A Fourier-Domain Analysis of BHEX for Photon-Ring Inference

Jonathan R. Landers

Independent researcher.

This project is inspired by published BHEX work and is not affiliated with the BHEX collaboration or mission team.

In recent work proposing the **Black Hole Explorer (BHEX)** mission, the BHEX authors argue that sufficiently long-baseline interferometric measurements can reveal, and in favorable regimes help isolate, the universal photon-ring signature of a black hole from the more complicated emission of the surrounding plasma [1, 2]. In practice, the measured visibility data contain both contributions, and the observational challenge is to determine whether a structured photon-ring signal can be reliably extracted from the mixture using long-baseline amplitude signatures or more explicit structured source-separation methods. This note develops one possible mathematical extension of that problem framing, inspired by the BHEX visibility program.

The natural story is two-stage. First, one can ask how far a simple long-baseline ring-oriented heuristic can be pushed before structured plasma leakage creates bias. Second, one can ask whether that same regime is still recoverable once the inverse problem is posed more explicitly. The first question leads to a mismatch bound between a direct amplitude-readout heuristic and a plasma-aware estimator. The second leads to a recoverability statement: the direct amplitude readout can lose clarity before ring non-recoverability is reached.

1 A first comparison: template mismatch versus plasma-aware projection

Lemma (Mismatch between heuristic ring detection and structured plasma-aware separation). Let $y \in \mathcal{H}$ be observed visibility data in a complex Hilbert space, decomposed as

$$y = \alpha g + p + \varepsilon,$$

where $g \in \mathcal{H}$ is a photon-ring template, $\alpha \in \mathbb{C}$ is the ring amplitude, $p \in \mathcal{P} \subset \mathcal{H}$ is a plasma component lying in a plasma subspace $\mathcal{P}$, and $\varepsilon \in \mathcal{H}$ is noise. Define the heuristic estimator

$$\hat{\alpha}_{\text{heur}} = \frac{\langle y, g\rangle}{\|g\|^2},$$

and the plasma-aware estimator

$$\hat{\alpha}_{\text{model}} = \frac{\langle (I-\Pi_{\mathcal{P}})y,\ (I-\Pi_{\mathcal{P}})g\rangle}{\|(I-\Pi_{\mathcal{P}})g\|^2},$$

where $\Pi_{\mathcal{P}}$ denotes the orthogonal projector onto $\mathcal{P}$. Then

$$\left|\hat{\alpha}_{\text{heur}} - \hat{\alpha}_{\text{model}}\right| \le \mu(g,\mathcal{P})\frac{\|p\|}{\|g\|} + \frac{\|\varepsilon\|}{\|g\|} + \frac{\|(I-\Pi_{\mathcal{P}})\varepsilon\|}{\|(I-\Pi_{\mathcal{P}})g\|},$$

where

$$\mu(g,\mathcal{P}) = \frac{\|\Pi_{\mathcal{P}} g\|}{\|g\|}$$

is the ring-plasma coherence.

Proof. First, by direct substitution,

$$\hat{\alpha}_{\text{heur}} = \frac{\langle \alpha g + p + \varepsilon,\ g\rangle}{\|g\|^2} = \alpha + \frac{\langle p, g\rangle}{\|g\|^2} + \frac{\langle \varepsilon, g\rangle}{\|g\|^2}.$$

Hence

$$\hat{\alpha}_{\text{heur}} - \alpha = \frac{\langle p, g\rangle}{\|g\|^2} + \frac{\langle \varepsilon, g\rangle}{\|g\|^2}.$$

Now define

$$g_\perp = (I-\Pi_{\mathcal{P}})g.$$

Since $p \in \mathcal{P}$, we have

$$(I-\Pi_{\mathcal{P}})p = 0.$$

Therefore

$$(I-\Pi_{\mathcal{P}})y = \alpha(I-\Pi_{\mathcal{P}})g + (I-\Pi_{\mathcal{P}})\varepsilon = \alpha g_\perp + (I-\Pi_{\mathcal{P}})\varepsilon.$$

Substituting into the plasma-aware estimator gives

$$\hat{\alpha}_{\text{model}} = \frac{\langle \alpha g_\perp + (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle}{\|g_\perp\|^2} = \alpha + \frac{\langle (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle}{\|g_\perp\|^2}.$$

So

$$\hat{\alpha}_{\text{model}} - \alpha = \frac{\langle (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle}{\|g_\perp\|^2}.$$

Subtracting the two estimators yields

$$\hat{\alpha}_{\text{heur}} - \hat{\alpha}_{\text{model}} = \frac{\langle p, g\rangle}{\|g\|^2} + \frac{\langle \varepsilon, g\rangle}{\|g\|^2} - \frac{\langle (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle}{\|g_\perp\|^2}.$$

Applying the triangle inequality,

$$\left|\hat{\alpha}_{\text{heur}} - \hat{\alpha}_{\text{model}}\right| \le \frac{|\langle p, g\rangle|}{\|g\|^2} + \frac{|\langle \varepsilon, g\rangle|}{\|g\|^2} + \frac{|\langle (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle|}{\|g_\perp\|^2}.$$

Using Cauchy–Schwarz on the last two terms,

$$\frac{|\langle \varepsilon, g\rangle|}{\|g\|^2} \le \frac{\|\varepsilon\|\,\|g\|}{\|g\|^2} = \frac{\|\varepsilon\|}{\|g\|},$$

and

$$\frac{|\langle (I-\Pi_{\mathcal{P}})\varepsilon,\ g_\perp\rangle|}{\|g_\perp\|^2} \le \frac{\|(I-\Pi_{\mathcal{P}})\varepsilon\|\,\|g_\perp\|}{\|g_\perp\|^2} = \frac{\|(I-\Pi_{\mathcal{P}})\varepsilon\|}{\|g_\perp\|}.$$

It remains to bound the plasma term. Since $p \in \mathcal{P}$,

$$\langle p, g\rangle = \langle p, \Pi_{\mathcal{P}} g\rangle,$$

so again by Cauchy–Schwarz,

$$|\langle p, g\rangle| = |\langle p, \Pi_{\mathcal{P}} g\rangle| \leq \|p\| \, \|\Pi_{\mathcal{P}} g\|.$$

Thus

$$\frac{|\langle p, g\rangle|}{\|g\|^2} \leq \frac{\|p\| \, \|\Pi_{\mathcal{P}} g\|}{\|g\|^2} = \left(\frac{\|\Pi_{\mathcal{P}} g\|}{\|g\|}\right) \frac{\|p\|}{\|g\|} = \mu(g, \mathcal{P}) \frac{\|p\|}{\|g\|}.$$

Combining the three bounds gives

$$\left|\hat{\alpha}_{\text{heur}} - \hat{\alpha}_{\text{model}}\right| \leq \mu(g, \mathcal{P}) \frac{\|p\|}{\|g\|} + \frac{\|\varepsilon\|}{\|g\|} + \frac{\|(I - \Pi_{\mathcal{P}})\varepsilon\|}{\|(I - \Pi_{\mathcal{P}})g\|}.$$

This proves the claim. □

Interpretation. This lemma makes precise the gap between a template-based long-baseline heuristic and a plasma-aware structured estimator in the same general spirit as the BHEX visibility-analysis program [2]. The mismatch is controlled by three interpretable quantities: ring-plasma coherence, plasma magnitude, and measurement noise. The more the photon-ring template overlaps realistic plasma structure, the more the heuristic estimator can be biased relative to a structured projection.

2 From mismatch to recoverability

The previous lemma quantifies the discrepancy between two estimators, but it does not yet answer the more interesting question: if the direct amplitude readout becomes hard to interpret, does recoverability necessarily disappear with it? The next proposition formalizes the stricter and more useful claim that the answer can be no. In other words, the heuristic readability limit can arrive before the source-separation limit.

Proposition (direct-readout limit before separation limit). Let the observed visibility be

$$y = \alpha g_\theta + q + \varepsilon,$$

in a Hilbert space $\mathcal{H}$, where:

- $g_\theta \in \mathcal{G}$ is the photon-ring visibility family, parameterized by ring shape θ,
- $\alpha \in \mathbb{C}$ is the ring amplitude,
- $q \in \mathcal{Q}$ is a structured nuisance term representing plasma, finite-width $n = 1$ distortion, scattering, or averaging artifacts,
- ε is noise.

Assume, as a simplified abstraction motivated by the BHEX visibility program, that a long-baseline amplitude heuristic estimates θ from the oscillatory structure of the **visibility amplitude** $|y|$ [2]. Assume also that a structured estimator is given by

$$(\hat{\alpha}, \hat{\theta}, \hat{q}) \in \arg \min_{\alpha, \theta, q \in \mathcal{Q}} \|y - \alpha g_\theta - q\|_{\mathcal{H}}^2 + \lambda R(q),$$

where R encodes nuisance structure. Finally, assume a local identifiability condition: for some $c > 0$,

$$\|\alpha g_\theta + q - \alpha' g_{\theta'} - q'\|_{\mathcal{H}} \geq c\Big(|\alpha - \alpha'| + d(\theta, \theta') + \|q - q'\|_{\mathcal{H}}\Big)$$

for all nearby triples (α, θ, q) and (α', θ', q').

Then there exists a nonempty regime in which the amplitude-based heuristic becomes biased or visually unstable, but the structured estimator still recovers θ stably.

Proof. Write

$$y = \alpha g_\theta + r, \qquad r := q + \varepsilon.$$

The heuristic extracts θ from extrema or oscillation spacing of $|y|$ on long baselines. It is therefore accurate only when

$$|y| = |\alpha g_\theta + r|$$

has nearly the same oscillatory pattern as $|\alpha g_\theta|$. But modulus is nonlinear, so even moderate perturbations can shift extrema and spacing. Pointwise,

$$\big||\alpha g_\theta + r| - |\alpha g_\theta|\big| \leq |r|,$$

yet small amplitude perturbation does **not** imply small perturbation of extrema locations. Near a shallow oscillation peak, an $O(\delta)$ perturbation can induce an $O(\delta/\kappa)$ peak-location error, where κ denotes local curvature. Thus the heuristic can lose direct-readout accuracy once structured contamination distorts the oscillatory pattern, even if the ring signal is still present.

By contrast, the structured estimator fits the full model $y = \alpha g_\theta + q + \varepsilon$. Under the identifiability inequality, standard stability of constrained least squares implies that if $\|\varepsilon\|_{\mathcal{H}} \leq \eta$, then

$$|\hat{\alpha} - \alpha| + d(\hat{\theta}, \theta) + \|\hat{q} - q\|_{\mathcal{H}} \leq \frac{C}{c} \eta$$

for some constant C, since any competing triple separated from the truth by more than $O(\eta)$ would violate the lower bound and increase the residual.

Therefore there is a regime where:

- the nuisance q is large enough to distort the amplitude oscillation and make direct peak-reading unreliable,
- but still small enough, and sufficiently structured, that the full inverse problem remains identifiable and stable.

Hence

$$\{\text{cleanly amplitude-readable}\} \subsetneq \{\text{structured-separation recoverable}\}$$

is possible. □

One-line interpretation. The BHEX heuristic is especially powerful when the ring is **cleanly readable** in amplitude space [2]; the structured estimator succeeds under the weaker condition that the ring remains **identifiable** after joint modeling of structured contamination.

3 Physical picture

At the image level, the distinction is intuitive. The observed black-hole image is typically dominated by a **broader plasma emission ring or crescent**, produced by synchrotron emission from hot plasma, while the **photon ring** is a much thinner structure at nearly the same diameter, set more directly by strong lensing geometry [2, 3]. This means the source naturally presents a broad, bright, plasma-dominated feature, while the measurement goal is to isolate a much thinner, more universal ring hidden within or beneath it. In Fourier space, that geometric difference becomes a problem of structured overlap: broad plasma emission and thin photon-ring oscillations can coexist on the same visibility domain, but need not be equally identifiable.

This is exactly why the source-separation perspective is useful as an extension alongside the BHEX program. The heuristic says: go where the photon ring should show itself most directly. The structured method says: even if that regime is not perfectly clean, the ring may still be recoverable provided its structure remains distinguishable from the nuisance family.

4 Connection to noise

The first lemma separates the mismatch into a **plasma leakage term** and two **noise terms**. The first noise term, $\|\varepsilon\|/\|g\|$, is the raw sensitivity of the heuristic estimator to measurement noise, while the second,

$$\frac{\|(I-\Pi_{\mathcal{P}})\varepsilon\|}{\|(I-\Pi_{\mathcal{P}})g\|},$$

shows that the structured method is governed by the component of the noise that survives plasma projection, normalized by the strength of the projected ring template. Thus the plasma-aware method reduces bias from source confusion, but its advantage is greatest when the projected ring signal remains sufficiently strong relative to the projected noise.

The proposition sharpens this picture. Noise and nuisance structure do not merely degrade the signal-to-noise ratio; they can deform the oscillatory pattern used by the heuristic before they destroy identifiability altogether. This is the conceptual gap between a direct long-baseline readout strategy and a more robust inverse-problem formulation.

5 Astrophysical relevance

The mathematical content becomes most compelling when the abstract nuisance class is tied to realistic plasma and propagation effects. A natural next step is therefore to estimate either the ring-plasma coherence $\mu(g, \mathcal{P})$ or a more general identifiability margin for nuisance classes informed by GRMHD simulations, radiative transfer calculations, scattering models, or empirical variability constraints [2, 3]. Such a result would connect the abstract bounds directly to the astrophysical conditions under which long-baseline photon-ring detection is expected to succeed, become biased, or remain recoverable only through structured separation.

In that sense, the BHEX heuristic can be viewed as an elegant front-end to the inference problem: it targets the regime where the photon ring more or less announces itself [2]. The source-separation formalism is a complementary robustness layer: it targets the broader regime in which the data are less cleanly readable, but the ring still leaves enough structured trace to be recovered.

6 Key interesting points

- The writeup develops a **structured source-separation abstraction in Fourier space** inspired by BHEX photon-ring inference [1, 2].
- The first lemma gives a **clean mismatch bound** between a direct amplitude-readout heuristic and a plasma-aware estimator, controlled by ring-plasma coherence, plasma magnitude, and noise.
- The new proposition makes the stronger point that the **heuristic readability limit can occur before the recoverability limit** in this abstraction, provided the ring remains identifiable under joint modeling.
- The physical picture is that the image naturally shows a **broad plasma ring or crescent**, while the scientific target is a **much thinner photon ring** at nearly the same diameter [2, 3].
- The most promising bridge to astrophysics is to estimate coherence or identifiability margins for **GRMHD-informed nuisance classes**, linking this abstract theory to real mission regimes anticipated by BHEX forecast work [2, 3].

References

[1] Black Hole Explorer Collaboration, *Black Hole Explorer: Motivation and Vision*, arXiv:2406.12917. https://arxiv.org/abs/2406.12917

[2] Black Hole Explorer Collaboration, *The Black Hole Explorer: Photon Ring Science, Detection and Shape Measurement*, arXiv:2406.09498. https://arxiv.org/abs/2406.09498

[3] Black Hole Explorer Collaboration, *Interferometric Inference of Black Hole Spin from Photon Ring Size and Brightness*, arXiv:2509.23628. https://arxiv.org/abs/2509.23628