

GPU-accelerated X-ray pulse profile modeling

Tianzhe Zhou^{1,*} and Chun Huang²

¹ Department of Physics, Tsinghua University, Beijing 100084, China

² Physics Department and McDonnell Center for the Space Sciences, Washington University in St. Louis, St. Louis, MO 63130, USA

Received 3 February 2026 / Accepted 11 March 2026

ABSTRACT

Aims. The aim of this work is to remove the accuracy-speed bottleneck in Bayesian pulse profile modeling (PPM) of thermal X-ray emission from rotation-powered millisecond pulsars. This would enable a reliable inference of the stellar mass, M , radius, R , and the equation of state of cold, dense matter for extreme hotspot geometries and higher fidelity forward models.

Methods. We developed, validated, and publicly released a GPU-accelerated X-ray PPM framework. We benchmarked the GPU implementation against established reference calculations across standard and extreme geometries. We quantified the performance on an RTX 4080. We also diagnosed numerical systematics associated with atmosphere-table interpolation near lookup boundaries using two targeted tests and mitigated them with a mixed-order interpolation scheme.

Results. Our framework reproduces established benchmarks to a $\sim 10^{-3}$ relative accuracy, including the case of extreme hotspot configurations that are difficult to resolve at production settings. At high fidelity, we reduced per-evaluation runtimes from minutes to 2–5 ms on an RTX 4080, corresponding to 10^3 – $10^4\times$ speedups, thereby making posterior exploration feasible at resolutions and model complexities that were previously impractical. We identified a systematic error near atmosphere-table interpolation boundaries and show that the proposed mixed-order interpolator substantially reduces this bias in diagnostic tests.

Conclusions. By coupling benchmark-level accuracy with millisecond-scale evaluations, the framework expands the accessible hotspot model space and mitigates key numerical systematics in X-ray PPM, strengthening mass-radius inference for current and future X-ray missions.

Key words. methods: numerical – stars: neutron – pulsars: general – X-rays: stars

1. Introduction

Rotation-induced modulations of emission from neutron-star surfaces known as X-ray pulse profiles encode the surface temperature distribution and the global parameters: mass, M , and radius, R . The emission is often dominated by one or more localized, hot regions (i.e., hotspots). Photons from these regions experience gravitational redshift and strong-field light bending, as well as special-relativistic Doppler boosting and time delays. These effects leave signatures in the energy-resolved waveform, encoding the star's compactness, our primary interest, as well as other global and geometric parameters (Pechenick et al. 1983; Beloborodov 2002; Poutanen & Beloborodov 2006; Morsink et al. 2007; Cadeau et al. 2007; Psaltis & Özel 2014; Watts et al. 2016; Nättilä & Pihajoki 2018; Poutanen 2020). Although this picture is conceptually clean, to model observed pulse profiles accurately, and to extract the hotspot configuration and mass-radius (M - R) relation, we must incorporate the realistic atmospheric beaming pattern, including composition and magnetic-field effects (Zavlin et al. 1996; Ho & Lai 2001, 2003; Suleimanov et al. 2012; Salmi et al. 2020, 2023), rotation-induced oblateness (Miller & Lamb 1998; Poutanen & Gierliński 2003; Morsink et al. 2007; AlGendy & Morsink 2014; Nättilä & Pihajoki 2018; Jakab & Morsink 2025), hotspot geometry (Poutanen & Beloborodov 2006; Viironen & Poutanen 2004; Bogdanov et al. 2008), interstellar absorption (Verner et al. 1996; Wilms et al. 2000), and the detailed instrument response

(Gendreau et al. 2016; Prigozhin et al. 2016; Remillard et al. 2022; Markwardt et al. 2024). With these ingredients and practical approximations established, reliable forward modeling procedures can be carried out across the full parameter space, enabling inference of M and R using the X-ray pulse profile modeling (PPM) technique (Bogdanov et al. 2019b, 2021; Miller et al. 2019; Riley et al. 2019; Miller et al. 2021; Riley et al. 2021; Salmi et al. 2024a; Choudhury et al. 2024a; Vinciguerra et al. 2024; Dittmann et al. 2024).

Thermal emission from neutron-star surfaces peaks in the soft X-ray band (~ 0.1 – 2 keV) (Zavlin et al. 1996; Potekhin et al. 2015) and, thus, obtaining precise pulse profile measurements requires high throughput at these energies, sub-microsecond time tagging, and good energy resolution. NASA's Neutron Star Interior Composition Explorer (NICER) provides this capability: a 0.2–12 keV bandpass, large effective area at soft X-rays, ~ 100 ns absolute timing, and $\sim 10^2$ eV energy resolution near 1 keV (Gendreau et al. 2016; Prigozhin et al. 2016). NICER's primary pulse profile targets are rotation-powered millisecond pulsars (MSPs), whose highly coherent spin permits phase-folding of long observations to build energy-resolved waveforms with high signal-to-noise ratio. Statistical uncertainties then decrease roughly as the square root of the total photon counts (or exposure time). Accordingly, NICER has concentrated on a small set of bright, well-characterized MSPs that would be suitable for running the PPM (Bogdanov et al. 2019a; Guillot et al. 2019). Looking ahead, the upcoming enhanced X-ray Timing and Polarimetry mission (eXTP), currently planned for launch in the early 2030s, is designed to combine large-area timing,

* Corresponding author: zhoutz22@mails.tsinghua.edu.cn

spectroscopy, and polarimetry, and is expected to make PPM an even more powerful probe of neutron-star dense matter (Zhang et al. 2025; Li et al. 2025; Watts et al. 2019).

The first landmark NICER PPM result came from PSR J0030+0451, an isolated millisecond pulsar with no prior radio mass measurement (Manchester et al. 2005; Lommen et al. 2006). Independent analyses of the same NICER dataset produced joint posteriors for the stellar mass and radius and favored nonantipodal emitting-region geometries, establishing the first NICER mass-radius constraints via PPM. Reanalyses have refined these constraints (Miller et al. 2019; Riley et al. 2019; Vinciguerra et al. 2024). NICER’s subsequent measurement for PSR J0740+6620, whose mass is tightly constrained by radio timing at $\approx 2 M_{\odot}$ (Cromartie et al. 2020; Fonseca et al. 2021), yielded radius measurements at the high-mass end from independent analyses (Miller et al. 2021; Riley et al. 2021), with subsequent updates using additional NICER data (Salmi et al. 2024a; Dittmann et al. 2024) placing strong constraints on the dense-matter equation of state and disfavoring very soft models that predict markedly smaller radii (e.g., Raaijmakers et al. 2021b; Zhang & Li 2021). In both sources, the inferred hotspot locations are inconsistent with a simple centered dipole, suggesting more complex (e.g., offset or multipolar) surface-field configurations (Bilous et al. 2019; Pétri et al. 2025; Kalapotharakos et al. 2021; Chen et al. 2020; Huang & Chen 2025). Recent and ongoing NICER campaigns have targeted bright MSPs such as PSR J0437–4715, whose radio timing yields a precise mass of $M = 1.418 \pm 0.044 M_{\odot}$, and PSR J0614–3329, thereby extending radius measurements toward the canonical $\sim 1.4 M_{\odot}$ mass scale (Reardon et al. 2016; Reardon et al. 2024; Choudhury et al. 2024a; Mauviard et al. 2025).

Together with tidal-deformability constraints from gravitational-wave events (Abbott et al. 2017; Abbott et al. 2018; Abbott et al. 2019), these results have energized the nuclear-astronomy community and motivated increasingly parametric and nonparametric equation-of-state inference frameworks that confront neutron-star structure with multimessenger data (e.g., De et al. 2018; Landry & Essick 2019; Raaijmakers et al. 2019, 2020; Raaijmakers et al. 2021a; Essick et al. 2020; Greif et al. 2019; Legred et al. 2021; Huth et al. 2022; Annala et al. 2022; Huang et al. 2024a,c,b; Rutherford et al. 2024; Zhou & Huang 2025; Cartaxo et al. 2026; Huang & Sourav 2025; Huang 2025a,b). Since NICER’s initial PPM results, the applications and methodological advancements of PPM have expanded substantially. Open-source frameworks, such as X-PSI, have made fully Bayesian PPM accessible and reproducible, providing flexible hotspot parameterizations (Riley et al. 2023a,b).

Reanalyses of PSR J0030+0451 have shown that the choice of hotspot maps and the treatment of instrumental and background systematics can shift the inferred M – R posteriors, especially when joint NICER+XMM constraints are imposed (Vinciguerra et al. 2024). This approach has since been extended to additional NICER targets, notably PSR J0437–4715 (Choudhury et al. 2024a), and recent studies of PSR J1231–1411 (Salmi et al. 2024b). In parallel, key physical ingredients continue to be refined and several formalisms now propagate polarization predictions suitable for cross-mission tests (Heyl & Shaviv 2000; Viironen & Poutanen 2004; Loktev et al. 2020; Salmi et al. 2025). Methodological progress also includes faster and more accurate relativistic ray-tracing schemes (Nättilä & Pihajoki 2018; Poutanen 2020). Overall, PPM remains a rapidly advancing field.

Despite major advances, several bottlenecks still limit current PPM analyses. First, degeneracies among hotspot geometry and stellar compactness can broaden posterior inferences (Poutanen & Beloborodov 2006; Lo et al. 2013; Psaltis & Özel 2014; Zhao et al. 2024; Güver et al. 2026; Bootsma et al. 2025). One response is to adopt physics-informed hotspot parameterizations that reduce ad hoc geometric degrees of freedom; such approaches have recently been applied to PSR J0030+0451 and PSR J0437–4715 (Huang & Chen 2025). It has also been shown by Lamb et al. (2009a,b); Miller & Lamb (2015) that for NICER-like data, several plausible sources of systematic error do not yield apparently acceptable fits with radius biases exceeding $1 - \sigma$.

Second, PPM is computationally intensive: accurate forward models must couple relativistic ray tracing, realistic atmospheric beaming, rotation-induced oblateness, and instrumental background and calibration systematics. Rigorous Bayesian inference (e.g., nested sampling or ensemble Markov chain Monte Carlo, i.e., MCMC) typically requires $\sim 10^6$ – 10^8 likelihood evaluations in a high-dimensional parameter space (Speagle 2020; Buchner 2021) and incorporating radio constraints or joint NICER+XMM data further increases the cost (Miller et al. 2021; Salmi et al. 2024a; Choudhury et al. 2024a). Third, there is an accuracy-speed trade-off: coarse angular, phase, and energy grids can bias the likelihood and, thus, the posteriors for certain hotspot configurations; whereas the near-converged resolutions needed for robust inference (i.e., to render numerical errors subdominant to Poisson counting fluctuation) make forward modeling computationally demanding and can limit the scale of large-scale Bayesian analyses (Choudhury et al. 2024b; Bogdanov et al. 2019b, 2021; Salmi et al. 2024b; Vinciguerra et al. 2024). Recent cross-code tests and case studies identify extreme geometries where insufficient resolution measurably alters the waveform and likelihood computation, motivating higher-resolution validation runs (Choudhury et al. 2024b). At present, the relatively large observational uncertainties from NICER mean that resolution choices do not yet have a visible effect on shifting the inferred mass-radius. Nonetheless, these differences are expected to become important for future missions such as eXTP, for which PPM is a key dense-matter science driver (Watts et al. 2019; Li et al. 2025). In practice, a single high-fidelity likelihood evaluation can take ~ 10 , s or more (depending on resolution and model complexity), so many analyses adopt simplified hotspot parameterizations and/or multistage workflows (coarse-resolution exploration followed by high-resolution verification) to keep compute costs tractable. The principal bottleneck remains computational speed.

In this work, we address these limitations by introducing a GPU-accelerated PPM framework written in C++/CUDA. The framework supports energy- and phase-resolved data. NICER-specific instrument response and background models are implemented as modular components that can be swapped for other observatories. From the outset, the forward model accommodates complex, physics-motivated surface-temperature distributions beyond circular, isothermal caps, while preserving high-fidelity physics. Conceptually, our forward model adopts the standard physical ingredients established in prior CPU frameworks (e.g., Riley et al. 2023b). Our contributions are: (i) a reengineered, GPU-optimized implementation of PPM for contemporary hardware; (ii) end-to-end reproducibility checks against existing benchmarks and codebases and a demonstration that the accuracy-speed trade-off can be broken by performing high-fidelity computation at unprecedented speed; and (iii)

identification of a source of systematic error introduced by interpolating atmosphere lookup tables near grid boundaries.

Both the CPU and GPU implementations are released as open source to provide a reproducible foundation for future studies. In validation tests, the framework reproduces published waveform benchmarks to within $\lesssim 10^{-3}$ relative error while delivering 10^3 – $10^4\times$ speedups over representative CPU baselines. This reduces a single likelihood evaluation from seconds-minutes to milliseconds on a single GPU, making full Bayesian inference, which was previously practical only on large clusters, feasible on personal hardware.

The increased computation speed enables high angular, phase, and energy resolution without a prohibitive cost, mitigating the resolution-induced likelihood biases noted in prior work (Choudhury et al. 2024b). Consequently, a much broader family of fine-detail hotspot maps becomes computationally tractable, including cases previously impractical due to compute budgets. We anticipate that this framework will facilitate more realistic geometry priors, wider parameter exploration, and multi-instrument or multiwavelength joint analyses in forthcoming PPM studies built on this codebase.

The paper is organized as follows. Section 2 summarizes the theoretical framework, highlights subtleties that affect numerical accuracy, presents a transparent computational recipe to facilitate full reproducibility, and details our GPU implementation. Section 3 provides the first set of validation tests, demonstrating that our CPU and GPU implementations reproduce theoretical benchmarks under standard OS assumptions (see Bogdanov et al. 2019b). Section 4 investigates interpolation within atmosphere tables, detailing our interpolation strategy and introducing two new test cases that underscore the need for careful treatment. Section 5 reports production-grade computations that incorporate instrument response, interstellar absorption, and realistic atmospheric beaming (i.e., all the essential ingredients of PPM forward computation) and shows that the framework remains fast and robust even for hotspot configurations previously regarded as impractical, thereby overcoming the long-standing accuracy-speed trade-off. Section 6 concludes with a brief summary and discussion. Appendices A–C supplement the main text with rederivations of some core formulas useful for PPM forward computation.

2. Model framework and computational recipe

In this section, we present the theoretical framework for X-ray PPM forward computation and provide a transparent, step-by-step recipe that covers the theoretical framework and implementation details to enable end-to-end reproducibility. The objective of PPM inference is to infer fundamental neutron-star properties from the energy- and phase-resolved X-ray waveforms. Given the statistical precision of modern data, we aim to limit forward-model inaccuracies to $\lesssim 1\%$ in phase-resolved flux where feasible.

The relevant physics combines general relativity (GR) and special relativity (SR). In GR, strong gravity near the neutron star (NS) alters photon trajectories and energies: gravitational light bending increases the visible surface fraction, revealing regions beyond the local limb; gravitational redshift reduces photon energies; and time-of-flight delays in curved spacetime introduce phase shifts. Rotation adds SR effects: Doppler boosting modulates the observed flux as a hotspot approaches or recedes from the observer, while rapid spin induces stellar oblateness. These effects are set by the global parameters

(M, R, ν_*) and by the temperature map on the neutron-star surface. Foundational works include Pechenick et al. (1983); Beloborodov (2002); Poutanen & Gierliński (2003); Miller & Lamb (1998); Poutanen & Beloborodov (2006); Morsink et al. (2007); Cadeau et al. (2007); Psaltis & Özel (2014); Watts et al. (2016); Nättilä & Pihajoki (2018); Poutanen (2020); Zavlin et al. (1996); Ho & Lai (2001, 2003); Suleimanov et al. (2012); Salmi et al. (2020, 2023); AlGendy & Morsink (2014); Viironen & Poutanen (2004); Bogdanov et al. (2008).

The theoretical waveform is then propagated through the astrophysical and instrumental layers that connect the source to the data, namely, interstellar absorption and the instrument response. We address each component explicitly in dedicated subsections.

Our approach builds on modeling practices summarized by Bogdanov et al. (2019b, 2021); Salmi et al. (2023) and the open-source X-PSI package (Riley et al. 2023b), which we use as standard benchmarks. We extend these frameworks in a self-contained manner by providing fully relativistic derivations and implementations, applying controlled approximations only when they improve efficiency without biasing inference.

For orientation, Section 2 is organized as follows. Section 2.1 defines the physical setup and notation (parameters and geometry). Section 2.2 specifies the spacetime metric and the oblate-surface model. Section 2.3 details photon propagation (light bending and time-of-flight delays). Section 2.4 incorporates the rotational SR Doppler effect. Section 2.5 models the surface emission and atmospheric treatment, integrating to obtain the observed flux. Section 2.6 discusses the default surface temperature map deployed in this work and the surface discretization strategy. Section 2.7 discusses interstellar absorption and instrument response. Section 2.8 provides the detailed road map of our computation recipe, and the GPU implementation is discussed in Section 2.9. The objective is a production-grade, ready-to-use codebase for future PPM studies.

2.1. Geometrical setup

We modeled a rotating neutron star of mass M , equatorial radius R_e , and spin frequency ν_* (angular frequency of $\Omega \equiv 2\pi\nu_*$). At the equator of a typical NS, the dimensionless surface speed is

$$\beta_{\text{eq}} \equiv \frac{v_{\text{eq}}}{c} \approx 0.1\text{--}0.3, \quad (1)$$

so both GR and SR effects are necessary for sub-percent-level waveform accuracy.

Photons emitted from the oblate stellar surface experience light bending, gravitational redshift, time-of-flight delays, and Doppler boosting. To place these ingredients in a coherent framework, we adopted a star-centered coordinate system with the observer at infinity along a fixed line of sight. The stellar surface was discretized into small elements (i.e., surface patches) and we traced photons from each patch to the observer, integrating the contributions over the surface. Rather than restricting emission to one or a few circular, isothermal spots, our algorithm is designed from the outset to accept full-surface temperature maps, making it suitable for nontrivial configurations (e.g., accretion-heated distributions or physics-motivated hotspot maps). This relaxes symmetry assumptions that are typical for spot-based models; consequently, we emphasize computational optimizations to ensure the tractability of full-surface maps.

Our geometric notation follows Loktev et al. (2020) for consistency with prior literature. A schematic of the setup is shown

Table 1. Unit vectors in the geometric setup.

Symbol	Description
\hat{k}	Direction from the star center to the distant observer.
\hat{r}	Direction from the star center to the emitting surface patch.
\hat{n}	Local surface normal at the emitting patch (coincides with \hat{r} for spherical stars).
\hat{k}_0	Initial photon direction at emission (before propagation).
\hat{z}	Stellar rotation axis.
$\hat{\beta}$	Local velocity direction of the emitting patch (tangent to the azimuthal circle).

Table 2. Angles in the geometric setup.

Symbol	Description
i	Inclination: angle between \hat{z} and \hat{k} .
θ	Colatitude: angle between \hat{z} and \hat{r} .
ϕ	Azimuth of \hat{r} in the star frame.
ψ	Angle between \hat{r} and \hat{k} , with $\cos \psi = \hat{r} \cdot \hat{k} = \sin i \sin \theta \cos \phi + \cos i \cos \theta$.
η	Angle between \hat{r} and \hat{n} (zero for spherical stars; nonzero for oblate stars).
α	Angle between \hat{r} and \hat{k}_0 (emission angle relative to the radial direction).
σ	Angle between \hat{n} and \hat{k}_0 (emission angle relative to the local normal).
ξ	Angle between $\hat{\beta}$ and \hat{k}_0 ; $\cos \xi \equiv \hat{\beta} \cdot \hat{k}_0 = -\sin \alpha \sin i \sin \phi / \sin \psi$.




Fig. 1. Demonstration of the geometric setup used in our X-ray PPM forward computation, shown in the lab-frame coordinate system $\hat{x} - \hat{y} - \hat{z}$. A surface patch (orange square) on the stellar surface at colatitude θ and rotational phase ϕ , viewed along the observer’s line-of-sight unit vector, \hat{k} . Definitions of all unit vectors are given in Table 1 and angle conventions are summarized in Table 2.

in Figure 1. For clarity, the key unit vectors and angles used in this work are summarized in Tables 1 and 2.

Physically, i denotes the observer’s inclination; (θ, ϕ) specify the location of an emitting surface element; α is the emission angle measured from the local radial direction in the near-star static frame; and ψ is the angle between the surface radial vector and the observer’s line of sight. In GR, regions with $\psi > \pi/2$ remain visible due to gravitational light bending; the maximum visible ψ is determined by the compactness (defined below).

For a spherical star the surface normal coincides with the radial unit vector, $\hat{n} = \hat{r}$; hence, $\eta = 0$ and $\sigma = \alpha$. For an oblate star, rotational flattening tilts the surface normal away from \hat{r} ; hence, $\eta \neq 0$. We quantify this with the dimensionless surface-slope parameter,

$$f \equiv \frac{1}{R} \frac{dR}{\sqrt{1-u}} \frac{dR}{d\theta}, \quad u \equiv \frac{R_S}{R}, \quad (2)$$

with $R_S \equiv 2GM/c^2$ the Schwarzschild radius. We have

$$\cos \eta = \hat{r} \cdot \hat{n} = \frac{1}{\sqrt{1+f^2}}, \quad \sin \eta = \frac{f}{\sqrt{1+f^2}}. \quad (3)$$

The emission angle with respect to the local normal is

$$\cos \sigma = \cos \alpha \cos \eta + \frac{\sin \alpha \sin \eta}{\sin \psi} \times (\cos i \sin \theta - \sin i \cos \theta \cos \phi). \quad (4)$$

These definitions lay the foundation for computing the observed flux. This subsection provides a robust geometric framework that connects the static surface description with the fully relativistic treatment used in our forward model.

2.2. Spacetime and oblate-Schwarzschild (OS) approximation

We go on to specify the exterior spacetime and the rotation-induced stellar shape. Throughout this work, we adopted the oblate-Schwarzschild (OS) approximation: we traced the photon geodesics in the Schwarzschild spacetime of mass, M , while representing the stellar surface as rotationally flattened using the empirical shape prescription of Morsink et al. (2007), as refined by AIGendy & Morsink (2014) and Silva et al. (2021). This extends the classic ‘‘Schwarzschild+Doppler’’ (S+D) treatment (Poutanen & Beloborodov 2006) by including surface oblateness, while neglecting frame dragging and the external quadrupole moment of the metric. The Schwarzschild exterior metric is expressed as

$$ds^2 = -\left(1 - \frac{R_S}{r}\right) c^2 dt^2 + \left(1 - \frac{R_S}{r}\right)^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2. \quad (5)$$

The g_{tt} component determines gravitational redshift and time dilation, while spatial curvature bends null geodesics. For a spherical star, the stellar radius is constant, $R = R_e$, while the comoving-frame area element is

$$dS' = R_e^2 d\cos\theta d\phi. \quad (7)$$

Rotation deforms the star because centrifugal support reduces the effective gravity near the equator. We modeled the radius profile as

$$\begin{aligned} R &= R_e (1 + o_2 \cos^2 \theta), \\ o_2 &= \bar{\Omega}^2 (-0.788 + 1.030 x), \end{aligned} \quad (8)$$

where o_2 is the shape coefficient and the dimensionless frequency, $\bar{\Omega}$, and compactness parameter, x , are defined as

$$\bar{\Omega} \equiv 2\pi\nu_* \sqrt{\frac{R_e^3}{GM}}, \quad x \equiv \frac{GM}{c^2 R_e}. \quad (9)$$

The quadratic term encodes polar flattening ($o_2 < 0$) and remains accurate for $\nu_* \lesssim 700$ Hz and $R_S/R_e \lesssim 1$, with fractional errors $\lesssim 1\%$ when higher-order terms are neglected. For oblate star, the corresponding surface area element is expressed as

$$dS' = R^2 \sqrt{1 + f^2} d\cos\theta d\phi, \quad (10)$$

with f defined in Equation (2). This ensures that the element measures the proper area of the tilted surface. The effective surface gravity varies with latitude due to the combined effects of oblateness and rotation. Following [AlGendy & Morsink \(2014\)](#); [Bogdanov et al. \(2019b\)](#), we have

$$\begin{aligned} g(\theta)/g_0 &= 1 + \sin^2\theta [\bar{\Omega}^2(-0.791 + 0.776x) \\ &\quad + \bar{\Omega}^4(-1.315x + 2.431x^2) + \bar{\Omega}^6(-1.172x)] \\ &+ \cos^2\theta [\bar{\Omega}^2(1.138 - 1.431x) \\ &\quad + \bar{\Omega}^4(0.653x - 2.864x^2) + \bar{\Omega}^6(0.975x) \\ &\quad - \bar{\Omega}^4(13.47x - 27.13x^2)] \\ &+ |\cos\theta| \bar{\Omega}^4(13.47x - 27.13x^2). \end{aligned} \quad (11)$$

Here, the reference gravity g_0 of the corresponding spherical star is

$$g_0 = \frac{GM}{R_e^2} \frac{1}{\sqrt{1 - R_S/R_e}}. \quad (12)$$

At the equator ($\theta = \pi/2$), gravity is reduced, potentially affecting atmospheric structure and beaming, whereas at the poles it is enhanced.

The OS approximation is well suited to the millisecond pulsars targeted by NICER (typical spins $\nu_* \lesssim 400$ Hz) and delivers $\lesssim 1\%$ waveform accuracy in the tests we consider. Errors remain modest up to ~ 600 – 700 Hz, which motivates recent refinements of the OS approximation ([Jakab & Morsink 2025](#)). At higher spins ($\gtrsim 700$ – 800 Hz), frame dragging and the external metric's quadrupole terms become non-negligible, and Hartle–Thorne or fully relativistic rotating-star spacetime might be required. These choices fix the geometric setting for the rest of the calculation: null geodesics in Schwarzschild provide the mapping from the local emission direction and emission time at the surface to the observed angle and phase, naturally incorporating gravitational light bending and time-of-flight delays.

2.3. Photon propagation: light bending and time delays

The dominant GR effects relevant to PPM are gravitational light bending, gravitational redshift, and time-of-flight delays in the curved spacetime around the neutron star. Gravitational light bending deflects photon trajectories away from straight lines, increasing the visible fraction of the surface; regions with $\psi > \pi/2$ that would be hidden in Newtonian geometry can remain visible. Time-of-flight delays introduce phase shifts between photons emitted at different surface locations and directions, modifying the waveform shape and its harmonic content. Accurate treatment of these effects is essential because modeling errors can bias the inferred compactness $GM/(Rc^2)$ and hence the derived M and R , rather than merely inflating statistical uncertainties.

Gravitational light bending is quantified by the deflection angle ψ for a photon emitted at radius, R , with the emission angle, α , measured from the local radial direction. For outgoing rays ($0 \leq \alpha < \pi/2$), it is given by the Schwarzschild null-geodesic integral,

$$\psi(u, \alpha) = \int_R^\infty \frac{dr}{r^2} \left[\frac{1}{b^2} - \frac{1}{r^2} \left(1 - \frac{R_S}{r} \right) \right]^{-1/2}, \quad (13)$$

where the impact parameter is $b = R \sin\alpha / \sqrt{1-u}$ and $u = R_S/R$. This integral derives from the null geodesic equation in the Schwarzschild metric, reflecting how gravity pulls photons inward.

For inward emission ($\alpha > \pi/2$), the photon path passes through periastron before escaping via

$$\psi(u, \alpha) = 2\psi_{\max} - \psi(u, \pi - \alpha), \quad (14)$$

where $\psi_{\max} \equiv \psi(u_p, \pi/2)$ and the periastron compactness is $u_p = u/\rho$, with

$$\rho = \frac{-2 \sin\alpha}{\sqrt{3(1-u)}} \cos\left(\frac{1}{3} \left[\arccos\left(\frac{3u\sqrt{3(1-u)}}{2 \sin\alpha} \right) + 2\pi \right] \right). \quad (15)$$

Using Equations (13)–(15), we precomputed two-dimensional (2D) lookup tables that map $(u, \cos\psi) \mapsto (\cos\alpha, d\cos\alpha/d\cos\psi)$; these speed up ray tracing while retaining a high level of accuracy.

To calculate the time delay for a spherical star, we define zero delay for a photon emitted from the point directly beneath the observer (along \hat{k}). The gravitational and geometric delay for a photon emitted at the same radius R but at angle $\alpha < \pi/2$ is

$$\frac{c \Delta t_p}{R}(u, \alpha) = \frac{1}{R} \int_R^\infty \frac{dr}{1 - R_S/r} \left\{ \left[1 - \frac{b^2}{r^2} \left(1 - \frac{R_S}{r} \right) \right]^{-1/2} - 1 \right\}. \quad (16)$$

For an oblate star, an additional delay arises from the difference between the actual emission radius, R , and the equatorial radius, R_e ,

$$c \delta t(R) = R_e - R + R_S \ln\left(\frac{R_e - R_S}{R - R_S} \right). \quad (17)$$

For direct emission ($\alpha < \pi/2$), the total delay is $\Delta t = \Delta t_p + \delta t$. For inward emission ($\alpha > \pi/2$), the photon first moves inward and then outward; the total delay is

$$\begin{aligned} \frac{c \Delta t}{R}(u, \alpha) &= \frac{c}{R} \left[2 \Delta t_p(u_p, \pi/2) - \Delta t_p(u, \pi - \alpha) \right. \\ &\quad \left. + 2 \delta t(R_p) - \delta t(R) \right], \end{aligned} \quad (18)$$

where $R_p = \rho R$ is the periastron radius. Precomputing $(u, \cos \alpha) \mapsto (c \Delta t/R)$ yields fast, stable phase corrections for rotating stars. For a detailed explanation and discussion about the derivation above, we refer to [Salmi et al. \(2018, Section 2.1 and Appendix A\)](#). Interpolating from these three tables during likelihood evaluations avoids repeated geodesic integrations and is therefore essential for Bayesian inference, which typically requires $\sim 10^6$ – 10^8 forward-model evaluations.

2.4. Rotation: Relativistic Doppler effect

We incorporate rotation by superimposing SR surface effects on the full GR propagation described above: we Lorentz-boosted the emitted photons in comoving frame into the local static frame on the oblate surface and then trace them along Schwarzschild null geodesics (e.g., [AIGendy & Morsink 2014; Bogdanov et al. 2019b](#)). For millisecond pulsars with spins $\nu_* \sim 300$ Hz (typical equatorial $\beta_{\text{eq}} \equiv v/c \sim 0.1$), this ‘‘OS + Doppler’’ treatment provides an excellent approximation, whereas higher spins may necessitate Hartle–Thorne or fully relativistic rotating-star metrics. Special-relativistic Doppler boosting generates pulse asymmetry and higher harmonics, which are crucial for recovering the hotspot geometry on the neutron-star surface.

The surface velocity, accounting for the gravitational-redshift factor in the metric, is

$$\beta = \frac{2\pi\nu_* R \sin \theta}{c\sqrt{1-u}}, \quad \gamma = \frac{1}{\sqrt{1-\beta^2}}, \quad (19)$$

where the $\sqrt{1-u}$ factor accounts for the gravitational-redshift correction to the rotating frequency.

The relativistic Doppler factor

$$\delta = \frac{1}{\gamma(1-\beta \cos \xi)}, \quad (20)$$

modifies the observed photon energy, E_{obs} , in static frame relative to the emitted energy in the comoving frame, E_{emit} , as $E_{\text{obs}} = \delta E_{\text{emit}}$. This factor depends on the velocity, β , of the emitting region and the angle, ξ , of the velocity vector with respect to the emitting direction. It tends to amplify energy and intensity for photons from approaching patches ($\cos \xi > 0$) and diminishes them for receding ones. The angles in the comoving frame (where the patch is at rest) are transformed via Doppler aberration,

$$\mu = \cos \sigma' = \delta \cos \sigma, \quad \cos \alpha' = \delta \cos \alpha. \quad (21)$$

Quantities with a prime (') are measured in the comoving frame. A detailed derivation of the above result is presented in [Appendix A](#).

The rapid rotation of the neutron star means that the emission from a surface hotspot will be continuously Doppler-shifted throughout the star’s rotational phase. As the hotspot moves across the stellar disk towards the observer, its emission is progressively blueshifted, reaching maximum blueshift as it approaches the line of sight. Conversely, as it recedes, the emission becomes increasingly redshifted. This modulation of photon energy directly impacts the observed flux and the shape of the pulse profile. Therefore, any realistic model of a pulse profile must incorporate the relativistic Doppler effect to correctly interpret the observed light curves and accurately constrain fundamental neutron-star parameters. These SR effects couple directly to the intrinsic atmospheric beaming, as discussed in the next section.

2.5. Surface emission and atmospheric beaming

When photons escape from surface hotspots, a simple starting point is to model the emergent spectrum as a blackbody in the comoving frame. For quantitative inference, however, physically motivated neutron-star atmosphere models are required. Thin (likely accreted) hydrogen or helium layers produce angle- and energy-dependent specific intensities (‘‘beaming’’) and spectra that deviate from a blackbody spectrum, typically appearing harder and exhibiting energy- and composition-dependent limb darkening or brightening. For rotation-powered millisecond pulsars with magnetic fields $B \sim 10^8$ – 10^9 G, weakly magnetized hydrogen atmospheres are commonly adopted; whereas strongly magnetized models are required at higher fields. We followed this atmosphere-based approach throughout, drawing on established calculations and recent updates ([Zavlin et al. 1996; Ho & Lai 2001, 2003; Suleimanov et al. 2012; Salmi et al. 2020, 2023; Choudhury et al. 2024b](#)).

According to ([Bogdanov et al. 2019b, Equation \(31\)](#)), the number flux of photons can be calculated as

$$df(E) = \sqrt{1-u} \delta^3 \mu \frac{d \cos \alpha}{d \cos \psi} \frac{I(E', \mu)}{E} \times \frac{\gamma R^2 \sqrt{1+f^2} d \cos \theta d\phi}{D^2}, \quad (22)$$

where df denotes the photon count rate per unit detector area and per unit observed photon energy; I is the atmospheric specific intensity; E and E' denote the observed and emitted photon energies, respectively; and D is the distance between the star and the observer. A detailed derivation of this expression is provided in [Appendix C](#). Integrating over all visible surface elements, $d \cos \theta d\phi$, yields the full pulse profile.

For realistic beaming, we followed the procedure of [Bogdanov et al. \(2021\)](#). The characteristic field strength at which magnetic effects become non-negligible in the atmosphere is $B_0 \equiv e^3 m_e^2 c / \hbar^3 \simeq 2 \times 10^9$ G. Magnetic effects are negligible in this context because surface fields are smaller than B_0 . We also fixed the composition and ionization. After accretion ceases, gravitational settling stratifies the outer layers, leaving the lightest element at the photosphere; for recycled MSPs, this is typically hydrogen, with helium as a plausible variant. Over the NICER band and at $T_{\text{eff}} \sim (0.5\text{--}2) \times 10^6$ K, differences in angular beaming between H and He are modest (a few percent below 1 keV and at most tens of percent near 2 keV). However, at grazing angles ($\mu \lesssim 0.5$, $\sigma' \gtrsim 60^\circ$) these differences can become substantial (see [Salmi et al. 2023](#)). We therefore adopted a pure, nonmagnetic H atmosphere as the baseline and treat a pure He atmosphere as a systematic variant. At these temperatures the neutral fraction is small and for $\mu \gtrsim 0.5$ fully ionized opacities reproduce the beaming of partially ionized solutions to within a few percent over 0.5–2 keV. Accordingly, we use fully ionized models for $I_\nu(\mu)$. We note that non-local thermodynamic equilibrium (NLTE) and Compton corrections are also negligible in this regime (sub-percent below 2 keV). These choices keep the beaming physics self-consistent with the magnetic-field assumption above and isolate composition and ionization effects to small, quantifiable systematics in PPM inference. Throughout, we adopt the standard assumption of a deep-heated atmosphere in which external energy is deposited below the photosphere. Alternative heating geometries can alter the emergent spectrum and beaming (e.g., [Bauböck et al. 2019; Salmi et al. 2020](#)) and full consistency with the magnetosphere would require modeling

the return-current particle velocities and their stopping column depths, which we defer to future work for improvement.

To compute realistic beaming $I_\nu(\mu)$, we did not solve the radiative transfer during PPM inference. Instead, we used precomputed, high-resolution lookup tables of $\log(I T_{\text{eff}}^{-3})$ as a function of $(\log T_{\text{eff}}, \log g, \mu, \log(E/kT_{\text{eff}}))$. The hydrogen grid spans $\log T_{\text{eff}} \in [5.1, 6.8]$ (helium: $[5.1, 6.5]$), $\log g \in [13.7, 14.7]$, $\mu \in [10^{-6}, 1]$ with extra nodes near $\mu \approx 0, 1$, and $\log(E/kT_{\text{eff}}) \in [-1.3, 2.0]$. Production analyses use the fully ionized H grid. An identically sampled He grid is propagated as a composition systematic. For the full details of this treatment, we refer to [Bogdanov et al. \(2021\)](#) for further discussion.

With the atmosphere model established, the near-surface radiative physics is specified. Two additional components complete the forward model used for data comparison: (i) the surface temperature map and its numerical discretization on the oblate star; and (ii) the propagation to the detector, including interstellar absorption and the instrument response. The next two subsections develop these ingredients in detail. We now transition from the GR+SR forward model to its observational implementation, which converts model photon fluxes into measured counts.

2.6. Surface temperature maps and surface discretizations

Unlike most published PPM studies, which, for computational reasons, typically discretize only parametrized hotspot regions, our framework natively discretizes the entire stellar surface temperature map. Prior work has largely employed one- or two-spot models with simple geometries (e.g., circular caps) and single or dual temperature components. While effective for many applications, such prescriptions cannot capture spatially complex temperature structures within a spot or extended patterns across the surface. Full-surface discretization, in contrast, accommodates nontrivial temperature distributions and global morphologies, which are required in several physical scenarios (e.g., accreting neutron stars with extended heating patterns and nonuniform temperature distribution; see [Das et al. 2022](#); [Das et al. 2025](#)).

In our codebase, the primary input is a user-specified full-surface map $T_{\text{eff}}(\theta, \phi)$ defined on an oblate surface mesh in colatitude, θ , and azimuth, ϕ . The mesh resolution is user-controlled and can be increased until interpolation errors are sub-percent in the phase–energy domain.

For benchmarking with results from the literature, we also provide parameterizations that reproduce standard geometric hotspot models used in previous PPM works. For example, a single circular cap is defined by its center, (θ_0, ϕ_0) , angular radius, ζ , and spot temperature, T_{spot} ; all mesh elements with angular separation $\leq \zeta$ from (θ_0, ϕ_0) are assigned T_{spot} , while elements outside the cap are set to 0, indicating that no photons are emitted. The same strategy extends to antipodal two-cap configurations and other predefined shapes.

Our default physics-motivated generator for whole-surface temperature maps follows [Huang & Chen \(2025\)](#), which derives return-current heating patterns from a force-free magnetosphere with an off-centered dipole. In that study, the implementation used the X-PSI ([Riley et al. 2023a,b](#)) full-surface module and achieved per-map generation times of order one second on a CPU, sufficient for static tests and small-scale inference but not optimized for large-scale PPM inference and for testing different sampler setups for robustness. In the present framework, we adopt this physics-motivated hotspot module as the default for hotspot configurations in future PPM inference. For the full algorithmic details and validation of the hotspot configuration, see [Huang & Chen \(2025\)](#).

2.7. Interstellar absorption and instrument response

Photons emerging from the stellar surface are attenuated by interstellar absorption, which we model as a multiplicative factor applied to the intrinsic spectrum. Accounting for this energy-dependent attenuation is essential for linking theoretical pulse profiles to the observed data and enabling reliable parameter inference. In this work, we model line-of-sight attenuation with a multiplicative transmission factor applied to the emitted photon spectrum,

$$f_{\text{ISM}}(E; N_{\text{H}}) = \exp[-\sigma_{\text{ISM}}(E)N_{\text{H}}] f_0(E), \quad (23)$$

where $f_0(E)$ is the source-frame spectral flux, N_{H} is the hydrogen column, and $\sigma_{\text{ISM}}(E)$ is the energy-dependent cross section from the `tbnew` prescription (modern gas-phase abundances and grain physics) ([Wilms et al. 2000](#)). For efficiency, we use a precomputed attenuation vector on the NICER energy grid at a fiducial $N_{\text{H}}^{\text{ref}}$ and exploit the linear scaling in the exponent to obtain transmission for arbitrary N_{H} . The MSPs targeted for PPM have very low columns ($\sim 10^{20} \text{ cm}^{-2}$), so extinction is small over 0.3–2 keV and, being phase independent, has negligible impact on geometric inferences; we nevertheless carry N_{H} as a nuisance parameter and propagate it through the likelihood.

Another component that connects the theoretical pulse profile to the observed data is the instrument response. Here, our instrument response follows the same strategy implemented in previous NICER works as a standard pipeline. Comparison to NICER data is performed in detector space by convolving the attenuated model spectrum with the instrument response,

$$C(N) = \int_0^\infty R(N, E) f_{\text{ISM}}(E; N_{\text{H}}) dE, \quad (24)$$

where $C(N)$ is the predicted count rate in channel N (counts s^{-1}) and $R(N, E)$ is the full response (RMF×ARF), namely, the probability of detecting a photon of energy, E , in channel, N , multiplied by the effective area. Energy-resolved pulse profiles are generated by evaluating $f_0(E)$ from the atmospheric beaming tables for each visible surface element and phase bin, applying the ISM transmission, and convolving with $R(N, E)$. We verified our forward model by reproducing standard XSPEC results for simple spectra with the same response files and by cross-checking independent implementations. All production fits use the identical RMF+ARF set and channelization as the data to ensure consistency.

2.8. Computation recipe

In this section, we outline the step-by-step computational procedure for generating a phase-resolved spectrum from a rotating neutron star. The process (i) discretizes the stellar surface, (ii) precomputes gravitational light bending, (iii) computes the observed flux from each surface element, and (iv) convolves the result with models for interstellar absorption and the instrument response.

2.8.1. Surface discretization

To accurately model the emission from the star’s surface, we first divide it into a finite number of small patches or pixels. For this, we employ the `HEALPix` (Hierarchical Equal Area isoLatitude Pixelization) scheme ([Górski et al. 2005](#)). This method is particularly advantageous because it divides the sphere into patches of equal area, ensuring that each patch contributes proportionally

to the total flux and simplifying the geometric calculations. The resolution of this grid is controlled by the parameter N_{side} . Let N_θ be the number of rings and $i_\theta = 1, \dots, N_\theta$ the ring index. Each ring is further discretized into N_ϕ azimuthal bins with $i_\phi = 1, \dots, N_\phi$. The center coordinates (θ, ϕ) of each surface patch are computed from $(i_\theta, N_\theta, i_\phi, N_\phi)$ according to the chosen discretization scheme.

For $1 \leq i_\theta < N_{\text{side}}$ and $1 \leq i_\phi \leq 4i_\theta$, we have

$$\cos \theta(i_\theta, i_\phi, N_{\text{side}}) = 1 - \frac{i_\theta^2}{3N_{\text{side}}^2}, \quad (25)$$

$$\phi(i_\theta, i_\phi, N_{\text{side}}) = \frac{\pi}{2i_\theta} \left(i_\phi - \frac{1}{2} \right). \quad (26)$$

For $N_{\text{side}} \leq i_\theta \leq 2N_{\text{side}}$ and $1 \leq i_\phi \leq 4N_{\text{side}}$, we have

$$\cos \theta(i_\theta, i_\phi, N_{\text{side}}) = \frac{4}{3} - \frac{2i_\theta}{3N_{\text{side}}}, \quad (27)$$

$$\phi(i_\theta, i_\phi, N_{\text{side}}) = \frac{\pi}{2N_{\text{side}}} \left(i_\phi - \frac{(i_\theta - N_{\text{side}} + 1) \bmod 2}{2} \right). \quad (28)$$

The southern hemisphere is a mirror image of the northern hemisphere, i.e., for $2N_{\text{side}} < i_\theta \leq 4N_{\text{side}} - 1$ and $1 \leq i_\phi \leq 4 \min(4N_{\text{side}} - i_\theta, N_{\text{side}})$, we have

$$\cos \theta(i_\theta, i_\phi, N_{\text{side}}) = -\cos \theta(4N_{\text{side}} - i_\theta, i_\phi, N_{\text{side}}), \quad (29)$$

$$\phi(i_\theta, i_\phi, N_{\text{side}}) = \phi(4N_{\text{side}} - i_\theta, i_\phi, N_{\text{side}}). \quad (30)$$

In the HEALPix discretization scheme, the surface of the sphere is partitioned into $12N_{\text{side}}^2$ equal-area patches, so each patch covers a solid angle of $d \cos \theta d\phi = \pi/(3N_{\text{side}}^2)$.

In modeling complex emission geometries, we identify the surface patches that belong to the emitting region and designate them as active patches, whose total number is N_{active} . This discretization can introduce a small scaling bias at sharp boundaries because the pixelization may slightly over- or under-represent the true emitting area. A more careful edge treatment could mitigate this bias and potentially reduce computational load, but given that performance is satisfactory under the current setup, we defer such refinements to a future code release.

2.8.2. Lensing table generation

As discussed in Section 2.3, computing gravitational light bending and the associated time-of-flight delays for every surface patch is computationally expensive. To accelerate the calculation, we precompute three lookup tables: (i) the cosine of the emission angle, $\cos \alpha$; (ii) the lensing factor, $d \cos \alpha / d \cos \psi$; and (iii) the time delay, Δt .

These tables are indexed by grids defined over the key physical parameters. We use a helper function to define linearly spaced arrays:

$$\mathbf{A} = \text{Linspace}(x_{\text{min}}, x_{\text{max}}, x_{\text{num}}), \quad (31)$$

to denote an array \mathbf{A} of size x_{num} , whose elements are given by

$$\mathbf{A}[i] = x_{\text{min}} + (x_{\text{max}} - x_{\text{min}}) \frac{i}{x_{\text{num}} - 1}, \quad (32)$$

where $i = 0, \dots, x_{\text{num}} - 1$. The grids for the compactness (u), the cosine of the final bending angle ($\cos \psi$), and the cosine of the emitting angle ($\cos \alpha$) are defined as

$$\mathbf{u} = \text{Linspace}(u_{\text{min}}, u_{\text{max}}, N_u),$$

$$\mathbf{c}_\psi = \text{Linspace}(c_{\psi, \text{min}}, c_{\psi, \text{max}}, N_{c_\psi}),$$

$$\mathbf{c}_\alpha = \text{Linspace}(c_{\alpha, \text{min}}, c_{\alpha, \text{max}}, N_{c_\alpha}).$$

For all calculations, we adopted the following empirically chosen values, selected to ensure all lookups remain within bounds for typical neutron-star parameters:

$$u_{\text{min}} = 0.3, \quad u_{\text{max}} = 0.5, \quad N_u = 256,$$

$$c_{\psi, \text{min}} = -0.9, \quad c_{\psi, \text{max}} = 1.0, \quad N_{c_\psi} = 512,$$

$$c_{\alpha, \text{min}} = -0.2, \quad c_{\alpha, \text{max}} = 1.0, \quad N_{c_\alpha} = 512.$$

For $\alpha < \pi/2$,

$$\psi(u, \alpha) = \frac{2 \sin \alpha}{\sqrt{1-u}} \int_0^1 \frac{x dx}{\sqrt{\cos^2 \alpha + x^2 q}}, \quad (33)$$

$$q = [2 - x^2 - u(1 - x^2)^2 / (1 - u)] \sin^2 \alpha, \quad (34)$$

and for $\alpha > \pi/2$,

$$\psi(u, \alpha) = 2\psi(u/\rho, \pi/2) - \psi(u, \pi - \alpha). \quad (35)$$

We created two 2D arrays, \mathbf{C}_α and \mathbf{LF} , both of a size (N_u, N_{c_ψ}) , to store precomputed lookup tables for the cosine of the local emission angle $\cos \alpha$ and the lensing factor $d \cos \alpha / d \cos \psi$. For each $\mathbf{u}[i_u]$ in the array \mathbf{u} , we set $\alpha_{\text{cur}} = 0$, calculated $\psi_{\text{cur}} = \psi(\mathbf{u}[i_u], \alpha_{\text{cur}})$, and appended $(\alpha_{\text{cur}}, \psi_{\text{cur}})$ to a list $\{(\alpha, \psi), \dots\}$. We then increased α_{cur} by $d\alpha = \pi/2048$ (an empirically chosen constant; it can be decreased for higher accuracy) until $\psi_{\text{cur}} > \arccos c_{\psi, \text{min}}$. With a simple transformation, we converted the list $\{(\alpha, \psi), \dots\}$ to $\{(\cos \psi, \cos \alpha), \dots\}$; using numerical differentiation, we obtain the list $\{(\cos \psi, d \cos \alpha / d \cos \psi), \dots\}$. Then, for each $\mathbf{c}_\psi[i_{c_\psi}]$ in the array \mathbf{c}_ψ , we interpolated $\mathbf{c}_\psi[i_{c_\psi}]$ within these two lists to obtain the corresponding $\mathbf{C}_\alpha[i_u, i_{c_\psi}]$ and $\mathbf{LF}[i_u, i_{c_\psi}]$. Next, we moved on to the calculation of the time delay. For $\alpha < \pi/2$,

$$\frac{c\Delta t_p}{R}(u, \alpha) = \frac{2 \sin^2 \alpha}{1-u} \int_0^1 \frac{x dx}{\sqrt{\cos^2 \alpha + x^2 q} (1 + \sqrt{\cos^2 \alpha + x^2 q})}. \quad (36)$$

For $\alpha > \pi/2$,

$$\frac{c\Delta t_p}{R}(u, \alpha) = 2 \frac{c\Delta t_p}{R}(u/\rho, \pi/2) - \frac{c\Delta t_p}{R}(u, \pi - \alpha). \quad (37)$$

We used a 2D array \mathbf{TD} of size (N_u, N_{c_α}) to store the lookup table of time delay as a function of (u, c_α) . For each grid point $(\mathbf{u}[i_u], \mathbf{c}_\alpha[i_{c_\alpha}])$, we directly calculated the time delay using the integral above and stored the result, expressed as

$$\mathbf{TD}[i_u, i_{c_\alpha}] = \frac{c\Delta t_p}{R}(\mathbf{u}[i_u], \arccos \mathbf{c}_\alpha[i_{c_\alpha}]). \quad (38)$$

The light-bending integral (Equation (33)) and the time-delay integral (Equation (36)) are given in Salmi et al. (2018, Appendix A).

2.8.3. Number flux calculation

This step computes the observed flux from each surface patch and assembles the pulse profile. We used a widely employed “shift-and-add” technique: rather than simulating every active patch across all rotation phases, we first computed a detailed light curve for a single pseudo-patch located at $\phi = 0$ on a given colatitude ring θ . This template light curve was then phase-shifted and added to the total to account for the actual longitudes of the active patches on that ring.

For each ring indexed by i_θ at a colatitude, θ , we determined which patches are active and collect the azimuthal angles, ϕ , of the active patches into an array, ϕ_a , of a length N_{ϕ_a} . If $N_{\phi_a} = 0$, the ring contains no active patches and is skipped. Otherwise, we prepared four arrays to store the pseudo-patch light-curve data at $\phi = 0$. These were reused later when accumulating contributions from the active patches:

$$\mathbf{F}[i_{\text{fp}}], \quad \mathbf{z}[i_{\text{fp}}], \quad \boldsymbol{\varphi}_o[i_{\text{fp}}], \quad \boldsymbol{\mu}[i_{\text{fp}}], \quad (39)$$

where $i_{\text{fp}} = 0, \dots, N_{\text{fp}}$ and N_{fp} is the number of fine rotation-phase points, this parameter controls precision. When the pseudo-patch rotates to azimuth $\phi(i_{\text{fp}}) = 2\pi i_{\text{fp}}/N_{\text{fp}}$, we perform the following for each $i_{\text{fp}} \in [0, N_{\text{fp}}]$:

1. Initialize invisible-range indices $i_{\text{inv_min}} = N_{\text{fp}}$ and $i_{\text{inv_max}} = 0$.
2. Compute $\cos \psi = \sin i \sin \theta \cos \phi + \cos i \cos \theta$.
 - Use the lookup table to obtain $\cos \alpha$ and the lensing factor $d \cos \alpha / d \cos \psi$.
 - Evaluate $\cos \sigma$ from Equation (4):

$$\cos \sigma = \cos \alpha \cos \eta + \frac{\sin \alpha \sin \eta}{\sin \psi} \quad (40)$$

- $\times (\cos i \sin \theta - \sin i \cos \theta \cos \phi)$.
- If $\cos \sigma < 0$, the point is invisible (and if $i_{\text{fp}} = 0$, the entire ring is invisible). Update $i_{\text{inv_min}} = \min(i_{\text{inv_min}}, i_{\text{fp}})$ and $i_{\text{inv_max}} = \max(i_{\text{inv_max}}, i_{\text{fp}})$. Otherwise, set

$$\mathbf{F}[i_{\text{fp}}] = \sqrt{1-u} \delta^3 \cos \sigma' \frac{d \cos \alpha}{d \cos \psi} \frac{\gamma R^2 \sqrt{1+f^2} d \cos \theta d \phi}{D^2},$$

$$\mathbf{z}[i_{\text{fp}}] = \frac{1}{\delta \sqrt{1-u}},$$

$$\boldsymbol{\varphi}_o[i_{\text{fp}}] = \frac{\phi}{2\pi} + v_* \Delta t,$$

$$\boldsymbol{\mu}[i_{\text{fp}}] = \delta \cos \sigma. \quad (41)$$

If a ring is partially visible, we must handle the phase segment where it is hidden. For these invisible points ($i_{\text{inv_min}} \leq i_{\text{fp}} \leq i_{\text{inv_max}}$), we set the flux and apparent emission angle to zero. To ensure continuity and avoid numerical artifacts, the redshift and observed phase are smoothly filled using linear interpolation between the values at the points of disappearance ($i_{\text{inv_min}} - 1$) and reappearance ($i_{\text{inv_max}} + 1$), namely,

$$\mathbf{F}[i_{\text{fp}}] = 0, \quad \boldsymbol{\mu}[i_{\text{fp}}] = 0, \quad (42)$$

$$\mathbf{z}[i_{\text{fp}}] = \mathcal{L}(i_{\text{fp}}; i_{\text{inv_min}} - 1, i_{\text{inv_max}} + 1, \mathbf{z}[i_{\text{inv_min}} - 1], \mathbf{z}[i_{\text{inv_max}} + 1]), \quad (43)$$

$$\boldsymbol{\varphi}_o[i_{\text{fp}}] = \mathcal{L}(i_{\text{fp}}; i_{\text{inv_min}} - 1, i_{\text{inv_max}} + 1, \boldsymbol{\varphi}_o[i_{\text{inv_min}} - 1], \boldsymbol{\varphi}_o[i_{\text{inv_max}} + 1]). \quad (44)$$

where \mathcal{L} denotes linear interpolation:

$$\mathcal{L}(x; x_0, x_1, y_0, y_1) = y_0 + (y_1 - y_0) \frac{x - x_0}{x_1 - x_0}. \quad (45)$$

After preparing these calculations, we used the “shift-and-add” technique to accumulate all active patches’ contributions to the total waveform. We defined the phase points to be output as an array, $\boldsymbol{\varphi}_{\text{out}}$, of a size $N_{\boldsymbol{\varphi}_{\text{out}}}$, and the observed photon energies as an array, \mathbf{E}_o , of a size $N_{\mathbf{E}_o}$, and the total output photon number flux as a matrix, \mathbf{F}_t , of a specific size $(N_{\boldsymbol{\varphi}_{\text{out}}}, N_{\mathbf{E}_o})$. For each $\phi = \phi_a[i_{\phi_a}]$ with $i_{\phi_a} = 0, \dots, N_{\phi_a} - 1$, each $\varphi = \boldsymbol{\varphi}_{\text{out}}[i_{\varphi_{\text{out}}}]$ with $i_{\varphi_{\text{out}}} = 0, \dots, N_{\boldsymbol{\varphi}_{\text{out}}} - 1$, and each $E = \mathbf{E}_o[i_{E_o}]$ with $i_{E_o} = 0, \dots, N_{\mathbf{E}_o} - 1$. Next, we can compute

$$\varphi_p = ((\varphi + \frac{\phi}{2\pi} + 1 - \boldsymbol{\varphi}_o[0]) \bmod 1) + \boldsymbol{\varphi}_o[0], \quad (46)$$

$$\mathbf{F}_p = \text{Interp}(\varphi_p; \boldsymbol{\varphi}_o, \mathbf{F}),$$

$$\mathbf{z}_p = \text{Interp}(\varphi_p; \boldsymbol{\varphi}_o, \mathbf{z}), \quad (47)$$

$$\boldsymbol{\mu}_p = \text{Interp}(\varphi_p; \boldsymbol{\varphi}_o, \boldsymbol{\mu}),$$

where $\text{Interp}(x_0; \mathbf{x}, \mathbf{y})$ denotes building an interpolation function from arrays \mathbf{x} and \mathbf{y} and evaluating it at x_0 . The contribution to the total flux is found by scaling, \mathbf{F}_p , with the specific intensity of the atmospheric emission, $I(z_p E, \boldsymbol{\mu}_p)$, which depends on the redshifted energy and the apparent emission angle. This contribution is then added to the final light-curve matrix, \mathbf{F}_t ,

$$\mathbf{F}_t[i_{\varphi_{\text{out}}}, i_{E_o}] += \mathbf{F}_p \frac{I(z_p E, \boldsymbol{\mu}_p)}{E}. \quad (48)$$

2.8.4. Interstellar absorption and instrument response

The photon number flux calculated above is a theoretical quantity. To compare it with real data, we must model its interaction with the interstellar medium (ISM) and the response of the detector. We can define logarithmically spaced arrays using the helper function

$$\mathbf{A} = \text{Logspace}(x_{\text{min}}, x_{\text{max}}, x_{\text{num}}), \quad (49)$$

which represents an array \mathbf{A} of length x_{num} whose elements are given by

$$\mathbf{A}[i] = \exp\left(\ln x_{\text{min}} + (\ln x_{\text{max}} - \ln x_{\text{min}}) \frac{i}{x_{\text{num}} - 1}\right), \quad (50)$$

where $i = 0, \dots, x_{\text{num}} - 1$.

Throughout all calculations, we use the grid

$$\mathbf{E}_o = \text{Logspace}(0.1 \text{ keV}, 5.215 \text{ keV}, N_{\mathbf{E}_o} = 128). \quad (51)$$

We modeled the absorption by the ISM using a precomputed table from the `tbnew` model. The table provides the attenuation fraction as a function of energy for a reference hydrogen column density ($N_{\text{H}}^{\text{ref}} = 4 \times 10^{18} \text{ cm}^{-2}$). We used `col1` to represent the energy column in the table and `col2` to represent the attenuation-fraction column. Our simulation uses $N_{\text{H}} = 2 \times 10^{19} \text{ cm}^{-2}$. Assuming the optical depth scales linearly with N_{H} , the new attenuation is the original value raised to the power of the ratio $N_{\text{H}}/N_{\text{H}}^{\text{ref}}$ (i.e., = 5). We created an array `ATT` by interpolating the table to our energy grid and applying this scaling; `ATT` is of a size denoted as $N_{\mathbf{E}_o}$ and elements expressed as

$$\mathbf{ATT}[i_{E_o}] = \text{Interp}(\mathbf{E}_o[i_{E_o}]; \text{col1}, \text{col2})^5. \quad (52)$$

Next, we convolve the signal with the instrumental response of the detector (e.g., NICER). This requires two standard data files: an ARF file and an RMF file. The ARF file provides the effective area of the detector across different energies and contains three columns that we denote E_l , E_h , and A_e all with size N_e . The RMF file describes how photons of a given true energy are recorded in different detector output channels and contains a matrix, \mathbf{R} , of a specific size (N_e, N_c) . We multiplied the effective area from the ARF with the RMF to obtain a fine-resolution response matrix, \mathbf{R}_f , with a specific size (N_e, N_c) :

$$\mathbf{R}_f[i_e, i_c] = A_e[i_e] \times \mathbf{R}[i_e, i_c]. \quad (53)$$

We defined the array, e_i , with a length, N_{E_o} , whose elements are given by

$$e_i[j] = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{otherwise.} \end{cases} \quad (54)$$

Define the downsampling matrix \mathbf{DS} with size (N_{E_o}, N_e) by

$$\mathbf{DS}[i_{E_o}, i_e] = \int_{E_l[i_e]}^{E_h[i_e]} \text{Interp}(x; E_o, e_{i_{E_o}}) dx. \quad (55)$$

Combining ISM attenuation, \mathbf{ATT} , the downsampling matrix, \mathbf{DS} , and the fine instrument response, \mathbf{R}_f , yields the final, comprehensive response matrix, \mathbf{RSP} . This matrix directly converts the theoretical, energy-resolved photon flux into the expected counts per detector channel. The resulting \mathbf{RSP} has a specific size (N_{E_o}, N_c) and is given by

$$\mathbf{RSP}[i_{E_o}, i_c] = \sum_{i_e=0}^{N_e-1} \mathbf{ATT}[i_{E_o}] \times \mathbf{DS}[i_{E_o}, i_e] \times \mathbf{R}_f[i_e, i_c]. \quad (56)$$

2.9. GPU implementation

The NVIDIA GPU architecture is designed for massively parallel computation, executing thousands of threads simultaneously. The core processing unit is the Streaming Multiprocessor (SM), each of which contains hundreds of simpler arithmetic units known as CUDA cores. When a computational task is launched, it is organized into a grid of thread blocks; one or more blocks can be assigned to an SM for execution. The efficiency of this architecture depends on its memory hierarchy. All threads can access a large, high-capacity global memory, which is analogous to CPU random-access memory (RAM) but has high latency, making frequent access a primary performance bottleneck. To mitigate this, each SM is equipped with a small, extremely fast on-chip shared memory. This memory is private to the threads within a single block and acts as a programmable cache. In a typical GPU algorithm for a problem such as the one described below, a block of threads first cooperatively loads the necessary subset of data from global memory into this low-latency shared memory. All subsequent computations can then be performed by accessing the shared memory, reducing memory overhead and unlocking the full parallel processing power of the CUDA cores before writing the final results back to global memory.

In our waveform generation process (as described in Section 2.8), we treated each ring independently. In the GPU implementation, we assigned each ring's computation to a dedicated SM. For each SM, the four helper arrays \mathbf{F} , \mathbf{z} , $\boldsymbol{\varphi}_o$, and $\boldsymbol{\mu}$ were staged in shared memory, whereas the number-flux matrix, \mathbf{F}_t , the lensing lookup table, the (optional) numerical atmosphere

lookup table, and the \mathbf{RSP} matrix reside in global memory. When a kernel is launched, the host (CPU) provides all configuration data describing the task, including the star's mass, radius, and spin frequency (M, R, ν_*) ; the observer's inclination (i) ; the hotspot's shape, position, and size $(\{\theta_i, \phi_i, \zeta_i, \dots\})$; and the ring indices it operates on (i_θ, N_θ) . For each active surface patch, output phase, and output energy, the kernel computes the corresponding number flux and immediately writes the result to the matrix \mathbf{F}_t stored in global memory using Equation (48). After all tasks on the SMs were complete, we multiplied the matrix, \mathbf{F}_t , by the matrix, \mathbf{RSP} , which combines ISM absorption and the detector's instrument response. We then copy the final result from GPU global memory to CPU host memory for analysis and plotting.

3. OS approximation benchmarks

In this section, we describe how we tested the OS approximation and reproduced the theoretical benchmarks established in Bogdanov et al. (2019b) that were generated by the Illinois–Maryland (IM) group code. At this stage, no atmosphere model or instrument-specific modeling was included; all tests assumed blackbody emission. We therefore evaluated the code accuracy using blackbody spectra and prescribed beaming laws.

Following Bogdanov et al. (2019b), we verified both S+D and OS. The S+D scheme treats special-relativistic effects exactly at the surface, assuming a spherical star and using the exterior Schwarzschild metric; it is very fast and accurate for slowly rotating stars (see also Poutanen & Beloborodov 2006). The OS scheme incorporates the star's spin-induced oblateness while retaining an exterior Schwarzschild spacetime; it is now the standard in X-ray PPM. Direct comparisons against ray tracing in numerically computed rotating-star metrics show that OS waveforms agree with the exact results to within $\lesssim 0.1\%$ for spin frequencies $\nu_* \lesssim 300$ Hz, whereas S+D can incur percent-level errors at higher spins. Hence, we adopted OS as the default for our production runs (Bogdanov et al. 2019b; Morsink et al. 2007) and implemented the OS approximation in our GPU codebase.

We followed the hotspot geometries and spacetime choices in Bogdanov et al. (2019b). Specifically, we tested both point-like ($\zeta_{\text{spot}} = 0.01$ rad) and extended ($\zeta_{\text{spot}} = 1$ rad) circular hotspots. Beaming patterns included (i) isotropic emission and (ii) simple parametric laws $\propto \cos^2 \sigma$ and $\propto \sin^2 \sigma$ used for verification in Bogdanov et al. (2019b, Tables 3–4), where σ is the angle between the surface normal and the emission direction. In this section, we assume a canonical neutron-star mass $M = 1.4 M_\odot$, equatorial radius $R_e = 12$ km, distance $D = 200$ pc, and a surface temperature $kT = 0.35$ keV. We considered monochromatic pulse profiles evaluated at $E_{\text{obs}} = 1$ keV with a photon number flux of $1 \text{ cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$, matching the reference tests in Bogdanov et al. (2019b).

We implemented both CPU and GPU versions of our codebase. For validation, we compare our waveforms phase-by-phase against the IM reference waveforms used in Bogdanov et al. (2019b). Consistent with the NICER validation protocol, our target is a fractional accuracy better than 0.1% across phases under likelihood-relevant weighting, where phase bins contribute in proportion to their expected counts. In this sense, we achieve the target accuracy, with larger relative deviations confined to ingress/egress when the hotspot lies behind the stellar limb and the flux is orders of magnitude below the pulse maximum. Because Poisson likelihood contributions scale with the expected counts, these low-flux bins carry little statistical

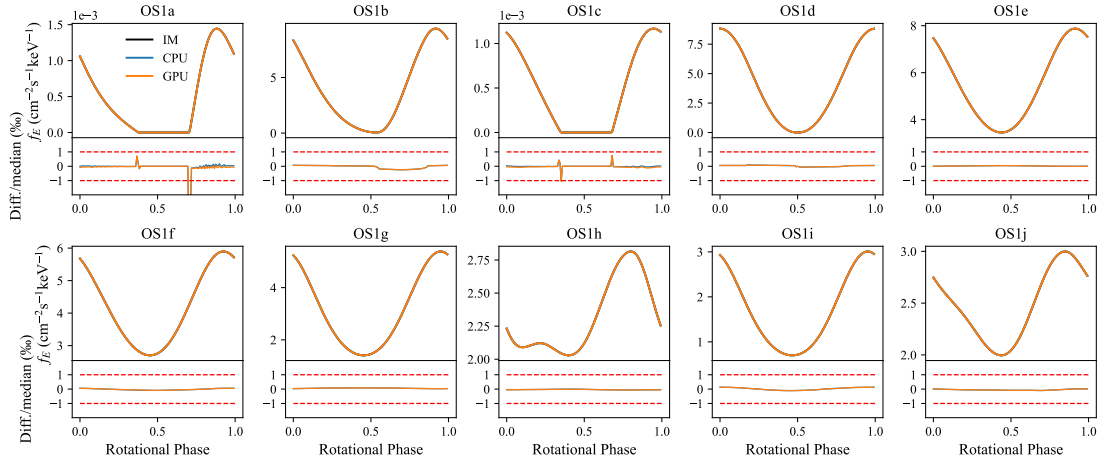


Fig. 2. Monochromatic ($E_{\text{obs}} = 1$ keV) waveforms under the oblate-Schwarzschild (OS) approximation for the OS1 test suite (subset shown), computed with our CPU (blue) and GPU (orange) implementations. The theoretical benchmarks from Bogdanov et al. (2019b) are shown in black. Residuals are taken with respect to the benchmark; red dashed lines indicate $\pm 0.1\%$. Test-case definitions follow Bogdanov et al. (2019b).

Table 3. Parameter values for atmosphere-interpolation tests..

Case name	Test 1	Test 2
Number of hotspots	1	1
Angular radius of hotspot (rad)	0.01	0.01
Colatitude of spot center θ_c (rad)	1.983	$\pi - 0.01$
Observer colatitude i (rad)	0.01	0.964
NS mass (M_{\odot})	1.4	1.4
NS equatorial radius (km)	12	12
ν at infinity (Hz)	600	600
Spectrum of emission	nsx_H_v200804	nsx_H_v200804
Temperature of emission (K)	10^6	10^6
Distance (pc)	150	150

weight in typical datasets. Moreover, the IM treatment about the stellar limb, as implemented in Bogdanov et al. (2019b), is a useful reference but not ground truth. Accordingly, we defer any certification of limb (small- μ) influence on the waveform computation to a dedicated stress test in Sect. 4. Since these systematic errors become significantly more consequential under realistic atmosphere model. These localized discrepancies have a negligible impact on inference because likelihood evaluations are generally insensitive to tiny, low-count portions of the profile.

Across all OS test cases (OS1a–OS1j; see Figure 2 for a subset), our CPU and GPU implementations reproduce the IM reference waveforms within the 0.1% target at essentially all phases outside eclipse ingress/egress. The few outliers occur precisely where the spot rotates out of and back into view; as noted above, these phases have minimal photon counts and do not affect practical parameter estimation. This level of agreement meets the precision standard required for theoretical PPM forward computation and provides a reliable foundation for extensions such as energy-resolved profiles with atmospheric beaming and instrument responses. In the next section, we incorporate a realistic atmosphere model and identify an important interpolation artifact in the widely adopted atmospheric lookup table. We then introduce two diagnostic test cases to quantify its impact and present a more detailed implementation that mitigates the artifact.

4. Test cases on atmosphere interpolation

In the previous section, we explain how we verified that our implementation reproduces published theoretical benchmarks under the OS approximation and that the resulting light curves agree with canonical test problems. Before turning to production-grade studies, we highlight another common source of systematic error that is missing from much of the existing literature: interpolation within precomputed neutron-star atmosphere tables.

To avoid repeatedly solving the radiative-transfer problem on the fly, modern forward models use predefined lookup tables of the emergent specific intensity. This step is essential for realistic PPM: replacing a blackbody with a physically motivated atmosphere is required for meaningful comparison with NICER observations. In practice, many groups use nonmagnetized atmosphere grids for hydrogen or helium compositions, often referred to as NSX-H and NSX-He, which tabulate $I_{\nu}(\mu; T_{\text{eff}}, g, E)$ as a function of photon energy E , emission-angle cosine $\mu = \cos \sigma'$, effective temperature T_{eff} , and surface gravity g (e.g., Ho & Lai 2001, 2003). In what follows, we examine how interpolation choices on these grids can affect the modeled waveforms, and we propose additional test cases designed to demonstrate the numerical behavior of interpolation on the NSX atmosphere lookup tables. In our experiments, evaluating the NSX atmosphere tables and performing the associated interpolation

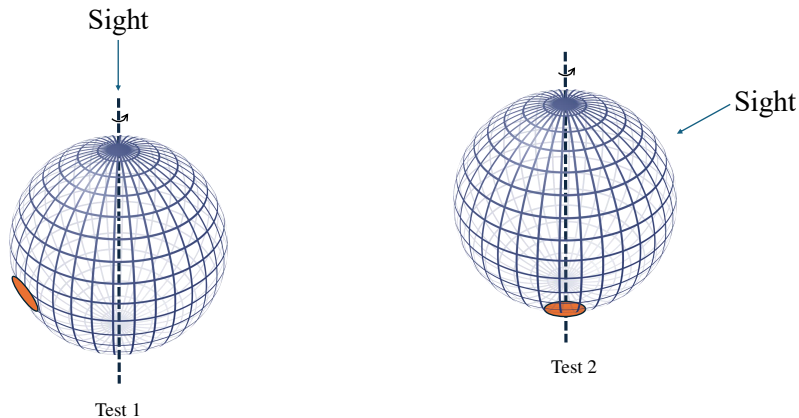


Fig. 3. Demonstration of two geometric test cases for atmosphere interpolation. Test 1: point-like hotspot on the southern hemisphere; the line of sight is aligned with the spin axis and lies above the rotational north pole. Test 2: point-like hotspot at the south pole; the observer is located in the northern hemisphere.

constitute the dominant computational cost in pulse profile calculations. More importantly, different interpolation schemes can introduce distinct numerical artifacts (e.g., overshooting near table boundaries), which may bias predicted fluxes and spectra quite significantly under our proposed test cases.

Several widely used codes adopt different interpolation strategies for these tables. As reported in Bogdanov et al. (2019b, 2021); Choudhury et al. (2024b), the IM code employs quadratic and quartic polynomial interpolation, X-PSI primarily uses cubic Lagrange interpolation, and the Alberta implementation relies on linear schemes (see Bogdanov et al. 2019b, 2021; Choudhury et al. 2024b for details). Because IM is not open source, its precise interpolation details are not publicly documented. In contrast, the publicly available X-PSI code uses cubic Lagrange interpolation for atmosphere lookups (Riley et al. 2023b; Bogdanov et al. 2021). Given the structure of the tabulated NSX H/He atmosphere grids (Ho & Lai 2001, 2003), these choices can introduce method-dependent artifacts, underscoring the need for careful treatment.

A commonly underappreciated numerical issue in atmosphere-table usage is polynomial overshoot near grid boundaries. Because the widely used nonmagnetized NSX H/He grids have finite spacing and regions of rapid curvature, selecting an interpolation scheme that avoids spurious oscillations while preserving physical constraints (e.g., non-negativity of the specific intensity) is nontrivial. In principle, as the grid becomes sufficiently dense, low-order (e.g., linear) interpolation is typically the most robust and economical. A comprehensive, controlled comparison of interpolation orders and monotone schemes, ideally alongside finer atmosphere grids or direct radiative-transfer recalculations, is deferred to future work.

In our experiments with NSX-H grids, cubic Lagrange interpolation can exhibit boundary instabilities (overshoot/undershoot) that depend on local spacing and curvature, particularly as the emission-angle cosine approaches an endpoint ($\mu \rightarrow 0$; grazing emission). To enforce non-negativity, the open-source X-PSI implementation clips any negative interpolants to zero.¹ While this safeguard prevents unphysical negative intensities, it still cannot compute the predicted flux downward properly in phases where many rays sample boundary regions.

Motivated by these considerations, we adopt a hybrid interpolation scheme in this work. For the temperature and gravity axes of the NSX-H table, $\log T_{\text{eff}}$ and $\log g$ (the NSX-H tables tabulate these quantities in logarithmic units), we employ linear interpolation. The grids are sufficiently dense, making linear interpolation both robust and accurate in practice. For the emission-energy axis, we use cubic Lagrange interpolation in $\log(E/kT_{\text{eff}})$; for the emission-angle cosine μ , we use cubic Lagrange interpolation in the interior but switch to linear interpolation within boundary zones at $\mu \rightarrow 0$. Because a cubic Lagrange interpolant uses four points, we designate the outermost four grid nodes on the small- μ end as “protected” regions in which linear interpolation is enforced.

In a preliminary numerical test, we directly interpolated the NSX-H atmosphere tables and examined the resulting specific intensities. We found that the boundary behavior as $\mu \rightarrow 0$ (i.e., near the stellar limb) is particularly susceptible to interpolation overshoot, occasionally producing unphysical negative intensities. This sensitivity arises from the steep angular dependence of the tabulated intensities close to the limb and the discrete structure of the NSX-H grids. To illustrate how such overshoot can affect otherwise typical observing configurations, we present two deliberately chosen test cases.

Both tests use a single circular hotspot with uniform temperature $T = 10^6$ K, representing the simplest hotspot configuration. The hotspot has an angular radius of 0.01 rad (i.e., a compact, highly localized patch), ensuring that viewing-geometry effects dominate, minimizing complex-shape-induced systematics and avoiding dilution from extended or complex emission regions. In both cases, the neutron-star mass and radius are fixed at $M = 1.4 M_{\odot}$ and $R = 12$ km. To probe rapid rotation while remaining within the OS regime, we set the spin frequency to 600 Hz. The same version of the NSX-H atmosphere tables is used throughout, and the source distance is fixed at 150 pc. The colatitude of the hotspot center is $\theta_c = 1.983$ rad for test 1 and $\theta_c = \pi - 0.01$ rad for test 2; the observer colatitude is $i = 0.01$ rad for test 1 and $i = 0.964$ rad for test 2. A complete specification of both configurations is summarized in Table 3, and their geometries are illustrated in Figure 3.

The two tests differ only in viewing geometry and hotspot location. They are constructed to create more phases when the spot approaches or recedes beyond the visible limb, driving rays with $\mu \rightarrow 0$, where unconstrained cubic interpolation tends to overshoot. Although these examples intentionally maximize the

¹ X-PSI version 3.1.

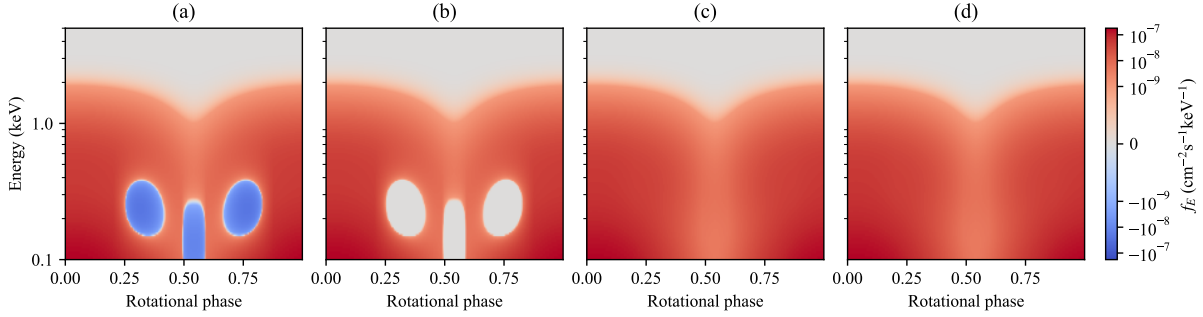


Fig. 4. Simulated phase–energy–resolved X-ray pulse profiles for interpolation test 1, computed using four interpolation schemes applied to the NSX tables: (a) cubic interpolation in all dimensions; (b) as in (a), with negative intensities clipped to zero; (c) cubic interpolation in all dimensions with linear interpolation enforced at table boundaries; (d) the baseline configuration used in this work, cubic along two axes and linear along the remaining two, with linear interpolation enforced at table boundaries.




Fig. 5. Simulated phase–energy–resolved X-ray pulse profiles for interpolation test 2, computed using four interpolation schemes applied to the NSX tables: (a) cubic interpolation in all dimensions; (b) as in (a), with negative intensities clipped to zero; (c) cubic interpolation in all dimensions with linear interpolation enforced at table boundaries; (d) the baseline configuration used in this work, cubic along two axes and linear along the remaining two, with linear interpolation enforced at table boundaries.

visibility of the effect, the underlying issue is generic: for many reasonable geometries, a rotating hotspot will periodically sample small- μ rays (as it grazes the limb or re-enters the line of sight), making careful treatment near $\mu \rightarrow 0$ essential.

Importantly, these setups are not exotic. Test 1 is consistent with a southern magnetic pole in a canonical inclined dipole or an off-centered dipole configuration. Test 2 is likewise plausible for an off-centered dipole whose south pole lies near the rotation pole. Variants such as a narrow ring hotspot at the same colatitude (e.g., arising from higher multipolar fields) can be analyzed analogously and would exhibit similar sensitivity to the limb region.

In each test case, we evaluated four strategies for interpolating the NSX table. Strategy (a) applies cubic interpolation uniformly across all four dimensions. Strategy (b) also uses cubic Lagrange interpolation but enforces a non-negativity safeguard by clipping negative values to zero. Strategy (c) employs cubic interpolation uniformly in all dimensions with linear interpolation enforced at the $\mu \rightarrow 0$ boundary. Strategy (d), our baseline in this work, uses linear interpolation in two dimensions and cubic interpolation in the remaining two, while reverting the nominally cubic dimensions to linear near the $\mu \rightarrow 0$ boundary to suppress overshoot. All results are produced with the same CPU codebase and a high-resolution configuration.

Figure 4 presents phase–energy–resolved X-ray pulse profiles for test 1. For a small viewing angle ($i = 0.01$), the profile exhibits mild asymmetry over spin phase. Under uniform cubic interpolation (panel a), we observe pronounced negative intensities (large undershoots) centered near phases ~ 0.25 , 0.50 , and 0.75 , predominantly in the soft X-ray band. These unphysical

negatives would severely corrupt the inferred pulse profile and any derived likelihood. For the X-PSI code (which we have not run ourselves), a preliminary replication (T. Salmi, priv. comm.) reports a smaller negative flux (occurring at a different location) than we find here. We have not independently verified the phase–energy locations of these occurrences and therefore do not draw conclusions about where in the profile they arise.

Introducing non-negativity clipping (panel b) removes the undershoots but creates extended zero-intensity “holes” that still end up biasing the likelihood computation by artificially suppressing flux. In contrast, uniform cubic interpolation with linear treatment at the edges (panel c) yields a qualitatively reasonable profile: by construction, linear schemes at the boundaries strongly limit overshoot and thus avoid negative intensities. Our baseline mixed scheme (panel d) likewise produces a physically consistent profile; the boundary reversion to linear in the nominally cubic dimensions effectively eliminates the negative regions, directly implicating cubic overshoot at the table edges of the μ axis as the root cause of these artifacts.

Test 2 (Figure 3) is intentionally close in observables to test 1, but it is physically distinct: the hotspot is fixed at the rotational south pole while the line of sight changes which longitudes are viewed as the star rotates. This geometry is well motivated by centered or off-centered dipolar fields heating neutron-star hotspots. The same qualitative failure modes reappear under uniform cubic interpolation, as shown in Figure 5: three phase-localized negative regions in the soft band, which become zeroed “holes” after clipping, leading to substantial underestimation of photon counts. As before, both the uniform linear scheme and

our mixed baseline recover stable, physically plausible profiles for this configuration.

Crucially, these cases are not isolated to a single demonstrated geometry. Small perturbations of the hotspot location, changes in mass–radius, or alternative hotspot shapes readily reproduce these negative regions when unconstrained cubic interpolation is used near table boundaries. Because Bayesian inference explores the full parameter space of a given model, it is effectively inevitable that such configurations will be encountered during inference. Unless interpolation is handled with care, mass–radius constraints risk hidden systematics and hence reduced credibility. We therefore recommend that existing pipelines incorporate these two test cases into their validation suites and adopt atmosphere-table interpolation strategies that control boundary overshoot (e.g., uniform linear or mixed schemes with boundary linearization, as implemented here).

After this detailed discussion of the implementation of the atmosphere model, we proceeded to production-grade tests. Specifically, we reproduced end-to-end computation cases that incorporate all necessary ingredients, such as our baseline atmosphere interpolation scheme, interstellar absorption, and convolution with the instrument response. These tests demonstrate that our treatment of the forward computation of the pulse profile is implemented correctly. With these components in place, the codebase is prepared for inference while achieving between three and four orders of magnitude speedup in pulse profile synthesis relative to existing CPU- and Python-based codebases.

5. Production-level computation result

In this section, we proceed by starting with theoretical benchmarks, which omit realistic radiative physics and other effects. We move on to production-grade tests that connect directly to observational data, which is done by incorporating a realistic neutron-star atmosphere, ISM absorption, and NICER’s instrument response. In this section, we compare the results of the CPU and GPU versions of the codebase. At previous section, we are not discussing the computation performance specifically, because the most time-consuming part of this PPM is the realistic atmosphere model table interpolation, so vast computation improvement is based on improving this atmosphere treatment strategy. Here in this section, we will compare both the modeling accuracy and the computation performance, using the benchmarks provided in Choudhury et al. (2024b). As we discuss in the previous section, we are using different strategy on surface discretizations, instead of using the discretization of only the hotspot region, we consider the whole star surface discretization as default option, so when comparing the modeling accuracy of the surface elements in the hotspot, we using the active grid number in the hotspot as the comparison. The goal of this section is to demonstrate that the extreme geometry tests of Choudhury et al. (2024b), including cases that are challenging to model accurately at standard resolutions in existing codes such as X-PSI and IM, can be computed at our standard setup with discrepancies well below the Poisson fluctuation level, while achieving 10^4 – 10^5 × speedups relative to the Ultra-resolution X-PSI runs. In this section, CPU tests were performed on an Intel i7-13700KF with single core and GPU tests on an NVIDIA GeForce RTX 4080 (16 GB).

Following Choudhury et al. (2024b), we adopt the same five “extreme” waveform tests to compare PPM treatments. Four tests use a single-temperature hot spot. All shapes are built only

from circular components, either emitting or masking, because some code in that test limited to circular primitives. Masking components simply hide the flux from the region they cover. We excluded the double-temperature test case from Choudhury et al. (2024b) because the compactness parameter was set to an extreme value, rendering the emission spot always visible through gravitational light bending and permitting multiple imaging—an effect not treated by our present code version. Complete parameter setup and the exact geometries (rings, crescents, and the bithermal configuration) are tabulated and illustrated in Choudhury et al. (2024b, Table 1 and Figure 1).

All tests share the same non-hotspot settings: distance $D = 150$ pc; observer inclination $i = 2.391$ rad measured from the rotational north pole (a viewing geometry under which the near pole remains visible across rotational phase), spin frequency at infinity $\nu = 300$ Hz, and neutral hydrogen column density $N_{\text{H}} = 2 \times 10^{19}$ cm $^{-2}$, and exposure time of 10^6 s. The inclination is chosen to mirror the viewing of a well-studied nearby millisecond pulsar PSR J0437–4715, and the spin frequency is set at 300 Hz because the OS approximation is accurate to $\leq 0.1\%$ at and below this rate (Bogdanov et al. 2019b), while discrepancies grow to a few percent by ~ 600 Hz (Cadeau et al. 2007; Jakab & Morsink 2025).

The four single-temperature tests use $M = 1.4 M_{\odot}$ and equatorial radius $R_e = 12.0$ km. Within the single-temperature set, two tests form nonconcentric rings by nesting a mask inside an emitter (with sizes varied to stress overlap logic and discretization), and two form elongated crescents by letting a mask protrude into an emitter to accentuate edge effects.

Atmosphere setup is uniform across tests. Each emitting component is modeled as a geometrically thin, fully ionized hydrogen atmosphere; specific intensities are drawn from NSX-H lookup tables (Ho & Lai 2001). The resulting spectra are then attenuated using the same N_{H} above. For any instrument-response handling and the full per-component quantities, we refer the reader to Choudhury et al. (2024b), whose setup we follow here.

5.1. Resolution tiers setup

We evaluate two resolution tiers designed for distinct purposes: (i) a standard tier that is sufficiently accurate for routine inference at practical cost and (ii) a high-resolution tier intended to match the fidelity of the “Ultra-style” reference settings in (Choudhury et al. 2024b). To isolate the effect of surface discretization, we vary only the active number of surface patches within each hotspot. All other precision controlled parameters are held fixed at high values comparable to those used in rigorous code-verification runs. This ensures that numerical errors associated with phase and energy discretization, along with per-patch ray statistics, are negligible. The remaining lever on accuracy and speed is the surface discretization.

Table 4 summarizes the two tiers detailed resolution setup. The “high-resolution” tier uses an active 256×256 surface tiling per hotspot, which is comparable to the surface discretization commonly used in “Ultra-level” resolution tier (e.g., in X-PSI). The “standard” tier uses 128×128 patches per hotspot and is intended to keep numerical systematics below typical Poisson counting uncertainties for representative NICER-like exposures, while offering a substantial speedup. Because our codebase differs in architecture and algorithms from other frameworks, a one-to-one mapping of resolution settings is not straightforward; by “comparable” we mean that our tiers achieve similar accuracy

Table 4. Numerical settings for the two resolution tiers in this work.

Resolution Tier	N_{active}	N_u	N_{c_ψ}	N_{c_α}	N_{fp}	N_{E_o}
std-resolution	128×128	256	512	512	192	128
high-resolution	256×256	256	512	512	192	128

Notes. Only the count of active surface patch N_{active} varies across tiers; all other precision parameters (the number of discretized compactness in lensing table N_u , the number of discretized $\cos \psi$ in lensing table N_{c_ψ} , the number of discretized $\cos \alpha$ in lensing table N_{c_α} , the number of fine phase N_{fp} , the number of down-sampling observed photon energy grid N_{E_o}) are held fixed.

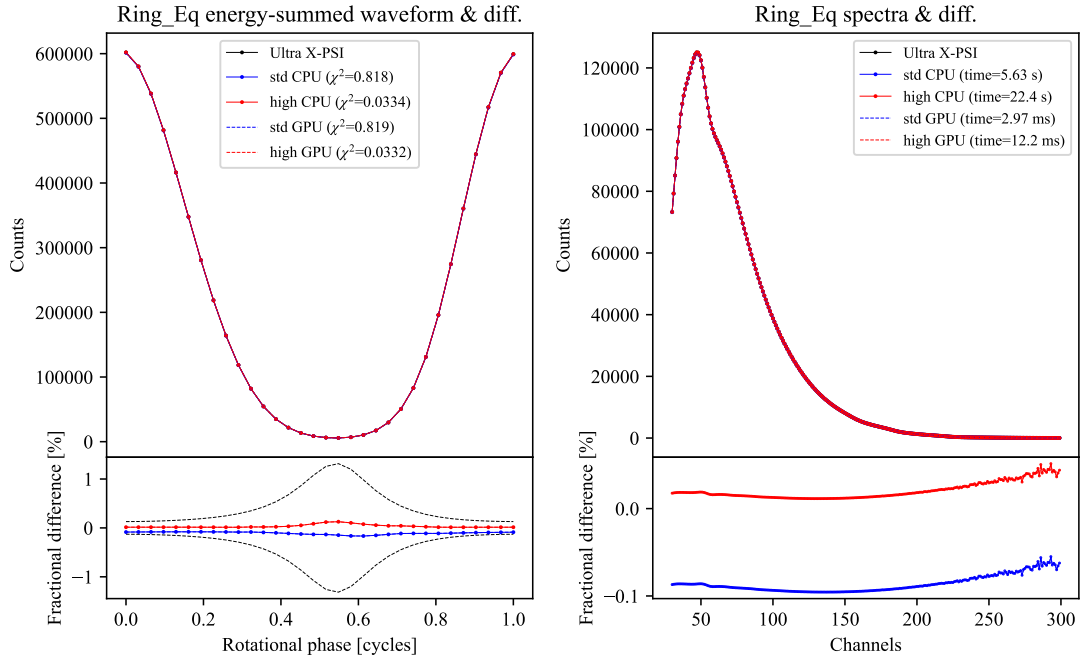


Fig. 6. Energy-summed waveforms (left) and phase-summed spectra (right) for the Ring-Eq case. The bottom subpanels show the fractional difference (in %) of each configuration relative to the X-PSI Ultra-resolution benchmark (black dotted). CPU results are plotted with solid lines: standard resolution (blue solid) and high resolution (red solid). GPU results are plotted with dashed lines: standard resolution (blue dashed) and high resolution (red dashed). In the left panel, the legend reports χ^2 as defined in Equation (57), computed jointly over energy and phase channels as a quantitative measure of the discrepancy between full energy-resolved pulse profile. In the right panel, the legend lists the runtime for each configuration. Gray dashed lines in the bottom subpanels indicate the expected Poisson fluctuations of the Ultra-resolution X-PSI benchmark.

to the production-grade benchmarks reported by Choudhury et al. (2024b).

In our test runs, both CPU and GPU implementations produce phase-resolved, energy-resolved waveforms under identical physics and numerical settings. For validation, we compare our waveforms to a trusted reference (eg. X-PSI Ultra-resolution setup in Choudhury et al. (2024b)) and require fractional residuals $\lesssim 10^{-3}$ (i.e., $\lesssim 0.1\%$) over phases where the flux is non-negligible. Small, localized deviations may occur near the photon number extremes. These phases carry very low counts and have negligible impact on likelihood evaluations used in pulse profile inference.

The reason we need to compare our CPU and GPU version code internally is because porting a ray-tracing pipeline from CPU to GPU is nontrivial. Memory layouts, parallel reduction order, and instruction-level differences can lead to small floating-point discrepancies even for the same numerical algorithm. While CPU implementations can reach either tier, end-to-end, high-resolution scans of the full parameter space are typically computationally prohibitive within reasonable computation time, whereas the GPU implementation enables production-grade inference at the high-resolution setting. The standard tier

of GPU setup is designed to be extremely fast on GPU but already become untractable on CPU even with taking advantage of modern C++ coding, while keeping numerical systematics subdominant to Poisson fluctuation for all these extreme hot-spot setup.

To demonstrate the robustness of our GPU code on the whole possible hotspot geometry parameter space, in the following subsections, we evaluate four geometries that stress the ray-tracing and integration pipeline: Ring-Eq, Ring-Polar, Crescent-Eq, and Crescent-Polar. Their precise definitions (spot shapes, locations, and masks) follow Choudhury et al. (2024b, Table 1 and Figure 1). Each case is discussed in turn using identical numerical algorithm and precision-controlled parameters so that differences reflect only hotspot geometry or resolution. To ensure full transparency and reproducibility, we have released the source code at GitHub repository². A versioned snapshot, including all scripts required to regenerate the figures for both CPU and GPU implementations at both resolution tiers, has been archived on Zenodo³.

² https://github.com/zhoutz/gpu_ppm

³ <https://doi.org/10.5281/zenodo.17315493>

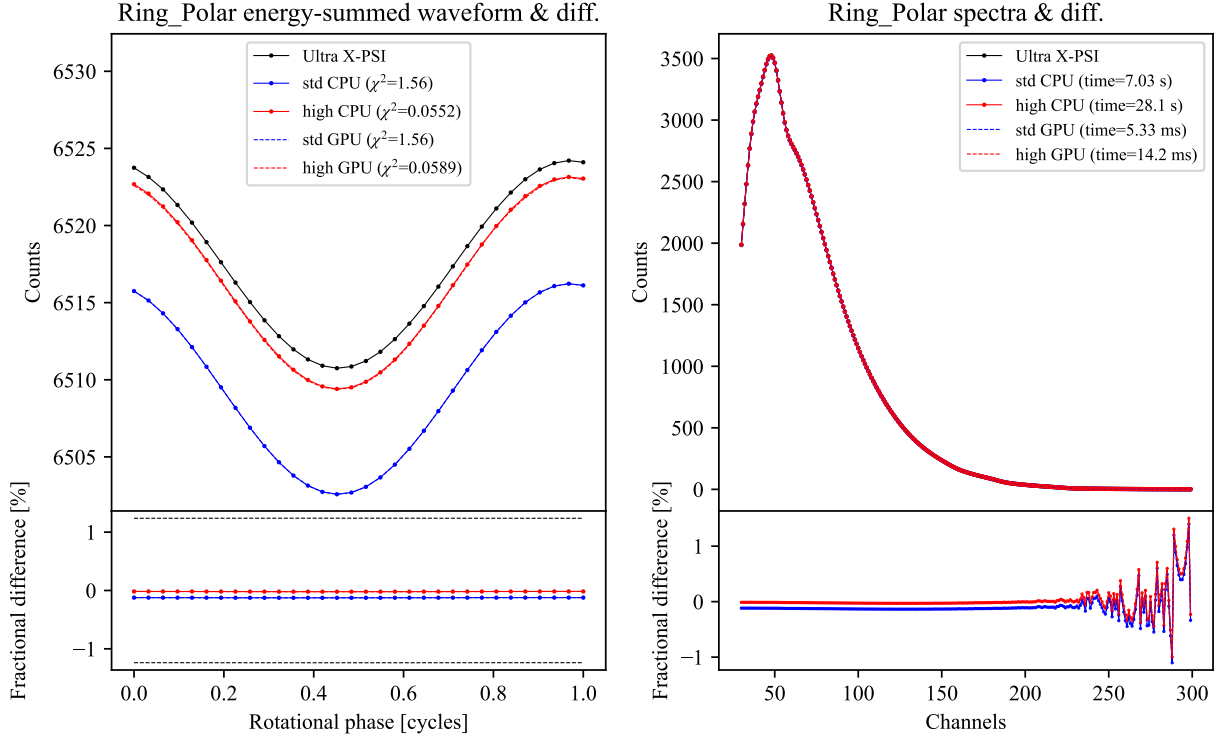


Fig. 7. Energy-summed waveforms (left) and phase-summed spectra (right) for the Ring-Polar case. The bottom subpanels show the fractional difference (in %) of each configuration relative to the X-PSI Ultra-resolution benchmark (black dotted). CPU results are plotted with solid lines: standard resolution (blue solid) and high resolution (red solid). GPU results are plotted with dashed lines: standard resolution (blue dashed) and high resolution (red dashed). In the left panel, the legend reports χ^2 as defined in Equation (57) computed jointly over energy and phase channels as a quantitative measure of the discrepancy between full energy-resolved pulse profile. In the right panel, the legend lists the runtime for each configuration. Gray dashed lines in the bottom subpanels indicate the expected Poisson fluctuations of the X-PSI Ultra-resolution benchmark.




Fig. 8. Energy-summed waveforms (left) and phase-summed spectra (right) for the Ring-Polar case. The bottom subpanels show the fractional difference (in %) of each configuration relative to the X-PSI Ultra-resolution benchmark (black dotted). CPU results are plotted with solid lines: standard resolution (blue solid) and high resolution (red solid). GPU results are plotted with dashed lines: standard resolution (blue dashed) and high resolution (red dashed). In the left panel, the legend reports χ^2 as defined in Equation (57) computed jointly over energy and phase channels as a quantitative measure of the discrepancy between full energy-resolved pulse profile. In the right panel, the legend lists the runtime for each configuration. Gray dashed lines in the bottom subpanels indicate the expected Poisson fluctuations of the X-PSI Ultra-resolution benchmark.




Fig. 9. Energy-summed waveforms (left) and phase-summed spectra (right) for the Ring-Polar case. The bottom subpanels show the fractional difference (in %) of each configuration relative to the X-PSI Ultra-resolution benchmark (black dotted). CPU results are plotted with solid lines: standard resolution (blue solid) and high resolution (red solid). GPU results are plotted with dashed lines: standard resolution (blue dashed) and high resolution (red dashed). In the left panel, the legend reports χ^2 as defined in Equation (57) computed jointly over energy and phase channels as a quantitative measure of the discrepancy between full energy-resolved pulse profile. In the right panel, the legend lists the runtime for each configuration. Gray dashed lines in the bottom subpanels indicate the expected Poisson fluctuations of the X-PSI Ultra-resolution benchmark.

5.2. Ring-Eq case

The Ring-Eq configuration is a ring-shaped hotspot centered on the stellar equator. It typically produces a single predominant pulse per rotation. Figure 6 compares energy-summed waveforms and phase-summed spectra across our two resolution tiers (standard and high) and across CPU and GPU backends, benchmarked against the X-PSI Ultra-resolution reference used in Choudhury et al. (2024b). The X-PSI Ultra-resolution configuration serves here (and throughout) as a high-fidelity standard for extreme geometries.

To quantify the agreement, we can compute a joint, 2D χ^2 statistic,

$$\chi^2 = K \sum_{i,j} \frac{(m_{ij} - d_{ij})^2}{m_{ij}}, \quad (57)$$

where m_{ij} and d_{ij} denote model and reference counts in energy bin i and phase bin j , respectively. The normalization factor K is defined as

$$K = \frac{10^6}{\sum_{i,j} m_{ij}}. \quad (58)$$

The normalization constant is included to enable fair comparisons between different waveforms. Without it, Equation (57) implies that the overall χ^2 would scale with the total photon counts. Accordingly, we normalize all computations to a fixed total of 10^6 counts, as suggested in Choudhury et al. (2024b). Under Poisson statistics and in the large-count limit, $-2 \Delta \ln L \approx$

χ^2 , so maintaining small χ^2 controls the corresponding change in log-likelihood ($\Delta \ln L \approx -\frac{1}{2}\chi^2$). Our goal is to demonstrate that both CPU and GPU implementations meet production-grade accuracy for pulse profile inference.

The energy-summed waveforms from all four runs (standard+high resolution \times CPU+GPU) agree with the benchmark to within the expected Poisson fluctuations. Differences cluster by resolution rather than by backend: increasing the surface-tiling density (from 128×128 to 256×256 patches per hotspot) uniformly reduces fractional residuals, while switching from CPU to GPU at fixed resolution has no visible effect on accuracy. Small deviations concentrate near flux minima (when the ring rotates partially out of view), where counts are lowest. These phases carry negligible statistical weight in likelihood evaluations.

The phase-summed spectra exhibit the same pattern. Curves at a given resolution overlap closely regardless of backend, and the high-resolution tier yields uniformly smaller absolute fractional differences. This behavior is consistent with surface discretization being the dominant residual error once the per-patch rays, phase sampling, and energy sampling are held fixed at high settings.

Runtime measurements are shown in the legend of Figure 6. Within our codebase, increasing the surface-patch count by a factor of four (standard \rightarrow high resolution) scales wall time by approximately the same factor (e.g., 5.63 s \rightarrow 22.4 s on CPU), as expected for a largely linear compute-to-patch workload. Even with careful modern C++ implementation, exhaustive parameter-space exploration at high resolution is computationally prohibitive on CPU, whereas the GPU implementation

achieves $\sim 2000\times$ speedups over our CPU runs at identical settings, enabling production-grade inference even under the high-resolution setup. For context, the quoted X-PSI Ultra-resolution runtime for comparable resolution is $\sim 1\text{--}5$ minutes; our GPU timings are 4–5 orders of magnitude faster while maintaining agreement at the $\chi^2 \sim 10^{-2}$ level.

5.3. Ring-Polar case

The Ring-Polar configuration is a small ring-shaped hotspot centered near the rotational pole. Because the emitting area is small, a spot at the same surface temperature produces fewer total counts than larger hotspot. Moreover, for typical neutron-star compactness (e.g., $M \approx 1.4 M_\odot$, $R_e \sim 12$ km), gravitational light bending keeps most of a polar ring at least partially in view over the rotation cycle for a wide range of observer inclinations, so the modulation amplitude is modest rather than strongly pulsed.

Figure 7 compares energy-summed waveforms and phase-summed spectra across our two resolution tiers (standard and high) and across CPU and GPU backends, using the X-PSI Ultra-resolution configuration as the high-fidelity reference. Agreement is quantified by the joint, 2D statistic χ^2 . In our runs, energy-summed waveforms from all configurations are consistent with the benchmark at the Poisson level: the standard tier yields $\chi^2 \approx 1.5$ for both CPU and GPU, while the high-resolution tier reduces this to $\chi^2 \approx 0.05$. As in the other cases, differences group by resolution rather than by backend, indicating that surface discretization, not CPU vs. GPU, is the dominant factor governing residuals once other precision settings are held fixed.

The phase-summed spectra show the same pattern. At energies where the counts are highest, curves overlap closely; in the highest-energy channels (e.g., channel indices ≥ 250 , where counts are lowest), fractional residuals can reach the $\sim 1\text{--}1.5\%$ level without changing the total χ^2 . This behavior is expected: low-count bins carry little statistical weight, so localized percent-level differences there have negligible impact on the likelihood computation.

Within a given backend, as in the Ring-Eq case, increasing the surface-patch count from 128×128 to 256×256 (a $\times 4$ increase) scales the Ring-Polar runtime by approximately the same factor, as expected for a patch-parallel workload. Absolute runtime for the Ring-Polar case is comparable to, or slightly longer than, that for the equatorial ring. While the CPU implementation meets the accuracy target, exhaustive parameter scans at either resolution tier are computationally prohibitive on CPU within practical budgets. The GPU implementation achieves $\sim 10^3\times$ speedups at identical settings, enabling production-grade Bayesian inference at very high fidelity.

5.4. Crescent-Eq case

In the Crescent-Eq configuration, the hotspot is an elongated crescent located near the stellar equator. Because the crescent spans a substantial azimuthal extent, it remains visible over a large fraction of the rotational phase and yields comparatively high counts. As noted by Choudhury et al. (2024b), this geometry can be slightly more sensitive to code implementations and resolution choices, with percent-level differences sometimes appearing at phases where the flux approaches zero; such localized discrepancies are expected when counts are minimal and carry little statistical weight.

In our experiments (Figure 8), the energy-summed waveforms agree with the X-PSI Ultra-resolution benchmark at the

$\lesssim 0.1\%$ level across phase, and increasing the surface-patch density from the standard tier (128×128 patches per hotspot) to the high-resolution tier (256×256) produces no visually discernible change. The joint 2D statistic χ^2 defined earlier likewise remains very small: at the standard tier we obtain $\chi^2 \approx 0.005$, while at the high-resolution tier we find $\chi^2 \approx 0.006$. Differences of this magnitude are consistent with floating-point and discretization effects rather than systematic modeling errors; at such small values, monotonic improvements in χ^2 with resolution are not guaranteed.

The phase-summed spectra show a similar pattern. Curves overlap closely in energy ranges with substantial counts, while the highest-energy bins can exhibit slightly larger fractional residuals without materially affecting the total χ^2 . Small backend-dependent trends (CPU vs. GPU) in these highest-energy bins are consistent with floating-point accumulation effects. At fixed resolution, results group by resolution rather than by hardware.

We attribute the robustness in this geometry, in part, to our full-surface discretization strategy: the entire stellar surface is finely meshed, and the crescent is selected by masking within that mesh. This approach represents sharp crescent edges cleanly across phase, reducing edge-alignment artifacts that can arise when only the hotspot boundary is discretized. While alternative pipelines employ different but valid tiling schemes (e.g., Bogdanov et al. 2019b, 2021), our results indicate that a full-surface mesh provides stable accuracy for elongated, edge-dominated spots.

Finally, the runtime remains highly favorable. On the GPU, the standard tier computes a single energy-resolved profile in ~ 3.4 ms, enabling production-grade scans. On the CPU, the same configuration requires ~ 7.16 s at the standard tier and ~ 27 s at the high-resolution tier. Thus, although the CPU implementation attains the desired accuracy, only the GPU implementation provides the required speedup to support large-scale Bayesian inference.

The crescent geometry is particularly relevant for future studies incorporating higher-order magnetic multipoles, which can naturally yield noncircular, elongated hotspots. For example, Huang & Chen (2025) show that, in an off-centered dipole magnetosphere, incorporating a volumetric return-current region modifies the polar-cap temperature map, adding a crescent-like, lower-temperature component. These effects are modest for a centered (aligned) dipole but become more pronounced as the dipole offset increases. Thus, testing this crescent configuration under different placements is extremely valuable for future physics-motivated, hotspot-based PPM.

5.5. Crescent-Polar case

In the Crescent-Polar configuration, we place the same elongated hotspot near the rotational north pole and adopt a near-limb viewing geometry, such that only a small portion of the emitting region is directly visible at any given phase. In this edge-dominated regime, most detected photons originate from a narrow strip along the spot boundary where projected area, light bending, and atmospheric beaming combine to produce a steep intensity gradient. Consequently, the energy-summed waveform becomes highly sensitive to numerical resolution specifically at the spot edge. Across different codes, and even within the same code at different resolution settings, this geometry can produce visibly different waveforms. Although the largest fractional differences tend to occur in low-count bins, the associated

mismatch metric, χ^2 (as defined in Equation (1) of Choudhury et al. (2024b) and here computed jointly over energy and phase channels), can still become large enough to change the likelihood by a factor $\exp(-\chi^2/2)$, potentially biasing inference if not controlled.

Figure 9 shows that our implementation closely matches the benchmark in this challenging configuration. Using both the CPU and GPU pipelines at the standard resolution tier yields $\chi^2 \approx 1.96$. In typical PPM analyses, a forward-model mismatch of order unity is negligible relative to the χ^2 variations that drive parameter updates and therefore has no visible impact on the inference. At the high-resolution tier, the discrepancy remains at $\chi^2 \approx 1.82$, indicating only marginal improvement in precision. The right panel shows that the spectral fractional differences decrease at high resolution, this is the primary source of the modest improvement from the standard to the high-resolution tier. By contrast, the left panel exhibits a persistent pattern in the fractional differences that generally stays below the Poisson fluctuation level; however, near half phase (the trough between the two peaks), the residuals reach up to $\sim 5\%$. This phase interval corresponds to when the crescent-shaped polar hotspot becomes faint and approaches the limb ($\mu \rightarrow 0$). As discussed in Section 4, cubic-interpolation overshoot in this regime, followed by the non-negativity clipping safeguard in X-PSI, leads to an underestimation of the local photon counts. While the clipping prevents the systematic from exceeding the Poisson fluctuation level, it nevertheless produces a real, phase-localized bias that is evident here and results in an overall underestimation of the counts. Comparing with Figure 7 of Choudhury et al. (2024b), the IM group’s results oscillate between under- and overestimating the photon counts relative to the X-PSI benchmark, suggesting that their atmosphere-model interpolation is sensitive to interpolation choices. In contrast, the Alberta codebase, which uses linear interpolation for the atmosphere model, is more robust at boundaries and does not suffer from overshoot-induced underestimation. Accordingly, it yields consistently higher counts at the limb. These comparisons reinforce that careful treatment of atmosphere-table interpolation is essential for reliable pulse profile computations. Consistent with our other tests, the phase-summed spectra indicate that the residual fractional differences are dominated by high-energy channels with very low counts.

From a performance standpoint, our code remains highly efficient in this case. On the CPU, the standard setup runs in ~ 1.63 s and the high-resolution setup in ~ 6.63 s. On the GPU, the corresponding run times are ~ 1.77 ms and ~ 6.9 ms, respectively, yielding speedups of order 10^3 relative to our own CPU implementation for both resolution tiers. These gains are achieved without emulators or machine-learning surrogates; we use high-fidelity forward calculations with algorithmic refinements executed on the GPU.

This case underscores an important methodological point. Because edge-on, elongated spots can exist in physically motivated scenarios (e.g., modestly off-centered dipoles can place volumetric hotspot regions near a rotational pole; see Figures 3a–b of Huang & Chen 2025), there is no compelling physical argument a priori to forbid such geometries. Reliable inference therefore requires numerical treatments that remain accurate in edge-dominated regimes. By breaking the traditional speed-versus-accuracy trade-off through a GPU implementation, our code provides one practical pathway to control χ^2 to the ~ 1 level even in this extreme configuration, while simultaneously achieving $\sim 10^3$ – 10^4 speedups.

Our CPU implementation uses IEEE 754 double-precision floating-point numbers (64-bit, FP64), whereas the current GPU implementation on an RTX 4080 employs single precision (32-bit, FP32) for all computations. Because our forward model involves no long time integration, cumulative round-off error is limited; moreover, we adopt numerically stable formulations to control FP32 error. In the current benchmarks at matched resolution, the FP32 GPU results are indistinguishable from the FP64 CPU results, indicating that FP32 is adequate for our present configuration. From an engineering perspective, enabling FP64 in our GPU codebase is straightforward. Users who run this code on scientific GPU clusters may optionally switch to FP64 to improve numerical precision.

6. Discussion and conclusion

This work presents a GPU-accelerated X-ray pulse profile modeling (PPM) codebase that directly addresses the typical accuracy-speed trade-off in current analyses. We first outline the theoretical framework of pulse profile forward modeling, rigorously verifying and deriving all formulas essential to PPM forward computation (with more details in Appendices A–C). We then provide a detailed computational recipe for both CPU and GPU implementations and clearly document the porting of the CPU code to the GPU to maximize reproducibility.

We verified our implementation by first reproducing the theoretical light-curve benchmarks of Bogdanov et al. (2019b), demonstrating accurate computation of light curves for specified spacetime and hotspot geometries under the oblate-Schwarzschild (OS) approximation. The Schwarzschild+Doppler (S+D) approximation is also supported, although OS remains our default production setting.

We further examined a commonly ignored source of systematic error that arises when moving from theoretical benchmarks to models with realistic atmospheres. Because of the current structure of the widely used NSX model, several codebases interpolated within these precomputed neutron-star atmosphere lookup tables using high-order polynomial interpolation. Using two representative viewing-geometry test cases, we demonstrate that cubic Lagrange interpolation can produce significant overshoot, yielding unphysical negative specific intensities. One approach to mitigation is to clip these negative values to zero. However, in the test cases presented here, both strategies (cubic Lagrange interpolation with or without clipping) can lead to underestimation of the total photon flux at many rotational phases, particularly in the soft X-ray band, depending on implementation. Preliminary X-PSI tests indicate smaller negative flux values (T. Salmi, priv. comm.) and the phase–energy locations of any deficits may differ across implementations.

Moreover, these new test cases are neither extreme nor exotic. They plausibly represent the emission geometry of a southern polar cap arising from either a canonical centered dipole or an off-centered dipole configuration. The region of parameter space that exhibits this overshoot behavior is potentially broad; modest variations of the pulse profile model parameters readily produce similar behavior across different settings. In fact, for many common viewing and hotspot geometries, rotation will bring portions of the hotspot into and out of the line of sight near the limb, where these interpolation artifacts are most likely to arise. The interpolation artifact identified here suggests that previously reported PPM inferences that relied on high-order polynomial interpolation without boundary-aware treatment could contain an uncontrollable systematic error. Reexamining those analyses

with a boundary-aware interpolation scheme, like the one proposed in this work, would be extremely valuable for assessing the impact on inferred parameters, especially M and R .

To mitigate interpolation artifacts, we adopted a mixed-order interpolation scheme. Within the interior of the atmospheric lookup tables, we used cubic Lagrange interpolation; but at the grazing-emission boundary, we reverted to linear interpolation to suppress overshoot. We also allow the interpolation order to differ across table dimensions, mixing linear and cubic along different axes to reduce computational cost while maintaining accuracy. The optimal configuration is table- and physics-dependent; constructing denser atmospheric grids would enable a more definitive validation of the interpolation scheme. A comprehensive evaluation is deferred to future work.

Building on these theoretical checks and the careful treatment of the atmosphere model, we turned to production-grade benchmarks, which include all of the necessary ingredients for PPM forward computation. We reproduced the suite of benchmarks in Choudhury et al. (2024b), incorporating a realistic atmospheric treatment, interstellar absorption, and the NICER instrument response. We implemented both CPU and GPU codes and examined two resolution configurations that differ primarily in the number of surface grid elements within the hotspot. Holding all other setup choices fixed allowed us to isolate how discretization influences runtime scaling while maintaining comparable physical fidelity.

The observed GPU performance indicates clear implications for inference, enabling higher-resolution forward models, broader parameter exploration, and more robust convergence diagnostics within modest computational budgets. With a modern C++ implementation, CPU runtimes at these extremely high resolutions, ranging from several seconds to tens of seconds per profile, remain prohibitive for full Bayesian analyses and robustness testing. In contrast, the GPU implementation presented here reduces per-profile computation time to the millisecond regime for the standard resolution and to ~ 10 ms for the extremely high-resolution setting on a single RTX 4080 GPU. Because fully converged posterior exploration can require 10^7 - 10^8 likelihood evaluations (depending on the problem's dimensionality), these timings make end-to-end inference feasible on a single GPU or modest GPU clusters.

Crucially, the speed improvement does not compromise accuracy. At the fine resolutions adopted here, our results closely match published benchmarks while retaining millisecond-scale runtimes. Even for geometries previously viewed as resolution-sensitive, our high-resolution GPU runs reproduce reference benchmarks, effectively removing the hard accuracy-speed bottleneck.

Beyond the common single- or two-temperature prescriptions used in many current inferences, our framework accepts arbitrary, user-specified global temperature maps on the stellar surface. This capability enables inference based on physics-motivated hotspots, such as those predicted by off-centered dipole or more general multipolar magnetic configurations (e.g., Huang & Chen 2025; Gralla et al. 2017; Lockhart et al. 2019), and naturally extends to accretion-powered millisecond pulsars, where global temperature distributions are more appropriate than bounded caps (Das et al. 2022; Das et al. 2025). Because purely geometric hotspot choices can introduce degeneracies and ad hoc modeling systematics, we advocate reducing arbitrary

geometry in favor of models tied to the magnetospheric structure of the neutron star (e.g., controlled multipolar expansions) and potentially enabling multiwavelength constraints. Related efforts (e.g., Kalapotharakos et al. 2021; Olmschenk et al. 2025; Pétri et al. 2025) typically assume uniform-temperature caps; observational and theoretical indications suggest, however, that polar-cap regions are structured, with nonuniform temperatures and subregions dominating the X-ray emission (Spitkovsky 2006; Bai & Spitkovsky 2010; Gralla et al. 2017; Lockhart et al. 2019). Our code is designed to accommodate such more complex hotspot models.

While the present work emphasizes fast and accurate forward modeling, the inference layer is a natural next step. Our GPU optimization extends beyond waveform synthesis to the entire likelihood evaluation: phase- and energy-resolved model counts, background treatment, and Poisson log-likelihood accumulation run on GPU and parallelize naturally across energies, phase bins on different parameters. We support two complementary paths for sampling: (i) a thin Python/C++ interface that interoperates with established CPU-side samplers; for instance, MULTINEST (Feroz et al. 2009), ULTRANEST (Buchner 2021), and DYNesty (Speagle 2020). We also kept the likelihood kernel resident on the GPU to minimize host-device overhead; and (ii) a fully GPU-resident parallel-tempered MCMC (PT-MCMC); and, resources permitting, a GPU-friendly nested-sampling engine) to handle multimodal posteriors and strong parameter correlations. Both modes can scale to multiple GPUs by sharding chains or live points across GPUs and batching likelihood calls within each device. Using this architecture together with physics-motivated temperature maps, we will reanalyze NICER sources, systematically varying samplers and settings to assess robustness and to deliver stable, reproducible constraints on neutron-star masses and radii.

Before proceeding to inference, we note an additional physical effect: multiple imaging due to strong-field gravitational light bending. For sufficiently compact stars, approximately when $GM/(Rc^2) \gtrsim 0.28$, photon trajectories from a given surface element can reach the observer along multiple paths, yielding secondary images. Our current framework is not yet equipped to model these higher-order images and, thus, implementing this capability for production-grade inference is deferred to future work. For context, current NICER constraints for PSR J0740+6620 imply a central-value compactness likely below this threshold, although portions of the allowed parameter space may reach it; in those more compact configurations, this multiple-imaging effect could become relevant.

In summary, the codebase presented here provides a modern computational framework for PPM: it achieves production-grade accuracy at millisecond-scale runtimes. Our results are aligned with the majority of previous community tests and we propose two new atmosphere-interpolation tests. Our work substantially breaks the accuracy-speed bottleneck of Bayesian PPM studies. In parallel with advancing reanalyses of existing observational sources with physics-motivated hotspot models, we invite the community to use our GPU PPM codebase to reexamine the robustness of previous studies under diverse inference setups. For transparency and pedagogical reasons, we also provide a CPU reference implementation with a clear structure and high readability. Taken together, these efforts advance pulse profile

modeling toward a more efficient, reliable, and reproducible regime.

Data availability

GPU-PPM Github repository: https://github.com/zhoutz/gpu_ppm. Zenodo repository of reproduction file of this work: <https://doi.org/10.5281/zenodo.17315493>.

Acknowledgements. H.C. acknowledges support from NSF Grant AST-2308111. The authors extend gratitude to Sharon Morsink, Tuomo Salmi for providing valuable Benchmark, help on clarifying the theoretical details about PPM, Alexander Chen about the GPU implementation suggestions and H.C. would like to thank Anna Watts, Alexander Dittmann, Cole Miller, Roger Romani, Wenbin Lu, Roger Blanford for insightful discussion and comments.

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *Phys. Rev. Lett.*, **119**, 161101
- Abbott, B., Abbott, R., Abbott, T., et al. 2018, *Phys. Rev. Lett.*, **121**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *Phys. Rev. X*, **9**, 011001
- AlGendy, M., & Morsink, S. M. 2014, *ApJ*, **791**, 78
- Annala, E., Gorda, T., Katerini, E., et al. 2022, *Phys. Rev. X*, **12**, 011058
- Bai, X.-N., & Spitkovsky, A. 2010, *ApJ*, **715**, 1282
- Bauböck, M., Psaltis, D., & Özel, F. 2019, *ApJ*, **872**, 162
- Beloborodov, A. M. 2002, *ApJ*, **566**, L85
- Bilous, A. V., Watts, A. L., Harding, A. K., et al. 2019, *ApJ*, **887**, L23
- Bogdanov, S., Grindlay, J. E., & Rybicki, G. B. 2008, *ApJ*, **689**, 407
- Bogdanov, S., Guillot, S., Ray, P. S., et al. 2019a, *ApJ*, **887**, L25
- Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019b, *ApJ*, **887**, L26
- Bogdanov, S., Dittmann, A. J., Ho, W. C. G., et al. 2021, *ApJ*, **914**, L15
- Bootsma, E., Vinciguerra, S., Watts, A. L., Kini, Y., & Salmi, T. 2025, *MNRAS*, **537**, 3769–3780
- Buchner, J. 2021, arXiv e-prints [arXiv:2101.09604]
- Buchner, J. 2021, *JOSS*, **6**, 3001
- Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, *ApJ*, **654**, 458
- Cartaxo, J., Huang, C., Malik, T., et al. 2026, *ApJS*, **282**, 33
- Chen, A. Y., Yuan, Y., & Vasilopoulos, G. 2020, *ApJ*, **893**, L38
- Choudhury, D., Salmi, T., Vinciguerra, S., et al. 2024a, *ApJ*, **971**, L20
- Choudhury, D., Watts, A. L., Dittmann, A. J., et al. 2024b, *ApJ*, **975**, 202
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *Nat. Astron.*, **4**, 72
- Das, P., Porth, O., & Watts, A. L. 2022, *MNRAS*, **515**, 3144
- Das, P., Salmi, T., Davelaar, J., Porth, O., & Watts, A. 2025, *ApJ*, **987**, 34
- De, S., Finstad, D., Lattimer, J. M., et al. 2018, *Phys. Rev. Lett.*, **121**, 091102
- Dittmann, A. J., Miller, M. C., Lamb, F. K., et al. 2024, *ApJ*, **974**, 295
- Essick, R., Landry, P., & Holz, D. E. 2020, *Phys. Rev. D*, **101**, 063007
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, *ApJ*, **915**, L12
- Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *SPIE Conf. Ser.*, **9905**, 99051H
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Gralla, S. E., Lupsasca, A., & Philippov, A. 2017, *ApJ*, **851**, 137
- Greif, S. K., Raaijmakers, G., Hebel, K., Schwenk, A., & Watts, A. L. 2019, *MNRAS*, **485**, 5363
- Guillot, S., Kerr, M., Ray, P. S., et al. 2019, *ApJ*, **887**, L27
- Güver, T., Psaltis, D., Özel, F., & Zhao, T. 2026, *ApJ*, **997**, 266
- Heyl, J. S., & Shaviv, N. J. 2000, *MNRAS*, **311**, 555
- Ho, W. C. G., & Lai, D. 2001, *MNRAS*, **327**, 1081
- Ho, W. C. G., & Lai, D. 2003, *MNRAS*, **338**, 233
- Huang, C. 2025a, *ApJ*, **978**, L14
- Huang, C. 2025b, *ApJ*, **985**, 216
- Huang, C., & Chen, A. Y. 2025, *ApJ*, **991**, 90
- Huang, C., & Sourav, S. 2025, *ApJ*, **983**, 17
- Huang, C., Raaijmakers, G., Watts, A. L., Tolos, L., & Providência, C. 2024a, *MNRAS*, **529**, 4650
- Huang, C., Tolos, L., Providência, C., & Watts, A. 2024b, *MNRAS*, **536**, 3262
- Huang, C., Malik, T., Cartaxo, J., et al. 2024c, arXiv e-print [arXiv:2411.14615]
- Huth, S., Pang, P. T. H., Tews, I., et al. 2022, *Nature*, **606**, 276
- Jakab, Z., & Morsink, S. 2025, *ApJ*, **994**, 163
- Kalopotharakos, C., Wadiasingh, Z., Harding, A. K., & Kazanas, D. 2021, *ApJ*, **907**, 63
- Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009a, *ApJ*, **706**, 417
- Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009b, *ApJ*, **705**, L36
- Landry, P., & Essick, R. 2019, *Phys. Rev. D*, **99**, 084049
- Legred, I., Chatziioannou, K., Essick, R., Han, S., & Landry, P. 2021, *Phys. Rev. D*, **104**, 063003
- Li, A., Watts, A. L., Zhang, G., et al. 2025, *Sci. China Phys., Mech. Astron.*, **68**, 119503
- Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, *ApJ*, **776**, 19
- Lockhart, W., Gralla, S. E., Özel, F., & Psaltis, D. 2019, *MNRAS*, **490**, 1774
- Loktev, V., Salmi, T., Näätä, J., & Poutanen, J. 2020, *A&A*, **643**, A84
- Lommen, A. N., Kipporn, R. A., Nice, D. J., et al. 2006, *ApJ*, **642**, 1012
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, **129**, 1993
- Markwardt, C., Arzoumanian, Z., Gendreau, K., Hare, J., & Nicer Team. 2024, in *AAS/High Energy Astrophysics Division*, **21**, 105.36
- Mauviard, L., Guillot, S., Salmi, T., et al. 2025, *ApJ*, **995**, 60
- Miller, M. C., & Lamb, F. K. 1998, *ApJ*, **499**, L37
- Miller, M. C., & Lamb, F. K. 2015, *ApJ*, **808**, 31
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, *ApJ*, **887**, L24
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2021, *ApJ*, **918**, L28
- Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, *ApJ*, **663**, 1244
- Näätä, J., & Pihajoki, P. 2018, *A&A*, **615**, A50
- Olmschenk, G., Broadbent, E., Kalopotharakos, C., et al. 2025, *ApJ*, **991**, 169
- Pechenick, K. R., Ftaclas, C., & Cohen, J. M. 1983, *ApJ*, **274**, 846
- Pétri, J., Guillot, S., Guillemot, L., et al. 2025, *A&A*, **701**, A39
- Potekhin, A. Y., Pons, J. A., & Page, D. 2015, *Space Sci. Rev.*, **191**, 239
- Poutanen, J. 2020, *A&A*, **640**, A24
- Poutanen, J., & Beloborodov, A. M. 2006, *MNRAS*, **373**, 836
- Poutanen, J., & Gierliński, M. 2003, *MNRAS*, **343**, 1301
- Prigozhin, G., Gendreau, K., Doty, J. P., et al. 2016, *SPIE Conf. Ser.*, **9905**, 990511
- Psaltis, D., & Özel, F. 2014, *ApJ*, **792**, 87
- Raaijmakers, G., Riley, T. E., Watts, A. L., et al. 2019, *ApJ*, **887**, L22
- Raaijmakers, G., Greif, S. K., Riley, T. E., et al. 2020, *ApJ*, **893**, L21
- Raaijmakers, G., Greif, S. K., Hebel, K., et al. 2021a, *ApJ*, **918**, L29
- Raaijmakers, G., Greif, S. K., Hebel, K., Hinderer, T., et al. 2021b, *ApJ*, **918**, L29
- Reardon, D. J., Hobbs, G., Coles, W., et al. 2016, *MNRAS*, **455**, 1751
- Reardon, D. J., Bailes, M., Shannon, R. M., et al. 2024, *ApJ*, **971**, L18
- Remillard, R. A., Loewenstein, M., Steiner, J. F., et al. 2022, *AJ*, **163**, 130
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *ApJ*, **887**, L21
- Riley, T. E., Watts, A. L., Ray, P. S., et al. 2021, *ApJ*, **918**, L27
- Riley, T. E., Choudhury, D., Salmi, T., et al. 2023a, *JOSS*, **8**, 4977
- Riley, T. E., Choudhury, D., Salmi, T., et al. 2023b, *JOSS*, **8**, 4977
- Rutherford, N., Mendes, M., Svensson, S., et al. 2024, *ApJ*, **971**, L19
- Salmi, T., Näätä, J., & Poutanen, J. 2018, *A&A*, **618**, A161
- Salmi, T., Suleimanov, V. F., Näätä, J., & Poutanen, J. 2020, *A&A*, **641**, A15
- Salmi, T., Vinciguerra, S., Choudhury, D., et al. 2023, *ApJ*, **956**, 138
- Salmi, T., Choudhury, D., Kini, Y., et al. 2024a, *ApJ*, **974**, 294
- Salmi, T., Deneva, J. S., Ray, P. S., et al. 2024b, *ApJ*, **976**, 58
- Salmi, T., Dorsman, B., Watts, A. L., et al. 2025, *MNRAS*, **538**, 2562
- Silva, H. O., Pappas, G., Yunes, N., & Yagi, K. 2021, *Phys. Rev. D*, **103**, 063038
- Speagle, J. S. 2020, *MNRAS*, **493**, 3132
- Spitkovsky, A. 2006, *ApJ*, **648**, L51
- Suleimanov, V., Poutanen, J., & Werner, K. 2012, *A&A*, **545**, A120
- Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *ApJ*, **465**, 487
- Viironen, K., & Poutanen, J. 2004, *A&A*, **426**, 985
- Vinciguerra, S., Salmi, T., Watts, A. L., et al. 2024, *ApJ*, **961**, 62
- Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, *Rev. Mod. Phys.*, **88**, 021001
- Watts, A. L., Yu, W., Poutanen, J., et al. 2019, *Sci. China Phys. Mech. Astron.*, **62**, 29503
- Wilms, J., Allen, A., & McCray, R. 2000, *ApJ*, **542**, 914
- Zavlin, V. E., Pavlov, G. G., & Shibano, Y. A. 1996, *A&A*, **315**, 141
- Zhang, N.-B., & Li, B.-A. 2021, *ApJ*, **921**, 111
- Zhang, S.-N., Santangelo, A., Xu, Y., et al. 2025, *Sci. China Phys. Mech. Astron.*, **68**, 119502
- Zhao, T., Psaltis, D., Özel, F., & Beklen, E. 2024, *ApJ*, submitted [arXiv:2412.12283]
- Zhou, T., & Huang, C. 2025, *ApJ*, submitted [arXiv:2504.08662]

Appendix A: Relativistic Doppler effect

To account for the relativistic Doppler effect, we consider two inertial reference frames: the lab frame, in which the stellar surface element is in motion, and the comoving frame, in which the element is at rest. Throughout this section, we set light speed $c = 1$. A prime (') denotes quantities measured in the comoving frame. The goal is to relate (i) the angle between the photon direction and the surface normal, and (ii) the solid-angle element, as seen in the two frames. These relations are what enter the specific-intensity transformations.

Suppose that, in the lab frame, a small emitting surface patch translates along the $+\hat{y}$ direction with speed v , while its outward surface normal is aligned with \hat{z} . A photon leaves the patch in direction \hat{k} . We project \hat{k} onto the x - z plane. Define k_{\perp} in x - z plane,

$$k_{\perp} = \hat{k} - \hat{k} \cdot \hat{y}, \quad \hat{k}_{\perp} = \frac{k_{\perp}}{|k_{\perp}|}. \quad (\text{A.1})$$

We then parametrize the direction by a polar angle ξ measured from $+\hat{y}$ direction and an azimuth ϕ measured in the x - z plane. With this choice,

$$\hat{k} \cdot \hat{y} = \cos \xi, \quad \hat{k}_{\perp} \cdot \hat{x} = \cos \phi, \quad \hat{k}_{\perp} \cdot \hat{z} = \sin \phi, \quad (\text{A.2})$$

so that the unit propagation vector has the component form

$$\hat{k} = (\sin \xi \cos \phi, \cos \xi, \sin \xi \sin \phi). \quad (\text{A.3})$$

Along the null ray, the spatial coordinates advance at light speed in the direction \hat{k} , hence

$$x = t \sin \xi \cos \phi, \quad y = t \cos \xi, \quad z = t \sin \xi \sin \phi, \quad (\text{A.4})$$

We now boost to the comoving frame with a standard Lorentz transformation along $+\hat{y}$:

$$t' = \gamma(t - vy), \quad y' = \gamma(y - vt), \quad x' = x, \quad z' = z, \quad (\text{A.5})$$

where the Lorentz factor is

$$\gamma = \frac{1}{\sqrt{1 - v^2}}. \quad (\text{A.6})$$

Next, we introduce the angle σ between the photon direction and the surface normal \hat{z} . This is the local emission angle relevant for atmosphere beaming, in our geometry,

$$\cos \sigma = \frac{z}{\sqrt{x^2 + y^2 + z^2}} = \sin \xi \sin \phi, \quad (\text{A.7})$$

and the same definition applied in the comoving frame gives σ' . Using the boosted coordinates and the null condition, we obtain

$$\begin{aligned} \cos \sigma' &= \frac{z'}{\sqrt{x'^2 + y'^2 + z'^2}} \\ &= \frac{\sin \xi \sin \phi}{\sqrt{(\sin \xi \cos \phi)^2 + [\gamma(y/t - v)]^2 + (\sin \xi \sin \phi)^2}} \\ &= \frac{\sin \xi \sin \phi}{\cos \sigma} \\ &= \frac{\sqrt{\sin^2 \xi + \gamma^2 (\cos \xi - v)^2}}{\cos \sigma} \\ &= \frac{\gamma \sqrt{(1 - v^2) \sin^2 \xi + \cos^2 \xi - 2v \cos \xi + v^2}}{\cos \sigma} \\ &= \frac{\gamma \sqrt{1 - 2v \cos \xi + v^2 \cos^2 \xi}}{\cos \sigma} \\ &= \frac{\gamma \sqrt{1 - v \cos \xi}}{\cos \sigma} \\ &= \delta \cos \sigma, \end{aligned} \quad (\text{A.8})$$

where we have identified the usual Doppler factor

$$\delta = \frac{1}{\gamma(1 - v \cos \xi)}. \quad (\text{A.9})$$

This expression shows explicitly how the emission angle with respect to the surface normal is modulated by the boost.

We now turn to the transformation of the solid angle. With our spherical coordinates we define

$$d\Omega = d \cos \xi d\phi, \quad d\Omega' = d \cos \xi' d\phi', \quad (\text{A.10})$$

Because the Lorentz boost is along y and leaves x and z unchanged, the azimuth in the x - z plane is unaffected:

$$\phi' = \arctan\left(\frac{z'}{x'}\right) = \arctan\left(\frac{z}{x}\right) = \phi, \quad (\text{A.11})$$

so that

$$d\phi = d\phi'. \quad (\text{A.12})$$

The polar cosine with respect to the boost axis is, by definition,

$$\cos \xi = \frac{y}{\sqrt{x^2 + y^2 + z^2}}, \quad (\text{A.13})$$

and its correspondence in comoving frame is

$$\cos \xi' = \frac{y'}{\sqrt{x'^2 + y'^2 + z'^2}} = \frac{\gamma(y/t - v)}{\gamma(1 - v \cos \xi)} = \frac{\cos \xi - v}{1 - v \cos \xi}. \quad (\text{A.14})$$

This is the standard formula for angles measured relative to the boost direction. Differentiating with respect to $\cos \xi'$ gives the Jacobian of the transformation:

$$\begin{aligned} d \cos \xi' &= \frac{d \cos \xi (1 - v \cos \xi) - (\cos \xi - v) (-v d \cos \xi)}{(1 - v \cos \xi)^2} \\ &= d \cos \xi \frac{1 - v^2}{(1 - v \cos \xi)^2} \\ &= \delta^2 d \cos \xi. \end{aligned} \quad (\text{A.15})$$

which immediately implies the solid-angle relation

$$d\Omega' = \delta^2 d\Omega. \quad (\text{A.16})$$

Appendix B: Time dilation approximation

In this section, we derive the mapping between emission time at a rotating stellar surface element and the reception time at a distant detector. The goal is to obtain a compact differential relation $dt_{\text{emit}} \leftrightarrow dt_{\text{receive}}$

Suppose a nongravitating spherical star which center locate at $(x, y, z) = (0, 0, 0)$ with mass M and equatorial radius R_e , a distant observer is located in the $+x$ direction. The location of a small beaming patch on the star surface while rotating we denote as $\mathbf{r}(t_{\text{emit}})$ (all t are measured in observer's frame). A photon emit from \mathbf{r} , that finally hit the plane $x = D$, where D is a large number represent the star-observer distance, will undergo 2 stages.

The 1st stage is, from the photon been emitted, to it arrive the plane $x = R$, we use Δt_1 to denote the the time elapsed in stage 1,

$$t_{\text{hitR}} = t_{\text{emit}} + \Delta t_1. \quad (\text{B.1})$$

Here, $x = R$ is a fiducial plane placed beyond the stellar surface along $+\hat{x}$.

The 2nd stage begins when photon hit $x = R$ and ends when photon hit plane $x = D$,

$$t_{\text{receive}} = t_{\text{hitR}} + \Delta t_2, \quad \Delta t_2 = \frac{D - R}{c}. \quad (\text{B.2})$$

Because the ray is taken to propagate parallel to $+\hat{x}$ once it leaves the near zone, Δt_2 is a constant independent of emission phase. This cleanly separates phase-dependent (Stage 1) from phase-independent (Stage 2) contributions.

In the 1st stage, we consider 2 consecutive emitted photon from the same patch, photon A and photon B, with emitting time relation

$$t_{\text{emit}}^B = t_{\text{emit}}^A + dt_{\text{emit}}, \quad (\text{B.3})$$

since the star is rotating, these 2 photon are not emit at the same location, so we have $\Delta t_1^A \neq \Delta t_1^B$. Since we are treating the non-gravitating case, we ignore the gravitational light bending effect, we assume the photon go straight line towards $+x$ direction. From geometry we can derive

$$\Delta t_1 = \frac{R - \hat{x} \cdot \mathbf{r}(t_{\text{emit}})}{c}, \quad (\text{B.4})$$

so

$$\frac{d(\Delta t_1)}{dt_{\text{emit}}} = \frac{-\hat{x} \cdot \dot{\mathbf{r}}(t_{\text{emit}})}{c} = -\beta \cos \xi, \quad (\text{B.5})$$

where $\cos \xi = \hat{x} \cdot \hat{\mathbf{v}}$ is the cosine of the angle between patch velocity and the photon emit direction, and

$$\beta = \frac{v}{c} = \frac{|\dot{\mathbf{r}}|}{c}. \quad (\text{B.6})$$

In stage 1, we get

$$\frac{dt_{\text{hitR}}}{dt_{\text{emit}}} = 1 + \frac{d(\Delta t_1)}{dt_{\text{emit}}} = 1 - \beta \cos \xi, \quad (\text{B.7})$$

which expresses the local time-compression/expansion between emission and crossing of the $x = R$ plane.

In stage 2, Δt_2 is constant, so

$$\frac{dt_{\text{receive}}}{dt_{\text{emit}}} = \frac{dt_{\text{hitR}}}{dt_{\text{emit}}} = 1 - \beta \cos \xi, \quad (\text{B.8})$$

and the total mapping from emission to reception, in the non-gravitating straight-line limit, is purely kinematic. When the surface moves toward the observer ($\cos \xi > 0$) the received cadence is compressed, when it recedes, it is dilated.

When considering a Schwarzschild gravitating star, the surface of star and the observer have different gravitational potential, so we have to include the gravitational redshift effect, which is

$$\frac{dt_{\text{receive}}}{dt_{\text{emit}}} = \frac{1}{\sqrt{1 - R_S/R}}, \quad (\text{B.9})$$

where R_S is the Schwarzschild radius associated with M and R .

Combine both, we have

$$dt_{\text{emit}} = \frac{\sqrt{1 - R_S/R}}{1 - \beta \cos \xi} dt_{\text{receive}}, \quad (\text{B.10})$$

which is the working relation used in the main text. The numerator represents gravitational time dilation, and the denominator represents kinematic compression/expansion from motion along the line of sight. In the limiting cases $\beta \rightarrow 0$ or $R_S/R \rightarrow 0$, Equation (B.10) correctly reduces to purely gravitational or purely kinematic effects, respectively. If both vanish, $dt_{\text{emit}} = dt_{\text{receive}}$ as expected.

Appendix C: Blackbody radiation and flux formula

We begin with the Planck's law of blackbody radiation,

$$B(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1}, \quad (\text{C.1})$$

where $B(\nu, T)$ is the spectral radiance of a blackbody in thermal equilibrium at temperature T , with unit power per (unit solid angle \times unit area normal to the propagation \times unit frequency). Express it in more precise form,

$$\frac{EdN}{dt d\Omega \cos \sigma dS d\nu} = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1}, \quad (\text{C.2})$$

where E is the photon energy, dN the photon count, $d\Omega$ the differential solid angle around the propagation direction, $\cos \sigma dS$ the projected surface area element, and $d\nu$ the frequency bin. Writing Planck's law in this fully differential form will let us transform variables and frames cleanly.

The advantage of using this form is that we can easily change variables, suppose we want to change ν to E with

$$E = h\nu, \quad dE = h d\nu, \quad (\text{C.3})$$

we directly plug them in and get

$$\frac{EdN}{dt d\Omega \cos \sigma dS dE} = \frac{2E^3}{h^3 c^2} \frac{1}{e^{E/kT} - 1}. \quad (\text{C.4})$$

This recasts the blackbody intensity in per-energy units, which is convenient for applications on detector responses that are tabulated versus E .

We use prime(') to denote quantities measured in the comoving frame of the emitting patch. For a small patch on the star surface emitting blackbody radiation, we rewrite the above equation in comoving frame,

$$\frac{E' dN'}{dt' d\Omega' \cos \sigma' dS' dE'} = I_{bb} \equiv \frac{2E'^3}{h^3 c^2} \frac{1}{e^{E'/kT'} - 1}. \quad (\text{C.5})$$

Here I_{bb} is the comoving specific intensity of a blackbody. The primes emphasize that both thermodynamic variables (T') and emission angles (σ') are defined in the instantaneous rest frame of the surface element. Our task below is to relate these comoving frame differentials to the ones measured by a distant observer.

We now calculate each term in the observer's frame. According to Doppler effect and gravitational redshift,

$$E = \delta \sqrt{1 - u} E'. \quad (\text{C.6})$$

The observed photon energy is boosted by the special-relativistic Doppler factor δ and reduced by the gravitational redshift factor $\sqrt{1 - u}$, with $u \equiv R_S/R$.

According to number conservation, each emit photon should arrive the observer,

$$dN = dN'. \quad (\text{C.7})$$

According to Equation (B.10),

$$dt' = \frac{\sqrt{1 - u}}{1 - \beta \cos \xi} dt = \sqrt{1 - u} \delta y dt. \quad (\text{C.8})$$

The comoving time element dt' relates to the observer frame time interval dt by the same factors derived in Appendix B: gravitational time dilation (numerator) and kinematic arrival-time compression/expansion (denominator).

According to Equation (A.16),

$$\begin{aligned} d\Omega' &= \delta^2 d\Omega = \delta^2 d \cos \alpha d\lambda \\ &= \delta^2 \frac{d \cos \alpha}{d \cos \psi} d \cos \psi d\lambda = \delta^2 \frac{d \cos \alpha}{d \cos \psi} \frac{dA}{D^2}, \end{aligned} \quad (\text{C.9})$$

where λ is the azimuthal angle around the outward radial vector, dA is the differential area of the detector, and D is the distance from star to the observer.

The first equality is the standard formula $d\Omega' = \delta^2 d\Omega$ from Appendix A. We then (i) resolve $d\Omega$ into polar and azimuthal differentials around radial vector, (ii) introduce the gravitational light-bending Jacobian $d \cos \alpha / d \cos \psi$ that maps the apparent angle ψ at infinity to local emission angle α of spherical surface, and (iii) relate $d \cos \psi d\lambda$ to detector's collecting area dA via $d \cos \psi d\lambda = dA / D^2$.

Combine all these together, we have

$$\frac{EdN}{dt dAdE} = \sqrt{1-u} I_{bb} \delta^3 \gamma \cos \sigma' \frac{d \cos \alpha}{d \cos \psi} \frac{dS'}{D^2}. \quad (\text{C.10})$$

This is the energy flux per unit detector area and per unit energy. The prefactors have transparent origins: δ (from E), δ^2 (from $d\Omega'$), and γ (from the dt' mapping) combine into $\delta^3 \gamma$, $\sqrt{1-u}$ accounts for gravitational redshift, the cosine factor uses the comoving local emission angle σ' , and the Jacobian encodes Schwarzschild gravitational lensing. The dS' / D^2 factor represent the differential area of the emitting patch and the stellar distance information.

In practice, we may substitute I_{bb} with a more general intensity function $I(E', \sigma')$ to account for non-blackbody and anisotropic atmosphere beaming, and we want the number flux rather than energy flux, so we divide both side by E and get

$$\frac{dN}{dt dAdE} = \sqrt{1-u} \delta^3 \gamma \cos \sigma' \frac{d \cos \alpha}{d \cos \psi} \frac{I(E', \sigma')}{E} \frac{dS'}{D^2}. \quad (\text{C.11})$$