Note on IXPE Statistics

Steven Ehlert, Steve O'Dell, Brian Ramsey, & Martin C. Weisskopf

September 14, 2022

1 Introduction

This note recommends practices for statistical treatment of IXPE results. It addresses detection and measurement of linear polarization using basic Gaussian statistics for the Stokes parameters, recognizing that more sophisticated approaches may be appropriate for addressing complex questions. Section 2 summarizes the recommendations. Remaining sections document these recommendations and contain material already familiar to many readers. Section 3 describes characterization of linear polarization using Stokes parameters (q, u) and using polarization parameters (Π, ψ); Section 4 sketches the analysis of x-ray polarimetry using XSPEC.

2 Summary

This section summarizes recommendations and suggestions for reporting basic statistical results from analyses of IXPE polarization data. Although some suggestions—e.g., choice of confidence intervals—are arbitrary, their goal is to standardize reporting of IXPE results. Derivations are based upon the most fundamental measurements made by IXPE—i.e., the normalized Stokes parameters q and u.

2.1 Key Variables

- The subscript 0 in q_0 and u_0 typically indicates "ground-truth" values, estimated by averaging q and u over a large number of events. Throughout this note, q_0 and u_0 denote either the "ground-truth" values or measured mean values, depending upon context.
- The symbol σ denotes the Gaussian standard deviation about q_0 and u_0. By virtue of the central limit theorem, the Stokes parameters q and u are independent, Gaussian distributed variables with equal standard deviations σ.
- The symbols Π and ψ denote the (linear) polarization degree and (electric) position angle, respectively. These are derived from q and u (§3.1). Similarly, Π_0 and ψ_0 denote either "ground-truth" or measured mean values, depending upon context.

2.2 Stokes parameters

Work with Stokes parameters (q, u) until polarization parameters (Π, ψ) are needed for reporting.

2.3 Significance indicator

The main indicator of statistical significance is $(\Pi_0/\sigma) = (\sqrt{q_0^2 + u_0^2}/\sigma)$. Indeed, as χ^2 is the relevant statistic, the range of a polarization parameter (q, u, Π, or ψ) at given confidence is proportional to σ.

2.4 Detection

Polarization detection at confidence C can be claimed if $\Pi > \mathrm{MDP}_C = \sqrt{-2\ln(1-C)}\ \sigma$. Suggested descriptions of a detection are "probable" ($C > 99\%$), "highly probable" ($C > 99.9\%$), and "secure" ($C > 99.99\%$). If the observation is split into independent data sets, use Equation 8 (§3.4).

2.5 Measurement

Polarization measurements should be reported using 2-D error contours (§2.5.1), especially when the result is not highly significant (see §3.4 for details). For secure detections of polarization, 1-D 1-sigma error bars (§2.5.2) may be presented instead. For tabulating polarization results, 1-D errors are typically required.

2.5.1 2-D error contours

Error contours on (q, u) or (Π, ψ) are based upon χ^2_2 and specified confidence levels $\{C\}$. Suggested confidence levels are $\{50\%, 90\%, 99\%, 99.9\%\}$. While (q, u) error contours are circles, (Π, ψ) error contours are complicated unless (Π_0/σ) is large. Consequently, (Π, ψ) error contours ("protractor" plots) should be presented, especially when the result is not highly significant. The suggested position-angle range for these plots is $(-90°, +90°)$ with respect to North.

2.5.2 1-D error bars

It is often useful to present the uncertainties for each of the polarization parameters **without regard to the particular values of the other parameters**. The distribution of values for a polarization parameter p—i.e., q, u, Π or ψ—are Gaussian distributed (or approximately so for ψ) when the detection is secure (i.e., when Π is several times larger than σ). The suggested reporting of such errors is the one-dimensional 1-sigma statistical confidence interval—i.e., the Gaussian standard deviation σ_Π for polarization parameter Π. Thus, the measurement results would be stated as $\Pi = \Pi_0 \pm \sigma_\Pi$, where $\sigma_q = \sigma_u = \sigma_\Pi = \sigma$ and $\sigma_\psi \approx \sigma/(2\sqrt{\Pi_0\Pi}) \rightarrow \sigma/(2\Pi_0)$.

3 Characterization

One may characterize linear polarization using normalized Stokes parameters $q = Q/I$ and $u = U/I$ or polarization degree Π and position angle ψ. The two Stokes parameters are "well-behaved" in that they are independent, normally distributed random variables that are straightforward to manipulate. Most of the interesting physical questions involve estimating Π_0 and ψ_0, which are the true values of the X-ray polarization degree and (electric) position angle, respectively. These are not known *a priori*, so we are interested in estimating their values from the IXPE measurements of q and u. **Thus, perform calculations using Stokes parameters q and u, then transform to Π and ψ, if desired.**

3.1 Transformation

The transformation between the Stokes parameters (q, u) and polarization parameters (Π, ψ) is similar to that between Cartesian (x, y) and polar (r, ϕ) coordinates, except that (Π, ψ) are defined only over the half-plane. Like polar coordinates (r, ϕ), polarization coordinates (Π, ψ) suffer a coordinate singularity at the origin.

$$\Pi = \sqrt{q^2 + u^2} \tag{1}$$

$$\psi = \frac{1}{2}\arctan(\frac{u}{q}) \tag{2}$$

$$q = \Pi\cos(2\psi) \tag{3}$$

$$u = \Pi\sin(2\psi) \tag{4}$$

3.2 Distribution

The Stokes parameters q and u are each normally distributed with means q_0 and u_0, respectively, and with standard deviations $\sigma_q = \sigma_u = \sigma$. Consequently, an error contour in (q, u) is a circle of radius ϵ centered on q_0 and u_0, where $(\epsilon/\sigma)^2$ is distributed as χ^2_2 — i.e, χ^2 on two degrees of freedom. The probability α that ϵ exceeds ϵ_C is then

$$\alpha(\epsilon > \epsilon_C) = 1 - C(\epsilon < \epsilon_C) = 1 - \mathrm{CDF}((\epsilon_C/\sigma)^2, 2) = \exp[-\frac{1}{2}(\frac{\epsilon_C}{\sigma})^2] \ , \tag{5}$$

where $C(\epsilon < \epsilon_C)$ is the confidence level that ϵ does not exceed ϵ_C. Here, $\mathrm{CDF}(\chi^2, \nu)$ is the cumulative distribution function of χ^2 on ν degrees of freedom, which is particularly simple for $\nu = 2$.

Inverting, the error contour radius for a given probability α (or confidence level C) is

$$\epsilon_C = \sqrt{-2\ln(\alpha)}\ \sigma = \sqrt{-2\ln(1-C)}\ \sigma \ . \tag{6}$$

Table 1 uses Equation 6 to calculate error radii for a large range of confidence levels (C).

α	C	(ϵ_C/σ)	χ^2_2
0.500000	0.500000	1.177	1.386
0.100000	0.900000	2.146	4.605
0.010000	0.990000	3.035	9.210
0.001000	0.999000	3.717	13.81
0.000100	0.999900	4.292	18.42
0.000010	0.999990	4.799	23.03
0.000001	0.999999	5.257	27.63

Table 1: Stokes error radii for various confidence levels.

3.3 Significance indicator

For linear polarization, the primary indicator of statistical significance is $(\Pi/\sigma) = \sqrt{q^2+u^2}/\sigma$. To the extent that errors are normally distributed, the confidence level for detection and the confidence intervals for measurement depend upon the ratio of the polarization parameter to its standard deviation—through the cumulative distribution function of χ^2 on an appropriate number of degrees of freedom.

Interestingly—albeit obviously—χ^2 tests of statistical significance do not depend upon the modulation factor. This also applies to other quantities that are ratios of polarization parameters or functions of such ratios: fractional errors (ϵ_Π/Π) for polarization degree, position angle (ψ), and position-angle error (ϵ_ψ).

3.4 Detection

The criterion for detection is based upon the minimum detectable polarization (MDP), which follows from the probability that a polarization $\Pi > \mathrm{MDP}_C$ is observed under the null hypothesis that the true polarization $\Pi_0 = 0$. This, of course, is the special case of the error contour (§3.2) for $q_0 = u_0 = 0$, such that $\epsilon_C \to \mathrm{MDP}_C = \sqrt{-2\ln(1-C)}\ \sigma$. Inverting this relationship, the confidence level for observing a polarization less than Π if the true polarization $\Pi_0 = 0$ is

$$C(1) = \mathrm{CDF}((\Pi/\sigma)^2, 2) = 1 - \exp[-\frac{1}{2}(\frac{\Pi}{\sigma})^2] \ . \tag{7}$$

If the observation is split into multiple, independent data sets—e.g., binned in energy or time—the statistical significance of an observation in any one bin must be evaluated considering all available bins ("tries"). One way to calculate the confidence level that polarization is detected in any of J independent data sets (or bins), is to test against the null hypothesis that the true polarization $\Pi_{j,0} = 0$ in every bin j. Thus, an appropriate test is χ^2 on $2J$ degrees of freedom:

$$C(J) = \mathrm{CDF}(\Sigma_{j=1}^{J}(\Pi_j/\sigma_j)^2, 2J) \ , \tag{8}$$

where $\Pi_j = \sqrt{q_j^2 + u_j^2}$ and σ_j are the observed polarization and standard deviation in bin j. If the data are treated as a single bin ($J = 1$), the standard calculation for MDP_C is recovered.

Suggested adjectives for describing a detection are "probable" for $C > 99\%$, "highly probable" for $C > 99.9\%$, and "secure" for $C > 99.99\%$. An observation obtaining $C < 99\%$ is *not* a detection, although it might suggest that a detection is plausible with additional data. NB: Use Equation 7 for C only if the observation is treated as a single data set ($J = 1$); use Equation 8 when the data are divided into J bins.

3.5 Measurement

In terms of Stokes parameters (q, u), expressing the measurement of linear polarization is straightforward: $q = q_0 \pm \sigma$ and $u = u_0 \pm \sigma$, where σ is the standard deviation about the mean values (q_0, u_0) of the individually measured Stokes parameters. As 1-D errors in q and u are each normally distributed, 2-D errors (§3.2) are distributed as

$$\chi_2^2 = [(q - q_0)^2 + (u - u_0)^2]/\sigma^2 \ , \tag{9}$$

giving circular error contours (Equation 6) centered on (q_0, u_0).

Using Equations 3 and 4, it is straightforward to substitute into Equation 9 to express χ_2^2 using polarization parameters (Π, ψ):

$$\chi_2^2 = [(\Pi - \Pi_0)^2 + 4\Pi_0\Pi \sin^2(\psi - \psi_0)]/\sigma^2 \ . \tag{10}$$

In terms of polarization parameters (Π, ψ), expressing the measurement is complicated by the coordinate singularity at the origin, which excludes negative values for Π. If the (q, u) error contours at desired confidence levels enclose the origin or even come close, plot the 2-D error contours in (Π, ψ) coordinates—the "protractor" plot—to visualize the behavior of (Π, ψ).

Sufficiently far from the origin (which can be confirmed visually by looking at the 2-D error contour), Π and ψ are each nearly normally distributed, with standard deviations

$$\sigma_\Pi = \sigma \tag{11}$$

$$\sigma_\psi \approx \frac{1}{2}\left(\frac{\sigma}{\sqrt{\Pi_0\Pi}}\right) = \frac{\sigma}{2\Pi_0}\sqrt{\frac{\Pi_0}{\Pi}} \rightarrow \frac{\sigma}{2\Pi_0} \ . \tag{12}$$

As Equation 12 shows, one manifestation of the coordinate singularity is that the uncertainty in ψ diverges as $(\Pi_0\Pi/\sigma^2) \rightarrow 0$, rendering ψ essentially unconstrained near the origin.

For a secure detection, the 1-D measurements and their uncertainties can be presented as $\Pi = \Pi_0 \pm \sigma$ and $\psi = \psi_0 \pm \sigma_\psi \approx \psi_0 \pm \frac{1}{2}(\sigma/\Pi_0)$. Note that, as they average over the second parameter, one-dimensional uncertainties are smaller than extrema of two-dimensional error contours at the same confidence level.

3.5.1 2-D error contours

As discussed earlier (§3.2), a 2-D error contour for the Stokes plot (q, u) is a circle of radius $\epsilon_C = \sqrt{-2\ln(1 - C)}\ \sigma$, depending upon confidence level C (Equation 6 and Table 1). Mapped onto a polarization plot (Π, ψ), the error contour is asymptotically an ellipse for $(\Pi_0/\sigma) \gg 3$ but progressively more diffuse as $(\sqrt{\Pi_0\Pi}/\sigma) \rightarrow 0$. When expressing results in terms of polarization degree and position angle, it is always advisable to include a (Π, ψ) error-contour ("protractor") plot to clarify behavior near the origin, especially of the uncertainty in the position angle ψ. Even when the detection is secure, a contour plot can help identify values favored or disfavored by the data.

Use of (Π, ψ) contour plots facilitates visualization of the dependence of polarization upon energy, time, or position. However, use of (q, u) contour plots may better facilitate analysis—especially when (Π_0/σ) is not large. Suggested confidence levels for 2-D error contours—(q, u) or (Π, ψ)—are 50%, 90%, 99%, and 99.9% (Table 1). Using 68.27% as a confidence level for 2 degrees of freedom can be confusing, as it is the 1-sigma confidence level for 1 degree of freedom.

The suggested position-angle range for the (Π, ψ) plots is $(-90^\circ, +90^\circ)$ East of North. If needed, individual quadrants—NW $(-90^\circ, 0)$ or NE $(0, +90^\circ)$—may be displayed.

3.5.2 1-D error bars

Having extolled the virtues of Stokes parameters (§3), reporting the best value for q and u and the statistical precision for each is straightforward:

$$q = q_0 \pm \sigma_q = q_0 \pm \sigma \tag{13}$$

$$u = u_0 \pm \sigma_u = u_0 \pm \sigma \ , \tag{14}$$

where σ is the Gaussian standard deviation. As the 1-D statistical errors are Gaussian, there is no point in providing a "2-sigma" or "3-sigma" error or some 1-D confidence interval. Just give the 1-D 1-sigma error—i.e., σ. For secure detections, report measurements of Π and ψ as best value and 1-D 1-sigma statistical precision (§3.5, Equations 11 and 12):

$$\Pi = \Pi_0 \pm \sigma_\Pi = \Pi_0 \pm \sigma \tag{15}$$

$$\psi = \psi_0 \pm \sigma_\psi = \psi_0 \pm \frac{\sigma}{2\sqrt{\Pi_0 \Pi}} = \psi_0 \pm \frac{\sigma}{2\Pi_0}\sqrt{\frac{\Pi_0}{\Pi}} \to \psi_0 \pm \frac{\sigma}{2\Pi_0} \ . \tag{16}$$

Based upon inspection of the 2-D error contours, it seems feasible to use the above equations when $\Pi_0 > \mathrm{MDP}_{99} = 3.035 \times \sigma$ ("probable" detection). Conveniently, this ensures measurement of Π_0 at better than 3 sigma.

Conversely, for nondetections (i.e., for $\Pi_0 < \mathrm{MDP}_{99}$), measurement of polarization at better than 3-sigma is not possible. Nonetheless, one may wish to quote a best value and error or upper limit to the polarization degree, while recognizing that the position angle ψ_0 is too poorly constrained to report. As $\Pi \geq 0$, the 1-D error may become asymmetric:

$$\Pi = \Pi_0{}^{+\sigma}_{-\Pi_0} \tag{17}$$

when $\Pi_0 < \sigma$. In such cases, it's perhaps better to quote an upper limit for Π computed from $\Pi < \Pi_0 + 3\sigma$ for a 3-sigma upper limit or perhaps a 99%-confidence (1 degree of freedom) upper limit.

4 Spectropolarimetric fitting with XSPEC

XSPEC provides an extremely useful forward-modeling engine for analyzing x-ray spectropolarimetric data. It convolves a parameterized model spectrum with the instrument response to predict an observed spectrum for each Stokes parameter (I, Q, U), which it compares to the data to obtain the set of best-fit parameters that minimize the fit statistic (typically χ^2 for IXPE observations). After running `fit` to determine the best-fit parameters, use the command `error` to estimate 1-D confidence intervals (Table 2) for specified parameters, based upon the delta statistic ($\Delta\chi^2_1$) for 1 parameter of interest.

$\Delta\chi^2_1$	C
1.0	0.6827 (1σ)
2.706	0.9000
4.000	0.9545 (2σ)
6.635	0.9900
9.000	0.9973 (3σ)

Table 2: $\Delta\chi^2_1$ and corresponding confidence intervals C on a single parameter (i.e., 1 degree of freedom).

To examine confidence regions for m parameters of interest, use the command `steppar` to compute the fit statistic in an m-dimensional grid. Then run `plot contour` to generate error contours for any two parameters (e.g., Π and ψ) in the grid, based upon confidence limits for 2 degrees of freedom.

Internally, XSPEC correctly manipulates polarization data using Stokes parameters (q, u). This is the case even if the polarization fitting parameters and their errors are expressed externally as polarization degree and position angle (Π, ψ). Uncertainties obtained using the XSPEC `err` command may be reported as a 1-D error or as a specified confidence interval, depending upon the input arguments. As an example, the 1-sigma (68.3%) confidence interval for parameter 5 (Π) in a model fit is computed using the command

```
err 1.0 5 .
```

Assuming that the best-fit value was previously determined to be $\Pi = 0.045$ (i.e., 4.5%), the output of this command will look like

```
5 0.030120 0.0601224 (-0.0151152,0.0150072) ,
```

where the output numbers are parameter ID, 68.3% lower limit, 68.3% upper limit, and (in parentheses) lower and upper limits minus the best-fit value. The result and 1-sigma confidence interval is therefore reported as $\Pi = 4.5 \pm 1.5\%$.

To calculate a 99% confidence interval for an example where there is no detection and the best-fit value of $\Pi = 1.3\%$, the input command changes only slightly to account for the different confidence interval:

```
err 6.635 5 .
```

The output looks like

```
5 0 0.0826826 (-0.0131152,0.0695674) ,
```

where now the values correspond to the 99% confidence interval ($\chi^2_1 = 6.635$). For this case, the data constitute a non-detection, and the 99% confidence upper-limit of the polarization degree is $\Pi < 8.2\%$. Note that there are no differences in presentation between these two cases other than the fact that the lower confidence limit is equal to the physical limit of $\Pi = 0$ in the latter example. As noted previously, these 1-D errors are smaller than extrema of the corresponding 2-D confidence contours and independent of errors in the other parameter.