Im}(G_j) \cdot \cos(2\pi j \cdot t) \\ &+ \mathrm{Re}(G_j) \cdot \sin(2\pi j \cdot t)], \end{aligned} \tag{8}$$

where $\boldsymbol{\theta} = (\mathrm{Re}(G_{-\lfloor (L-1)/2 \rfloor}), \ldots, \mathrm{Re}(G_{\lfloor L/2 \rfloor}), \mathrm{Im}(G_{-\lfloor (L-1)/2 \rfloor}), \ldots, \mathrm{Im}(G_{\lfloor L/2 \rfloor}))$. The inverse Fourier transformation guarantees that $(s_x(l/L; \boldsymbol{\theta}), s_y(l/L; \boldsymbol{\theta})) = (x_l, y_l)$ for each l and thus guarantees that $s(t; \boldsymbol{\theta}, L)$ becomes arbitrarily close to C as L grows large. The high-frequency terms in Equation (8) capture local noise-driven fluctuations, and thus we typically truncate them and include only terms with $j = -\lfloor (J-1)/2 \rfloor, \cdots, \lfloor J/2 \rfloor$, for some J such that $2 \leqslant J < L$, i.e., we replace $s(t; \boldsymbol{\theta}, L)$ with $s(t; \boldsymbol{\theta}, J)$.

Geometrically, each pair of summands in Equation (8) with frequency $\pm j$ represents an ellipse with major semiaxis $|G_{-j}| + |G_j|$ and minor semiaxis $||G_{-j}| - |G_j||$, and orientation angle $\frac{1}{2}[\tan^{-1}(\frac{\mathrm{Im}(G_{-j})}{\mathrm{Re}(G_{-j})}) + \tan^{-1}(\frac{\mathrm{Im}(G_j)}{\mathrm{Re}(G_j)})]$. The tuple $\langle G_{-j}, G_j \rangle$ is referred to as the "Fourier descriptor pair" (FDP) with frequency j, with the low-frequency FDP describing the global information and the high-frequency FDP describing the local information of C. In particular, we use the FD with frequency 0 to quantify the shape's centroid, and the FDP with frequency 1 to quantify the shape's size and eccentricity. This is a significant advantage of the FD parameterization. Since the parameterization is geometrically meaningful, quantifying uncertainty in terms of the parameterization is also geometrically meaningful.

3.2. Overparameterization

The FD generally depends on the choice of (x_0, y_0), as the choice of this value affects the other $L-1$ evenly space points used to compute the descriptors. The "shift property" of the DFT means that shifting the signal $\boldsymbol{g}$ by $-k_s$ position modifies the coefficients from G_j to $\tilde{G}_j = e^{ij\frac{2\pi k_s}{L}}G_j$. Hence, the value of $\boldsymbol{\theta}$ in Equation (8) is not unique, and there is ambiguity in the parameterization of the curve defining the segment. Motivated by J. C. Ye et al. (2000), we constrain the FDs by normalizing G_j to $\tilde{G}_j$ for each j. Specifically, we force $\tilde{G}_1 = |G_1|$, which is equivalent to shifting the phase of G_1 by $-\psi_1$, where

$$\psi_1 = \begin{cases} \tan^{-1}\frac{\mathrm{Im}(G_1)}{\mathrm{Re}(G_1)}, & \text{if} \quad \mathrm{Im}(G_1) \geqslant 0, \\ \left(\pi + \tan^{-1}\frac{\mathrm{Im}(G_1)}{\mathrm{Re}(G_1)}\right), & \text{if} \quad \mathrm{Im}(G_1) < 0. \end{cases} \tag{9}$$

It follows from the shift property of the DFT that

$$\tilde{G}_j = \begin{cases} e^{-ij\cdot\psi_1} \cdot G_j, & \text{for } 0 \leqslant j \leqslant \lfloor L/2 \rfloor, \\ e^{ij\cdot\psi_1} \cdot G_j, & \text{for } (L - \lfloor (L-1)/2 \rfloor \leqslant j \leqslant L-1. \end{cases} \tag{10}$$

In the following sections, we assume that $\boldsymbol{G}$ is normalized in this way.

[10] $\lfloor x/y \rfloor$ denotes the floor of x/y, i.e., the largest integer smaller than x/y.

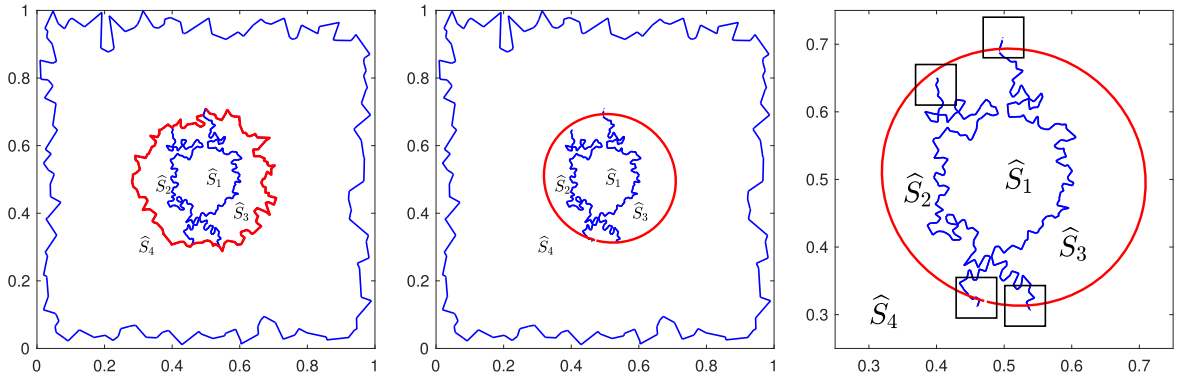


Figure 2. The potential issue of estimating J using a full segmentation, $\hat{\boldsymbol{S}}$, with $K > 2$. This segmentation is taken from a simulation study detailed in Section 5.1. Left panel: the blue and red lines represent the boundaries of segments in $\hat{\boldsymbol{S}}$. The boundary of interest, C_1, was extracted from $\widetilde{\boldsymbol{S}}$ and is highlighted in red. Middle panel: replacing the C_1 with $\hat{s}(J)$ within $\hat{\boldsymbol{S}}$ (with $J = 3$). Right panel: detail of the middle panel. The replacement of C_1 with $\hat{s}(J)$ causes multiple discontinuities in the boundaries of $\widehat{\mathcal{S}}_1$, $\widehat{\mathcal{S}}_2$, $\widehat{\mathcal{S}}_3$, and $\widehat{\mathcal{S}}_4$. These are highlighted with black boxes. These discontinuities mean that the middle panel is not segmented, and its log-likelihood function is undefined.

4. Statistical Inference

We turn now to the primary contribution of this paper: the extraction of a single boundary from an SRGonG-complied event-list segmentation (ESRGonG) and the quantification of the uncertainty of this boundary (Auto-BUQ).

4.1. ESRGonG: The Merge-to-two Step and Boundary Extraction

To isolate the boundary of a single extended source within an SRGonG segmentation, $\hat{\boldsymbol{S}}(\hat{K})$, we propose to merge segments to form a reduced segmentation of predetermined size, $\widetilde{K} < \widehat{K}$. To accomplish this, we use a "forced-merge" operation that applies a beam merger step to $\hat{\boldsymbol{S}}(\hat{K})$ and arrives at the reduced segmentation, $\widetilde{\boldsymbol{S}}(\widetilde{K})$, by maximizing the profile log-likelihood over all potential sub-sequence segmentations with $\widetilde{K}$ segments. In this article, we are primarily focused on the case where $\widetilde{K} = 2$, and thus refer to the "merge-to-two step." We apply a merge-to-two step to obtain $\widetilde{\boldsymbol{S}} = (\widetilde{\mathcal{S}}_1, \widetilde{\mathcal{S}}_2)$. Let $\widetilde{\mathcal{S}}_1$ be the central segment; we assume that $\widetilde{\mathcal{S}}_1$ has a single boundary.

The boundaries of a segmentation are extracted using the Voronoi diagram and its dual Delaunay triangulation. If two photons located in different segments are adjacent in the Delaunay triangulation, the edge shared by their Voronoi cells is taken as part of their common boundary. The boundary of each segment is constructed by sequentially connecting such a set of edges.

4.2. Estimate of the Object Boundary

In this section, we use segmentation and parameterization tools to obtain a parametric estimate of an object's spatial boundary, which is essentially the first step of Auto-BUQ.

We assume that the true boundary, $s^\star$, takes the same parameterized functional form as in Equation (8) and can be written $s^\star = s(t; \boldsymbol{\theta}^\star, J^\star)$, where $\boldsymbol{\theta}^\star$ and $J^\star$ are the "true" values of $\boldsymbol{\theta}$ and J. To be sure that we can estimate the complete boundary, we assume that $s(t; \boldsymbol{\theta}^\star, J^\star)$ is fully contained in $\mathcal{F}$.

In order to compute the estimate, we apply ESRGonG (Section 4.1) to obtain the full SRGonG segmentation $\hat{\boldsymbol{S}} = (\hat{\mathcal{S}}_1, \ldots, \hat{\mathcal{S}}_K)$ and the two-region segmentation $\widetilde{\boldsymbol{S}} = (\widetilde{\mathcal{S}}_1, \widetilde{\mathcal{S}}_2)$, as well as the single boundary of $\widetilde{\mathcal{S}}_1$, denoted C_1. We further assume that $\widetilde{\mathcal{S}}_1$ shares this boundary with $\widetilde{\mathcal{S}}_2$ and that $\widetilde{\mathcal{S}}_2$ completely surrounds $\widetilde{\mathcal{S}}_1$.

Our estimate of the boundary is as follows:

$$\hat{s}(J) = \begin{bmatrix} s_x(t; \hat{\boldsymbol{\theta}}, J) \\ s_y(t; \hat{\boldsymbol{\theta}}, J) \end{bmatrix} = V_J(t)^T \hat{\boldsymbol{\theta}}, \quad \text{for} \quad t \in [0, 1], \tag{11}$$

where $\hat{\boldsymbol{\theta}} \in \mathbb{R}^{2J-1}$ represents the normalized FD parameters computed with a set of L points, $\{(x_l, y_l)\}$, evenly spaced along C_1; i.e., $\hat{\boldsymbol{\theta}}$ consists of the real and imaginary parts of the J lowest-frequency terms in $\boldsymbol{G}$,

$$\begin{aligned} \hat{\boldsymbol{\theta}} = (&\mathrm{Re}(G_{-\lfloor (J-1)/2 \rfloor}), \cdots, \mathrm{Re}(G_{\lfloor J/2 \rfloor}), \\ &\mathrm{Im}(G_{-\lfloor (J-1)/2 \rfloor}), \cdots, \mathrm{Im}(G_0), \mathrm{Im}(G_2), \\ &\cdots, \mathrm{Im}(G_{\lfloor J/2 \rfloor}))^T \in \mathbb{R}^{2J-1}, \end{aligned} \tag{12}$$

where $\boldsymbol{G}$ is the Fourier spectrum of C_1 of length L. Because $\hat{\boldsymbol{\theta}}$ is obtained from the Fourier spectrum of C_1, once C_1 is determined, the estimate of the boundary depends only on the truncation cut off, J. Thus, we denote the estimate $\hat{s}(J)$ to emphasize its dependence on J. In the conventional Fourier parameterization, FDs are paired (C. C. Lin & R. Chellappa 1987; W. Burger & M. J. Burge 2016), i.e., the complete tuple $\langle G_{-j}, G_j \rangle$ is included in Equation (8). Following this convention, we assume $J^\star$ is an odd integer $\geqslant 3$. If the user does not wish to pre-determine the value of J (e.g., based on the complexity of the boundary[11]), J can be estimated, as described in Section 4.3.

4.3. Estimating J via the Bayesian Information Criterion

We estimate J by maximizing its profile log-likelihood. Because the likelihood depends on J through the segmentation it induces, we let $\mathcal{S}_{\hat{s}(J)}$ denote the interior region of $\hat{s}(J)$ and let $\boldsymbol{S}_{\hat{s}(J)} = (\mathcal{S}_{\hat{s}(J)}, \mathcal{F} \setminus \mathcal{S}_{\hat{s}(J)})$[12] be the partition of $\mathcal{F}$ into the interior and exterior of $\hat{s}(J)$. The profile log-likelihood of J can be expressed as $\mathcal{L}_{\text{profile}}(K = 2, \boldsymbol{S}_{\hat{s}(J)} | \boldsymbol{X}, n)$ as specified in Equation (4). We formulate the log-likelihood with $K = 2$ because we cannot substitute the parameterized boundary, $\hat{s}(J)$ for the ESRGonG boundary, C_1 within a segmentation with $K > 2$. This is because there would be discontinuities between an SRGonG segmentation and $\hat{s}(J)$; see Figure 2. Without a

[11] When $J = 3$, $s(t; \boldsymbol{\theta}, J)$ always describes an ellipse.

[12] $A \backslash B = \{x | x \in A \text{and} x \notin B\}$.

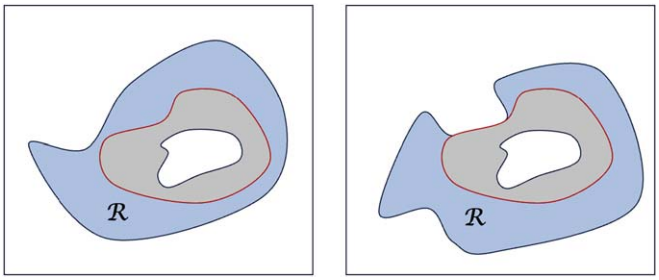


Figure 3. Illustration of enclosure assumption. The red curve is the object boundary $s(t, \boldsymbol{\theta}^\star, J^\star)$. The intensities over both blue and gray segments are assumed to be constant. Left panel: a subset $\mathcal{R}$ that satisfies the enclosure assumption. Right panel: a subset $\mathcal{R}$ that violates the enclosure assumption.

consistent segmentation, the log-likelihood is not defined. Thus, we evaluate the log-likelihood with $K = 2$. To accommodate the assumption that the Poisson intensities over both $\mathcal{S}_{\hat{s}(J)}$ and $\mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)}$ are constant, we compute $\mathcal{L}_{\text{profile}}$ using only photons recorded within a subset of $\mathcal{F}$, where the following holds.

Enclosure Assumption. There exists a subset $\mathcal{R} \subseteq \mathcal{F}$ that encloses $s^\star$, where photon arrivals can be appropriately modeled as an inhomogeneous Poisson process with an intensity function that is constant over both $\mathcal{S}^{\mathcal{R}}_{s^\star} = \mathcal{S}_{s^\star} \cap \mathcal{R}$ and $\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{s^\star} = (\mathcal{F}\backslash\mathcal{S}_{s^\star}) \cap \mathcal{R}$. By "encloses," we mean for all t, $s^\star = s(t, \boldsymbol{\theta}^\star, J^\star)$ is in the interior of $\mathcal{R}$.

Figure 3 illustrates two segmentations, one following the enclosure assumption and one violating it. We discuss the pitfalls of estimating J when the enclosure assumption is violated in Appendix B. Because $s^\star$ is unknown, we approximate $\mathcal{R}$ with the minimal union of segments in $\hat{\boldsymbol{S}}$, the full SRGonG segmentation, that fully encloses C_1, i.e.,

$$\mathcal{R} \approx \min_{\bigcup_k \widehat{\mathcal{S}}_k} \text{such that } C_1 \subset \text{interior}\left(\bigcup_k \widehat{\mathcal{S}}_k\right). \tag{13}$$

This is equivalent to the union of all $\widehat{\mathcal{S}}_k$ that share a boundary with C_1.

Under the enclosure assumption, we compute the profile log-likelihood of J using only the $n_{\mathcal{R}}$ photons observed in $\mathcal{R}$, i.e.,

$$\begin{aligned}\mathcal{L}^*_{\text{profile}}(J|\boldsymbol{X}_{\mathcal{R}}, n_{\mathcal{R}}) = \text{Num}(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}) \quad \log\left(\frac{\text{Num}(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)})}{\text{Area}(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)})}\right) \\ +\text{Num}(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}) \quad \log\left(\frac{\text{Num}(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)})}{\text{Area}(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)})}\right) \\ -n_{\mathcal{R}} - \log \quad n_{\mathcal{R}}!,\end{aligned} \tag{14}$$

where $\boldsymbol{X}_{\mathcal{R}} = \{\boldsymbol{x}_i : \boldsymbol{x}_i \in \mathcal{R}\}$. Because $\mathcal{L}^*_{\text{profile}}(J|\boldsymbol{X}_{\mathcal{R}}, n_{\mathcal{R}})$ is an increasing function of J, simply maximizing $\mathcal{L}^*_{\text{profile}}$ leads to overfitting of $s(t, \boldsymbol{\theta}, J)$. To avoid this, we introduce a BIC term to penalize model complexity. In particular, we use

$$\text{BIC}(J) = -2\mathcal{L}^*_{\text{profile}}(J|\boldsymbol{X}_{\mathcal{R}}, n_{\mathcal{R}}) + m_{\text{pseg}}(J)\log n_{\mathcal{R}}, \tag{15}$$

where $m_{\text{pseg}}(J)$ denotes the number of free parameters as a function of J, i.e., $m_{\text{pseg}}(J) = 2J + 1$, where $2J - 1$ corresponds to the free parameters in $\boldsymbol{\theta}$ given in Equation (12), and the remaining two account for the intensity parameters of $\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}$ and $\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}$. Consequently, the BIC estimate of J is

$$\hat{J}_{\text{BIC}} = \underset{J\in[3,L],\text{ with } J \text{ odd}}{\arg\min} \text{BIC}(J). \tag{16}$$

BIC is known for its statistical consistency in model selection problems,[13] and we illustrate its performance in Section 5.2.

The profile log-likelihood given in Equation (14) is closely related to that in Equation (4). Both encode the same Poisson model with its piecewise constant intensity function, with Equation (14) representing a "local fit" in that it removes events outside $\mathcal{R}$. We use Equations (3) and (4) to estimate the number of segments K, while setting the value of m_{seg} to represent the number of parameters required to model each segment. In Equations (14) and (15), on the other hand, we fix $K = 2$ when fitting J to control the complexity of the boundary. The two BIC penalties in Equations (3) and (15) are the same up to how the number of parameters are counted in Km_{seg} and $m_{\text{pseg}}(J)$, respectively.

4.4. Construct the Global Confidence Region through Bootstrapping

In this section, we introduce the second step of Auto-BUQ: the construction of an asymptotic global confidence region for the true object boundary $s^\star$. In order to construct such a global confidence region, we assume

$$\hat{\boldsymbol{\theta}} \overset{\text{asy}}{\sim} N(\boldsymbol{\theta}^\star, \Sigma_{\boldsymbol{\theta}}), \tag{17}$$

and thus from Equation (11) that

$$s(t; \hat{\boldsymbol{\theta}}, J) \overset{\text{asy}}{\sim} N(V_J(t)^T\boldsymbol{\theta}^\star, \Sigma_s(t)), \tag{18}$$

where $\Sigma_s(t) = V_J(t)^T\Sigma_{\boldsymbol{\theta}}V_J(t)$, and $\overset{\text{asy}}{\sim}$ indicates an asymptotic distribution. We follow the method proposed in J. C. Ye et al. (2000) to construct a small-sized global confidence region for $s(t, \boldsymbol{\theta}^\star, J^\star)$. For a fixed $t \in [0, 1]$, the smallest elliptical confidence region, $\hat{U}_\gamma(t)$, for $s(t, \boldsymbol{\theta}^\star, J^\star)$ with asymptotic confidence level $\alpha \in [0, 1]$, i.e.,

$$\Pr(s(t, \boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma(t)) = \alpha \tag{19}$$

under the asymptotic distribution of $s(t; \hat{\boldsymbol{\theta}}, J)$ is

$$\begin{aligned}\hat{U}_\gamma(t) = \{x \in \mathbb{R}^2 : (x - s(t; \hat{\boldsymbol{\theta}}, J))^T \\ \times \Sigma_s(t)^{-1}(x - s(t; \hat{\boldsymbol{\theta}}, J)) \quad \leqslant \gamma^2\},\end{aligned} \tag{20}$$

where γ^2 is the α quantile of a χ^2_2 distribution with two degrees of freedom,[14] and $J = J^\star$. The global confidence region for the

[13] BIC is known for its consistency in model selection problems, but may suffer from underfitting (K. P. Burnham & D. R. Anderson 2002). The Akaike information criterion (AIC) provides an alternative penalty term that could be used in place of BIC when estimating J. Both BIC and AIC measures are available in the implementation code, though we focus in this work on validating the BIC scheme. The estimate of J under the AIC penalty is

$$\hat{J}_{\text{AIC}} = \underset{J\in[3,L],\text{ with } J \text{ odd}}{\arg\min} -2\mathcal{L}^*_{\text{profile}}(J|\boldsymbol{X}_{\mathcal{R}}, n_{\mathcal{R}}) + 2m_{\text{pseg}}(J),$$

where $m_{\text{pseg}}(J)$ and $\mathcal{L}^*_{\text{profile}}$ are as in Equation (15). The AIC allows for a more flexible fitting of the boundary, but may favor overly complex fits; i.e., AIC is prone to overfitting. As the complexity of the true boundary increases, perhaps characterized by additional bulbs or sharp corners, FD requires more parameters, making the AIC a preferable choice. It is important to note, however, that complex shapes often entail wider global confidence regions as more parameters lead to larger model variance. In practice, this means an upper limit constraint on J is required, possibly guided by domain-specific information.

[14] I.e., $\gamma^2 = \chi^2_2(\alpha)$, where $\chi^2_2(\cdot)$ is the inverse cumulative distribution function of a chi-squared random variable with 2 degrees of freedom.

entire curve, $\hat{U}_\gamma$, is the union of the set of elliptical confidence regions for each t, i.e.,

$$\hat{U}_\gamma \triangleq \bigcup_{t\in[0,1]} \hat{U}_\gamma(t). \tag{21}$$

We must address two differences between our scenario and that of J. C. Ye et al. (2000). First, the value of $J^\star$ is predetermined in J. C. Ye et al. (2000) but not here. Second, J. C. Ye et al. (2000) was able to estimate $\Sigma_{\boldsymbol{\theta}}$ from multiple observations, while we have only one observed data set. We address the first issue by estimating $\hat{J}$ using the model selection procedure described in Section 4.3 and the second via a bootstrap method.

To estimate $\Sigma_{\boldsymbol{\theta}}$ via the bootstrap, we first generate N_B replicate data sets, $\boldsymbol{X}^{(N_b)}$, for $N_b = 1, \ldots, N_B$. Our procedure for doing this is described in Section 4.5. With the replicate data sets in hand, we apply ESRGonG to each data set and derive $\widetilde{\boldsymbol{S}}^{(N_b)}$ with two segments for each. Let $\boldsymbol{G}^{(N_b)}$ denote the Fourier spectrum of segment boundary $C_1^{(N_b)}$ extracted from $\widetilde{\boldsymbol{S}}^{(N_b)}$. Fixing J at $\hat{J}$, $\boldsymbol{\theta}^{(N_b)}$ consists of the $2\hat{J}-1$ lowest-frequency components in $\boldsymbol{G}^{(N_b)}$. We estimate $\Sigma_{\boldsymbol{\theta}}$ with

$$\widetilde{\Sigma}_{\boldsymbol{\theta}} = \frac{1}{N_B - 1}\sum_{N_b=1}^{N_B} (\boldsymbol{\theta}^{(N_b)} - \bar{\boldsymbol{\theta}})(\boldsymbol{\theta}^{(N_b)} - \bar{\boldsymbol{\theta}})^T$$
$$\text{with } \bar{\boldsymbol{\theta}} = \frac{1}{N_B}\sum_{N_b=1}^{N_B} \boldsymbol{\theta}^{(N_b)} \in \mathbb{R}^{2\hat{J}-1}. \tag{22}$$

The confidence region for $s^\star$ is constructed by replacing Σ_θ in Equation (20) with $\widetilde{\Sigma}_\theta$. The procedure is summarized in an algorithm, and we evaluate its performance through numerical experiments with extended sources of various shapes in Section 5.2. In Appendix D.1, we use Monte Carlo data to estimate the true crossing-level probability and compare it with the theoretical bounds proposed in J. C. Ye et al. (2000).

4.5. Bootstrapping

There are numerous nonparametric bootstrapping techniques to reconstruct spatial point patterns. Some widely used procedures include tiling (e.g., P. Hall 1985) as well as its many variants and extension (e.g., J. M. Loh & M. L. Stein 2004). An adaptive reconstruction method was proposed by A. Tscheschel & D. Stoyan (2006) that iteratively updates the points based on certain summary characteristics of the data. All of these methods, however, assume the spatial data are isotropic or aim to estimate some statistics depending on local patterns in a homogeneous way, which does not apply in our case. Thus, we propose two bootstrap procedures, one parametric and one nonparametric, that work in different scenarios.

4.5.1. Fourier-descriptors-based Parametric Bootstrap

The parametric bootstrap utilizes the FDs and the piecewise inhomogeneous Poisson assumption. More specifically, we assume that the intensity over both $\mathcal{S}_{s^\star}$, the segment enclosed by $s^\star$, and $\mathcal{F}\backslash\mathcal{S}_{s^\star}$ are constant and estimate the underlying intensity function as

$$\hat{\lambda}(\boldsymbol{x}) = \frac{\text{Num}(\mathcal{S}_{\hat{s}(J)})}{\text{Area}(\mathcal{S}_{\hat{s}(J)})}\mathbb{1}\{\boldsymbol{x} \in \mathcal{S}_{\hat{s}(J)}\} + \frac{\text{Num}(\mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)})}{\text{Area}(\mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)})}\mathbb{1}\{\boldsymbol{x} \in \mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)}\}. \tag{23}$$

Setting $\hat{\Lambda}$ equal to the integral of $\hat{\lambda}(\boldsymbol{x})$ over $\mathcal{F}$, i.e., $\hat{\Lambda} = \text{Num}(\mathcal{F})$, we obtain a bootstrap replicate, $\boldsymbol{X}^{(N_b)}$, by sampling from the Poisson process. I.e., we obtain an independent sample of size n from the density $\hat{\lambda}(\boldsymbol{x})/\hat{\Lambda}$.

4.5.2. Nonparametric Bootstrap

The parametric bootstrapping procedure generates piecewise constant Poisson data, which may not fully reflect the variance of the original Poisson process. This becomes particularly problematic in real data that are observations of diffuse sources. Therefore, we propose an alternative nonparametric solution. We first sample n photons with replacement from the observed event list, $\boldsymbol{X}$, and then create the Voronoi diagram from the unique photons that were selected. If an observed photon was selected $k > 1$ times, we independently place k uniformly distributed photons in its Voronoi cell.[15]

4.6. Empirical Quantities

With the bootstrap methods developed in Section 4.5, we can also quantify the uncertainty of ESRGonG without applying the parameterization step. Specifically, we first apply SRGonG and ESRGonG to the bootstrapped data $\boldsymbol{X}^{(N_b)}$ and compute $\widetilde{\boldsymbol{S}}^{(N_b)} = (\widetilde{\mathcal{S}}_1^{(N_b)}, \widetilde{\mathcal{S}}_2^{(N_b)})$ as described in Section 4.1, and then measure the uncertainty of the ESRGonG-generated boundary with the empirical distribution of the area and flux of the source segment $\widetilde{\mathcal{S}}_1^{(N_b)}$.

4.6.1. Area

We estimate the size of the object with $\text{Area}(\widetilde{\mathcal{S}}_1)$, and consider the empirical cumulative distribution function obtained from the bootstrapped data.

$$\hat{F}_{\text{Area}}(t) = \frac{1}{N_B}\sum_{N_b=1}^{N_B} \mathbb{1}\{\text{Area}(\widetilde{\mathcal{S}}_1^{(N_b)}) \leqslant t\}. \tag{24}$$

4.6.2. Flux

We estimate the flux as the counts within the segment of interest as $\text{Num}(\widetilde{\mathcal{S}}_1)$, which is the number of photons observed inside the segmented object boundary of $\widetilde{\boldsymbol{S}}$. The bootstrap estimate of the empirical cumulative distribution function of $\text{Num}(\widetilde{\mathcal{S}}_1)$ is

$$\hat{F}_{\text{Flux}}(t) = \frac{1}{N_B}\sum_{N_b=1}^{N_B} \mathbb{1}\{\text{Num}(\widetilde{\mathcal{S}}_1^{(N_b)}) \leqslant t\}, \tag{25}$$

where $\text{Num}(\widetilde{\mathcal{S}}_1^{(N_b)})$ represents the number of photons in the bootstrap sample N_b that are inside the bootstrap segment derived by ESRGonG.

[15] Alternatively, we could independently sample photon coordinates uniformly within Voronoi cells for the entire bootstrap sample, including cases where $k = 1$.

Table 2
Expected Source Counts (μ_{SC}, μ_{SH}) of Each Setting of the Simulation in Section 5.1, Specified in Terms of S/N, μ_B, and c

S/N	μ_B; c	c = 1.4	c = 2	c = 2.8
3	100	(6, 11)	(8, 8)	(11, 5)
3	500	(10, 18)	(14, 14)	(20, 8)
3	2500	(20, 37)	(29, 28)	(40, 17)
10	100	(39, 72)	(56, 55)	(78, 33)
10	500	(50, 92)	(71, 71)	(99, 43)
10	2500	(81, 149)	(115, 115)	(161, 69)
30	100	(319, 593)	(456, 456)	(638, 274)
30	500	(335, 621)	(478, 478)	(669, 287)
30	2500	(398, 739)	(569, 568)	(796, 341)

4.6.3. Background-subtracted Flux

In the presence of background, $\mathrm{Num}(\widetilde{\mathcal{S}}_1)$ will include background events and thus overestimate the flux attributable to the source. If we assume that the background is constant and $\widetilde{\mathcal{S}}_2^{(N_b)}$ includes only background, we can estimate the background intensity per unit area as $\mathrm{Num}(\widetilde{\mathcal{S}}_2)/\mathrm{Area}(\widetilde{\mathcal{S}}_2)$ and the expected background count in $\widetilde{\mathcal{S}}_1^{(N_b)}$ as $\mathrm{Area}(\widetilde{\mathcal{S}}_1)\mathrm{Num}(\widetilde{\mathcal{S}}_2)/\mathrm{Area}(\widetilde{\mathcal{S}}_2)$. Thus, the empirical cumulative distribution function of the adjusted estimator is

$$\hat{F}_{\mathrm{Flux}}(t) = \frac{1}{N_B} \times \sum_{N_b=1}^{N_B} \mathbb{1}\left\{\mathrm{Num}(\widetilde{\mathcal{S}}_1^{(N_b)}) - \frac{\mathrm{Area}(\widetilde{\mathcal{S}}_1^{(N_b)})\mathrm{Num}(\widetilde{\mathcal{S}}_2^{(N_b)})}{\mathrm{Area}(\widetilde{\mathcal{S}}_2^{(N_b)})} \leqslant t\right\}. \quad (26)$$

5. Simulations

5.1. Simulation Study: The Effects of Contrast and S/N on ESRGonG

In this section, Monte Carlo simulated data with various contrasts and S/Ns are generated to evaluate the performance of ESRGonG. We consider the scenario of a circular extended source surrounded by a halo. The expected photon count from a uniformly distributed background is denoted by μ_B, and those of the circular source and halo are denoted by μ_{SC} and μ_{SH}, respectively. The field of view, $\mathcal{F}$, is assumed to be $[0, 1] \times [0, 1]$ and the areas of the circular extended source and the halo are fixed at 0.03 and 0.09. Let $\mu_S = \mu_{SC} + \mu_{SH}$ be the total expected source count. The S/N is defined as $S/N = \mu_S/\sqrt{\mu_S + 0.12\mu_B}$, and the contrast is defined by $c = 4\mu_{SC}/\mu_S$. Our experiments are organized into a $3 \times 3 \times 3$ design, where the first factor, μ_B, takes values 100, 500, and 2500, representing small, medium, and high background noise; the second factor, S/N, takes values 3, 10, and 30, representing a small, medium, and high S/N; and the third factor, c, takes values 1.4, 2, and 2.8, representing small, medium, and high contrast. For each of the background and the two source objects, the photons are simulated from a Poisson process constructed with the given means, that is μ_B, μ_{SC} and μ_{SH}, respectively. The total number of photons is fixed to be $\mu_B + \mu_{SC} + \mu_{SH}$, and the photon locations within each region are uniformly distributed. The values of (μ_{SC}, μ_{SH}) of each simulation setting appear in Table 2. An illustrative figure of the extended source can be found in Figure 4.

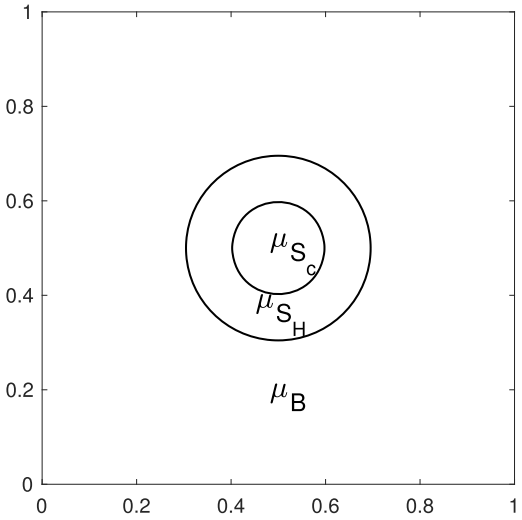


Figure 4. Illustration of the simulations in Section 5.1. An extended circular shape of area 0.03 is surrounded by a halo of area 0.09.

We generate $N_M = 500$ Monte Carlo replicate data sets for each of the 27 simulation settings and apply ESRGonG to each data set. Specifically, we first initialize the seeds on regular grids. For $\mu_B = 100$, 500, and 2500, we use 7×7, 10×10 and 10×10 grids with seed size = 1, 3, and 5, respectively. We then apply SRGonG with a random merger step to each simulated data set, specifying $M_C = 3$ and $M_R = 1000$. Lastly, we apply the beam merger step with $M_B = 20$ to the SRGonG segmentation and extract the boundary C_1 from the merged two-region segmentation.

We summarize the performance of ESRGonG in estimating the source radius. When $J = 2$, the radius of the fitted boundary is $|G_1|$. Thus, Figure 5 examines the distribution of $|G_1|$ in each of the 27 simulation settings. For some of these settings, such as low background brightness and low S/N, the BIC criteria frequently result in all subsegments being merged into one, i.e., $\widehat{K} = 1$, and there is no source in the fitted model. Even if $\widehat{K} > 1$, there may not be an identifiable source in the two-region segmentation, $\widetilde{\boldsymbol{S}}$, if there is no such $\widetilde{\mathcal{S}}_1$ whose boundary does not intersect with the boundary of $\mathcal{F}$ frame. When the source is identified (i.e., the boundary of $\widetilde{\mathcal{S}}_1$ does not intersect with $\mathcal{F}$ frame), we can extract its boundary, C_1.

We then apply DFT to each extracted C_1 and summarize the distribution of $|G_1|$ in Figure 5. Under the low- and medium-contrast levels, we expect ESRGonG to identify the inner circular extend source along with its halo. (These combine to form a circular shape with a radius ≈0.2.) Under the high-contrast setting, however, we expect ESRGonG to only identify the circular extended source. (This is a circular shape with radius ≈0.1.) Figure 5 illustrates that although $|G_1|$ fluctuates under small S/N, most of the time ESRGonG identifies a segment with $0.1 < |G_1| < 0.2$. The segmentation results are less ambiguous under the middle and high S/Ns levels.

5.2. Simulation Study: Performance of Auto-BUQ

We introduce three extended sources of varying complexity to illustrate the performance of ESRGonG and Auto-BUQ, including the FD boundary estimates, model selection techniques,

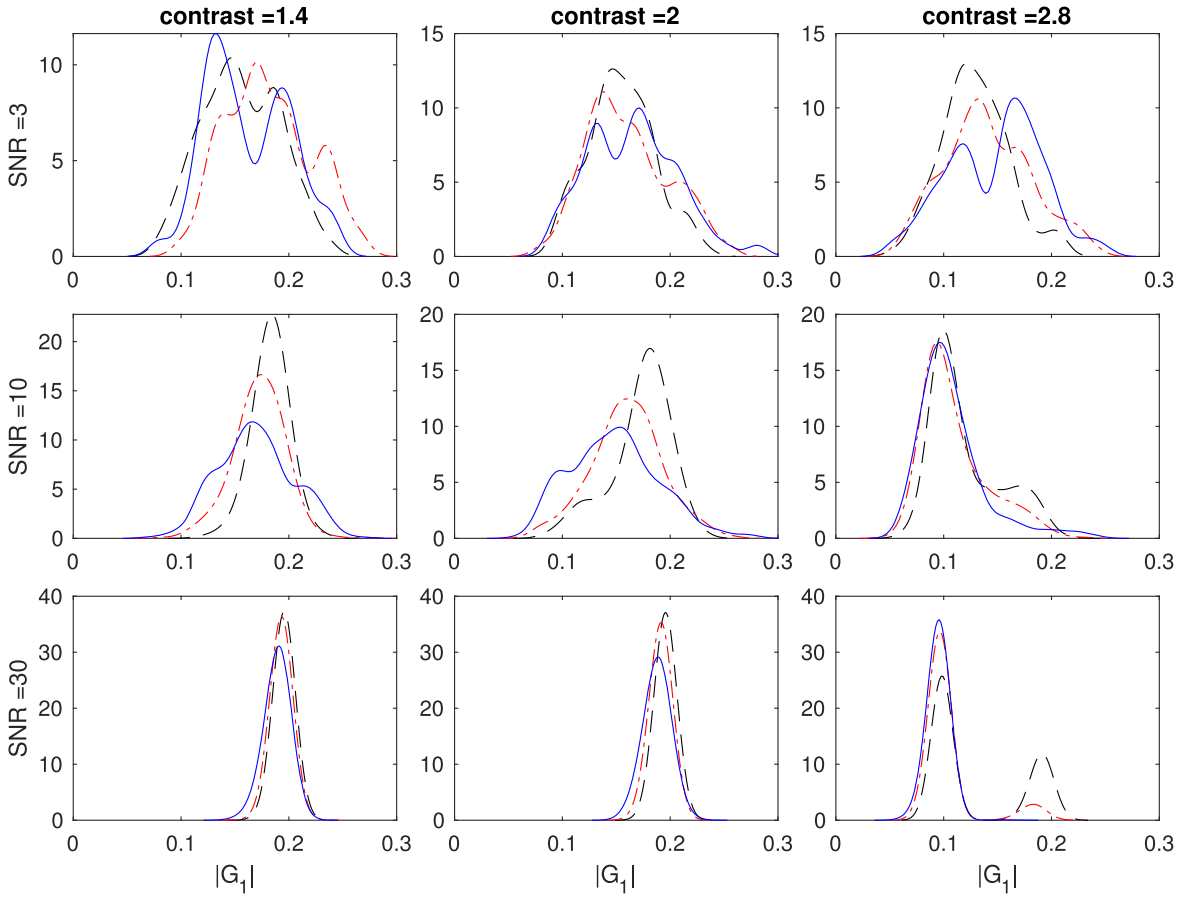


Figure 5. Simulation study for ESRGonG. The distribution of G_1 computed from the segmented object is plotted. The black dashed line, red dotted–dashed line and blue solid lines represent the frequency for $\mu_B = 100$, 500, and 2500, respectively.

and parametric bootstrap method for constructing a global confidence region.

The shapes of the extended sources that we consider appear in the first column of Figure 6, specifically, an ellipse, a bridge, and a spiral; these shapes can be constructed with 3, 7, and 11 FDs, respectively. Each extended source has an area of 0.2. For each source, we simulate $N_M = 300$ replicate data sets each with 1300 photons. The photons are randomly split between source and background with probability $3/13$ of being assigned to the extended source. Finally, the source and background photons are randomly distributed (uniformly) across their respective regions. This results in a surface brightness contrast between the source and the background of $1.5\times$.

We apply ESRGonG to each simulated data set, with $M_R = 1000$, $M_C = 3$ used in SRGonG randomized merger step and $M_B = 20$ in the forced-merge step and using a 9×9 regular seed grid, initialized with seed size at 5. To visualize the results, we randomly select one simulated data set from each of the three extended-source shape settings. The first column of Figure 6 shows these simulated data sets along with the corresponding ground truth. The second and third columns represent the ESRGonG segmentation results and the segment boundary obtained from the segmentation. Auto-BUQ results, including the FD estimates of the boundary obtained with the same three simulations, appear in Figure 7. The first column plots the FD boundaries, $s(t; \hat{\boldsymbol{\theta}}, \hat{J}_{\mathrm{BIC}})$, along with the simulated data (with $\hat{J}_{\mathrm{BIC}}$ computed as in Equation (16)), and the second and third columns illustrate the 39% and 68% global confidence regions, respectively, all constructed with $N_B = 200$ resampled parametric bootstrap data sets. We choose these two confidence levels as 39% corresponds to $\gamma = 1$, which straightforwardly illustrates the width of the band, and 68% is the standard confidence level that is commonly used in statistics tests.

Figure 8 summarizes the sampling distribution of $\hat{J}_{\mathrm{BIC}}$. BIC excels in modeling simple shapes but, in some cases, tends to favor overly simplistic models. This can result in underfitting, especially in cases involving intricate shapes like the spiral example shown in the third row of Figure 6. We suggest using BIC when the underlying extended object is not expected to exhibit sharp corners. As object complexity increases, characterized by additional bulbs or sharp corners, the FD parameterization requires more parameters.

6. Astronomical Applications

Here, we illustrate the FD parameterization and related analysis on two nearby galaxies with contrasting morphological features. The data sets used are listed in Table 3, and include Chandra and XMM-Newton observations of NGC 2300 and an XMM-Newton observation of Arp 299.

To apply ESRGonG, we first convert each data set to a $[0, 1] \times [0, 1]$ field of view by linearly transforming their sky coordinates as recommended by SRGonG implementation. The conversion parameters are given in Table 3; the center coordinate is mapped to (0.5, 0.5).

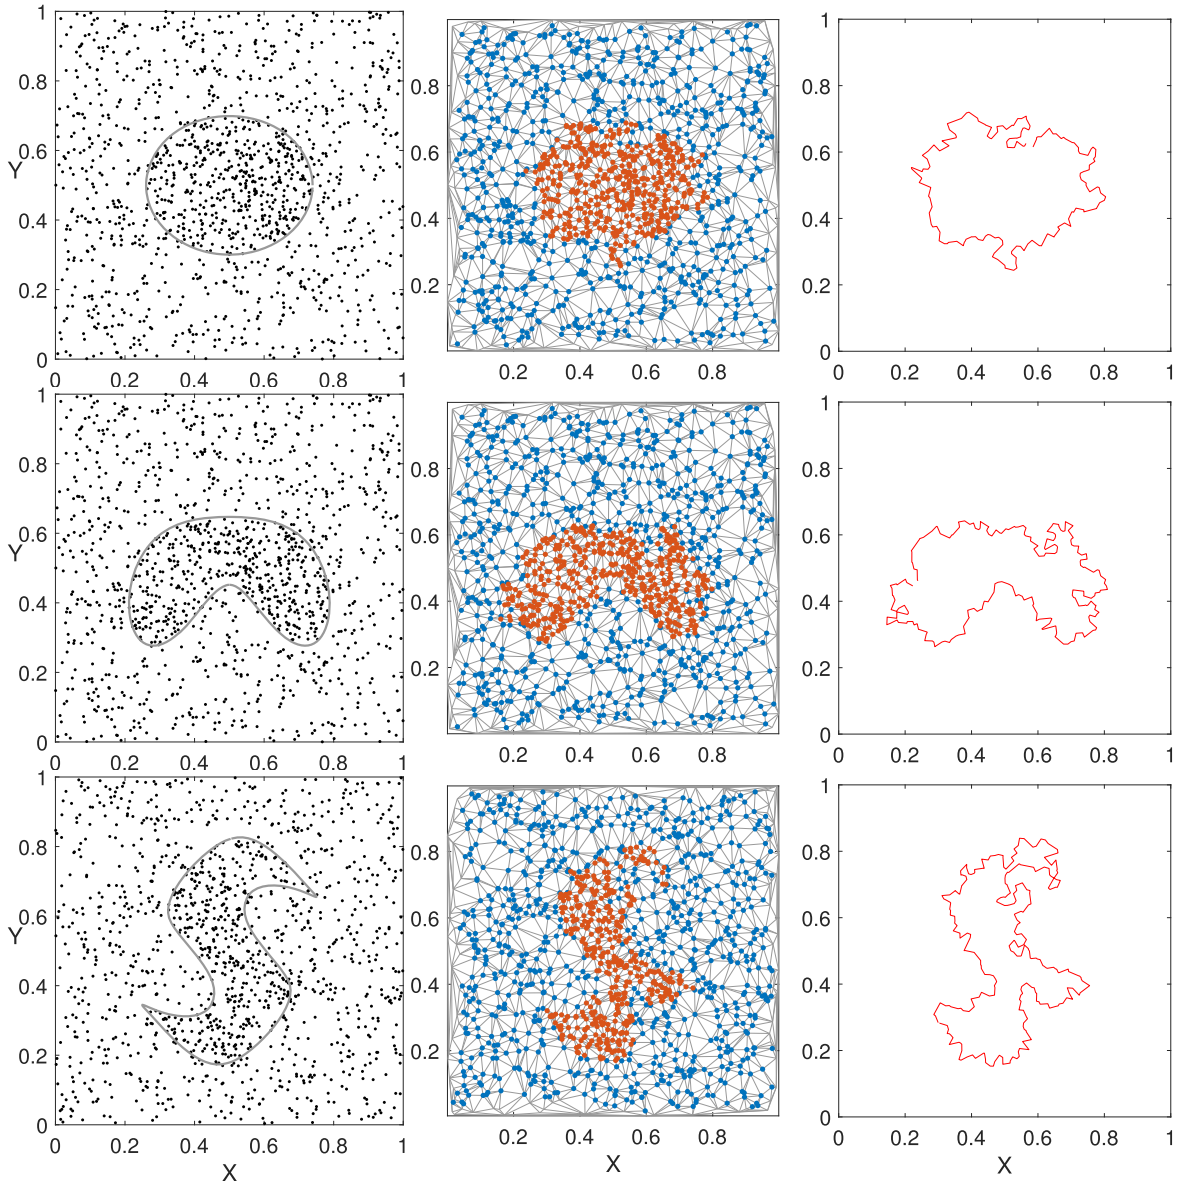


Figure 6. Simulation experiments with various extended sources. Each extended source has a size of 0.2, and the expected number of photons from the background and the source is 1000 and 300, respectively. Left panel: simulated observations whose ground truth of source boundaries is modeled with FD parameterization with $J = 3, 7, 11$ from top to bottom. Middle panel: `ESRGonG` two-region segmentation results. Right panel: segment's boundaries extracted from the segmentation.

In the `ESRGonG` steps, we set the seeds via Voronoi-area stratified sampling (M. Fan et al. 2023), construct the over-segmented event list, and apply the hybrid merging process with $M_s = 150$ randomized merger steps, and apply additional beam merger step with $M_B = 20$ to derive the two-region segmentation. We repeat the procedure with 24 combinations of the Voronoi-area stratified sampling seeding parameters, specifically, six values for the number strata, i.e., $m_{\rm R} = [5,8,11,14,17,20]$ crossed with four values for the minimum number of photons required for a Voronoi-area seed to be accepted, i.e., $m_{\rm vorthr} = [5,10,15,20]$. We choose the ($m_{\rm R}$, $m_{\rm vorthr}$) pair with the optimum penalized profile log-likelihood (given in Equation (4)) for its two-region segmentation. We then implement `Auto-BUQ`, including the estimation of J via BIC and the generation of N_B nonparametric bootstrap data sets as described in Section 4.3 and 4.5.2. By applying `Auto-BUQ` on each of the N_B resampled data sets, we obtain the global confidence region for the source boundary as described in Algorithm 1.

6.1. NGC 2300

NGC 2300 is an early-type galaxy residing in the eponymous NGC 2300 group. The surface brightness of its X-ray emission is fairly smooth and follows a β profile (D. S. Davis et al. 1996). This makes it an excellent test case for identifying the boundary of extended sources such as early-type galaxies, galaxy groups, or higher-redshift galaxy clusters. Since it has been observed with both Chandra and XMM-Newton, it allows us to compare `Auto-BUQ` for two data sets with different spatial resolutions, pixellations, effective areas, and background levels.

We show the result of application of `SRGonG` (M. Fan et al. 2023) to both the Chandra and XMM-Newton data sets of

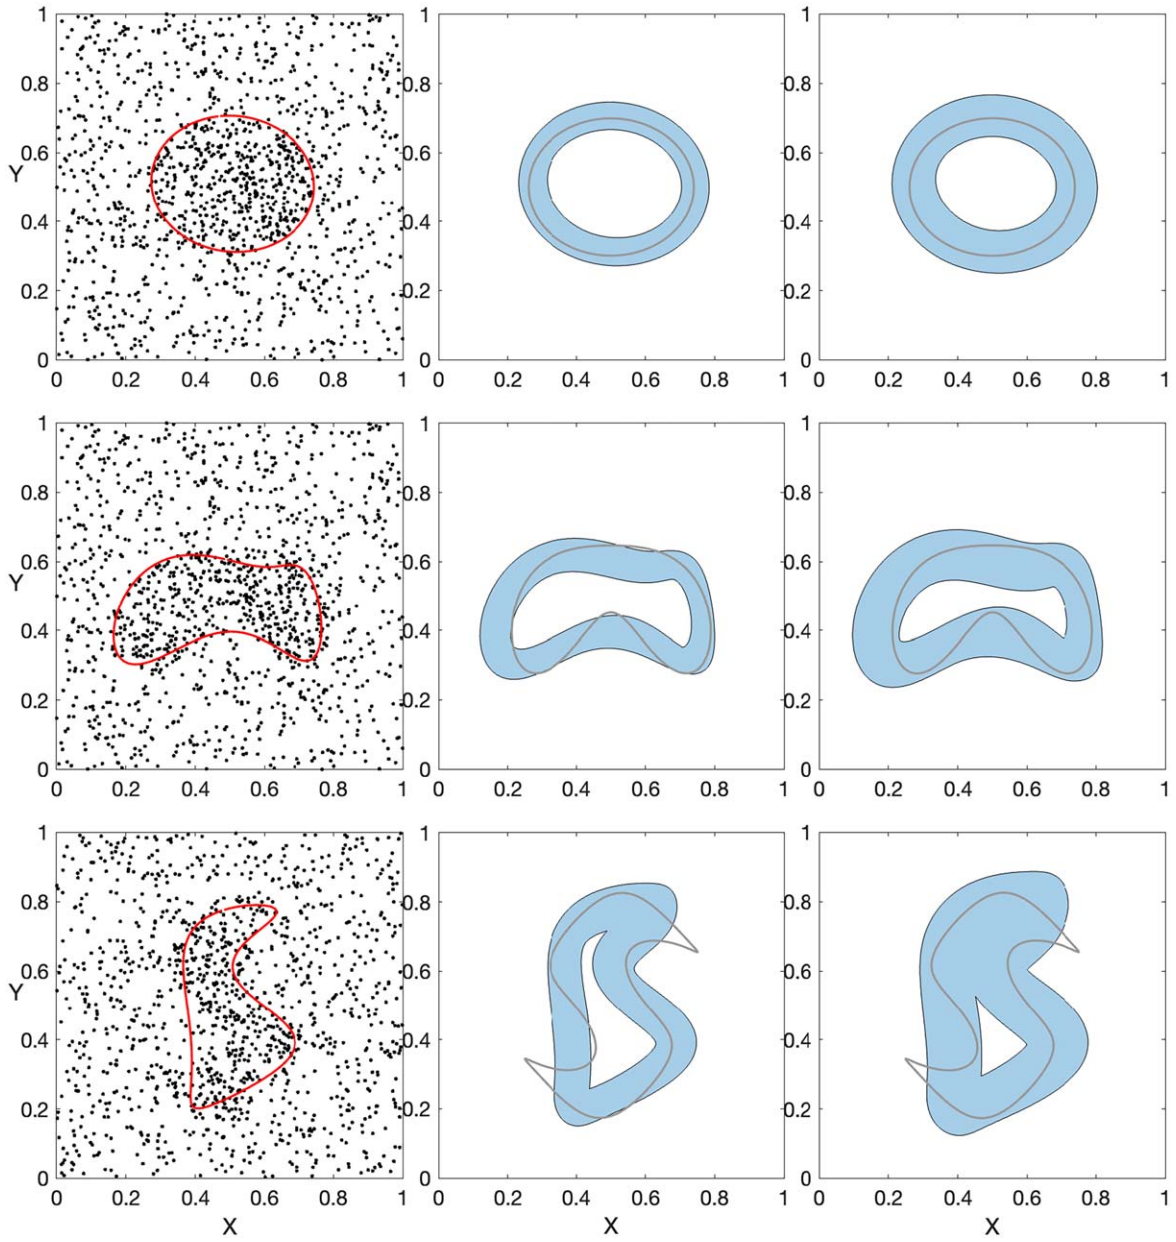


Figure 7. Simulation experiments on ellipse, bridge, and spiral shape. Left panels: estimated extended source boundary $s(t; \hat{\boldsymbol{\theta}}, \hat{J}_{\mathrm{BIC}})$ with $\hat{J}_{\mathrm{BIC}} = 3, 7, 7$ for the three shapes, respectively. Middle panel: global confidence region $\hat{U}_\gamma$ with $\gamma = 1$, corresponding to $\alpha \approx 0.39$. Right panel: global confidence region $\hat{U}_\gamma$ with $\gamma = 1.51$, corresponding to $\alpha \approx 0.68$.

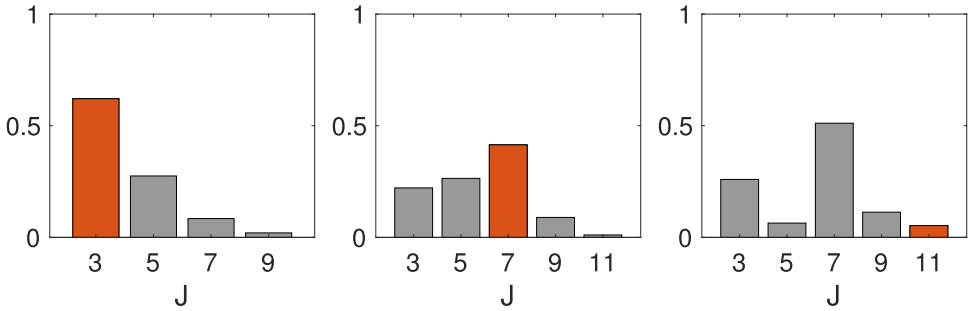


Figure 8. Simulation studies with various shapes of extended sources. The histogram of BIC fitted value $\hat{J}$ for 300 replicates with the true number of J highlighted.

Table 3
Data Sets Used to Illustrate the Application of Auto-BUQ

Object	Instrument	Observation ID	Dimensions of the Field of View $\mathcal{F}$			Counts in $\mathcal{F}$	Exposure Time
			Center (pix)	R.A. (width)	Decl. (height)		
NGC 2300	Chandra/ACIS-S	04968	(3415.0, 3830.0)	347″	335″	11750	45.6 ks
NGC 2300	XMM/MOS1	0022340201	(26215.1, 28185.2)	240″	240″	8800	54.2 ks
Arp 299	XMM/MOS1	0112810101	(25160.5, 27843.5)	352″	352″	6450	21.6 ks

NGC 2300 in Figure 9. The analysis of the Chandra data set is shown in the upper row and that of XMM-Newton is shown in the lower row. Both data sets are filtered to limit the photon energies to the range 0.5–8 keV. The observed photon locations are shown for both data sets in the left column, zoomed into a window of $120'' \times 120''$. The Voronoi tessellation and the adopted seeds, obtained using Voronoi-area stratified sampling (see M. Fan et al. 2023) are shown in the second column, and the SRGonG segmentations in the third column, obtained for $M_S = 80$ and 53, and resulting in $\widehat{K} = 7$ and five segments, respectively, for Chandra and XMM-Newton. The right column shows, as the red curve, the boundary of the segment selected for ESRGonG analysis after merging to two segments. The Auto-BUQ analyses are shown in Figure 10: the left column shows the FD parameterization as the smooth red curve superposed over the segment of interest shaded in blue; the middle column shows the original segment boundary in red superposed on uncertainty bands for $\gamma = 1, 2$ (corresponding to $\alpha = 0.39, 0.89$; see Figure 7) shaded in purple; and the third column shows 10 example bootstrap sample segments illustrating the uncertainty in the boundary.

Algorithm 1. Construct global confidence region for $s^\star$

Require: Coordinates of observed photons $\mathbf{X}$ of size n; γ that corresponds to confidence level of coverage.
Ensure: $\hat{U}_\gamma$, global confidence region of extended object boundary $s(t, \boldsymbol{\theta}^\star, J^\star)$, $t = [0, 1]$.
1: Derive $\hat{\mathbf{S}}$ (full segments) and $\tilde{S}$ (two segments) by applying a segmentation algorithm to input data $\boldsymbol{X}$. Extract curve C_1 from $\tilde{S}$.
2: Construct an estimator $\hat{\boldsymbol{\theta}}$ with length $2\hat{J} - 1$ by applying DFT to C_1 and a model selection method that estimates J.
3: **For** $N_b \leftarrow 1$ to N_B **do**
4: Generate bootstrap data $\boldsymbol{X}^{(N_b)}$ of size n with suitable bootstrapping procedure.
5: Derive $\tilde{\mathbf{S}}^{(N_b)}$ and extract $C_1^{(N_b)}$ by applying ESRGonG to $\mathbf{X}^{(N_b)}$. Derive $\boldsymbol{\theta}^{(N_b)}$ with length $2\hat{J} - 1$ by applying DFT to $C_1^{(N_b)}$.
6: end for
7: Estimate $\widetilde{\Sigma}_{\boldsymbol{\theta}}$ from $\boldsymbol{\theta}^{(N_b)}$, $N_b = 1,\ldots, B$. Construct $\hat{U}_\gamma(t)$ over $t \in [0, 1]$.
8: **return** $\hat{U}_\gamma \triangleq \bigcup_{t\in[0,1]} \hat{U}_\gamma(t)$.

We show the bootstrap uncertainties on the FD parameters in Figure 11, with the centroid positioning shown in the left column and the semiaxes in the right column. In all cases, the red dots indicate the best-fit value, and the blue crosses indicate bootstrap samples, whose spread is a measure of the uncertainty of the FD components. Notice that the difference in the locations of the centroids is $\approx(7\farcs5)$, consistent with the bootstrap uncertainties (see Table 4). Similarly, estimates of the semimajor and semiminor axes of NGC 2300, as estimated from the $+1$ and -1 components of the FDs, are consistent between the Chandra and XMM-Newton observations (see Table 4).

The astronomical consequences of the Auto-BUQ analysis are described in Figure 12, where the areas enclosed by the estimated boundaries (left column) and the corresponding enclosed background-subtracted counts (right column) are computed and compared for Chandra and XMM-Newton (see Table 4). We emphasize here that our results are not based on setting a threshold on an iso-brightness contour, but on adaptively merging segments that result in the largest contrast between the average brightness of the segments. Notice that, despite the differences in PSF sizes, background levels, and instrument sensitivities, the estimated boundaries enclose similar areas in both the Chandra and XMM-Newton data. Our algorithm finds a similar boundary for NGC 2300 in both data sets. We attribute this to both data sets having similar data quality. Thus, while a consistent determination of boundaries is not expected in observations with strongly variable characteristics, over the spatial scale and spectral response of the two instruments, we are able to obtain similar estimates of the diffuse boundary of NGC 2300 as distinguishable from the background in both cases. The effect of the higher background in XMM-Newton is seen in the form of a looser upper bound on the estimate; the higher background makes it easier for the segment boundary to expand into the area of the background due to local fluctuations. The observed ratio of count rates within the segments, $\frac{\text{XMM}}{\text{Chandra}} = 1.36 \pm 0.25$, is consistent with the values obtained for typical spectra associated with NGC 2300: assuming an absorption column of $N_{\rm H} = 4 \times 10^{20}\,\text{cm}^{-2}$, along with a power law with index $\Gamma = 1.4$–2.7, or a thermal plasma model with temperature $\text{kT} = 0.9$ keV (see D. S. Davis et al. 1996), we find relative model count rate ratios of $\frac{\text{XMM}}{\text{Chandra}} \approx 1.04 - 1.24$.

6.2. Arp 299

Arp 299 is a nearby interacting system of galaxies, also studied in detail with Chandra and XMM-Newton. In contrast to NGC 2300, it consists of two main components that are connected and a few discrete X-ray sources (A. Zezas et al. 2003; L. Ballo et al. 2004; K. Anastasopoulou et al. 2016). However, its X-ray emission is not symmetric and exhibits complex features. Chandra identifies multiple pointlike X-ray sources (see, e.g., K. Anastasopoulou et al. 2016), the analysis of which is beyond the scope of this work. Instead, we apply ESRGonG and Auto-BUQ to the XMM-Newton data as a demonstration of estimating source extent despite the presence of resolvable complex morphology.

We carry out a similar analysis for Arp 299 as we do for NGC 2300 in Section 6.1. The SRGonG and ESRGonG analysis steps are shown in Figure 13, using $M_S = 82$ seeds and resulting in $\widehat{K} = 13$ segments. The Auto-BUQ analysis

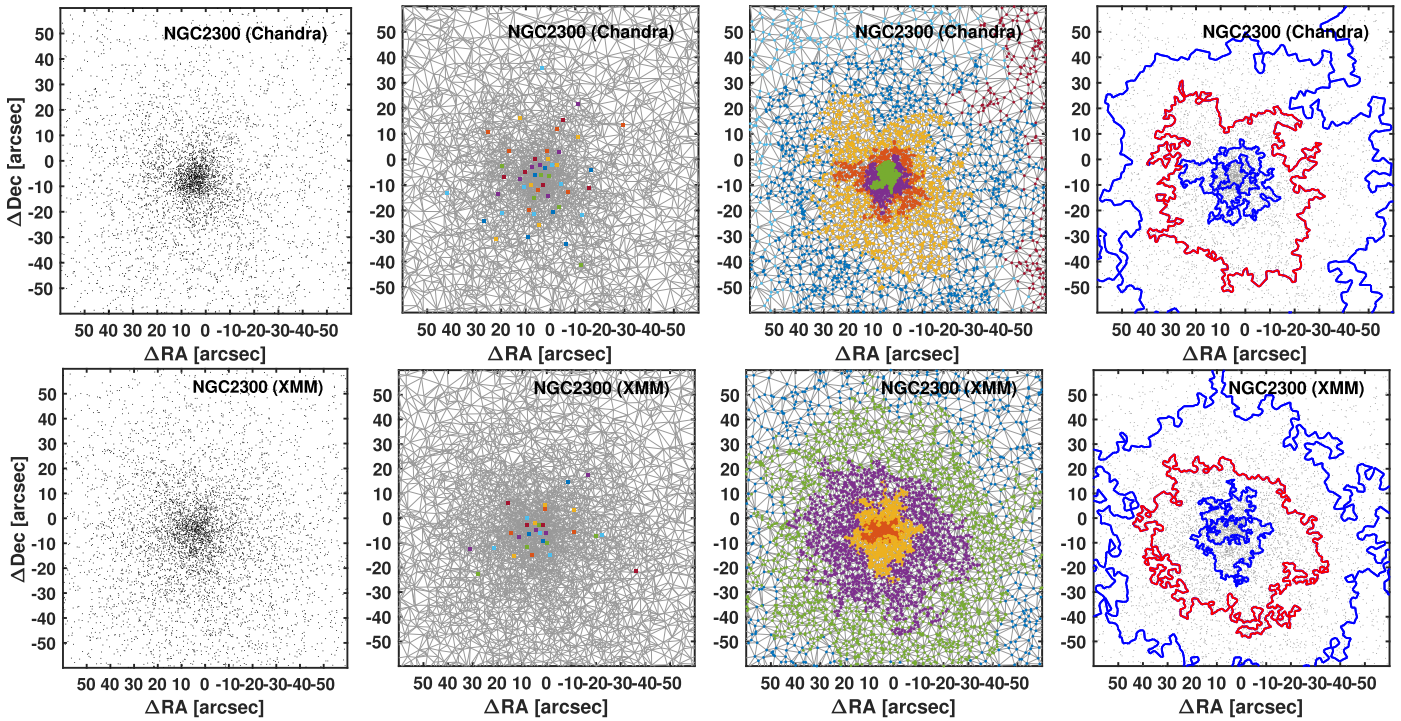


Figure 9. Representation of the application of SRGonG on NGC 2300 for Chandra (top panels) and XMM-Newton (bottom panels) data, each centered at (R.A., decl.) = (7:32:16.0517, +85:42:39.101). Plots in succeeding columns show different stages of the SRGonG analysis (Auto-BUQ steps are shown in Figure 10). Leftmost panels: scatter plots of the photons from ObsID 4968 (Chandra; ≈12,000 counts) and OBS_ID 0022340201 (XMM; ≈9000 counts), both filtered to include photon energies in the range 0.5–8 keV. Center-left panels: seed specification for SRGonG via nearest-neighbor labeling methods proposed in M. Fan et al. (2023). The hyperparameter pair (m_{R}, m_{vorthr}) is tuned via a grid search to give the maximum penalized profile log-likelihood after merging to two segments. We adopt $M_S = 80$ ($m_{\mathrm{R}} = 14$, $m_{\mathrm{vorthr}} = 15$) for Chandra and $M_S = 53$ ($m_{\mathrm{R}} = 14$, $m_{\mathrm{vorthr}} = 20$) for XMM-Newton. Center-right panels: raw SRGonG segmentation $\widehat{\boldsymbol{S}}$, with segments $\widehat{K} = 7$ for Chandra and $\widehat{K} = 5$ for XMM-Newton. Rightmost panels: boundaries $\widehat{\mathcal{S}}_k$ of all segments $\widehat{\boldsymbol{S}}$. The boundary curve C_1 extracted from the ESRGonG two-region segment $\widetilde{\boldsymbol{S}}$ is highlighted in red.

steps are illustrated in Figure 14. The centroid and semiaxes uncertainties are depicted in Figure 15, and the distributions of the area covered by the galaxies and the associated background-subtracted counts are in Figure 16. The estimated values and uncertainties are also reported in Table 4.

As illustrated in Figures 13 and 14, the merging of segments until only two are left successfully separates Arp 299 from the background and allows us to estimate the area covered by the system. The area of the system is ≈3230 arcsec^2, significantly more compact than the area of ≈12,000 arcsec^2 suggested by the D_{25} diameter of 122″ (G. de Vaucouleurs et al. 1991; H. G. j. Corwin et al. 1994). This indicates that integrated flux and mass estimates based on X-ray data, which rely on the nominal D_{25} values, are likely contaminated by background emission. Note further that the distribution of the areas in Figure 16 is multimodal, with the mode at smaller areas preferentially selecting the inner segment during bootstrapping. This suggests that the X-ray isophote of the Arp 299 system is more compact than the optical. This analysis shows the power of SRGonG and Auto-BUQ in determining the outline of sources with complex shapes, aiding their photometric analysis without relying on apertures defined in other parts of the spectrum.

7. Summary

We have developed a principled method to automatically determine the boundary of extended high-energy astronomical objects, and compute a measure for uncertainty on the boundary. The scheme minimizes human oversight and decision-making, resulting in estimates of boundaries that objectively characterize the sources. It is designed to work on photon event lists from diffuse objects, which may have complex internal structures, but are not disjoint, and are fully contained within the field of view.

The structure of the method is shown in Figure 17. We adopt and extend the SRGonG method (Seeded Region Growing on Graph; M. Fan et al. 2023), which generates well-defined segments (see review in Section 2.1) in X-ray and γ-ray photon event lists. Additionally, we incorporate a randomized merging step to minimize the effects of the greedy merging implemented in SRGonG (see Section 2.2). Next, a beam merging step is applied to the SRGonG segmentation to merge segments until only two are left: the source and a background (see Section 4.1). This allows us to obtain a well-defined outer boundary of the extended source, and forms the core of the ESRGonG (Boundary of Segmented Events) algorithm. Note that in principle any of the segment boundaries resulting from SRGonG may be analyzed using the subsequent steps, but that requires human intervention for selecting a particular segment and for defining unique branch-free segments to traverse its boundary, and thus we limit our investigations to the simpler case of an isolated diffuse structure whose boundary can be automatically extracted.

In order to quantify the uncertainty on the boundary thus obtained, we model it using FDs (see review in Section 3.1 and implementation in Section 4.2). We estimate the number of components in the FD via BIC (see Section 4.3), which allows us to construct global confidence regions via bootstrapping: the event lists are bootstrapped by sampling with

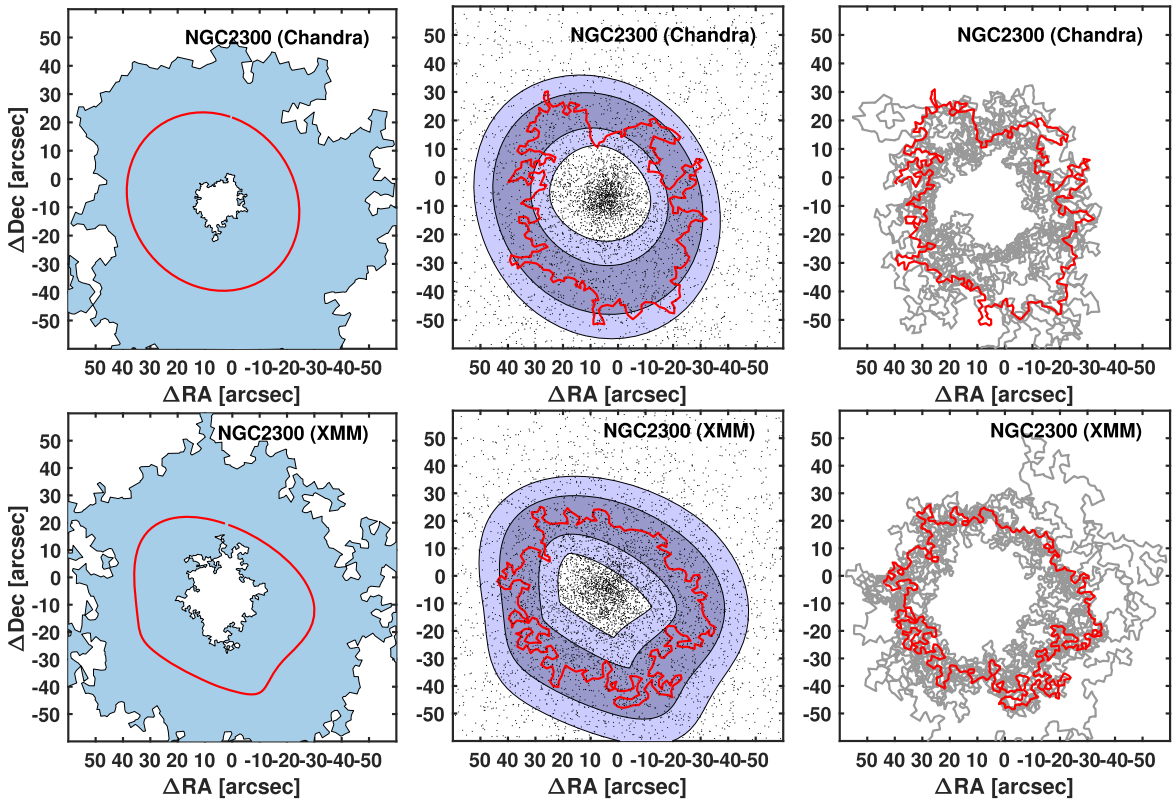


Figure 10. As in Figure 9, but illustrating the steps in `Auto-BUQ` analysis, showing the parameterization and uncertainty quantification of the boundary of the source. Left panels: the red curve shows the FD parameterization of C_1 with $\hat{J}$ estimated by BIC ($\hat{J}_{\rm BIC} = 3$ for Chandra, and $\hat{J}_{\rm BIC} = 7$ for XMM-Newton). The blue region shows $\mathcal{R}$, the smallest region that fully encloses C_1 identified within $\hat{S}$. Middle panels: $\hat{U}_\gamma$, the FD-based global confidence region constructed using nonparametric bootstrap data. The purple band shows $\hat{U}_1$; the light-purple band shows $\hat{U}_2$; and the red curve shows C. Right panels: illustrations of `ESRGonG` applied to 10 nonparametric bootstrap data sets. The red curve shows C_1, and the gray curves show $C_1^{(N_b)}$ for $N_b = 1, \ldots, 10$.

replacement, and `ESRGonG` is run on each instance to extract a boundary from the forced two-region segmentation (see Sections 4.4 and 4.5). We use the bootstrap results to construct astronomically useful quantities like the centroid and semiaxes of the source, as well as the area covered and the counts enclosed (Section 4.6). These tasks together form the `Auto-BUQ` algorithm.

We validate the algorithm via simulations and demonstrate that we obtain robust results in a wide variety of situations, and obtain realistic uncertainty regions even for complex shapes (Section 5). We then apply it to several real data sets: first, for a simple galaxy NGC 2300 observed with both Chandra and XMM-Newton where we demonstrate that `Auto-BUQ` obtains consistent estimates of the size and brightness of the galaxy (Section 6.1), and second, for the galaxy merger Arp 299 observed with XMM-Newton where we obtain a well-defined compact outer boundary in X-rays despite the presence of significant internal structure (Section 6.2).

Our current work is designed with a specific focus on conducting statistical inference for the external boundary of a single extended source within the field of view $\mathcal{F}$. We assume that the external boundary of this extended source is fully contained within $\mathcal{F}$. While it is recognized that FD-based parameterization is inherently well suited for addressing closed contours (W. Burger & M. J. Burge 2013), it is pertinent to note that this parameterization framework can also be adapted to accommodate a partial two-dimensional shape parameterization, as discussed in C. C. Lin & R. Chellappa (1987). One avenue of future research expansion entails the extension of our uncertainty quantification methodology to scenarios where the underlying extended source is only partially contained within the field of view. In such cases, inference methodology can be applied to complete the boundary first, as proposed by C. C. Lin & R. Chellappa (1987), and then quantify the uncertainty of the inferred shape.

In this paper, we extend the greedy algorithms proposed in M. Fan et al. (2023) to improve the quality of the resulting segmentation. An immediate question is whether it is possible to avoid a greedy search and obtain a global optimizer of the penalized profile log-likelihood in (3). J. D. Scargle et al. (2013) developed an innovative, dynamic programming strategy for a similar optimization involving one-dimensional time series data. M. J. Way et al. (2010) applied this strategy to two-dimensional image data via a translation to a one-dimensional orderable list of Voronoi cells based on their area. This list is segmented by grouping cells based on their area with no constraint on the resulting segments being connected. Disconnected segments, however, can be separated into subsegments that are connected, thus forming the final connected segmentation. This final setup increases the number of segments and, in our setting, would break the link between the BIC penalty and the ultimate complexity of the segmentation. This leads to various challenges. For example, the merge-to-two step could result in more than two segments. Nonetheless, exploring how the highly efficient algorithm of

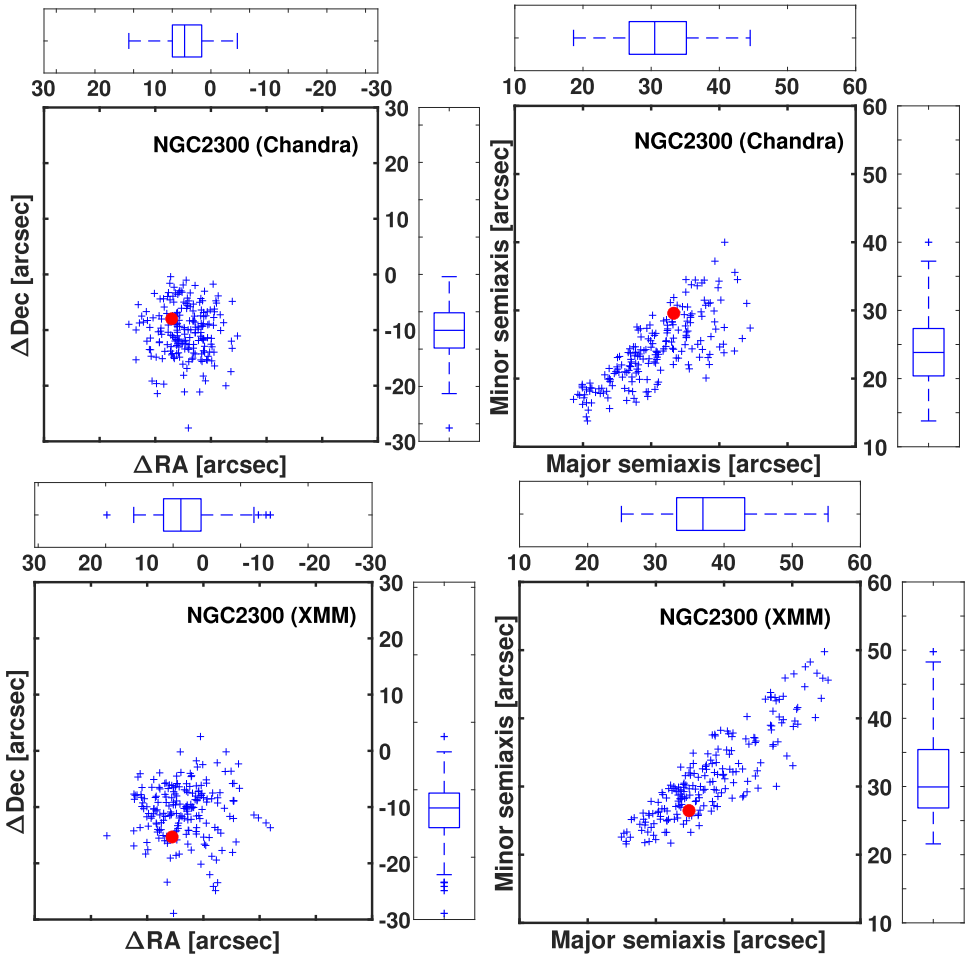


Figure 11. Quantification of the uncertainty on the first two FD components of the boundary of NGC 2300 as in Figure 10. Scatter plots of the centroid (left panels) and semiaxes (right panels) obtained for both data (red dots) and bootstraps (blue crosses) are shown. The box plots along the sides show the marginalized uncertainties of the bootstrapped distributions, with the bars indicating quartiles (the 50% quartile is marked in red, and the 25%–75% range is marked by the box), and the whiskers covering the full range of the bootstraps. Left panels: centroid of C_1, $(\mathrm{Re}(G_0),\ \mathrm{Im}(G_0))$ (red dot) and $(\mathrm{Re}(G_0^{(N_b)}),\ \mathrm{Im}(G_0^{(N_b)}))$ (blue crosses). Right panels: semiaxes (major semiaxis along the abscissa and minor semiaxis along the ordinate) for the data C_1, $(|G_{-1}| + |G_1|,\ ||G_{-1}|-|G_1||)$ (red dots) and for bootstraps $C_1^{(N_b)}$, $(|G_{-1}^{(N_b)}| + |G_1^{(N_b)}|,\ ||G_{-1}^{(N_b)}|-|G_1^{(N_b)}||)$ (blue crosses).

Table 4
Estimated Parameters of FD Fits

Data Set	Centroid[a] (RA,Dec)	Semiaxis[a] Major	Semiaxis[a] Minor	Area[a,b] (arcsec^2)	Counts[a,b] (counts)
Chandra/NGC 2300	(7:32:22.35 ± 3.″83, +85:42:31.13 ± 4.″58)	33.″3 ± 5.″99	29.″6 ± 4.″95	$3198^{<3430}_{>1555}$	$2668^{<2720}_{>2120}$
XMM/NGC 2300	(7:32:21.02 ± 5.″07, +85:42:23.80 ± 4.″81)	34.″9 ± 7.″19	26.″5 ± 6.″32	$3154^{<5660}_{>2590}$	$3093^{<3795}_{>2865}$
XMM/Arp 299	(11:28:32.35 ± 4.″06, +58:33:49.2 ± 2.″73)	32.″7 ± 5.″27	30.″6 ± 5.″02	$3231^{<4235}_{>2080}$	$3434^{<3655}_{>3045}$

Notes.
[a] Estimates, with uncertainties calculated via bootstrapping.
[b] Asymmetric (>)16% and (<)84% quantile bounds.

J. D. Scargle et al. (2013) might be applied in ESRGonG is an important area of future work.

While we acknowledge that real-world scenarios often involve the presence of multiple diffusive objects, our approach offers a foundational framework for statistical inference on individual objects, laying the groundwork for future developments in handling multiple segments within a unified statistical framework. One of the principal challenges impeding the development of an inference tool capable of handling multiple segments lies in the difficulty of aligning segmentations across various bootstrap data sets. Our observations indicate that, particularly in scenarios characterized by low S/Ns, the SRGonG algorithm may result in the segmentation of a single diffusive region into multiple segments or merge it with other segments. This introduces ambiguity in identifying segments that estimate the same region for further analysis. Addressing this challenge necessitates enhancing the robustness of the SRGonG algorithm to manage such scenarios effectively.

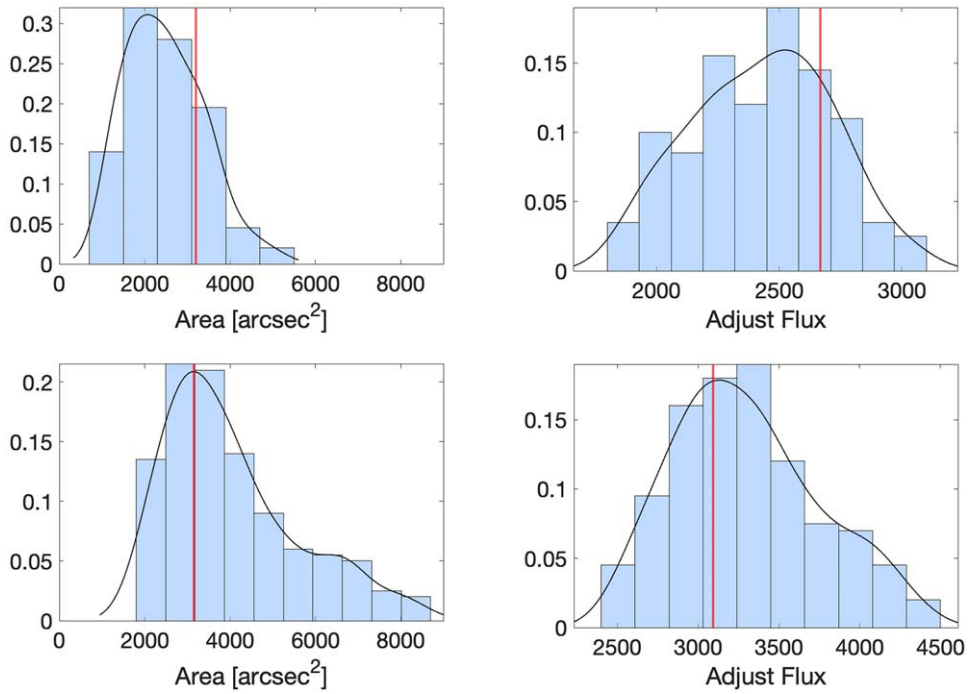


Figure 12. Quantification of the uncertainty in astronomically relevant estimates for NGC 2300, based on `Auto-BUQ` analysis shown in Figure 10 (top: Chandra; bottom: XMM-Newton). All panels show the frequency histogram based on bootstraps (blue shaded bars), the smoothed density (black curves), and the value estimated from the data (thin red vertical line). Left panels: the distributions of the area covered by NGC 2300, Area($\widetilde{\mathcal{S}}_1^{(N_b)}$). Right panels: the distribution of the background-subtracted counts contained within the selected areas, Num($\widetilde{\mathcal{S}}_1^{(N_b)}$) $-$ $\frac{\text{Area}(\widetilde{\mathcal{S}}_1^{(N_b)})\text{Num}(\widetilde{\mathcal{S}}_2^{(N_b)})}{\text{Area}(\widetilde{\mathcal{S}}_2^{(N_b)})}$. The estimated counts are consistent between the two observations (see the text).

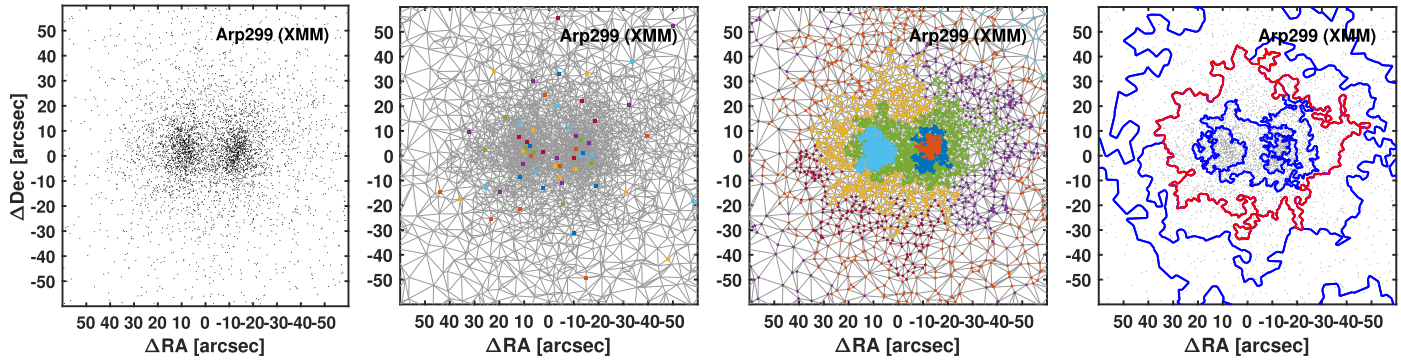


Figure 13. The same as in Figure 9, `SRGonG` analysis for an XMM-Newton observation of Arp 299. The center of the field of view is at (R.A., decl.) = (11:28:33.0292, +58:33:43.942). The adopted hyperparameter is $M_S = 82$ ($m_R = 17$,$m_{\text{vorthr}} = 15$) and results in $\widehat{K} = 13$ segments.

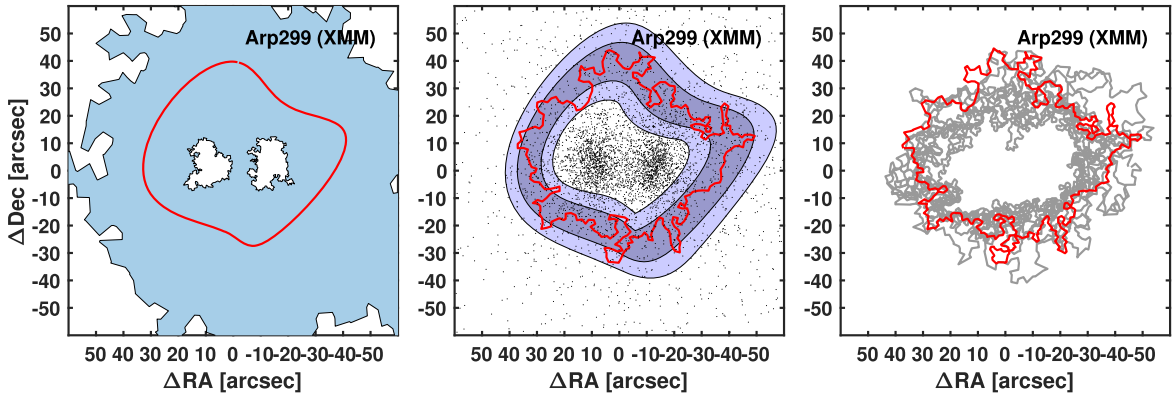


Figure 14. The same as in Figure 10, for `Auto-BUQ` analysis of XMM-Newton data set of Arp 299.

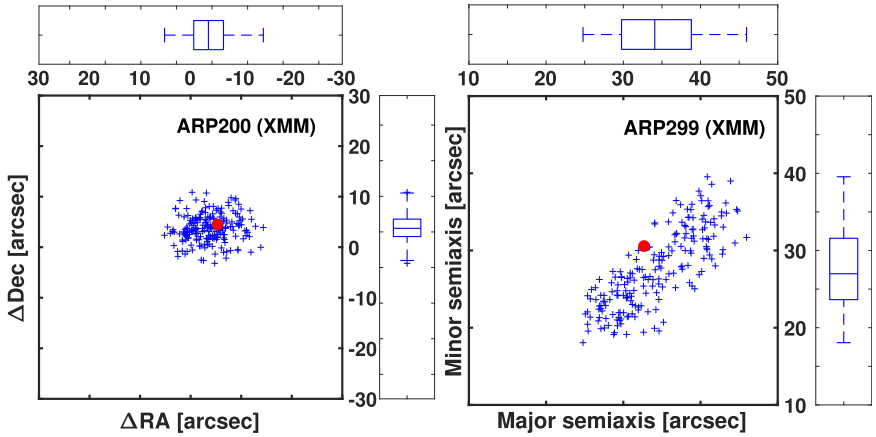


Figure 15. The same as in Figure 11, quantifying the uncertainty on centroid and semiaxes for XMM-Newton data set of Arp 299.

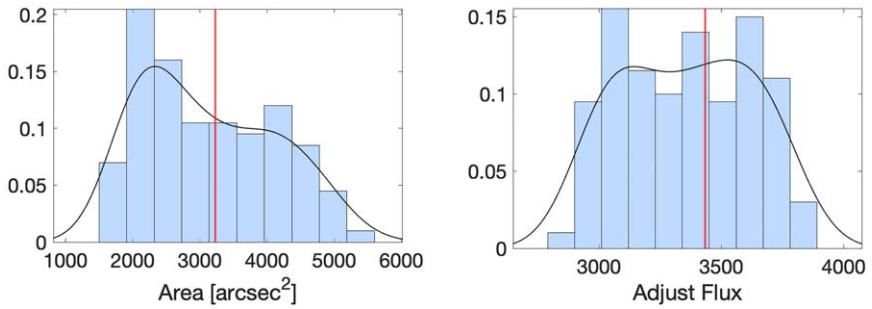


Figure 16. The same as in Figure 12, quantifying the uncertainty on the area and enclosed background-subtracted counts for XMM-Newton data set of Arp 299.

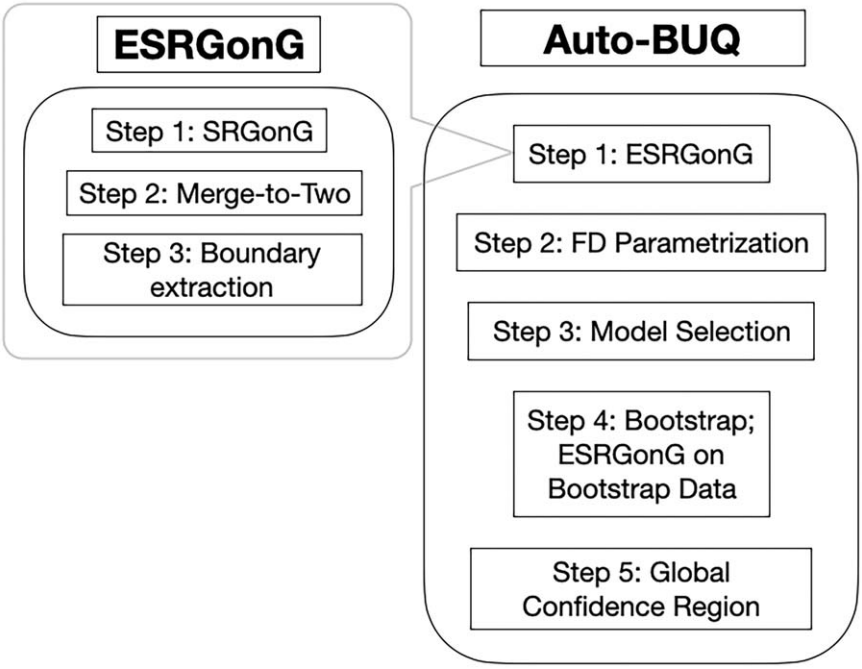


Figure 17. Flowchart of `Auto-BUQ`, showing how the expanded version of `SRGonG` (labeled `ESRGonG`) is incorporated into the algorithm. `ESRGonG` is used to estimate a boundary, and `Auto-BUQ` is used to quantify the uncertainty on the boundary.

Acknowledgments

This work was conducted under the auspices of the International CHASC AstroStatistics Center. CHASC is supported by NSF grants DMS-21-13397, DMS-21-13605, DMS-21-13615, and DMS-22-10388; by the UK Engineering and Physical Sciences Research Council [EP/W015080/1]; and by NASA APRA grant 80-NSSC21-K0285. D.V.D and A.Z. were also supported in part by a Marie-Skodowska-Curie RISE grant (H2020-MSCA-RISE-2019-873089) provided by the European Commission. V.L.K. acknowledges support from NASA contract to the Chandra X-ray

Center NAS8-03060. We thank CHASC members for many helpful discussions. This paper uses Chandra data sets, obtained by the Chandra X-ray Observatory, contained in the Chandra Data Collection (CDC) "359" doi:10.25574/cdc.359. It also uses XMM-Newton data sets, which are available via XMM-*Newton* Science Archive (XSA).

Facility: Chandra, XMM-Newton.

Software: MATLAB, `SRGonG` (M. Fan et al. 2023), `Auto-BUQ` https://github.com/jujWang96/UQ-BoSE (J. Wang 2025).

Appendix A
The Choice of M_s

In our investigation of the algorithms proposed to merge the over-segmented event list, we have found that as K decreases, the magnitude of change in the unpenalized profile log-likelihood in Equation (4) due to a merger operation exhibits a pronounced amplification. This is illustrated in Figure 18. This observation motivated our development of an ad hoc method for setting M_s, the number of random/beam merging steps in the hybrid merging process, so that we achieve a balance between computation time and segmentation quality by setting M_s sufficiently large to generate a diverse set of segmentations while avoiding unnecessary exploration by keeping M_s within reasonable bounds.

The setting of turning point grounds in the evaluation of this log-likelihood variation. In other words, we propose avoiding the randomized and/or beam merger steps until the greedy merger step yields a significant increase in the unpenalized profile log-likelihood. Operationally, we evaluate the change in profile log-likelihood, denoted as $-2\Delta\mathcal{L}_{\text{profile}}(K) = -2(\mathcal{L}_{\text{profile}}(K-1, \boldsymbol{S}(K-1)|\boldsymbol{X}) - \mathcal{L}_{\text{profile}}(K, \boldsymbol{S}(K)|\boldsymbol{X}))$ at each iteration of the greedy merger step. We select M_s as the maximum K value satisfying $-2\Delta\mathcal{L}_{\text{profile}}(K) > \delta$, where δ is a hyperparameter, which typically should be chosen within the range [1, 2]. In Figure 18, we plot $-2\Delta\mathcal{L}_{\text{profile}}(K)$ to illustrate the value of M_s chosen for the Arp 299 XMM data and the sensitivity of the number of segments on the choice of the value for δ. Notice that the dependence of M_s on δ is weak. Our ad hoc selection yields a value of $M_s = 132$ for $\delta = 2$, a choice that is not far from our default setting of 150.

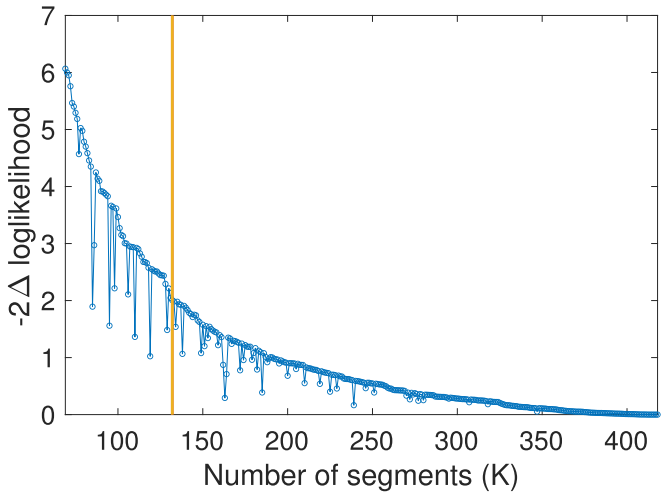


Figure 18. Ad hoc selection procedure of M_s for Arp 299 XMM-Newton data. For illustrative purposes, we choose the hyperparameter pair (m_{R}, m_{vorthr}) for the seed specification to be (17,5), leading to the largest number of seed $M_S = 418$ among all pairs. Blue line with dots: $-2\Delta\mathcal{L}_{\text{profile}}(K)$ vs. K; Yellow line: maximum K value (132) such that $-2\Delta\mathcal{L}_{\text{profile}}(K) > 2$.

Appendix B
Model Selection Pitfall When Enclosure Condition Is Violated

In this section, we will construct a simple example to illustrate the potential issue of fitting J when the enclosure assumption is violated. Consider the example in the left panel of Figure 19. We parameterize the object boundary C, highlighted by the red curve, with FD and aim to fit the number of FD with the BIC. The enclosure assumption is violated because partial C is not an interior point of $\mathcal{R}$, the blue segment. In the right panel of Figure 19, we display the parameterized function $s(t; \hat{\theta}, J)$ with $J = 2$. Notably, the parameterized curve now intersects with the boundary of $\mathcal{R}$ and creates an extra region. The presence of extra regions introduces ambiguity during the derivation of the profile log-likelihood function for the corresponding segmentation due to the undefined intensity within these additional segments.

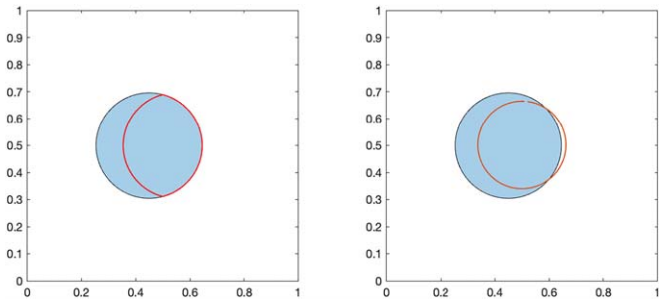


Figure 19. Illustration of pitfall caused when enclosure assumption is violated. Left panel: the red curve shows the object boundary C. The blue segment shows a construction of $\mathcal{R}$ that violates the enclosure assumption, as partial C is not interior point of $\mathcal{R}$. Right panels: the red curve shows the parameterized function $s(t; \hat{\theta}, J)$ with $J = 2$. Its intersection with the boundary of $\mathcal{R}$ creates an extra region.

Appendix C
Matrix Form of Fourier Parameterization

This section provides the construction of $V_J(t)$, outlined in Equation (11) for any $2 \leqslant J \leqslant L$.

$$V_J(t)^T = \begin{bmatrix} a(t)\, c(t) \\ b(t)\, d(t) \end{bmatrix} \tag{C1}$$

where the elements of the matrix $V_J(t)$ are themselves row matrices of the form

$$\begin{aligned}
a(t) &= [\cos(-2\pi(\lfloor (J-1)/2 \rfloor t) \ \cdots \ \cos(-2\pi t)\ 1\ \cos(2\pi t) \\
&\quad \cdots \cos(2\pi \lfloor J/2 \rfloor t)] \in \mathbb{R}^J \\
b(t) &= [\sin(-2\pi(\lfloor (J-1)/2 \rfloor t) \ \cdots \ \sin(-2\pi t)\ 1\ \sin(2\pi t) \\
&\quad \cdots \sin(2\pi \lfloor J/2 \rfloor t)] \in \mathbb{R}^J \\
c(t) &= [-\sin(-2\pi(\lfloor (J-1)/2 \rfloor t) \cdots \\
&\quad - \sin(-2\pi t) \ - 1 \ - \sin(4\pi t) \\
&\quad \cdots -\sin(2\pi \lfloor J/2 \rfloor t)] \in \mathbb{R}^{J-1} \\
d(t) &= [\cos(-2\pi(\lfloor (J-1)/2 \rfloor t) \ \cdots \ \cos(-2\pi t)\ 1\ \cos(4\pi t) \\
&\quad \cdots \cos(2\pi \lfloor J/2 \rfloor t)] \in \mathbb{R}^{J-1}.
\end{aligned}$$

Appendix D
Lower Bounds for the Global Confidence Level of FD-based Confidence Region

In this section, we delve into the examination of the global confidence level of FD-based confidence region constructed using Algorithm 1, i.e., $\Pr(s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma, \forall\, t)$. We introduce two lower bounds for $\Pr(s(t;\boldsymbol{\theta}, J^\star)\in\hat{U}_\gamma, \forall\, t)$ and their empirical estimations when the exact value is not directly accessible due to the absence of the true covariance $\Sigma_{\boldsymbol{\theta}}$. Additionally, we conduct a comparative analysis of the empirical global confidence level of $\hat{U}_\gamma$ across various selections of γ within the framework of three simulation scenarios, as outlined in Section 5.2. To evaluate the performance of the procedure in Section 4.4, we simulate Monte Carlo data to estimate the true crossing-level probability and compare it with the estimated theoretical bounds. We generate synthetic data with various shapes of extended source $s(t;\boldsymbol{\theta}^\star, J^\star)$, whose complexity is measured by $J^\star$. For a fixed γ, we construct a global confidence region $\hat{U}_\gamma^{(N_m)}$ for each of the N_M simulated data and estimate the true level-crossing probability, $1-\Pr(s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma, \forall\, t)$ by empirical level-crossing probability,

$$P_e = \frac{1}{N_M} \times \sum_{N_m=1}^{N_M} (1 - \mathbb{1}\{s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma^{(N_m)}, \forall\, t\}). \tag{D1}$$

In the construction of $\hat{U}$, we consider two scenarios, where J is true $J^\star$, and J is estimated with the BIC. We denote the empirical level-crossing probability as P_e and P_e^{BIC} when $J=J^\star$ and $J=\hat{J}_{\mathrm{BIC}}$, respectively. We compare the P_e and P_e^{BIC} with the gamma bound and the estimated level-crossing bound.

D.1. Empirical Estimation of the Bounds

The exact value of $\Pr(s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma, \forall\, t)$ is not accessible, but following the analysis in J. C. Ye et al. (2000), we can derive two lower bounds, the incomplete gamma bound and level-crossing bound, for the probability when $J=J^\star$ in the construction of $\hat{U}$. The incomplete gamma bound is written as:

$$\Pr(s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma, \forall\, t\in[0,1]) \geqslant \Gamma\left(\frac{2J^\star-1}{2}, \frac{\gamma^2}{2}\right). \tag{D2}$$

The level-crossing bound is written as:

$$\Pr(s(t;\boldsymbol{\theta}^\star, J^\star)\in\hat{U}_\gamma, \forall\, t\in[0,1]) \geqslant 1-(1+\kappa_o\gamma)\exp\left(-\frac{\gamma^2}{2}\right), \tag{D3}$$

where κ_o is given by

$$\kappa_o \triangleq \int_0^1 \left(\frac{\zeta(t)}{2\pi}\right)^{1/2} dt, \tag{D4}$$

and suppose $\Sigma_{\boldsymbol{\theta}}^{1/2} V_{J^\star}(t)$ has full column rank and let $A=\Sigma_{\boldsymbol{\theta}}^{1/2} V_{J^\star}(t) V_{J^\star}(t)^T \Sigma_{\boldsymbol{\theta}}^{1/2}$,

$$\zeta(t) = \mathrm{trace}(P_A^{\perp}\Sigma_{\boldsymbol{\theta}}^{1/2}\dot{V}_{J^\star}\Sigma_s^{-1}\dot{V}_{J^\star}\Sigma_{\boldsymbol{\theta}}^{1/2}), \tag{D5}$$

where $\dot{V}_{J^\star}$ denotes the derivative of $V_{J^\star}(t)$ with respect to t, and $P_A^{\perp}$ denotes the projection on the orthogonal complement of column span of A. To approximate the level-crossing bound, we estimate the $\Sigma_{\boldsymbol{\theta}}$ in equation (D5) with N_M Monte Carlo simulated data, i.e.,

$$\hat{\Sigma}_{\boldsymbol{\theta}} = \frac{1}{N_M}\sum_{m=1}^{N_M}(\hat{\boldsymbol{\theta}}^{(N_m)}-\bar{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}}^{(N_m)}-\bar{\boldsymbol{\theta}})^T, \tag{D6}$$

where $\bar{\boldsymbol{\theta}} = \frac{1}{N_M}\sum_{N_m=1}^{N_M}\hat{\boldsymbol{\theta}}^{(N_m)}$, and $\hat{\boldsymbol{\theta}}^{(N_m)}\in\mathbb{R}^{2J^\star-1}$ denote the truncated FD parameters of segment's boundary of one set of N_M simulated observations. Let $\hat{\kappa}_0$ denote the estimated κ_0 in Equation (D4); we then have the following two upper bounds for the level-crossing probability P_e:

$$P_{LC} = (1+\hat{\kappa}_o\gamma)\exp\left(-\frac{\gamma^2}{2}\right), \tag{D7}$$

$$P_G = 1-\Gamma\left(\frac{2J^\star-1}{2}, \frac{\gamma^2}{2}\right). \tag{D8}$$

Furthermore, a useful comparison with P_{LC} and P_G is given by:

$$P_l = \exp\left(-\frac{\gamma^2}{2}\right). \tag{D9}$$

According to J. C. Ye et al. (2000), P_l could be an approximation to P_e if the pointwise probability that $s(t;\boldsymbol{\theta}^\star, J^\star)$ lies outside $\hat{U}_\gamma(t)$ was approximately equal to P_e.

D.2. Simulation Results

We use the same simulation setting as in Section 5.2. For each scenario (shape corresponding to different true $J^\star$), we generate 300 replicates and compute the estimated level-crossing bound and gamma bound for each γ value. For each replicate, we construct the global confidence region described in Algorithm 1.

The empirical level-crossing probability and their (estimated) bounds are displayed in Figure 20 for each scenario. We denote the empirical level-crossing probability using BIC model selection by P_e^{BIC} and that using the true number of FD as P_e. The simulation results show that under this particular simulation setting, the empirical level-crossing probability is bounded by both the gamma bound and estimated level-crossing bound from above when $\hat{U}_\gamma$ is constructed with true $J^\star$. For both the ellipse and bridge cases, $\hat{U}_\gamma$ exhibits a high global confidence level. The P_e^{BIC} value is slightly higher than P_e when $J^\star=11$ due to the underfitting of BIC. The performance of $\hat{U}_\gamma$ degenerates under the spiral case mainly due to the challenge of fitting the hook-like shape at both ends of the spiral.

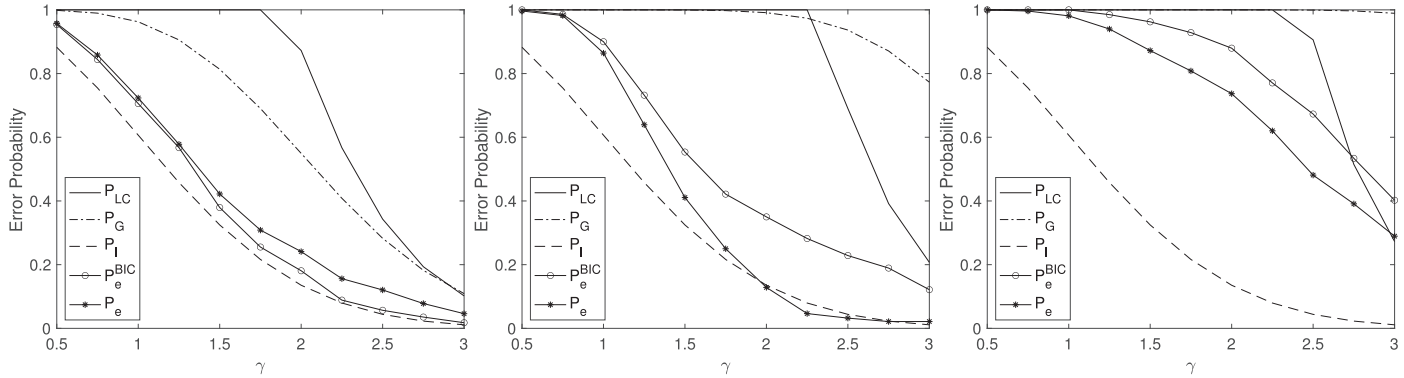


Figure 20. Comparisons of the Monte Carlo simulated empirical level-crossing probability P_e^{BIC} and P_e with the gamma bound P_G, estimated level-crossing bound P_{LC} and P_I for the shape in Figure 6. Left panel: $J^\star = 3$; middle panel: $J^\star = 7$; right panel: $J^\star = 11$).

ORCID iDs

Vinay L. Kashyap https://orcid.org/0000-0002-3869-7996
Thomas C. M. Lee https://orcid.org/0000-0001-7067-405X
David A. van Dyk https://orcid.org/0000-0002-0816-331X
Andreas Zezas https://orcid.org/0000-0001-8952-676X

References

Adams, R., & Bischof, L. 1994, ITPAM, 16, 641
Anastasopoulou, K., Zezas, A., Ballo, L., & Della Ceca, R. 2016, MNRAS, 460, 3570
Ballo, L., Braito, V., Della Ceca, R., et al. 2004, ApJ, 600, 634
Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
Burger, W., & Burge, M. J. 2013, Principles of Digital Image Processing, Vol. 3 (Berlin: Springer-Verlag)
Burger, W., & Burge, M. J. 2016, The Discrete Fourier Transform in 2D (London: Springer London), 481
Burnham, K. P., & Anderson, D. R. 2002, Advanced Issues and Deeper Insights (New York: Springer), 267
Byers, S. D., & Raftery, A. E. 2002, in Spatial Cluster Modelling, ed. A. B. Lawson & D. G. T. Denison (Boca Raton, FL: Chapman and Hall/CRC), 109
Chatterjee, S., Umbreit, S., Fregeau, J. M., & Rasio, F. A. 2013, MNRAS, 429, 2881
Corwin, H. G. j., Buta, R. J., & de Vaucouleurs, G. 1994, AJ, 108, 2128
Davis, D. S., Mulchaey, J. S., Mushotzky, R. F., & Burstein, D. 1996, ApJ, 460, 601
de Vaucouleurs, G., de Vaucouleurs, A., Corwin, H. G. J., et al. 1991, Third Reference Catalogue of Bright Galaxies, Vol. 1 (New York: Springer)
Ding, Y., Liu, J., Xu, X., et al. 2020, PMLR, 121, 156, https://proceedings.mlr.press/v121/ding20a.html
Efron, B. 1979, AOS, 7, 1
Efron, B., Hastie, T., Johnstone, I., & Tibshirani, R. 2004, AOS, 32, 407
Fan, M., Wang, J., Kashyap, V. L., et al. 2023, AJ, 165, 66
Gal, Y., & Ghahramani, Z. 2015, arXiv:1506.02142
Hall, P. 1985, Stoch. Process. Their Appl., 20, 231
Hamu Goldberg, H. B., Mushkin, J., Raviv, T. R., & Sochen, N. 2018, in 25th IEEE International Conf. on Image Processing (ICIP) (New York: IEEE), 734
Hubert, L., & Arabie, P. 1985, J. Classif., 2, 193
Kohl, S. A. A., Romera-Paredes, B., Meyer, C., et al. 2018, arXiv:1806.05034
Lin, C. C., & Chellappa, R. 1987, in IEEE Transactions on Pattern Analysis and Machine Intelligence (New York: IEEE), 686
Loh, J. M., & Stein, M. L. 2004, Statistica Sinica, 14, 69, http://www.jstor.org/stable/24307180
Mallat, S. G., & Zhang, Z. 1993, ITSP, 41, 3397
Mayer, J., Khairy, K., & Howard, J. 2010, AmJPh, 78, 648
Rand, W. M. 1971, JASA, 66, 846
Scargle, J. D. 2002, in AIP Conf. Proc. 617, Bayesian Inference and Maximum Entropy Methods in Science and Engineering (Melville, NY: AIP), 163
Scargle, J. D., Norris, J. P., Jackson, B., & Chiang, J. 2013, ApJ, 764, 167
Tillmann, C., & Ney, H. 2003, J. Comput. Linguist., 29, 97
Tscheschel, A., & Stoyan, D. 2006, CSDA, 51, 859
van Dyk, D., Connors, A., Esch, D. N., et al. 2006, BayAn, 1, 189
Wang, J. 2025, Auto-BUQ: Uncertainty Quantification for the Boundaries of Segmented Events, v1, Zenodo, doi:10.5281/zenodo.15128852
Way, M. J., Gazis, P. R., & Scargle, J. D. 2010, ApJ, 727, 48
Ye, J. C., Bresler, Y., & Moulin, P. 2000, ITIT, 46, 1881
Zezas, A., Ward, M. J., & Murray, S. S. 2003, ApJL, 594, L31