

Detecting new signals under background mismodeling

Sara Algeri^{*}*School of Statistics, University of Minnesota, Minneapolis, Minnesota 55455, USA*

(Received 25 June 2019; published 7 January 2020)

Searches for new astrophysical phenomena often involve several sources of nonrandom uncertainties that can lead to highly misleading results. Among these, model uncertainty arising from background mismodeling can dramatically compromise the sensitivity of the experiment under study. Specifically, overestimating the background distribution in the signal region increases the chances of missing new physics. Conversely, underestimating the background outside the signal region leads to an artificially enhanced sensitivity and a higher likelihood of claiming false discoveries. The aim of this work is to provide a unified statistical strategy to perform modeling, estimation, inference, and signal characterization under background mismodeling. The method proposed allows one to incorporate the (partial) scientific knowledge available on the background distribution and provides a data-updated version of it in a purely nonparametric fashion without requiring the specification of prior distributions on the unknown parameters. Applications in the context of dark matter searches and radio surveys show how the tools presented in this article can be used to incorporate nonstochastic uncertainty due to instrumental noise and to overcome violations of classical distributional assumptions in stacking experiments.

DOI: [10.1103/PhysRevD.101.015003](https://doi.org/10.1103/PhysRevD.101.015003)

I. INTRODUCTION

When searching for new physics, a discovery claim is made if the data collected by the experiment provide sufficient statistical evidence in favor of the new phenomenon. If the background and signal distributions are specified correctly, this can be done by means of statistical tests of the hypothesis, upper limits, and confidence intervals.

The problem. In practice, even if a reliable description of the signal distribution is available, providing accurate background models may be challenging, as the behavior of the sources that contribute to it is often poorly understood. Some examples include searches for nuclear recoils of weakly interacting massive particles over electron recoils backgrounds [1,2], searches for gravitational-wave signals over non-Gaussian backgrounds from stellar-mass binary black holes [3], and searches for a standard model-like Higgs boson over prompt diphoton production [4].

Unfortunately, model uncertainty due to background mismodeling can significantly compromise the sensitivity of the experiment under study. Specifically, overestimating the background distribution in the signal region increases the chances of missing new physics. Conversely, underestimating the background outside the signal region leads to an artificially enhanced sensitivity, which can easily result in false discovery claims. Several methods have been proposed in literature to address this problem (see, e.g., [5–7]). However, to the best of the author’s knowledge,

none of the methods available provides a unified strategy to (i) assess the validity of existing models for the background, (ii) fully characterize the background distribution, (iii) perform signal detection even if the signal distribution is not available, (iv) characterize the signal distribution, and (v) detect additional signals of new unexpected sources.

Goal. The aim of this work is to integrate modeling, estimation, and inference under background mismodeling and provide a general statistical methodology to perform (i)–(v). As a brief overview, given a source-free sample and the (partial) scientific knowledge available on the background distribution, a data-updated version of it is obtained in a purely nonparametric fashion without requiring the specification of prior distributions on the unknown parameters. At this stage, a graphical tool is provided in order to assess if and where significant deviations between the true and the postulated background distributions occur. The “updated” background distribution is then used to assess if the distribution of the data collected by the experiment deviates significantly from the background model. Also in this case, it is possible to assess graphically how the data distribution deviates from the expected background model. If a source-free sample is available, or if control regions can be identified, the solution proposed does not require the specification of a model for the signal; however, if the signal distribution is known (up to some free parameters), the latter can be used to further improve the accuracy of the analysis and to detect the signals of unexpected new sources. Finally, the method can easily be adjusted to cover situations in which a source-free sample or control

^{*}salgeri@umn.edu

regions are not available, the background is unknown, or incorrectly specified, but a functional form of the signal distribution is known.

The key of the solution. The statistical methodologies involved rely on the novel *LP approach to statistical modeling* first introduced by Mukhopadhyay and Parzen in 2014 [8]. As it will become clearer later on in the paper, the letter *L* typically denotes robust nonparametric methods based on quantiles, whereas *P* stands for polynomials [[9], Suppl. S1]. This approach allows the unification of many of the standard results of classical statistics by expressing them in terms of quantiles and comparison distributions and provides a simple and powerful framework for statistical learning and data analysis. The interested reader is directed to [10–14] and references therein, for recent advancements in mode detection, nonparametric time series, goodness-of-fit on prior distributions, and large-scale inference using an LP approach.

Organization. Section II is dedicated to a review of the main constructs of LP modeling. Section III highlights the practical advantages offered by modeling background distributions using an LP approach. Section IV introduces a novel LP-based framework for statistical inference. Section V outlines the main steps of a data-scientific approach for signal detection and characterization. In Sec. VI, the methods proposed are applied in the context of dark matter searches where the goal is to distinguish γ -ray emissions due to dark matter from those due to pulsars. In Sec. VII, the tools discussed are applied to a simulation of the Fermi Large Area Telescope γ -ray telescope, and it is shown how upper limits and Brazil plots can be constructed by means of comparison distributions. Section VIII is dedicated to model denoising. Section IX presents an application to data from the NRAO VLA Sky astronomical survey and discusses a simple approach to assess the validity of distributional assumptions on the polarized intensity in stacking experiments. A discussion of the main results and extensions is proposed in Sec. X.

II. LP APPROACH TO STATISTICAL MODELING

The *LP approach to statistical modeling* [8] is a novel statistical approach that provides an ideal framework to simultaneously assess the validity of the scientific knowledge available and fill the gap between the initial scientific belief and the evidence provided by the data. Sections II A–II C below introduce the LP modeling framework, whereas Sec. III discusses how the problem of background mis-modeling can be formulated under this paradigm.

A. The skew- G density model

Let X be a continuous random variable with cumulative distribution function (CDF) and probability density function (PDF) $F(x)$. Since F is the true distribution of the data, it is typically unknown. However, suppose a suitable CDF

$G(x)$ is available, and let $g(x)$ be the respective PDF. In order to understand if G is a good candidate for F , it is convenient to express the relationship among the two in a concise manner.

The *skew- G density model* [8,10] is a universal representation scheme that allows one to express any PDF $f(x)$ as

$$f(x) = g(x)d(G(x); G, F), \quad (1)$$

where $d(u; G, F)$ is called *comparison density* [15] and it is such that

$$d(u; G, F) = \frac{f(G^{-1}(u))}{g(G^{-1}(u))} \quad \text{with } 0 \leq u \leq 1, \quad (2)$$

with $u = G(x)$ and $G^{-1}(u) = \inf\{x: G(x) \geq u\}$ denoting the “postulated” quantile function of X . The comparison density is the PDF of the random variable $U = G(X)$; whereas its CDF is given by

$$D(u) = F(G^{-1}(u)) = \int_0^u d(v; G, F) \partial v, \quad (3)$$

and it is called *comparison distribution*.

Practical remarks. Equations (2) and (3) are of fundamental importance to understand the power of a statistical modeling approach based on the comparison density. Specifically, $d(u; G, F)$ allows one to “connect” any given PDF g to the true PDF f through the quantile transformation G^{-1} of u . Furthermore, $g \equiv f$ if and only if $d(u; G, F) = 1$ for all $u \in [0, 1]$; i.e., U is uniformly distributed over the interval $[0, 1]$. Whereas, if $g \neq f$, $d(u; G, F)$ models the departure of the true density $f(x)$ from the postulated model $g(x)$. Consequently, an adequate estimate of $d(u; G, F)$ not only leads to an estimate of the true $f(x)$ based on (1) but it also allows one to identify the regions where $f(x)$ deviates substantially from $g(x)$.

B. LP skew- G series representation

Denote with $L_2[0, 1]$ the Hilbert space of square integrable functions on the unit interval with respect to the measure G . A complete, orthonormal basis of functions in $L_2[0, 1]$ can be constructed considering powers of $G(x)$, i.e., $G(x), G^2(x), G^3(x), \dots$, and adequately orthonormalized via the Gram-Schmidt procedure [10]. The resulting bases can equivalently be expressed as normalized shifted Legendre polynomials,¹ namely $Leg_j(u)$, with $u = G(x)$.

¹Classical Legendre polynomials are defined over $[-1, 1]$; here, their “shifted” counterpart over the range $[0, 1]$ is considered. The first three normalized shifted Legendre polynomials are $Leg_0(u) = 1$, $Leg_1(u) = \sqrt{12}(u - 0.5)$, and $Leg_2(u) = \sqrt{5}(6u^2 - 6u + 1)$.

Under the assumption that (2) is a square integrable function on $[0, 1]$, i.e., $d \in L_2[0, 1]$, we can then represent $d(u; G, F)$ via a series of $\{Leg_j(u)\}_{j \geq 0}$ polynomials, i.e.,

$$d(u; G, F) = 1 + \sum_{j>0} LP_j Leg_j(u) \quad (4)$$

with coefficients $LP_j = \int_0^1 Leg_j(u) d(u; G, F) \partial u$. The representation in (4) is called *LP skew-G series representation* [10].

C. LP density estimate

Let x_1, \dots, x_n be a sample of independent and identically distributed (i.i.d.) observations from X . Observations from U are given by $u_1 = G(x_1), \dots, u_n = G(x_n)$. The LP_j coefficients in (4) can then be estimated via

$$\widehat{LP}_j = \frac{1}{n} \sum_{i=1}^n Leg_j(u_i). \quad (5)$$

Alternatively, in virtue of (3), the estimates \widehat{LP}_j can also be specified as

$$\widehat{LP}_j = \int_0^1 Leg_j(u) \partial \tilde{D}(u) = \int_0^1 Leg_j(u) \partial \tilde{F}(G^{-1}(u)), \quad (6)$$

where \tilde{F} and \tilde{D} denote the empirical distribution of the samples x_1, \dots, x_n and u_1, \dots, u_n , respectively.

The moments of the \widehat{LP}_j are

$$E[\widehat{LP}_j] = LP_j, \quad V(\widehat{LP}_j) = \frac{\sigma_j^2}{n},$$

and $Cov(\widehat{LP}_j, \widehat{LP}_k) = \frac{\sigma_{jk}}{n}, \quad (7)$

where $\sigma_j^2 = \int_0^1 (Leg_j(u) - LP_j)^2 d(u; G, F) \partial u$ and $\sigma_{jk}^2 = \int_0^1 (Leg_j(u) - LP_j)(Leg_k(u) - LP_k) d(u; G, F) \partial u$. When $f \equiv g$, the equalities in (7) reduce to

$$E[\widehat{LP}_j] = 0, \quad V(\widehat{LP}_j) = \frac{1}{n}, \quad \text{and} \quad Cov(\widehat{LP}_j, \widehat{LP}_k) = 0 \quad (8)$$

for all $j \neq k$. Derivations of (7) and (8) are discussed in the Appendix A.

If (4) is approximated by the first $M + 1$ terms,² an estimate of the comparison density is given by

²Recall that the first normalized shifted Legendre polynomial is $Leg_0(u) = 1$.

$$\hat{d}(u; G, F) = 1 + \sum_{j=1}^M \widehat{LP}_j Leg_j(u), \quad (9)$$

with variance

$$V[\hat{d}(u; G, F)] = \sum_{j=1}^M \frac{\sigma_j^2}{n} Leg_j^2(u) + 2 \sum_{j<k} \frac{\sigma_{jk}}{n} Leg_j(u) Leg_k(u). \quad (10)$$

See Appendix B for more details on the derivation of (10). Finally, the standard error of $\hat{d}(u; G, F)$ corresponds to the square root of (10), with σ_j^2 and σ_{jk} estimated by their sample counterpart, i.e.,

$$\hat{\sigma}_j^2 = \frac{1}{n} \sum_{i=1}^n (Leg_j(u_i) - \widehat{LP}_j)^2,$$

$$\hat{\sigma}_{jk} = \frac{1}{n} \sum_{i=1}^n Leg_j(u_i) Leg_k(u_i) - \widehat{LP}_j \widehat{LP}_k.$$

Finally, in virtue of the skew-G density model in (1) we can estimate $f(x)$ as

$$\hat{f}(x) = g(x) \hat{d}(G(x); G, F). \quad (11)$$

Since each $Leg_j(u)$ is a polynomial function of the random variable U , each \widehat{LP}_j estimate can be expressed as a linear combination of the first j sample moments of U , e.g.,

$$\widehat{LP}_2 = \frac{1}{n} \sum_{i=1}^n Leg_2(u_i) = \sqrt{5}(6\hat{\mu}_2 - 6\hat{\mu}_1 + 1),$$

where $\hat{\mu}_2 = \frac{1}{n} \sum_{i=1}^n u_i^2$, $\hat{\mu}_1 = \frac{1}{n} \sum_{i=1}^n u_i$. Therefore, the truncation point M can be interpreted as the order of the highest moment considered to characterize the distribution of U . (The reader is directed to Sec. IV C for a discussion on the choice of M .)

D. The bias variance trade-off

In order to understand how good (9) is in estimating $d(u; G, F)$ we consider the mean integrated squared error (MISE) of $\hat{d}(u; G, F)$, i.e.,

$$MISE = E \left[\int_0^1 (\hat{d}(u; G, F) - d(u; G, F))^2 \partial u \right] \quad (12)$$

$$= \sum_{j=1}^M \frac{\sigma_j^2}{n} + \sum_{j>M} LP_j^2, \quad (13)$$

where the first term in (13) corresponds to the integral of the (10) over $[0, 1]$; whereas the second term corresponds to the integrated squared bias (IBS), i.e.,

$$\text{IBS} = \int_0^1 (E[\hat{d}(u; G, F)] - d(u; G, F))^2 \partial u. \quad (14)$$

Interestingly, the latter can also be specified as

$$\text{IBS} = \int_0^1 \left(\frac{f(x) - g(x)}{g(x)} \right)^2 \partial u - \sum_{j=1}^M LP_j^2 \quad (15)$$

(see derivations in Appendix B). The first term on the right-hand side of (15) is particularly important in understanding the role played by g in obtaining a reliable estimate of f . Specifically, the closer g is to f the lower the bias of $\hat{d}(G(x); G, F)$ and $\hat{f}(x)$ in (11).

Practical remarks. Equation (13) implies that larger values of M do not necessarily lead to better estimates of $d(u; G, F)$. Specifically, when $n \rightarrow \infty$, the first term in (13) tends to zero. However, for large values of M , more and more terms to contribute to it and thus increasing M may lead to a substantial inflation of the variance in (10). Conversely, the bias is not affected by sample size, and it can be controlled by either choosing g sufficiently close to f [see (15)] and/or increasing M while preserving a good bias-variance trade-off.

Further remarks. Equation (6) implies that the estimator in (9) relies on the empirical distribution of the sample observed by means of the \widehat{LP}_j estimates. Therefore, an estimator of the comparison density based entirely on the empirical CDF can be expressed by setting $M = n - 1$ in (9). However, as discussed in this section, while this would reduce the bias, it would also increase the variance drastically. Therefore, for $M < n - 1$, the estimator in (9) not only leads to a reduction of the variance but, in virtue of (15), its bias is mitigated when the postulated model g is sufficiently close to the true PDF of the data f .

III. DATA-DRIVEN CORRECTIONS FOR MISSPECIFIED BACKGROUND MODELS

Let $\mathbf{x}_B = (x_1, \dots, x_N)$ be a sample of observations from control regions or the result of Monte Carlo simulations, where we expect no signal to be present. Hereafter, we refer to \mathbf{x}_B as the *source-free sample*. Therefore, \mathbf{x}_B can be used to “learn” the unknown PDF of the background, namely $f_b(x)$, and obtain an estimate for it via (11).

Despite the true background model being unknown, suppose that a candidate PDF, namely $g_b(x)$, is available. The candidate model $g_b(x)$ can be specified from previous experiments or theoretical results or can be obtained by fitting specific functions (e.g., polynomial or exponential) to \mathbf{x}_B . If $g_b(x)$ does not provide an accurate description of $f_b(x)$, the sensitivity of the experiment can be strongly affected.

Consider, for instance, a source-free sample of $N = 5000$ observations whose true (unknown) distribution corresponds to the tail of a Gaussian with mean 55 and width 15 over the range $[0, 50]$, i.e.,

$$f_b(x) = \frac{e^{-\frac{1}{2}\left(\frac{x-55}{15}\right)^2}}{k_{fb}} \quad (16)$$

with $k_{fb} = \int_0^{50} e^{-\frac{1}{2}\left(\frac{x-55}{15}\right)^2} \partial x$. Suppose that a candidate model for the background is obtained by fitting a second-degree polynomial on the source-free sample and adequately normalizing it in order to obtain a proper PDF, i.e.,

$$g_b(x) = \frac{9.52 - 2.22x + 0.15x^2}{k_{gb}} \quad (17)$$

with $k_{gb} = \int_0^{50} [9.52 - 2.22x + 0.15x^2] \partial x$. For illustrative purposes, assume that the distribution of the signal is a Gaussian centered at 25, with width 4.5 and PDF

$$f_s(x) = \frac{e^{-\frac{1}{2}\left(\frac{x-25}{4.5}\right)^2}}{k_{fs}} \quad (18)$$

with $k_{fs} = \int_0^{50} e^{-\frac{1}{2}\left(\frac{x-25}{4.5}\right)^2} \partial x$. The histogram of the source-free sample along with (16)–(18) is shown in Fig. 1. At the higher end of the spectrum, the postulated background (red dashed line) underestimates the true background distribution (green solid line). As a result, using (17) as the

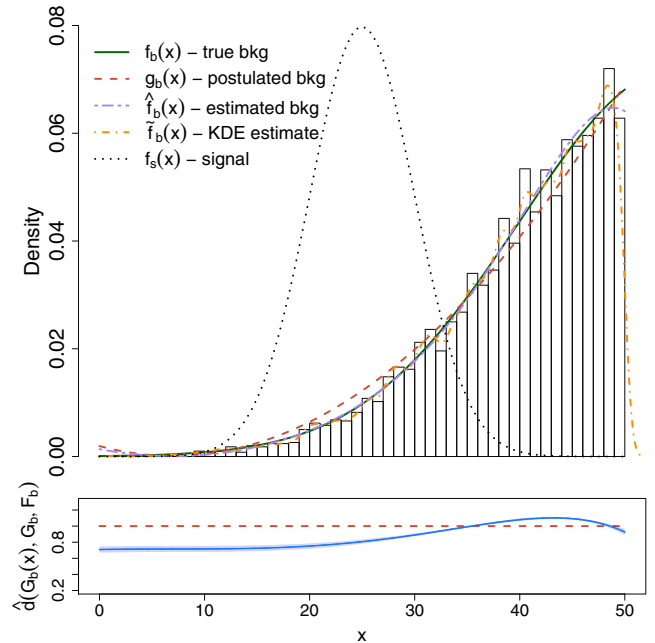


FIG. 1. Upper panel: histogram of a source-free sample simulated from the tail of a Gaussian with mean 55 and width 15 (green solid line). The candidate background distribution is given by the best fit of a second-degree polynomial (red dashed line), and it is updated using the source-free data by means of (19) (purple dot-dashed line). The Kernel density estimator of f_b is also displayed for comparison (orange dot-dashed line). Bottom panel: comparison density estimate (blue solid line) plotted on the x scale and respective standard errors (light blue area).

background model increases the chance of false discoveries in this region. Conversely, at the lower end of the spectrum, $g_b(x)$ underestimates $f_b(x)$, reducing the sensitivity of the analysis. For the sake of comparison, a Kernel density estimate (orange dot-dashed line) has been computed by selecting the bandwidth parameter as recommended in [16]. The latter exhibits substantial bias at the boundary and appears to overfit the data sample.

It is important to point out that the discrepancy of $f_b(x)$ from $g_b(x)$ is typically due to the fact that the specific functional form imposed (in our example, a second-degree polynomial) is not adequate for the data. Thus, changing the values of the fitted parameters (or assigning priors to them) is unlikely to solve the problem. However, it is possible to “repair” $g_b(x)$ and obtain a suitable estimate of $f_b(x)$ by means of (11). Specifically, $f_b(x)$ can be estimated via

$$\hat{f}_b(x) = g_b(x)\hat{d}(G_b(x); G_b, F_b), \quad (19)$$

where $\hat{d}(G_b(x); G_b, F_b)$ is the comparison density estimated via (9) on the sample $G_b(x_1), \dots, G_b(x_N)$, whereas F_b and G_b are the true and the postulated background distributions, with PDFs as in (16) and (17), respectively.

In our example, choosing $M = 2$ (see Sec. IV C), we obtain

$$\begin{aligned} \hat{d}(G_b(x); G_b, F_b) &= 1 + 0.063Leg_1[G_b(x)] \\ &\quad - 0.082Leg_2[G_b(x)], \end{aligned} \quad (20)$$

where $Leg_1[G_b(x)]$ and $Leg_2[G_b(x)]$ are the first and second normalized shifted Legendre polynomials evaluated at $G_b(x)$. Notice that, by combining (19) and (20), we can easily write the background model using a series of shifted Legendre polynomials. This may be especially useful when dealing with complicated likelihoods and for which a functional form is difficult to specify.

The upper panel of Fig. 1 shows that the “calibrated” background model in (19) as a purple dot-dashed line matches almost exactly the true background density in (16) (green solid line). The plot of $\hat{d}(G_b(x); G_b, F_b)$ in the bottom panel of Fig. 1 provides important insights on the deficiencies of (17) as a candidate background model. Specifically, the magnitude and the direction of the departure of $\hat{d}(G_b(x); G_b, F_b)$ from one corresponds to the estimated departure of $f_b(x)$ from $g_b(x)$ for each value of x . Therefore, if $\hat{d}(G_b(x); G_b, F_b)$ is below one in the region where we expect the signal to occur, using $\hat{f}_b(x)$ in place of $g_b(x)$ increases the sensitivity of the analysis. Conversely, if $\hat{d}(G_b(x); G_b, F_b)$ is above one outside the signal region, the use of $\hat{f}_b(x)$ instead of $g_b(x)$ prevents false discoveries.

Notice that in this article we only consider continuous data. In this respect, the goal is to learn the model of the background considered as a continuum and no binning is

applied. Therefore, the histograms presented here are only a graphical tool used to display the data distribution and are not intended to represent an actual binning of the data.

IV. LP-BASED INFERENCE

When discussing the skew- G density model in (1), we have witnessed that $f \equiv g$ if $d(u; G, F) = 1$ for all $u \in [0, 1]$. Additionally, the graph of $\hat{d}(u; G, F)$ provides an exploratory tool to understand the nature of the deviation of $f(x)$ from $g(x)$. This section introduces a novel inferential framework to test the significance of the departure of $f(x)$ from $g(x)$. Specifically, our goal is to test the hypotheses

$$H_0: d(u; G, F) = 1 \quad \text{for all } u \in [0, 1]$$

vs

$$H_1: d(u; G, F) \neq 1 \quad \text{for some } u \in [0, 1]. \quad (21)$$

First, an overall test, namely the *deviance test*, is presented. The deviance test assesses if $f(x)$ deviates significantly from $g(x)$ anywhere over the range of x considered. Second, adequate confidence bands are constructed in order to assess where significant departures occur.

A. The deviance test

Recall that the LP_j coefficients in (4) specify as $LP_j = \int_0^1 Leg_u(d) d(u; G, F) \partial u$. Consequently, by orthogonality of the $\{Leg_j(u)\}_j > 0$ polynomials and $Leg_0(u) = 1$, when H_0 in (21) is true, all the LP_j coefficients are equal to zero, including the first M of them. We can then quantify the departure of $\hat{d}(u; G, F)$ from one by means of the *deviance statistics* [13] which specifies as $\sum_{j=1}^M \widehat{LP}_j^2$. If the deviance is equal to zero, we may expect that g is approximately equivalent to f ; hence, we test

$$H_0: \sum_{j=1}^M LP_j^2 = 0 \quad \text{vs} \quad H_1: \sum_{j=1}^M LP_j^2 > 0 \quad (22)$$

by means of the test statistic

$$D_M = n \sum_{j=1}^M \widehat{LP}_j^2. \quad (23)$$

It can be shown [10] that, as $n \rightarrow \infty$,

$$\sqrt{n} \widehat{LP}_j \xrightarrow{d} N(0, 1), \quad (24)$$

where \xrightarrow{d} denotes convergence in distribution, and thus, under H_0 , D_M is asymptotically χ_M^2 distributed. Hence, an asymptotic p -value for (22) is given by

$$P(D_M > d_M) \xrightarrow{n \rightarrow \infty} P(\chi_M^2 > d_M), \quad (25)$$

where d_M is the value of D_M observed on the data.

Practical remarks. Notice that H_1 in (22) implies H_1 in (21). Similarly, H_0 in (21) implies H_0 in (22); however, the opposite is not true in general since there may be some nonzero LP_j coefficients for $j > M$. Therefore, even when choosing M small may lead to a conservative, but yet valid, inference.

B. Confidence bands

The estimator in (9) only accounts for the first $M + 1$ terms of the polynomial series in (4). Therefore, $\hat{d}(u; G, F)$ is a biased estimator of $d(u; G, F)$. Specifically, as discussed in Sec. IID, the integrated bias is given by $\sum_{j>M} LP_j^2$, whereas, as shown in Appendix B, the bias at a given point u is given by $\sum_{j>M} LP_j \text{Leg}(u)$. It follows that, when the bias is large, confidence bands based on $\hat{d}(u; G, F)$ are shifted away from the true density $d(u; G, F)$.

Despite the bias not being easily quantified in the general setting, it follows from (8) that, when H_0 in (21) [and consequently H_0 in (22)] is true, both the bias at a point u and the integrated bias are equal to zero. Thus, we can exploit this property to construct reliable confidence bands under the null. Specifically, the goal is to identify c_α , such that

$$\begin{aligned} 1 - \alpha &= P(-c_\alpha \leq \hat{d}(u; G, F) - 1 \leq c_\alpha, \text{ for all } u \in [0, 1] | H_0) \\ &= P(\max_u |\hat{d}(u; G, F) - 1| \leq c_\alpha | H_0), \end{aligned} \quad (26)$$

where α is the desired significance level.³

If the bias determines where the confidence bands are centered, the distribution and the variance of $\hat{d}(u; G, F)$ determine their width. As discussed in Sec. IIC [see (8)], under H_0 in (21), the \widehat{LP}_j estimates have mean zero, have variance $\frac{1}{n}$, and are uncorrelated with one another. Therefore, when $f \equiv g$, the standard error of $\hat{d}(u; G, F)$ corresponds to the square root of (10) with $\sigma_j^2 = 1$ and $\sigma_{jk}^2 = 0$, i.e.,

$$SE[\hat{d}(u; G, F) | H_0] = \sqrt{\sum_{j=1}^M \frac{1}{n} \text{Leg}_j^2(u)}. \quad (27)$$

Additionally, Eq. (24) implies that $\hat{d}(u; G, F)$ is asymptotically normally distributed, hence

³In astrophysics, the statistical significance α is often expressed in terms of the number of σ deviations from the mean of a standard normal, namely σ . For instance, a 2σ significance corresponds to $\alpha = 1 - \Phi(2) = 0.0227$, where $\Phi(\cdot)$ denotes the CDF of a standard normal.

$$\frac{\hat{d}(u; G, F) - 1}{\sqrt{\sum_{j=1}^M \frac{1}{n} \text{Leg}_j^2(u)}} \xrightarrow{d} N(0, 1) \quad (28)$$

as $n \rightarrow \infty$, for all $u \in [0, 1]$, under H_0 .

We can then construct approximate confidence bands under H_0 that satisfy (26) by means of tube formulas (see [[17], Chap. 5] and [18]), i.e.,

$$\left[1 - c_\alpha \sqrt{\sum_{j=1}^M \frac{1}{n} \text{Leg}_j^2(u)}, 1 + c_\alpha \sqrt{\sum_{j=1}^M \frac{1}{n} \text{Leg}_j^2(u)} \right], \quad (29)$$

where c_α is the solution of

$$2(1 - \Phi(c_\alpha)) + \frac{k_0}{\pi} e^{-0.5c_\alpha^2} = \alpha, \quad (30)$$

with $k_0 = \sqrt{\sum_{j=1}^M [\frac{\partial}{\partial u} \text{Leg}_j(u)]^2}$. If $\hat{d}(u; G, F)$ is within the bands in (29) over the entire range $[0, 1]$, we conclude that there is no evidence that f deviates significantly from g anywhere over the range considered and at confidence level $1 - \alpha$. Conversely, we expect significant departures to occur in regions where $\hat{d}(u; G, F)$ lies outside the confidence bands.

Practical remarks. Notice that, under H_0 in (22), the $\hat{d}(u; G, F)$ is an unbiased estimator of $d(u; G, F)$, regardless of the choice of M . This implies that the confidence bands in (29) are only affected by the variance and asymptotic distribution of $\hat{d}(u; G, F)$ under H_0 .

C. Choice of M

The number of \widehat{LP}_j estimates considered determines the level of ‘‘smoothness’’⁴ of $\hat{d}(u; G, F)$, with smaller values of M leading to smoother estimates. The deviance test can be used to select the value M , which maximizes the sensitivity of the analysis according to the following scheme:

- (i) Choose a sufficiently large value M_{\max} .
- (ii) Obtain the estimates $\widehat{LP}_1, \dots, \widehat{LP}_{M_{\max}}$ as in (5).
- (iii) For $m = 1, \dots, M_{\max}$:
calculate the deviance test p -value as in (25), i.e.,

$$p(m) = P(\chi_m^2 > d_m) \quad (31)$$

$$\text{with } d_m = n \sum_{j=1}^m \widehat{LP}_j^2.$$

⁴As an anonymous referee correctly pointed out, $\hat{d}(u; G, F)$ is always smooth as it is constructed as a series of infinitely differentiable functions. In statistics, however, the word ‘‘smoothness’’ is often used to indicate the flexibility of the estimator considered or, in other words, its degrees of freedom. Often, this is quantified in terms of magnitude of the second derivative of the function considered. Despite the abuse of terminology, throughout the manuscript we will refer to the latter definition of smoothness.

(iv) Choose M such that

$$M = \operatorname{argmin}_m \{p(m)\}. \quad (32)$$

1. Adjusting for postselection

As any data-driven selection process, the scheme presented above affects the distribution of (9) and can yield to overly optimistic inference [19,20]. Despite this aspect often being ignored in practical applications, the correct coverage can be guaranteed only if adequate corrections are implemented.

The issues arising in the context of postselection inference can be interpreted in terms of the looks-elsewhere effect [21,22] where one has to adjust the inference for the fact that, in practice, many different models have been considered and, consequently, many different tests have been conducted for the sake of assessing the goodness of fit.

In our setting, the number of models under comparison is typically small ($M_{\max} \leq 20$); therefore, postselection inference can be easily adjusted by means of Bonferroni's correction [23]. Specifically, the adjusted deviance p -value is given by

$$M_{\max} \cdot P(\chi_M^2 > d_M), \quad (33)$$

where M is the value selected via (32), whereas confidence bands can be adjusted by substituting c_α in (29), with $c_{\alpha, M_{\max}}$ satisfying

$$2(1 - \Phi(c_{\alpha, M_{\max}})) + \frac{k_0}{\pi} e^{-0.5c_{\alpha, M_{\max}}^2} = \frac{\alpha}{M_{\max}}. \quad (34)$$

Practical remarks. As noted in Sec. II, the estimate (9) involves the first M sample moments of U ; therefore, M_{\max} can be interpreted as the order of the highest moment that we expect to contribute in discriminating the distribution of U from uniformity. Notice that, in addition to the inflation of the variance of (9), when M is large, the computation of normalized shifted Legendre of higher order may face numerical instability (see Sec. VIII B). Therefore, as a rule of thumb, M_{\max} is typically chosen ≤ 20 . Finally, Steps i–iv aim to select the approximant based on the first most significant M moments, while excluding powers of higher order. A further note on model denoising is given in Sec. VIII.

V. A DATA-SCIENTIFIC APPROACH TO SIGNAL SEARCHES

The tools presented in Secs. II and IV provide a natural framework to simultaneously

- (a) assess the validity of the postulated background model and, if necessary, update it using the data (Sec. III);
- (b) perform signal detection on the physics sample;
- (c) characterize the signal when a model for it is not available.

Furthermore, if the model for the signal is known (up to some free parameters), it is possible to

- (d) further refine the background or signal distribution;
- (e) detect hidden signals from new unexpected sources.

Notice that, since Bonferroni's correction leads to an upper bound for the overall significance, the resulting coverage will be higher than the nominal one. Alternatively, approximate postselection confidence bands and inference can be constructed using Monte Carlo and/or resampling methods and repeating the selection process at each replicate.

Tasks (a)–(e) can be tackled in two main phases. In the first phase, the postulated background model is “calibrated” on a source-free sample in order to improve the sensitivity of the analysis and reduce the risk of false discoveries. The second phase focuses on searching for the signal of interest and involves both a nonparametric signal detection stage and a semiparametric stage for signal characterization. Both phases and respective steps are described in detail below and summarized in Algorithm 1.

Algorithm 1. A data-scientific signal search

INPUTS: source-free sample \mathbf{x}_B ;
 postulated background distribution $g_b(x)$;
 physics sample \mathbf{x} .
If available: signal distribution, $f_s(x, \theta_s)$.

PHASE A: background calibration

Step 1: Estimate $\hat{d}(u; G_b, F_b)$ on $\mathbf{u}_B = G_b(\mathbf{x}_B)$ and test (21) via deviance test and comparison density (CD) plot.

Step 2: **if** $F_b \neq G_b$, set $\hat{f}_b(x) = g_b(x)\hat{d}(u; G_b, F_b)$;
else set $\hat{f}_b(x) = g_b(x)$.

PHASE B: signal search

Stage 1: nonparametric signal detection

Step 3: set $g(x) = \hat{f}_b(x)$.

Step 4: estimate $\hat{d}(u; G, F)$ on $\mathbf{u} = G(\mathbf{x})$ and test (21) via deviance test and CD plot.

Step 5: **if** $G \neq F$, claim *evidence in favor of the signal* and go to Step 6;
else set $\hat{f}(x) = g(x)$, claim that *no signal is present* and stop.

Stage 2: semiparametric signal characterization

Step 6: **if** $f_s(x, \theta_s)$ given, fit $g_{bs}(x)$ in (38);
else use the CD plot of $\hat{d}(u; G, F)$ and the theory available to specify/fit a suitable model for $f_s(x, \theta_s)$ and fit $g_{bs}(x)$ in (38).

Step 7: estimate $\hat{d}(u; G_{bs}, F)$ on $\mathbf{u} = G_{bs}(\mathbf{x})$ and test (21) via deviance test and CD plot.

Step 8: **if** $G_{bs} \neq F$, claim *evidence of unexpected signal* and use the CD plot of $\hat{d}(u; G_{bs}, F)$ and the theory available to further investigate the nature of the deviation from G_{bs} ;
else go to Step 9.

Step 9: compute $\hat{d}(u; G, F)$ as in (39) and use it to refine $\hat{f}_b(x)$ or $f_s(x, \hat{\theta}_s)$ as in (40). Go back to Step 3.

A. Background calibration

As discussed in Sec. III, deviations of $\hat{d}(G_b(x); G_b, F_b)$ from one suggest that a further refinement of the candidate background model g_b is needed. However, as M increases, the deviations of $\hat{d}(G_b(x); G_b, F_b)$ from one may become more and more prominent while the variance inflates. Thus, it is important to assess if such deviations are indeed significant. In order to address this task, the analysis of Sec. III can be further refined in light of the inferential tools introduced in Sec. IV.

For the toy example discussed in Sec. III, we have seen that g_b overestimates f_b in the signal region and underestimates it at the higher end of the range considered (Fig. 1). We can now assess if any of these deviations are significant by implementing the deviance test in (23)–(25), whereas, to identify where the most significant departures occur, we construct confidence bands under the null model as in (29), i.e., assuming that no “update” of g_b is necessary.

The results are collected in the *comparison density plot* or *CD plot* presented in Fig. 2. First, a value $M = 2$ has been selected as in (32), and the respective deviance test (adequately adjusted via Bonferroni) indicates that the deviation of f_b from g_b is significant at a 6.430σ significance level (adjusted p -value of 6.397×10^{-11}). Additionally, the estimated comparison density in (20) lies outside the 2σ confidence bands in the region $[0, 50]$ where the signal specified in (18) is expected to occur. Hence, using (19) instead of (17) is recommended in order to improve the sensibility of the analysis in the signal region.

Important remarks on the CD plot. When comparing different models for the background or when assessing if the data distribution deviates from the model expected when no signal is present, it is common practice to visualize the results of the analysis by superimposing the models




FIG. 2. Deviance test and CD plot for the source-free sample. The comparison density is estimated via (20) (solid blue line), whereas its standard error (light blue area) is computed as the squared root of the estimate of the variance in (10). Finally, confidence bands have been constructed around one (grey areas) via (29) with c_α replaced by $c_{\alpha, M_{\max}}$ in (34). The notation $\geq 2\sigma$ is used to highlight that Bonferroni’s correction has been applied to adjust for postselection inference, leading to an increase of the nominal coverage.

under comparison to the histogram of the data observed on the original scale (e.g., upper panel of Fig. 1). This corresponds to a data visualization in the density domain. Conversely, the CD plot (e.g., Fig. 2) provides a representation of the data in the quantile domain, which offers the advantage of connecting the true density of the data with the quantiles of the postulated model [see (2) and (3)]. Consequently, the most substantial departures of the data distribution from the expected model are magnified, and those due to random fluctuations are smoothed out (see also Sec. VII B). Furthermore, the deviance tests and the CD plot together provide a powerful goodness-of-fit tool and exploratory that, conversely from classical methods such as Anderson-Darling [24] and Kolmogorov-Smirnov [25], not only allow one to test if the distributions under comparison differ, but they also allow one to assess *how* and *where* they differ. As a result, the CD plot can be used to characterize the unknown signal distribution (see Sec. V B 2) and to identify exclusion regions (e.g., Case I in Sec. V B 1).

As an additional advantage, the deviance test appears to enjoy a higher detection power than classical approaches. This aspect is highlighted in Table I where several methods for goodness of fit or two-samples comparisons are implemented, along with the deviance test, for all the cases discussed in Sec. V.

Reliability of the calibrated background model. The size N of the source-free sample plays a fundamental role in the validity of $\hat{f}_b(x)$ as a reliable background model. Specifically, the randomness involved in (19) only depends on the \widehat{LP}_j estimates. If N is sufficiently large, by the strong law of large numbers,

$$P(\lim_{N \rightarrow \infty} \widehat{LP}_j = LP_j) = 1.$$

Therefore, despite the variance of $\hat{f}_b(x)$ becoming negligible as $N \rightarrow \infty$, one has to account for the fact that $\hat{f}_b(x)$ leads to a biased estimate of $f_b(x)$ when $f_b \neq g_b$ (see Sec. II D). For sufficiently smooth densities, a visual inspection is often sufficient to assess if $\hat{d}(u; G_b, F_b)$ [and, consequently, $\hat{f}_b(x)$] provides a satisfactory fit for the data, whereas for more complex distributions the effect of the bias can be mitigated considering larger values of M and model denoising (see Sec. VIII A).

B. Signal search

1. Nonparametric signal detection

The background calibration phase allows the specification of a well tailored model for the background, namely $\hat{f}_b(x)$, which simultaneously integrates the initial guess, g_b , and the information carried by the source-free data sample. Hereafter, we disregard the source-free data sample and focus on analyzing the physics sample.

TABLE I. Comparison of deviance test and classical inferential tools. The first two columns report the p -values of Anderson-Darling [24] and Cramer-von Mises [25] goodness-of-fit tests obtained assuming as theoretical distribution the same G indicated in Secs. VA and VB for the calibration phase, cases I, II, and III, respectively. The raw deviance p -values and their postselection adjusted counterparts are reported in the third columns. Finally, the fourth and fifth columns report, respectively, the Kolmogorov-Smirnov [25] and Wilcoxon rank sum [26] tests used to compare directly the physics samples in Cases I, II, and III with the source-free sample used in Sec. VA.

Sample	Goodness-of-fit test p -values			Two-samples test p -values	
	Anderson-Darling	Cramer-von Mises	Deviance (adjusted)	Kolmogorov-Smirnov	Wilcoxon Rank Sum
Calibration	1.2×10^{-7}	4.2×10^{-7}	3.2×10^{-12} (6.4×10^{-11})
Case I	0.7776	0.7711	0.2657(>1)	0.9248	0.5487
Case II	4.6×10^{-7}	8.2×10^{-11}	9.0×10^{-33} (1.8×10^{-31})	1.9×10^{-13}	4.5×10^{-12}
Case III	4.6×10^{-7}	2.6×10^{-10}	2.6×10^{-28} (5.2×10^{-27})	2.1×10^{-15}	2.2×10^{-16}

Under the assumption that the source-free sample has no significant source contamination, we expect that, if the signal is absent, both the source-free and the physics samples follow the same distribution. Therefore, the calibrated background model, $\hat{f}_b(x)$, plays the role of the postulated distribution for the physics sample, i.e., the model that we expect the data to follow when no signal is present; hence, we set $g(x) = \hat{f}_b(x)$.

Let $f(x)$ be the (unknown) true PDF of the physics sample that may or may not carry evidence in favor of the source of interest. When no model for the signal is specified, it is reasonable to consider any significant deviation of f from g as an indication that a signal of unknown nature may be present. In this setting, similar to the background calibration phase, we can construct deviance tests and CD plots to assess if and where significant departures of f from g occur. Two possible scenarios are considered—a physics sample that collects only background data (Case I) and a physics sample of observations from both background and signal (Case II).

Case I: Background only.—Let \mathbf{x} be a physics sample of $n = 1300$ observations whose true (unknown) PDF $f(x)$ is equivalent to $f_b(x)$ in (16). We set

$$g(x) = \hat{f}_b(x) = g_b(x)\hat{d}(G_b(x); G_b, F_b), \quad (35)$$

where $g_b(x)$ and $\hat{d}(G_b(x); G_b, F_b)$ are defined as in (17) and (20), respectively. The resulting CD plot and deviance test are reported in the left panel of Fig. 3.

When applying the scheme in Sec. IV C with $M_{\max} = 20$, none of the values of M considered leads to significant results; therefore, for the sake of comparison with Case II below, we choose $M = 4$. Not surprisingly, the estimated comparison density approaches one over the entire range and lies entirely within the confidence bands. This suggests that the true distribution of the data does not differ significantly from the model which accounts only for the background. Similarly, the deviance test leads to very low significance (adjusted p -value > 1); hence, we conclude that our physics sample does not provide evidence in favor of the new source.

Case II: Background + signal.—Let \mathbf{x} be a physics sample of $n = 1300$ observations whose true (unknown) PDF $f(x)$ is equal to $f_{bs}(x)$ in (36),

$$f_{bs}(x) = (1 - \eta)f_b(x) + \eta f_s(x) \quad (36)$$

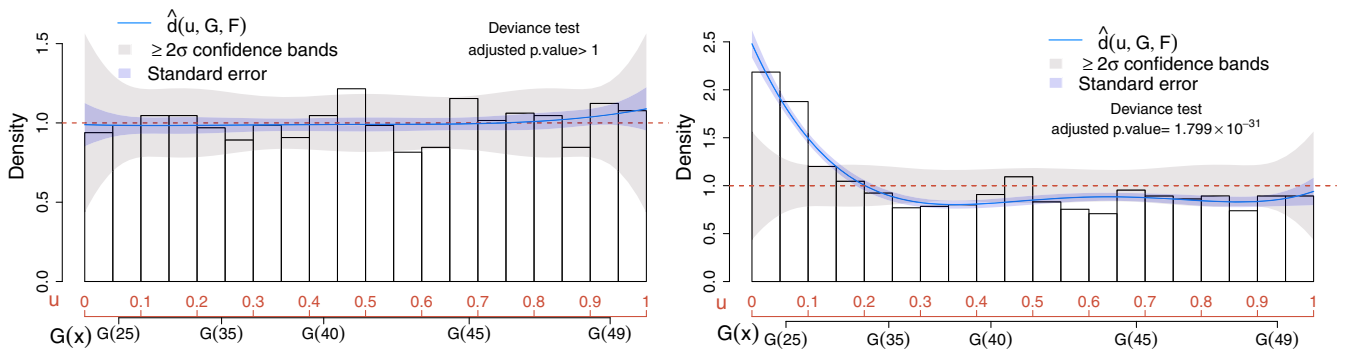


FIG. 3. Deviance test and CD plots for Case I where no signal is present (left panel) and Case II where the signal is present (right panel). In both cases, the postulated distribution G corresponds to the CDF of the calibrated background model in (35). For the sake of comparison, $d(u; G, F)$ has been estimated via (9) with $M = 4$ for both samples.




FIG. 4. Histogram of a physics sample of $n = 1300$ observations from both background and signal and with the PDF as in (36) (grey solid line). The true density has been estimated semiparametrically, as in (38) (pink dashed line), whereas the nonparametric estimates of $f(x)$ have been computed as in (11), by plugging in the $\hat{d}(G(x); G, F)$ estimates obtained with $M = 4$ (blue dot-dashed line) and $M = 9$ (black dotted line).

with $f_b(x)$ and $f_s(x)$ defined as in (16) and (18), respectively, and $\eta = 0.15$. The histogram of the data and the graph of $f_{bs}(x)$ are plotted in Fig. 4. As in Case I, we set $g(x)$ as in (35).

The CD plot and deviance test in the right panel of Fig. 3 show a significant departure of the data distribution from the background-only model in (35). The maximum significance of the deviance is achieved at $M = 4$, leading to a rejection of the null hypothesis at a 11.611σ significance level (adjusted p -value = 1.799×10^{-31}). The CD plot shows a prominent peak at the lower end of the spectrum; hence, we conclude that there is evidence in favor of the signal, and we proceed to characterize its distribution as described in Sec. V B 2.

2. Semiparametric signal characterization

The signal detection strategy proposed in Sec. V B 1 does not require the specification of a distribution for the signal. However, if a model for the signal is known (up to some free parameters), the analysis can be further refined by providing a parametric estimate of the comparison density and assessing if additional signals from new unexpected sources are present.

Case IIa: Background + (known) signal.—Assume that a model for the signal, $f_s(x, \theta_s)$, is given, with θ_s being a vector of unknown parameters. Since the CD plot in the right panel of Fig. 3 provides evidence in favor of the signal, we expect the data to be distributed according to the PDF,

$$\hat{f}_{bs}(x) = (1 - \eta)\hat{f}_b(x) + \eta f_s(x, \theta_s), \quad 0 \leq \eta \leq 1, \quad (37)$$

where $\hat{f}_b(x)$ is the calibrated background distribution in (35) and η and θ_s can be estimated via maximum likelihood (ML). Letting $\hat{\eta}$ and $\hat{\theta}_s$ be the ML estimates of η and θ_s , respectively, we specify

$$g_{bs}(x) = (1 - \hat{\eta})\hat{f}_b(x) + \hat{\eta}f_s(x, \hat{\theta}_s) \quad (38)$$

as a postulated model. For simplicity, let f_s be fully specified as in (18); we construct the deviance test and the CD plot to assess if (38) deviates significantly from the true distribution of the data. The scheme in Sec. IV C has been implemented with $M_{\max} = 20$, and none of the values of M considered led to significant results. The CD plot and deviance test for $M = 4$ are reported in the upper left panel of Fig. 5. Both the large p -value of the deviance test (adjusted p -value > 1) and the CD plot suggest that no significant deviations occur; thus, Eq. (38) is a reliable model for the physics sample.

Moreover, we can use (38) to further refine our $\hat{f}_b(x)$ or $f_s(x, \hat{\theta}_s)$ distributions. Specifically, we first construct a semiparametric estimate of $d(G(x); G, F)$, i.e.,

$$\hat{\hat{d}}G(x); G, F = (1 - \hat{\eta}) \frac{\hat{f}_b(x)}{f_s(x, \hat{\theta}_s)} + \hat{\eta}, \quad (39)$$

and rewrite

$$\begin{aligned} \hat{f}_b(x) &= \frac{\hat{f}_b(x)\hat{\hat{d}}G(x); G, F - \hat{\eta}f_s(x, \hat{\theta}_s)}{(1 - \hat{\eta})}, \\ \hat{f}_s(x) &= \frac{\hat{f}_b(x)\hat{\hat{d}}G(x); G, F - (1 - \hat{\eta})\hat{f}_b(x)}{\hat{\eta}}. \end{aligned} \quad (40)$$

In the upper right panel of Fig. 5, the true comparison density (grey dashed line) of our physics sample is compared with its semiparametric estimate computed as in (39) (pink dashed line) with $f_s(x, \hat{\theta}_s) = f_s(x)$ in (18). The graphs of two nonparametric estimates of $d(u; G, F)$ computed via (9) with $M = 4$ and $M = 9$ (blue dot-dashed line and black dotted line), respectively, are added to the same plot. Not surprisingly, incorporating the information available on the signal distribution drastically improves the accuracy of the analysis. The semiparametric estimate matches $d(u; G, F)$ almost exactly, whereas both nonparametric estimates show some discrepancies from the true comparison density. All the estimates suggest that there is only one prominent peak in correspondence of the signal region.

When moving from the comparison density domain to the density domain in Fig. 4, the discrepancies between the nonparametric estimates and the true density $f(x)$ are substantially magnified. Specifically, when computing (9)

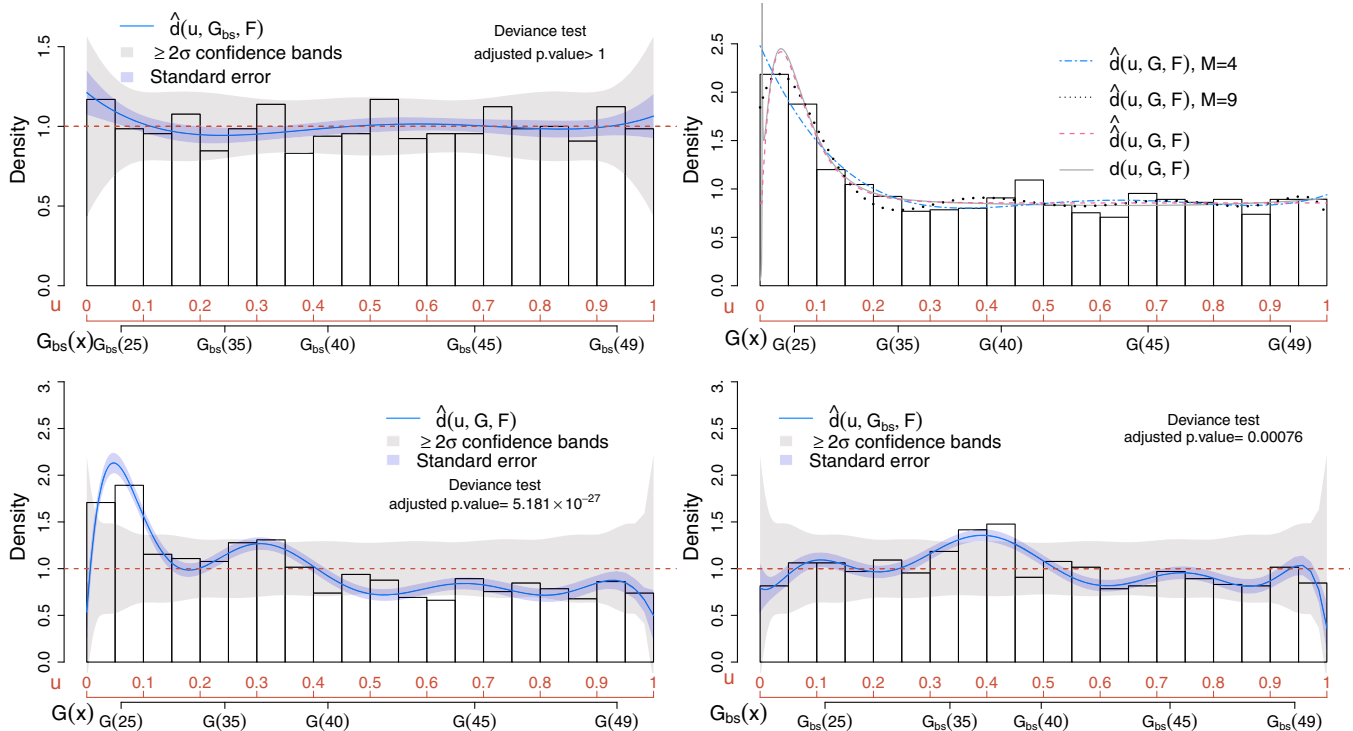


FIG. 5. Upper panels: Deviance test and CD plots for Case IIa where the signal is present (right panel), and the postulated distribution G_{bs} corresponds to the CDF of the estimated background + signal model in (38) with $\hat{\eta} = 0.146$. The comparison density estimate has been obtained considering $M = 3$. Bottom panels: Deviance test and CD plots for Case III where, in addition to the signal of interest, an additional resonance is present. The data are first analyzed considering the background-only PDF in (17) as the postulated model (left panel). The analysis is then repeated by assuming the fitted background + signal model in (38) as the postulated distribution (right panel). Both estimates of the comparison density in the left and right panels have been computed as in (9) with $M = 9$.

and (11) with $M = 4$ (blue dot-dashed line), the height signal peak is underestimated, whereas, when choosing $M = 9$, the $\hat{f}(x)$ exhibits high bias at the boundaries⁵ (dotted black line).

Case IIb: Background + (unknown) signal.—When the signal distribution is unknown, the CD plot of $\hat{d}(u; G, F)$ can be used to guide the scientist in navigating across the different theories on the astrophysical phenomenon under study and specify a suitable model for the signal, i.e., f_s . The model proposed can then be validated, as in Case IIa, by fitting (38) and constructing deviance tests and CD plots.

At this stage, the scientist has the possibility to iteratively query the data and explore the distribution of the signal by assuming different models. A viable signal characterization is achieved when no significant deviations of $\hat{d}(u, G_{bs}, F)$ from one are observed (e.g., see upper left panel of Fig. 5). Notice that a similar approach can be followed also in the background calibration stage (Sec. VA) to provide a parametric characterization of the background distribution.

⁵Boundary bias is a common problem among nonparametric density estimation procedures (e.g., [17], Chap. 5, Chap. 8). When aiming for a nonparametric estimate of the data density $f(x)$, solutions exist to mitigate this problem (e.g., [27]).

Case III: Background + (known) signal + unexpected source.—The tools proposed so far can also be used to detect signals from unexpected sources whose PDFs are, by design, unknown.

Suppose that the physics sample x contains $n = 1300$ observations whose true (unknown) PDF $f(x)$ is equal to $f_{bsh}(x)$,

$$f_{bsh}(x) = (1 - \eta_1 - \eta_2)f_b(x) + \eta_1 f_s(x) + \eta_2 f_h(x), \quad (41)$$

where $f_h(x)$ is the PDF of the unexpected signal and we assume its distribution to be normal with the center at 37 and width 1.8. Let $f_b(x)$ and $f_s(x)$ be defined as in (16) and (18), respectively, and let $\eta_1 = 0.15$ and $\eta_2 = 0.1$.

We can start with a nonparametric signal detection stage by setting $g(x) = \hat{g}_{bs}(x)$ in (35), with f_s defined as in (18) and $\hat{\eta}$ estimated via maximum likelihood estimate (MLE). The respective CD plot and deviance tests are reported in the bottom left panel of Fig. 5.

Choosing $M = 9$, as in (32), both the CD plot and deviance test indicate a significant departure from the expected background-only model and a prominent peak is observed in correspondence of the signal of interest centered around 25. A second but weaker peak appears to be right on the edge of our confidence bands, suggesting

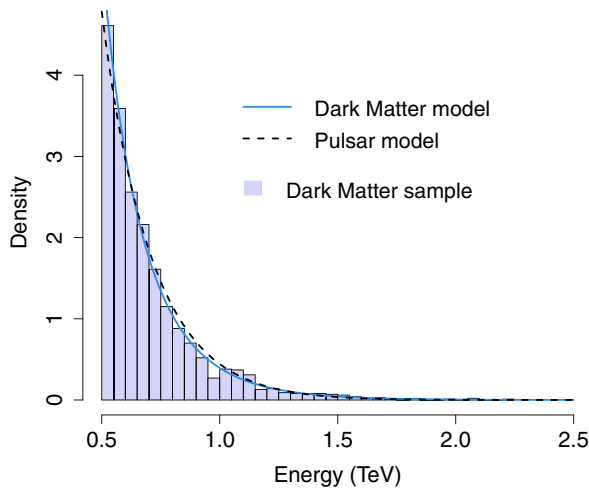
the possibility of an additional source. At this stage, if f_s was unknown, we could proceed with a semiparametric signal characterization as in Case IIb. Whereas assuming that the distribution of the signal of interest is known and given by (18), we fit (38), aiming to capture a significant deviation in correspondence of the second bump. This is precisely what we observe in the bottom right panel of Fig. 5. Here the estimated comparison density deviates from (35) around 35, providing evidence in favor of an additional signal in this region. We can then proceed as in Case IIb by exploring the theories available and/or collecting more data to further investigate the nature and the cause of the unanticipated bump.

VI. SIGNAL DETECTION WITHOUT CALIBRATION SAMPLE AND MODEL SELECTION

There are situations where a source-free sample is simply not available and thus the calibration phase in Sec. VA cannot be implemented. The tools described in Secs. II and IV can, however, still be applied in order to perform signal detection and goodness of fit when a model for the signal is known, up to some free parameters. In this framework, we expect the data to either come only from the signal (with at most some negligible background contamination) or only from the background.

In order to illustrate how to proceed in this setting, we consider a dark matter search where the postulated model for dark matter γ -ray emissions is the one of [[28], Eq. (29)], i.e.,

$$g_{\text{DM}}(y) = \frac{0.73M_\chi^{1.5}}{yk_{M_\chi}} \exp\left\{-7.8\frac{y}{M_\chi}\right\} \quad (42)$$



with $y \in [0.5, 5]$ teraelectron volt (TeV), $M_\chi \in [0.5, 5]$ TeV, and k_{M_χ} is a normalizing constant. The goal is to show that, when considering a background-only sample, the method proposed correctly rejects (42) as a suitable model for the data; whereas, when considering a dark matter sample, the dark matter model in (42) is “accepted.”

To further increase the complexity of the problem, we consider a situation where the background sample corresponds to γ -ray emissions due to a pulsar, with distribution

$$g_{\text{PS}}(y) = \frac{1}{yk_\tau} \exp\left\{-\left(\frac{y}{y_0}\right)^\tau\right\}, \quad (43)$$

with $y_0 = 0.5$, $y \in [0.5, 5]$ TeV, $\tau > 0$, and k_τ is a normalizing constant. Notice that, as discussed in [29], distinguishing γ -ray emissions due to pulsars from those due to dark matter is a particularly challenging task. The histograms of the two datasets considered are shown in Fig. 6; the overlapping curves correspond to the best fit of the models in (42) and (43) on each sample. Interestingly, for both samples, Eqs. (42) and (43) provide a very similar fit to the data; hence the importance of correctly selecting the most adequate model or excluding the dark matter hypothesis when observing emissions due to pulsars.

The upper panels of Fig. 7 display the CD plots obtained by setting $g = g_{\text{DM}}$ in (42) as a postulated model and comparing it with the distribution of the dark matter sample (upper left panel) and of the background pulsar sample (upper right panel). Remarkably, the CD plots and the adjusted deviance tests correctly lead to the conclusion that the distribution of the dark matter sample does not deviate significantly from (42), whereas the distribution of the pulsar sample does deviate substantially from (42) and the deviance test (adequately adjusted for postselection

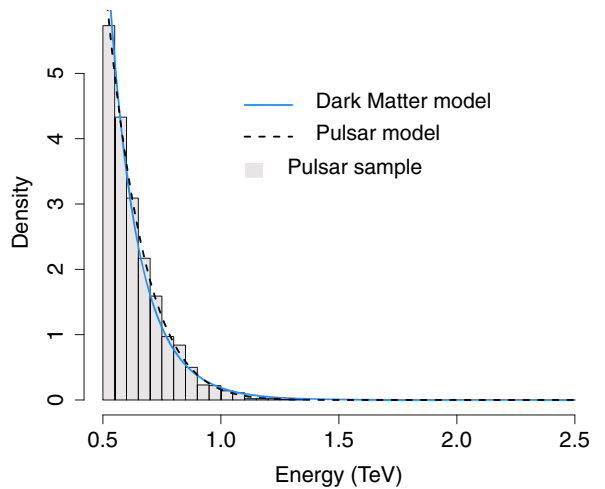


FIG. 6. Dark matter and pulsar samples. The left panel corresponds to the histogram of a sample of 2000 observations simulated from the model in (42) with $M_\chi = 2.5$. The right panel corresponds to the histogram of a sample of 2000 observations simulated from the model in (43) with $\tau = 2$. The best fit of (42) and (43) are also reported as blue solid line and black dashed lines, respectively, on top of each histogram.

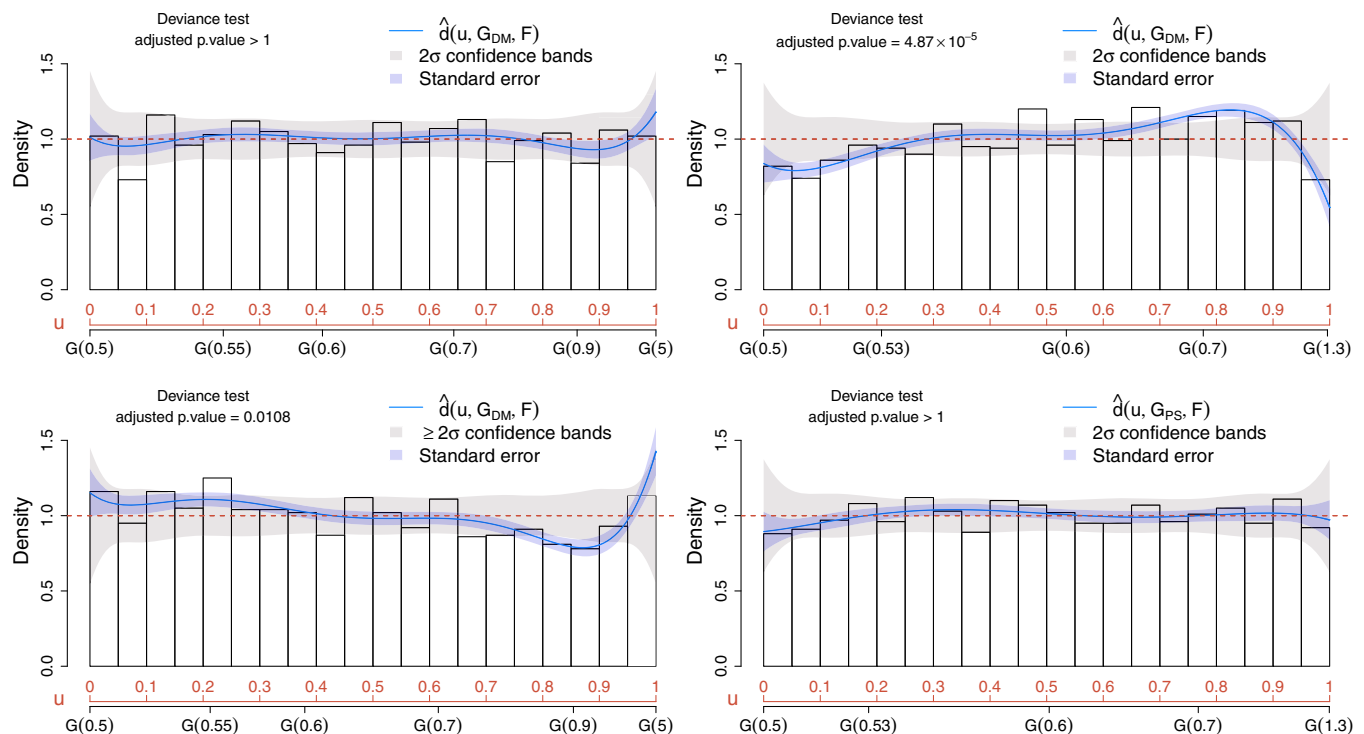


FIG. 7. CD plots and deviance tests for dark matter and pulsar samples. The upper left panel displays the CD plot and deviance test for the dark matter sample with $g = g_{DM}$ in (42). The upper right panel compares the distribution of the pulsar sample with the model in (42). The lower left panel displays the CD plot and deviance test for the dark matter sample with $g = g_{PS}$ in (43). The lower right panel displays the CD plot and deviance test for the pulsar sample with $g = g_{PS}$ in (43). For the plots on the left, the size of the basis selected is $M = 6$, whereas, for the plots on the right, $M = 5$.

inference) rejects the dark matter model with 3.897σ significance (adjusted p -value of 4.870×10^{-5}). Notice that, in both cases, we are ignoring the information regarding the pulsar distribution and the only inputs considered are the data and the signal model in (42).

Finally, when incorporating the knowledge of the pulsar distribution in (43) into the analysis, one can select between the models in (42) and (43) by constructing additional CD plots and deviance tests for both samples and setting $g = g_{PS}$ in (43). The results are shown in lower panels of Fig. 7. As expected, the dark matter model is rejected (lower left panel) with 2.297σ significance (adjusted p -value of 0.0108), whereas the pulsar model is accepted (lower right panel).

VII. BACKGROUND MISMODELING DUE TO INSTRUMENTAL NOISE AND UPPER LIMITS CONSTRUCTIONS

When conducting real data analyses one has to take into account that the data generating process is affected by both statistical and nonrandom uncertainty due to the instrumental noise. As a result, even when a model for the background is known, the data distribution may substantially deviate from it due to the smearing introduced by the detector (e.g., [30]). In order to account for the instrumental

error affecting the data, it is common practice to consider folded distributions where the errors due to the detector are often modeled assuming a normal distribution or estimated via nonparametric methods (e.g., [31,32]). In Sec. VII A, it is shown how the same approach described in Secs. V A and V B 1 can be used to assess if the instrumental error is negligible and, when not, how to update the postulated background model in order to incorporate the instrumental noise. Section VII B discusses upper limit constructions by means of comparison distributions.

A. Modeling the instrumental error

The data considered come from a simulated observation by the Fermi Large Area Telescope [33] with realistic representations of the effects of the detector and present backgrounds [21,34]. The Fermi-LAT is a pair-conversion γ -ray telescope on board the earth-orbiting Fermi satellite. It measures energies and images γ rays between about a 100 MeV and several TeV. The goal of the analysis is to assess if the data could result from the self-annihilation of a dark matter particle. Let the distribution of the astrophysical background be a power law, i.e.,

$$g_b(x) = \frac{1}{k_\phi x^{\phi+1}}, \quad (44)$$




FIG. 8. Histograms of simulated Fermi-LAT samples. The left panel displays the histogram of a source-free simulated Fermi-LAT sample of $N = 35, 157$ observations, whereas the black dashed line corresponds to the best fit of the power-law model in (44). The right panel shows the histogram of two simulated Fermi-LAT physics samples of $n = 200$ observations. The grey histogram corresponds to the background-only sample, whereas the blue histogram corresponds to the dark matter signal sample.

where k_ϕ is a normalizing constant and $x \in [1, 35]$ giga electron volt (GeV). Equation (44) corresponds to the distribution we would expect the background to follow if there was no smearing of the detector. The left panel of Fig. 8 shows the histogram of a source-free sample of 35,157 i.i.d. observations from a power-law distributed background source with index 2.4 [i.e., $\phi = 1.4$ in (44)] and contaminated by instrumental errors of unknown distribution.

In order to assess if (44) is a suitable distribution for these data, we proceed by fitting (44) via maximum likelihood and setting it as a postulated background distribution. The best fit of (44) is displayed on the left panel of Fig. 8 as a black dashed line.

We proceed estimating $d(G_b(x); G_b, F_b)$ and f_b as in (9) and (19), respectively, with $M = 4$ (chosen as in Sec. IV C). The deviance test and CD plot are reported in the left panel of Fig. 9 and suggest that significant departures from the fitted power-law model occur. This

implies that the instrumental error is not negligible, and thus, in order to account for it, we consider (45) as “calibrated” background density the model

$$\hat{f}_b(x) = \frac{1}{k_{\hat{\phi}} x^{\hat{\phi}+1}} (1 + 0.027 \text{Leg}_1[G_b(x)] - 0.067 \text{Leg}_2[G_b(x)] + 0.026 \text{Leg}_3[G_b(x)] - 0.045 \text{Leg}_4[G_b(x)]), \quad (45)$$

where $G_b(x)$ is the CDF of (44) and $\hat{\phi} = 1.359$ is the ML estimate of ϕ in (44).

For the sake of comparison, the same analysis has been repeated considering 35,157 i.i.d. observations from a power-law background source with index 2.4, without instrumental error. The respective CD plot and deviance test are shown on the right panel of Fig. 9 and indicate that the power-law model in (44), with ϕ replaced by its MLE (i.e., $\hat{\phi} = 1.391$), provides a good fit for the data; i.e., the instrumental error is, in this case, absent or negligible.

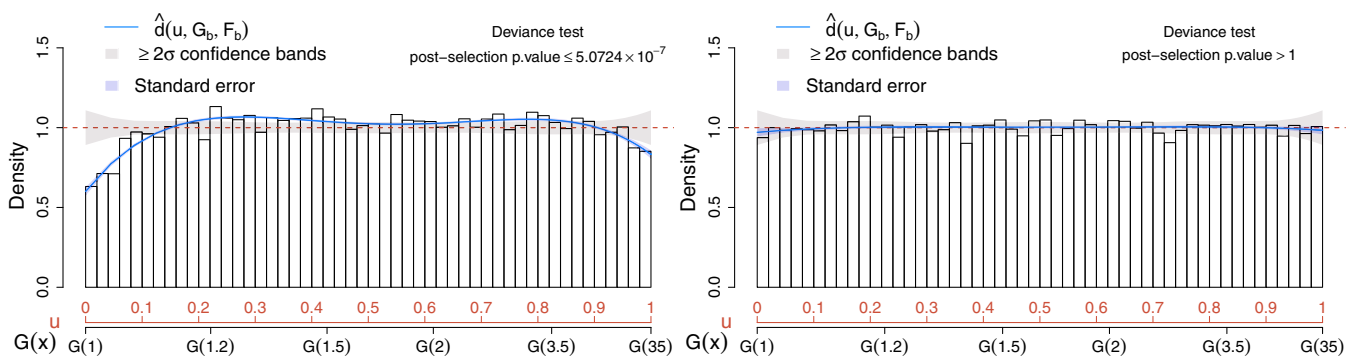


FIG. 9. Deviance test and CD plot for the source-free simulated Fermi-LAT sample including the instrumental error (left panel) and simulated sample with no instrumental error (right panel). In both cases $M = 4$.

B. Signal detection and upper limit construction

Once a calibrated background distribution is obtained, we proceed with the signal detection phase by setting $g(x) = \hat{f}_b(x)$ in (45). Similar to Sec. VB 1, two physics samples are given; one containing 200 observations from the background source distributed, as in (44), and the other containing 200 observations from a dark matter emission. The signal distribution from which the data have been simulated is the PDF of γ -ray dark matter energies in [[28], Eq. (28)] with $M_\chi = 3.5$. Both physics samples include the contamination due to the instrumental noise with unknown distribution. The respective histograms are shown in the right panel of Fig. 8.

The selection scheme in Sec. IV C suggests that no significant departure from (45) occurs on the background-only physics sample, whereas, for the signal sample, the strongest significance is observed at $M = 3$; therefore, for the sake of comparison, we choose $M = 3$ in both cases. The respective deviance tests and CD plots are reported in Fig. 10. As expected, the upper left panel of Fig. 10 shows a flat estimate of the comparison density on the background-only sample. Conversely, the upper right panel of Fig. 10 suggests that an extra bump is present over the [2, 3.5] region with 3.318σ significance (adjusted p -value = 4.552×10^{-4}). As in (39), it is possible to proceed with

the signal characterization stage (see Sec. VB 2); however, in this setting, one has to account for the fact that also the signal distribution must include the smearing effect of the detector.

As an anonymous referee pointed out, it is important to discuss how the upper limits and Brazil plot can be constructed via LP modeling and how they relate to the constructs discussed so far in this manuscript. Indeed, the confidence bands reported in the CD plots are themselves upper limits. Specifically, in the signal detection framework of Sec. VB 1, the confidence bands in (29) are constructed assuming that there is no signal in the data. Specifically, they correspond to the regions where the comparison density estimator is expected to lie, at $1 - \alpha$ confidence level, if the data include background-only events. Conversely, any deviation from the confidence bands characterizes the quantiles of the distribution where the data distribution does not conform with the one postulated under the assumption that no signal is present.

When the interest is in identifying areas of the search region where deviations from the background model occur, one can exploit the fact that $u = G(x)$, and thus upper limits and classical ‘‘Brazil plots’’ based on the comparison density can be obtained by plotting (20) and the respective confidence bands in (29) as a function of x . This is shown for our Fermi-LAT example in the bottom panels of Fig. 10.

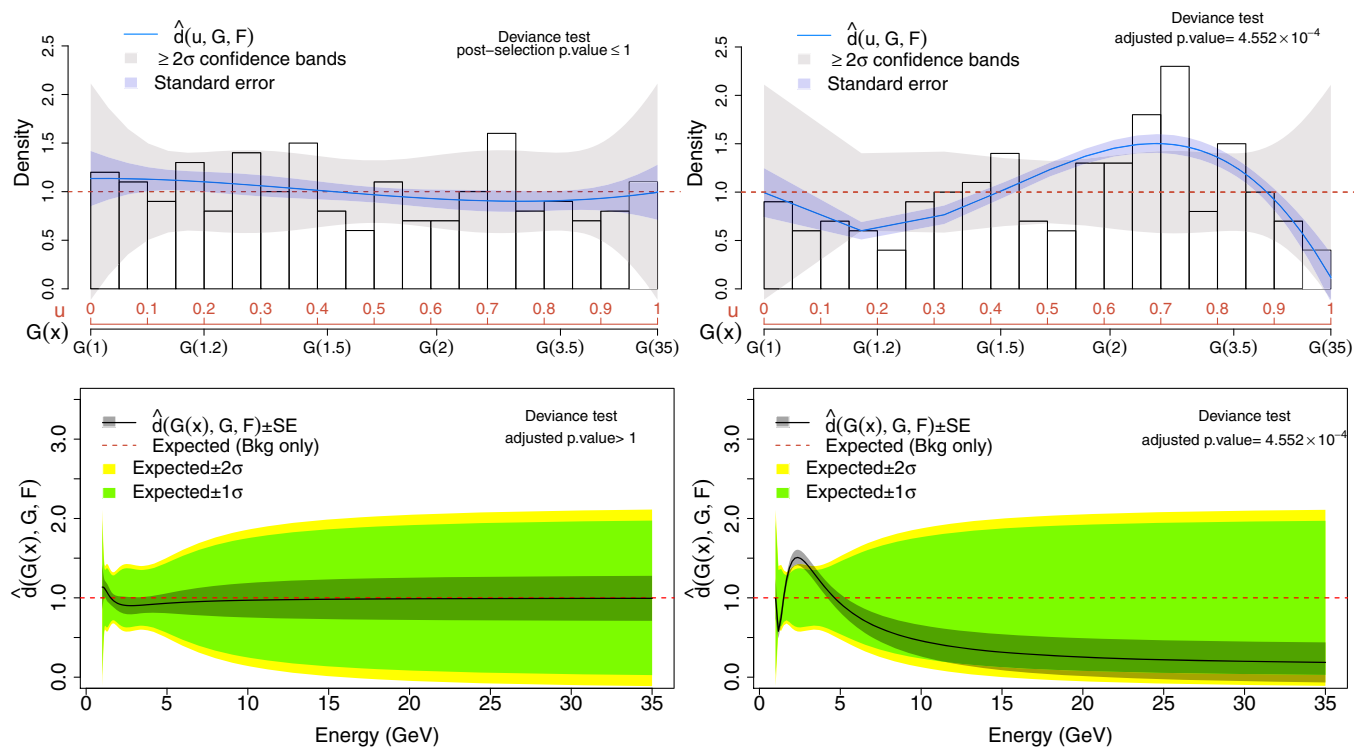


FIG. 10. Deviance test, CD plots, and Brazil plots for the simulated Fermi-LAT background-only sample of size 200 (left panels) and the simulated Fermi-LAT dark matter signal sample of size 200 (right panels). In both cases, the postulated distribution G corresponds to the CDF of the calibrated background model in (45). For the sake of comparison, $d(u; G, F)$ has been estimated via (9), with $M = 3$ in both cases.

Indeed, the upper and bottom panels in Fig. 10 carry essentially the same information in two different domains. Specifically, the CD plots display the departure of f from g in the quantile domain, whereas the Brazil plots show the same differences in the frequency domain. For signal detection purposes, the bottom panels may be preferred to identify the location where substantial deviations among the background and signal model occur, whereas the CD plots are more suitable for goodness-of-fit purposes as they provide a simultaneous visualization of the differences occurring at each quantile of the distribution.

VIII. MODEL DENOISING

As discussed in Sec. II D, the choice of M affects the resulting estimator of $d(u; G, F)$ in terms of both bias and variance. When dealing with complex background distributions, a large value of M may be necessary to reduce the bias of the estimated comparison density. At the same times, however, a large value of M leads to an inflation of the variance. In other words, considering a basis of M shifted Legendre polynomials may lead to overfitting.

Practically speaking, overfitting leads to wiggly (i.e., nonsmooth) estimates, and thus one may overcome this limitation by attempting to denoise the estimator in (9). Section VIII A reviews the model-denoising approach proposed by [8,10], whereas Sec. VIII B briefly discusses inference and model selection in this setting. Finally, Sec. VIII C compares the results obtained with a full and a denoised solution on the examples of Sec. V.

A. Denoising

Let $\widehat{LP}_1, \dots, \widehat{LP}_M$ be the estimate of the first M coefficients of the expansion in (4). The most “significant” LP_j coefficients are selected by sorting the respective \widehat{LP}_j estimates so that

$$\widehat{LP}_{(1)}^2 \geq \widehat{LP}_{(2)}^2 \geq \dots \geq \widehat{LP}_{(M)}^2$$

and choosing the value $k = 1, \dots, M$ for which the AIC(k) index in (46) is maximum,

$$\text{AIC}(k) = \sum_{j=1}^k \widehat{LP}_{(j)}^2 - \frac{2k}{n}. \quad (46)$$

The AIC-denoised estimator of $d(u; G, F)$ is given by

$$\hat{d}^*(u; G, F) = 1 + \sum_{j=1}^{k_M^*} \widehat{LP}_{(j)} \text{Leg}_{(j)}(u), \quad (47)$$

where $\widehat{LP}_{(j)}$ is the estimate whose square is the j th largest among $\widehat{LP}_1^2, \dots, \widehat{LP}_M^2$, $\text{Leg}_{(j)}(u)$ is the respective shifted Legendre polynomial and

$$k_M^* = \underset{k}{\text{argmax}} \{ \text{AIC}(1), \dots, \text{AIC}(M) \}. \quad (48)$$

Practical remarks. Recall that the first M coefficients LP_j can be expressed as a linear combination of the first M moments of U . Thus, the AIC-denoising approach selects the LP_j coefficients that carry all the “sufficient” information on the first M moments of the distribution.

B. Inference after denoising

The deviance test can be used, as in Sec. IV C, to choose the size of the initial basis of M polynomials among M_{\max} possible models. Finally, the k_M^* largest coefficients are chosen by maximizing (46). This two-step procedure selects $\hat{d}^*(u; G, F)$ in (47) from a pool of $M_{\text{tot}} = M_{\max} + \frac{M(M-1)}{2}$ possible estimators. Therefore, the Bonferroni-adjusted p -value of the deviance test is given by

$$M_{\text{tot}} \cdot P(\chi_{k_M^*}^2 > d_{k_M^*}) \quad (49)$$

with $d_{k_M^*} = \sum_{j=1}^{k_M^*} \widehat{LP}_{(j)}^2$. Similarly, confidence bands can be constructed as

$$\left[1 - c_{\alpha, M_{\text{tot}}} \sqrt{\sum_{j=1}^{k_M^*} \frac{1}{n} \text{Leg}_{(j)}^2(u)}, 1 + c_{\alpha} \sqrt{\sum_{j=1}^{k_M^*} \frac{1}{n} \text{Leg}_{(j)}^2(u)} \right], \quad (50)$$

where $c_{\alpha, M_{\text{tot}}}$ is the solution of

$$2(1 - \Phi(c_{\alpha, M_{\text{tot}}})) + \frac{k_0}{\pi} e^{-0.5c_{\alpha, M_{\text{tot}}}^2} = \frac{\alpha}{M_{\text{tot}}}. \quad (51)$$

Practical remarks. Given the possibility of denoising our solution, one may legitimately wonder why not to consider a large value of M_{\max} , e.g., $M_{\max} = 100$, and then select $k_{M_{\max}}^*$ directly. In other words, why should we first implement the procedure in Sec. IV C and, only after, refine our estimator as in Sec. VIII A and not vice versa? There are two main reasons why such an approach is discouraged.

First of all, one has to take into account that ignoring the selection stage proposed in Sec. IV C, there is no guarantee that the resulting $k_{M_{\max}}^*$ would include all the \widehat{LP}_j terms that provide the strongest evidence in favor of H_1 in (22). Therefore, the resulting p -value can in principle be lower than the one in (49). Indeed, the AIC criterion in (46), aims to improve the fit of the estimator to the data, whereas the deviance selection criteria in (32) aim to maximize the power of the inferential procedure.

Second, choosing $M_{\max} = 100$ is computationally unfeasible with most of the standard programming languages such as R and PYTHON, and the numerical computation of (9) may easily lead to divergent or inaccurate results.

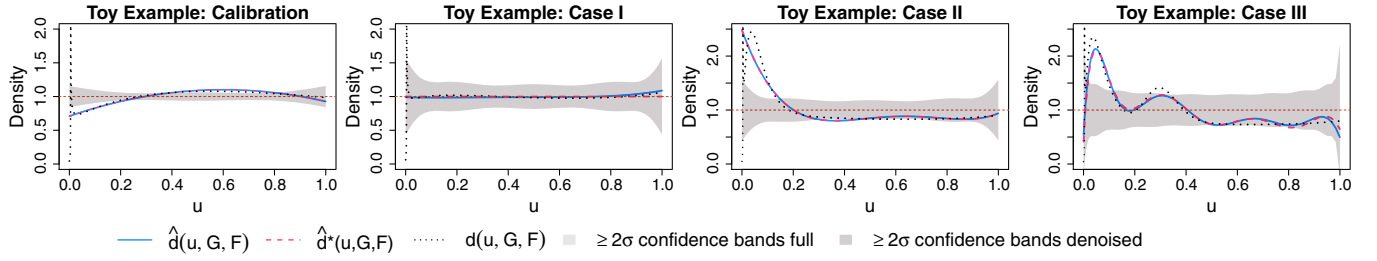


FIG. 11. Comparison of the estimators $\hat{d}(u; G, F)$ (blue solid lines) and $\hat{d}^*(u; G, F)$ (pink dashed lines) for the toy examples in Sec. V. The true comparison densities $d(u; G, F)$ are displayed as dotted black curves. The light gray areas correspond to the confidence bands of the full solution, whereas dark gray areas refer to the confidence bands of the denoised solution. In all the examples proposed, the latter is almost entirely overlapping with the former.

C. Comparing full and denoised solutions

Figure 11 compares the fit of the estimators $\hat{d}(u; G, F)$ and $\hat{d}^*(u; G, F)$ for the examples in Sec. V. For all the cases considered, M and k_M^* have been selected as in (32) and (48) (see second column of Table II). When no significance was achieved for any of the values of M considered, a small basis of $M = 3$ or $M = 4$ polynomials was chosen for the full estimator $\hat{d}(u; G, F)$, which was then further denoised in order to obtain $\hat{d}^*(u; G, F)$. Table II shows the results of the deviance tests of the full and the denoised solution for the examples in Sec. V. The unadjusted p -values and the Bonferroni-adjusted p -values are reported in the second and third columns, respectively. In half of the cases, $k_M^* = M$ and the estimators $\hat{d}(u; G, F)$ and $\hat{d}^*(u; G, F)$ overlap over the entire range $[0, 1]$. The inferential results were also approximately equivalent in the majority of the situations considered.

The main differences are observed in the analysis of the background-only physics sample (Case I). In this case, the deviance-selection procedure leads to nonsignificant results for all the values of M considered; the minimum p -value is observed at $M = 18$ (unadjusted p -value = 0.2657). In this

setting, the denoising process leads to $k_{18}^* = 2$ and the respective unadjusted p -value is 5.096×10^{-4} . This further emphasizes the importance of adjusting for model selection in order to avoid false discoveries. For modeling purposes and for the sake of comparison with the case where a signal is present, a basis of $M = 4$ was selected. Since the true distribution of the data is the same as the postulated one, the denoising process sets all the coefficients equal to zero ($k_M^* = 0$).

For Case III, only $k_M^* = 6$ out of $M = 9$ coefficients are selected when denoising (see Table II). Despite the right panel of Fig. 11 showing that the full and the denoised solutions are almost overlapping, the latter leads to an increased sensitivity (adjusted p -value = 2.496×10^{-28}) compared to the full solution (adjusted p -value = 5.181×10^{-27}).

These results suggest that the denoising approach can easily adapt to situations where a sparse solution is preferable (i.e., when only a few of the M coefficients LP_j are nonzero) without enforcing sparsity when many of the M coefficients considered are needed to adequately fit the data (e.g., bottom right panel of Fig. 11). From an inferential perspective, denoising can improve the sensitivity of the analysis; however, in order to avoid false discoveries, extra care needs to be taken when the deviance

TABLE II. Model selection and inference for the toy example in Sec. V. The second column reports the M and k_M^* values selected as in (32) and (48), respectively. The third column collects the unadjusted deviance p -values for the full and denoised solutions. The Bonferroni-adjusted p -values, computed as in (33) and (49), are reported in the fourth column. The correction terms applied correspond to $M_{\max} = 20$ for the full solution and $M_{\text{tot}} = M_{\max} + \frac{M(M-1)}{2}$ for the denoised solution.

	M, k_M^* selected	Method	Deviance p -values	Adjusted p -values
Toy example Calibration	$M = 2$	Full	$p(2) = 3.199 \times 10^{-12}$	$20 \times p(2) = 6.397 \times 10^{-11}$
	$k_2^* = 2$	Denoised	$p(2) = 3.199 \times 10^{-12}$	$21 \times p(2) = 6.717 \times 10^{-11}$
Toy example Case I	$M = 18$	Full	$p(18) = 0.2657$	$20 \times p(18) > 1$
	$k_{18}^* = 2$	Denoised	$p(2) = 5.096 \times 10^{-4}$	$21 \times p(2) = 0.0882$
Toy example Case II	$M = 4$	Full	$p(4) = 8.994 \times 10^{-33}$	$20 \times p(4) = 1.799 \times 10^{-31}$
	$k_4^* = 4$	Denoised	$p(4) = 8.994 \times 10^{-33}$	$21 \times p(4) = 2.338 \times 10^{-31}$
Toy example Case III	$M = 9$	Full	$p(9) = 2.590 \times 10^{-28}$	$20 \times p(9) = 5.181 \times 10^{-27}$
	$k_9^* = 6$	Denoised	$p(6) = 4.457 \times 10^{-30}$	$35 \times p(1) = 2.496 \times 10^{-28}$

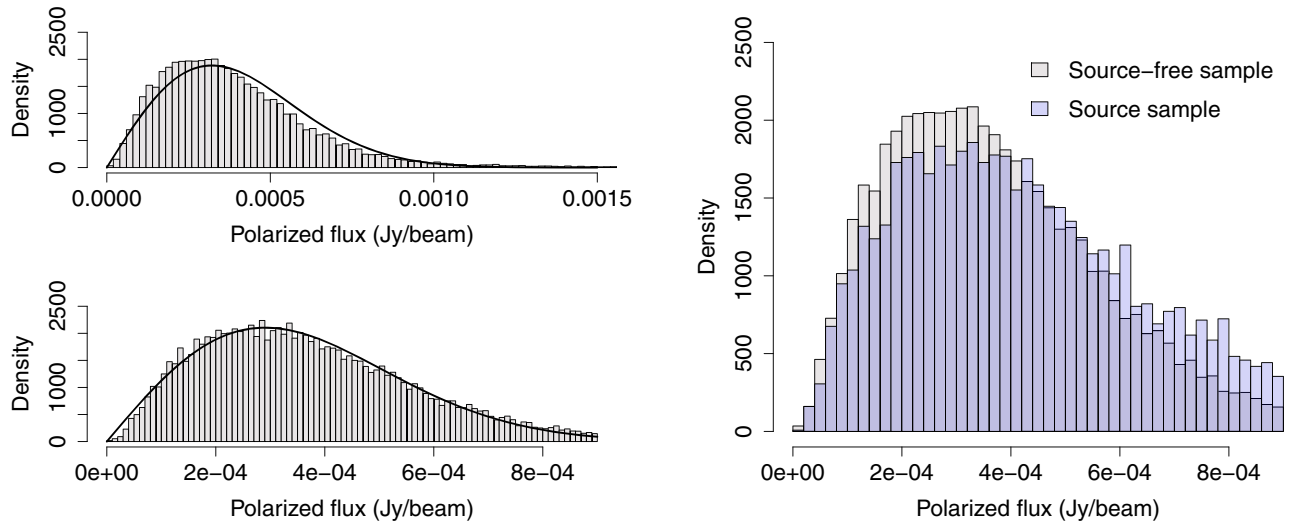


FIG. 12. Histograms of the NVSS samples. The left panels show the source-free sample with and without outliers (upper left and lower left panels, respectively). In both cases, the best fit of the Rayleigh model in (53) is displayed as a black curve. The right panel compares the source-free sample without outliers (grey histogram) with the source sample (blue histogram) truncated over the search area considered.

selection procedure leads to large p -values for all the M_{\max} models considered.

IX. AN APPLICATION TO STACKING EXPERIMENTS

In radio astronomical surveys, stacking techniques are often used to combine noisy images or “stacks” in order to increase the signal-to-noise ratio and improve the sensitivity of the analysis in detecting faint sources (e.g., [35–37]). In polarized signal searches, for instance, a faint population of sources is considered when the median polarized intensity observed over control regions differs significantly from the median of the region where the sources are expected to be present. In this context, under simplifying assumptions, the distribution of the intensity of the source polarization is often assumed to have Rice distribution, i.e.,

$$f(x) = \frac{x e^{-\frac{x^2 + \nu^2}{2\sigma^2}}}{k_\nu \sigma^2} \text{Bessel}\left(\frac{x\nu}{\sigma^2}\right), \quad (52)$$

where $\text{Bessel}(\cdot)$ denotes the Bessel function of the first kind of order zero and $k_\nu \sigma^2$ is a normalizing constant. Furthermore, Eq. (52) reduces to a Rayleigh PDF when no signal is present [38], i.e., when $\nu = 0$. Below, it is shown how the methods described in Secs. VA and VB 1 can be used to assess whether the Rayleigh distribution is a reliable model for the background and, when too simplistic, investigate the impact of incorrectly assuming a Rayleigh distribution on the reliability of the analysis.

The data considered come from the NRAO VLA Sky Survey (NVSS) [39]. The NVSS is an astronomical survey of the Northern hemisphere carried out by the Very Large

Array of the National Radio Astronomy Observatory. The NVSS has detected 1.8×10^6 sources in total intensity, but only 14% of these have reported a polarized signal peak greater than 3σ [37]. The original source-free sample contained 29 915 observations collected from four different control regions for each source with a brightness in total intensity between 0 and 0.0093 Jy/beam (see upper panel of Fig. 12). However, such a sample appears to contain several outliers that affect the data distribution, making it far from Rayleigh. A better Rayleigh fit is obtained when removing the outliers⁶ (see bottom left panel of Fig. 12). Since understanding the cause of these anomalous observations is beyond the scope of this manuscript, we proceed excluding them from the analysis, and we focus on assessing the validity of the Rayleigh assumption on the remaining 28 739 observations on the region $[0, 0.0009]$ Jy/beam. It has to be noted that the nominal noise in NVSS polarization is 0.00029 Jy/beam and we may expect as a reasonable threshold for the detection of one individual source to be 3 times the noise. Hence, a source sample of 6,220 observations has been selected from positions where compact radio sources with a brightness in total intensity between 0 and 0.0009 Jy/beam are known to be present. Both source-free and source samples are assumed to be i.i.d. The histograms of the source-free and signal samples considered are shown in the right panel of Fig. 12.

As first step, we fit a Rayleigh distribution (adequately truncated over the range $[0, 0.0009]$) on the source-free sample, i.e.,

⁶In statistics, an observation x_i is considered an outlier if $x_i < Q_{0.25} - 1.5[Q_{0.75} - Q_{0.25}]$ or $x_i > Q_{0.75} + 1.5[Q_{0.75} - Q_{0.25}]$ where $Q_{0.25}$ and $Q_{0.75}$ are the first and the third sample quartiles.




FIG. 13. Deviance test and CD plot for the NVSS source-free sample of size 28 739 compared to the Rayleigh distribution in (53).

$$g_b(x) = \frac{x e^{-x^2}}{k_{\hat{\sigma}^2}}, \quad (53)$$

where $k_{\hat{\sigma}^2}$ is a normalizing constant, $\hat{\sigma} = 0.0003$ is the ML estimate of the unknown parameter σ , and $x \in [0, 0.0009]$ Jy/beam. In order to assess if (53) provides a good fit for the data, we estimate the comparison density $d(G_b(x); G_b, F_b)$ by, first, selecting M as in (32) and then applying the AIC-based denoising approach described in Sec. VIII A. In this case, the denoised solution selects $k^* = 9$ out of $M = 10$ polynomial terms. The deviance tests and the CD plot in Fig. 13 suggest that, despite the fact that the median of the data coincides with the one of the Rayleigh model, overall, the latter does not provide a good fit for the distribution of the source-free sample. Specifically, the data distribution shows a higher right tail than one expected under the Rayleigh assumption, whereas the first quantiles are overestimated by the Rayleigh. Therefore, the researcher can either decide to use a more refined parametric model for the background or consider the calibrated background distribution of the form in (19), which in our setting specifies as

$$\begin{aligned} \hat{f}_b(x) = & \frac{x e^{-x^2}}{k_{\hat{\sigma}^2}} (1 - 0.018 \text{Leg}_1[G_b(x)] + 0.012 \text{Leg}_2[G_b(x)] \\ & + 0.052 \text{Leg}_3[G_b(x)] - 0.014 \text{Leg}_4[G_b(x)] \\ & + 0.047 \text{Leg}_5[G_b(x)] - 0.018 \text{Leg}_6[G_b(x)] \\ & + 0.031 \text{Leg}_7[G_b(x)] + 0.016 \text{Leg}_9[G_b(x)] \\ & - 0.015 \text{Leg}_{10}[G_b(x)]), \end{aligned} \quad (54)$$

where $G_b(x)$ is the CDF of (53).

The strategy described in Sec. V B 1 allows us to identify where significant differences between the control and source sample occur. In order to assess the effect of incorrectly assuming a Rayleigh background, we compare the distribution of the physics sample with both the Rayleigh and the calibrated background distribution in (54). Figure 14 reports deviance tests and CD plots obtained on the physics sample when setting $g(x) = \hat{f}_b(x)$ in (53) (left panel) and $g(x) = \hat{f}_b(x)$ in (54) (right panel). Both analyses provide strong evidence that the distribution of the physics sample differs significantly from the postulated models $\hat{f}_b(x)$ and $g_b(x)$, and the most substantial discrepancies occur on the right tail of the distribution. However, since the Rayleigh model underestimates the right tail of the background distribution (see Fig. 13), it leads to an artificially enhanced sensitivity in this region. The differences between the two CD plots are less prominent around the median expected under $\hat{f}_b(x)$ and $g_b(x)$ (i.e., in correspondence with $u = 0.5$ in both plots).

Figure 14 suggests that, for these data, assuming a background Rayleigh distribution would not substantially affect the results of a comparison between the source-free and signal sample based on the median. However, focusing solely on the median can strongly limit the overall sensitivity of the analysis since the major differences occur at the higher quantiles of the distribution. On the other

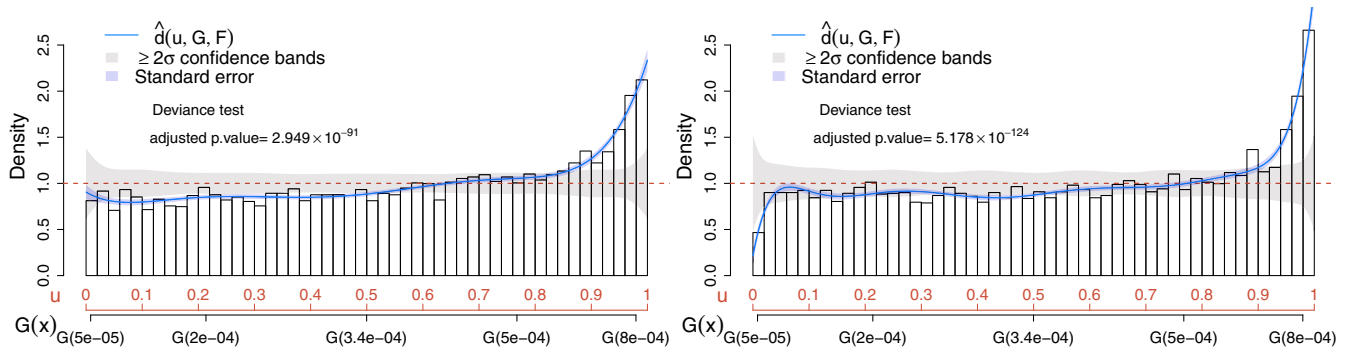


FIG. 14. Deviance tests and CD plots for the NVSS source sample assuming $g(x)$ to be the calibrated background model in (54) (left panel) and when letting $g(x)$ be the PDF of the truncated Rayleigh distribution in (53). In both cases the estimator of the comparison density has been denoised as described in Sec. VIII A. The values of M and k^* considered are $M = k^* = 6$ and $M = 9$, $k^* = 8$ for the estimators on the left and right panels, respectively.

hand, assuming a Rayleigh distribution for the background would artificially inflate the evidence in favor of the source. Specifically, the sigma significance of the deviance test obtained under the Rayleigh background assumption is 23.655σ (adjusted p -value = 5.178×10^{-124}), whereas the one obtained using (54) is 20.225σ (adjusted p -value = 2.948×10^{-91}).

Conversely, the calibrated background model in (54) allows us to safely compare the entire distribution of the polarized intensity in the source and control regions via CD plots and deviance tests without affecting the sensitivity of the analysis.

X. DISCUSSION

This article proposes a unified framework for signal detection and characterization under background mismodeling. From a methodological perspective, the methods presented here extend LP modeling to the inferential setting.

The solution discussed is articulated in two main phases: a calibration phase where the background model is “trained” on a source-free sample and a signal search phase conducted on the physics sample collected by the experiment. If a model for the signal is given, the method proposed allows the identification of hidden signals from new unexpected sources and/or the refining of the postulated background or signal distributions. Furthermore, the tools presented in this manuscript can easily be extended to situations where a source-free sample is not available and the background is unknown (up to some free parameters). As discussed in Sec. VI, however, in this setting the signal distribution is required to be known, and the physics sample is expected to contain only signal-like events; i.e., the background is almost completely reduced.

The theory of Sec. II D and the analyses in Sec. V have highlighted that, despite a fully nonparametric approach providing reliable inference, it may lead to unsatisfactory estimates when the postulated PDF g is substantially different from the true density f . In this setting, a semi-parametric stage can be performed in order to provide a reliable model for the data.

Each individual step in both the nonparametric and the semiparametric stages of Secs. V B 2 and V A provides useful scientific insights on the signal and background distribution. Hence, an automatized implementation of the steps of Algorithm 1 based solely on the p -values of the deviance tests is discouraged as it would lead to a substantial loss of scientific knowledge on the phenomena under study.

Finally, it is important to point out that, despite this article’s focus on the one-dimensional searches on continuous data, all the constructs presented in Secs. II and the deviance test in IV A also apply to the discrete case when considering i.i.d. events. More work is needed to extend these results and those of Sec. IV B to searches in multiple dimensions and when considering Poisson events with a functional mean. In the first case the difficulty mainly lies

in generalizing the constructs of Sec. IV to account for the dependence structure occurring across multiple dimensions. In the second case, the main challenge lies in identifying the equivalent of (1) to model the mean of the distribution, while incorporating the Poisson error.

A. Code availability

The LPBkg PYTHON package [40] and the LPBkg R package [41] allow the implementation of the methods proposed in this manuscript. Detailed tutorials on how to use the functions provided are also available at <http://salgeri.umn.edu/my-research>.

ACKNOWLEDGMENTS

The author thanks Jeroen Stil, who provided the NVSS datasets used in Sec. IX, and Lawrence Rudnick, who first recognized the usefulness of the method proposed in the context of stacking experiments. Conversations with Subhadeep Mukhopadhyay have been of great help when this work was first conceptualized. Discussions and e-mail exchanges with Charles Doss and Chad Shafer are gratefully acknowledged. Finally, the author thanks an anonymous referee whose feedback has been substantial to improve the overall quality of the paper.

APPENDIX A: MOMENTS OF THE \widehat{LP}_j ESTIMATES

Consider the general setting where $f \neq g$ and thus $d(u; G, F) \neq 1$ over $[0, 1]$. It follows that each u_i is independently and identically distributed with PDF $d(u; G, F)$; hence, all the expectations in $E[\widehat{LP}_j]$, $V(\widehat{LP}_j)$ and $Cov(\widehat{LP}_j, \widehat{LP}_k)$ are taken with respect to $d(u; G, F)$. Specifically,

$$\begin{aligned} E[\widehat{LP}_j] &= E\left[\frac{1}{n} \sum_{i=1}^n Leg_j(U_i)\right] \\ &= E[Leg_j(U)] \\ &= \int_0^1 Leg_j(u) d(u; G, F) \partial u = LP_j, \end{aligned}$$

where the second equality follows from the fact that each observed value u_i is a realization of a random variable U_i and each U_1, \dots, U_n is identically distributed as the random variable U , whose PDF is given by the comparison density $d(u; G, F)$. Notice that $d(u; G, F) = 1$ implies that $\int_0^1 Leg_j(u) d(u; G, F) \partial u = \int_0^1 Leg_j(u) \partial u = 0$, from which the first equivalence in (8) follows. Moreover,

$$\begin{aligned} V(\widehat{LP}_j) &= \frac{1}{n^2} V\left(\sum_{i=1}^n Leg_j(U_i)\right) \\ &= \frac{1}{n} V(Leg_j(U)) = \frac{\sigma_j^2}{n}, \end{aligned}$$

where $V(\text{Leg}_j(U)) = \int_0^1 (\text{Leg}_j(u) - LP_j)^2 d(u; G, F) \partial u = \sigma_j^2$. The second equality holds because of independence and identical distribution of each u_i . Notice that if $d(u; G, F) = 1$, $\sigma_j^2 = 1$ in virtue of the orthonormality of the $\text{Leg}_j(u)$ polynomials. Hence the second equivalence in (8) holds. Finally,

$$\begin{aligned} \text{Cov}(\widehat{LP}_j, \widehat{LP}_k) &= \text{Cov}\left(\frac{1}{n} \sum_{i=1}^n \text{Leg}_j(U_i), \frac{1}{n} \sum_{i=1}^n \text{Leg}_k(U_i)\right) \\ &= \frac{1}{n} \text{Cov}(\text{Leg}_j(U), \text{Leg}_k(U)) = \frac{\sigma_{jk}}{n}; \end{aligned}$$

also in this case, the second equality follows by independence and identical distribution of each u_i and

$$\begin{aligned} \text{Cov}(\text{Leg}_j(U), \text{Leg}_k(U)) &= \int_0^1 (\text{Leg}_j(u) - LP_j)(\text{Leg}_k(u) - LP_k) d(u; G, F) \partial u \\ &= \int_0^1 \text{Leg}_j(u) \text{Leg}_k(u) d(u; G, F) \partial u - LP_j LP_k \\ &= \sigma_{jk}. \end{aligned} \quad (\text{A1})$$

Because of the orthogonality of the $\text{Leg}_j(u)$, $\sigma_{jk} = 0$ when $d(u; G, F) = 1$, hence the third equivalence in (8).

APPENDIX B: BIAS, VARIANCE, AND MISE OF $\hat{d}(u; G, F)$

Given a point u over $[0, 1]$, the bias of (9) at u is

$$\text{Bias}[\hat{d}(u; G, F)] = E[\hat{d}(u; G, F)] - d(u; G, F) \quad (\text{B1})$$

$$= \sum_{j=1}^M E[\widehat{LP}_j] \text{Leg}_j(u) - \sum_{j>0} LP_j \text{Leg}_j(u) \quad (\text{B2})$$

$$= \sum_{j>M} LP_j \text{Leg}_j(u); \quad (\text{B3})$$

here (B2) follows from (4) and (9), whereas the integrated squared bias is

$$\text{IBS} = \int_0^1 \left(\sum_{j>M} LP_j \text{Leg}_j(u) \right)^2 \partial u \quad (\text{B4})$$

$$= \sum_{j>M} LP_j^2 \int_0^1 \text{Leg}_j^2(u) \partial u \quad (\text{B5})$$

$$+ 2 \sum_{M < j < k} LP_j LP_k \int_0^1 \text{Leg}_j(u) \text{Leg}_k(u) \partial u \quad (\text{B6})$$

$$= \sum_{j>M} LP_j^2, \quad (\text{B7})$$

where (B7) holds because of orthonormality of the $\text{Leg}_j(u)$ polynomials. Notice that

$$\text{IBS} = \sum_{j>M} LP_j^2 = \sum_{j>0} LP^2 - \sum_{j=1}^M LP^2 \quad (\text{B8})$$

$$= \int_0^1 (d(u; G, F) - 1)^2 \partial u - \sum_{j=1}^M LP^2 \quad (\text{B9})$$

$$\int_0^1 \left(\frac{f(G^{-1}(u)) - g(G^{-1}(u))}{g(G^{-1}(u))} \right)^2 \partial u - \sum_{j=1}^M LP_j^2, \quad (\text{B10})$$

where (B9) follows by Parseval's identity, whereas (B10) follows from (2).

The variance of (9) at a given point u is given by

$$V[\hat{d}(u; G, F)] = V\left(\sum_{j=1}^M \widehat{LP}_j \text{Leg}_j(u)\right) \quad (\text{B11})$$

$$= \sum_{j=1}^M \text{Leg}_j^2(u) V(\widehat{LP}_j) \quad (\text{B12})$$

$$+ 2 \sum_{j < k} \text{Leg}_j(u) \text{Leg}_k(u) \text{Cov}(\widehat{LP}_j, \widehat{LP}_k) \quad (\text{B13})$$

$$= \sum_{j=1}^M \frac{\sigma_j^2}{n} \text{Leg}_j^2(u) + 2 \sum_{j < k} \frac{\sigma_{jk}}{n} \text{Leg}_j(u) \text{Leg}_k(u). \quad (\text{B14})$$

By orthonormality of the polynomials $\text{Leg}_j(u)$, the integral of (B14) over $[0, 1]$ is

$$\int_0^1 V[\hat{d}(u; G, F)] \partial u = \sum_{j=1}^M \frac{\sigma_j^2}{n}; \quad (\text{B15})$$

also in this case, equality follows by orthonormality of the $\text{Leg}_j(u)$. Finally, the MISE is

$$\text{MISE}[\hat{d}(u; G, F)] = E\left[\int_0^1 (\hat{d}(u; G, F) - d(u; G, F))^2 \partial u\right] \quad (\text{B16})$$

$$= \int_0^1 E[(\hat{d}(u; G, F) - d(u; G, F))^2] \partial u \quad (\text{B17})$$

$$= \int_0^1 V[\hat{d}(u; G, F)] \quad (\text{B18})$$

$$+ (E[\hat{d}(u; G, F)] - d(u; G, F))^2 \partial u \quad (\text{B19})$$

$$= \sum_{j=1}^M \frac{\sigma_j^2}{n} + \sum_{j>M} LP_j^2, \quad (\text{B20})$$

where (B17) holds because of the Fubini-Tonelli theorem, whereas the last equality follows by (B7) and (B15).

- [1] E. Aprile *et al.*, *Phys. Rev. Lett.* **119**, 181301 (2017).
- [2] R. Agnese *et al.*, *Phys. Rev. D* **99**, 062001 (2019).
- [3] R. Smith and E. Thrane, *Phys. Rev. X* **8**, 021019 (2018).
- [4] A. M. Sirunyan *et al.*, *Phys. Lett. B* **793**, 320 (2019).
- [5] S. Yellin, *Phys. Rev. D* **66**, 032005 (2002).
- [6] N. Priel, L. Rauch, H. Landsman, A. Manfredini, and R. Budnik, *J. Cosmol. Astropart. Phys.* **05** (2017) 013.
- [7] P. D. Dauncey, M. Kenzie, N. Wardle, and G. J. Davies, *J. Instrum.* **10**, P04015 (2015).
- [8] S. Mukhopadhyay and E. Parzen, [arXiv:1405.2601](https://arxiv.org/abs/1405.2601).
- [9] S. Mukhopadhyay and K. Wang, [arXiv:1810.01724v2](https://arxiv.org/abs/1810.01724v2) [Biometrika (to be published)]; Supplement available at: http://unitedstatalgo.com/wp-content/uploads/2019/09/LPKsample_Supp.pdf.
- [10] S. Mukhopadhyay, *Electron. J. Stat.* **11**, 215 (2017).
- [11] S. Mukhopadhyay and D. Fletcher, *Sci. Rep.* **8**, 9983 (2018).
- [12] S. Mukhopadhyay and E. Parzen, *J. Risk Financ. Manage.* **11**, 37 (2018).
- [13] S. Mukhopadhyay, *Biometrics* **72**, 325 (2016).
- [14] S. Mukhopadhyay, *J. Nonparametr. Stat.* **30**, 1003 (2018).
- [15] E. Parzen, *Stat. Sci.* **19**, 652 (2004).
- [16] S. J. Sheather and M. C. Jones, *J. R. Stat. Soc. Ser. B* **53**, 683 (1991).
- [17] L. Wasserman, *All of Nonparametric Statistics* (Springer, New York, 2006).
- [18] R. S. Pilla, C. Loader, and C. C. Taylor, *Phys. Rev. Lett.* **95**, 230202 (2005).
- [19] X. Shen, H. C. Huang, and J. Ye, *J. Am. Stat. Assoc.* **99**, 751 (2004).
- [20] H. Leeb and B. M. Pötscher, *Econometric Theory* **21**, 21 (2005).
- [21] S. Algeri, D. A. van Dyk, J. Conrad, and B. Anderson, *J. Instrum.* **11**, P12010 (2016).
- [22] E. Gross and O. Vitells, *Euro. Phys. J. C* **70**, 525 (2010).
- [23] O. J. Dunn, *J. Am. Stat. Assoc.* **56**, 52 (1961).
- [24] T. W. Anderson and D. A. Darling, *Ann. Math. Stat.* **23**, 193 (1952).
- [25] D. A. Darling, *Ann. Math. Stat.* **28**, 823 (1957).
- [26] F. Wilcoxon, *Biometrics* **3**, 119 (1947).
- [27] S. Efromovich, *Nonparametric Curve Estimation: Methods, Theory, and Applications* (Springer, New York, 1999).
- [28] L. Bergström, P. Ullio, and J. H. Buckley, *Astropart. Phys.* **9**, 137 (1998).
- [29] E. A. Baltz, E. T. James, and L. W. Lawrence, *Astrophys. J. Lett.* **659**, L125 (2007).
- [30] L. Lyons, in *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding: CERN, Geneva, Switzerland, 2011* (2011), pp. 225–228, <https://cds.cern.ch/record/1306523/files/CERN-2011-006.pdf>.
- [31] V. Panaretos, in *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding: CERN, Geneva, Switzerland, 2011* (2011), pp. 229–239, <https://cds.cern.ch/record/1306523/files/CERN-2011-006.pdf>.
- [32] V. Blobel, in *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding: CERN, Geneva, Switzerland, 2011* (2011), pp. 240–251, <https://cds.cern.ch/record/1306523/files/CERN-2011-006.pdf>.
- [33] W. B. Atwood *et al.*, *Astrophys. J.* **697**, 1071 (2009).
- [34] S. Algeri, J. Conrad, and D. A. van Dyk, *Mon. Not. R. Astron. Soc. Lett.* **458**, L84 (2016).
- [35] S. Brown, A. Emerick, L. Rudnick, and G. Brunetti, *Astrophys. J. Lett.* **740**, L28 (2011).
- [36] R. L. White, D. J. Helfand, R. H. Becker, E. Glikman, and W. de Vries, *Astrophys. J.* **654**, 99 (2007).
- [37] J. M. Stil, B. W. Keller, S. J. George, and A. R. Taylor, *Astrophys. J.* **787**, 99 (2014).
- [38] J. F. L. Simmons and B. G. Stewart, *Astron. Astrophys.* **142**, 100 (1985).
- [39] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick, *Astron. J.* **115**, 1693 (1998).
- [40] Z. Zhang, LPBkg: Detecting new signals under background mismodelling, 2019, <https://pypi.org/project/LPBkg/>.
- [41] S. Algeri and H. Liu, LPBkg: Detecting new signals under background mismodelling, 2019, <https://CRAN.R-project.org/package=LPBkg>.