Evidence for a brief appearance of gamma-ray periodicity after a compact star merger

Reporter: Shuai Xie (谢帅)

2025.10.10 @SWIFAR, YNU

https://doi.org/10.1038/s41550-025-02649-w

Introduction

- Identifying whether the central engine of gamma-ray bursts (GRBs) is a black hole or a neutron star frequently lacks clear direct evidence (Berger 2014; Zhang 2018). Although the gravitational wave event GW170817 confirmed that the pre-merger binary system was composed of two neutron stars (Abbott et al. 2017), the characteristics of the resulting merger remnant are still undetermined.
- The recent identification of kilohertz quasi-periodic oscillations (QPOs) in two short GRBs (Chirenti et al. 2023), demonstrated through broad power spectra features, has renewed optimism in this field.
- GRB 230307A associated with a kilonova (Levan et al. 2024; Yang et al. 2024), this long-duration GRB strongly supports the compact star merger origin with a millisecond magnetar as its probable central engine (Sun et al. 2025). As the second brightest GRB ever detected, the observed minute-long duration and substantial photon counts provide an ideal case for investigating (quasi-)periodicity.

Data

- The observational data of GRB 230307A analysed in this study were obtained from the data archives of **GECAM-B** (Gravitational Wave High-energy Electromagnetic Counterpart All-sky Monitor satellite B, Li et al. 2021) and **GECAM-C** (Zhang et al. 2023) and from the GRB monitor onboard Fermi (**Fermi/GBM**, Meegan et al. 2009).
- A blind search for periodicity was conducted using event data from the three GECAM-B detectors with the smallest incident angles to the source. GECAM-B is prioritized because some of its detectors have **incident angles of less than 40°**. By contrast, GECAM-C and Fermi/GBM primarily have detectors with larger incident angles, typically greater than 50°. Furthermore, the Fermi/GBM **data experienced saturation issues** during certain time intervals(Dalessi, S. & Fermi GBM Team 2023), and the initial phase of the GECAM-C data was impacted by **high-energy particle interference**(Sun et al. 2025), rendering these datasets unsuitable for a comprehensive blind search across the entire prompt emission phase.

Blind Search

- we first applied a moving time window of 100 ms, which was shifted by 7 ms at each step, corresponding to the minimum variability timescale (MVT), to the GECAM-B event data of GRB 230307A. This resulted in the extrac tion of **7,287 valid 100-ms segments** (each with a signal-to-noise ratio, $\boldsymbol{SNR \geq 10}$). Each segment had count rates exceeding 4,000 photons per second and covered the energy range 22–10,053 keV.
- To account for energy-dependent periodicity, each segment was further divided into subsets based on energy channels, resulting in **629,314 subsets of event data**. These subsets formed our final sample for the subsequent analysis.
- Using each subset of event data, we computed the Rayleigh power across a trial frequency range of 500–2,500 Hz with a resolution of 1 Hz. Each subset yielded a candidate, defined by the maximum Rayleigh power obtained from 2,001 trial frequencies. The trial frequency corresponding to this maximum power was designated as the candidate frequency, whereas the maximum power itself was labelled as the candidate power for that subset.

- By conducting an identical blind search on the background noise data, a "false alarm threshold" was calculated to be $\boldsymbol{R_{false} = 15.81}$. Any candidate signal with a power below this threshold was considered to be random noise and was subsequently disregarded. After this filtering process, **115 candidate signals with powers exceeding the threshold were identified, distributed across 28 distinct frequencies.**
- For the signals at these 28 frequencies, the researchers calculated their **False Alarm Probability (FAP)**, which represents the probability that the signal could have arisen from random noise. This calculation accounted for the "look-elsewhere effect," a statistical correction that addresses the increased chance of finding a strong signal purely by coincidence due to the large number of frequencies searched.
- **Among all 28 frequencies, only the signals at 908 Hz and 909 Hz demonstrated extremely high statistical significance, with FAP values substantially lower than those of the other frequencies.** Further analysis revealed that **all data subsets producing these strong signals at both frequencies were concentrated within a single 160-millisecond time window.**

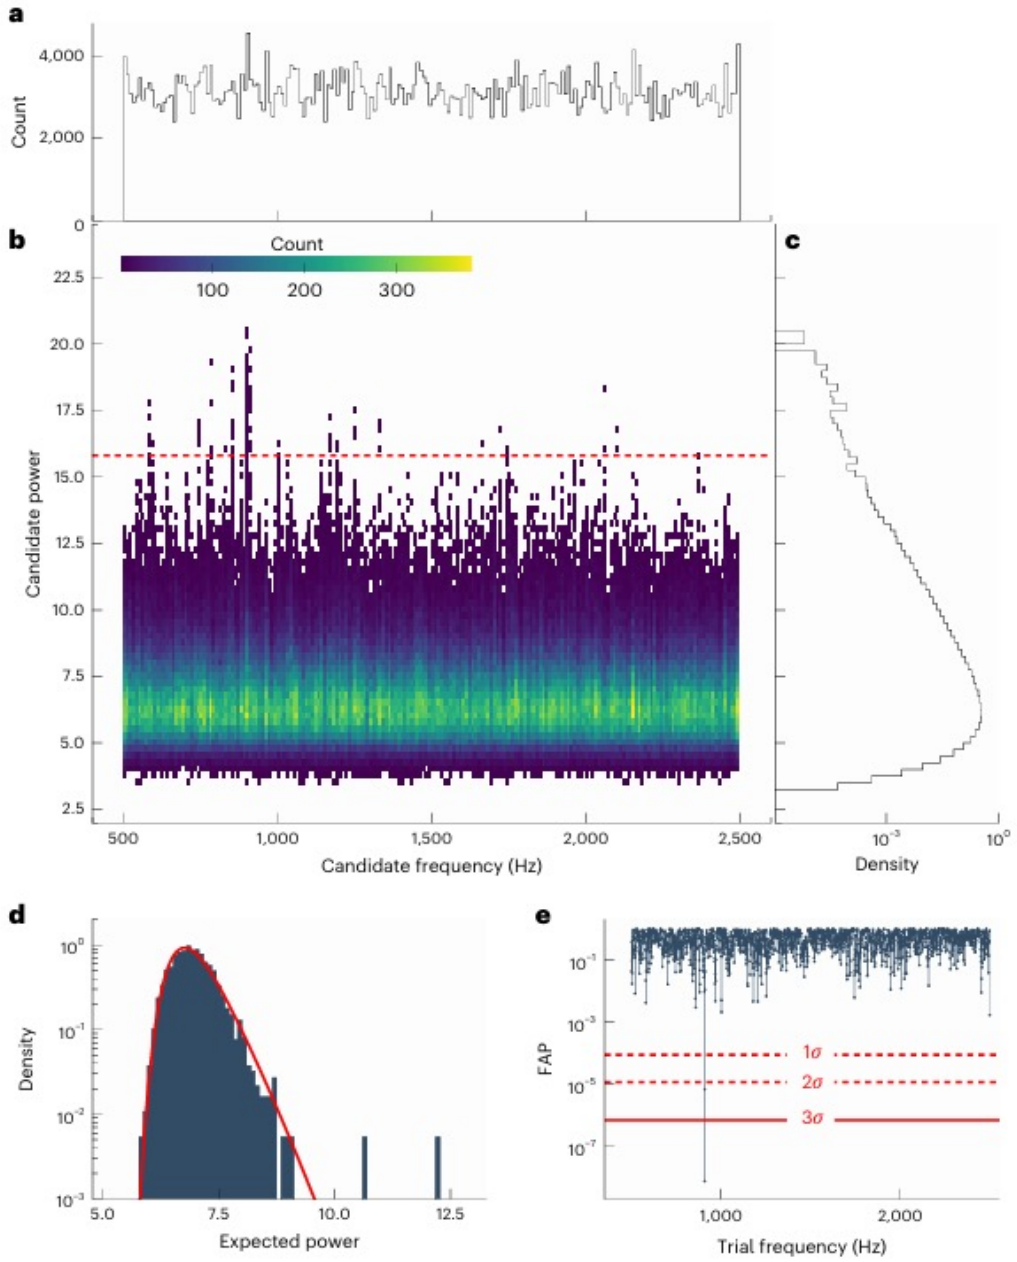

Results of blind search for periodicity in GRB 230307A.

a, c: Marginal distributions of the candidate frequencies (a) and powers (c).

b: The red dashed line (R_{false}=15.81) is a "threshold line," indicating the theoretical maximum power that random noise could likely achieve in such a large-scale search. We identified 115 candidate powers exceeding R_{false} at 28 distinct candidate frequencies.

d: Power Distribution of Pure Noise, the red curve shows the fitted probability density function for this distribution.

e: The FAP represents the probability that the signal strength observed at a given frequency is purely the result of random noise coincidence. The smaller this value, the more likely the signal is real. The red horizontal lines indicate the 1σ, 2σ, and 3σ significance levels.

- To accurately assess the statistical significance of the blind search, which spanned numerous overlapping time and energy segments, we first estimated the **total number of independent trials**. This figure was approximated as $(51s/100ms)\times(1/0.25)\times(2000Hz/10Hz) = 4.08\times10^5$, representing the product of the number of non-overlapping time segments, non-overlapping energy ranges, and effective independent frequency bins, respectively.
- Upon correcting the 28 candidate frequencies for this substantial number of trials, **only the 908 Hz signal emerged as statistically significant, with a trial-corrected False Alarm Probability (FAP) of $\approx \mathbf{1.45\times10^{-5}}$. In contrast, the 909 Hz signal was only marginally significant (FAP $\approx \mathbf{1.34\times10^{-2}}$), all other frequencies were consistent with noise (FAP $= \mathbf{1}$)**.
- Crucially, all significant signals detected around 908 and 909 Hz were temporally clustered within a single 160 ms window. This concentration suggests **a common origin from the same quasi-periodic oscillation (QPO)**. Consequently, the representative FAP for this QPO feature is defined as the minimum FAP across this frequency range, yielding a value of 1.45×10^{-5}.

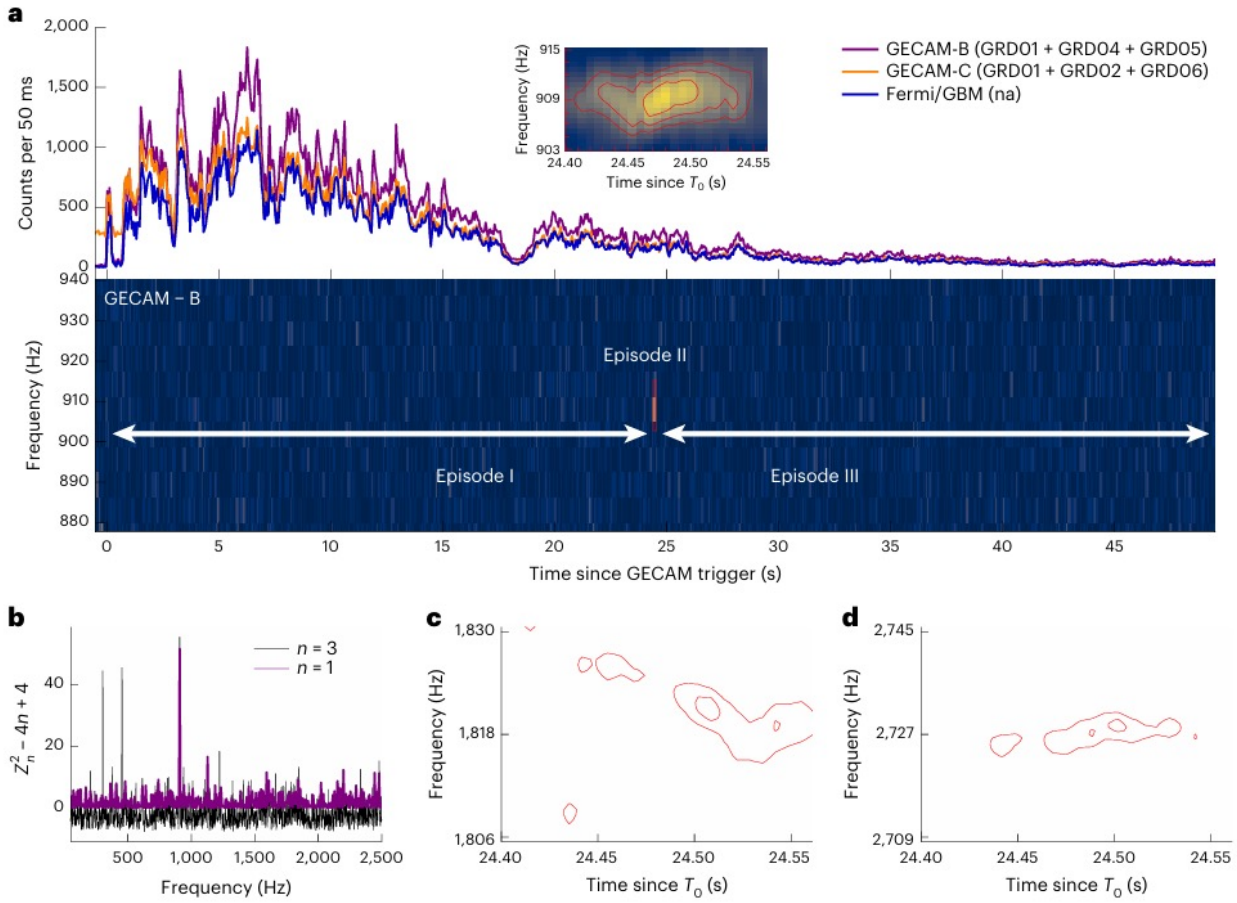

Occurrence and harmonics of the QPO signal from GRB 230307A

Panel a，Upper Panel (Light Curve), Lower Panel (Dynamical Power Spectra)

Panel b: This plot displays the results of the **H-test** applied to the signal within the Episode II time interval. At 909 Hz, the H-test statistic reaches its maximum peak, reconfirming that 909 Hz is the fundamental frequency of this periodic signal. Significant peaks also appear at approximately one-third (303 Hz) and one-half (455 Hz) of 909 Hz; these are known as subharmonics. Their presence greatly enhances the credibility that the signal originates from a **rotating physical entity**. The H-test result (H-statistic ≈ 55.32, with an optimal harmonic number of 3) indicates that, in addition to the fundamental frequency, at least two higher-order harmonics are present (Panel c: 1818 Hz; Panel d: 2727 Hz).

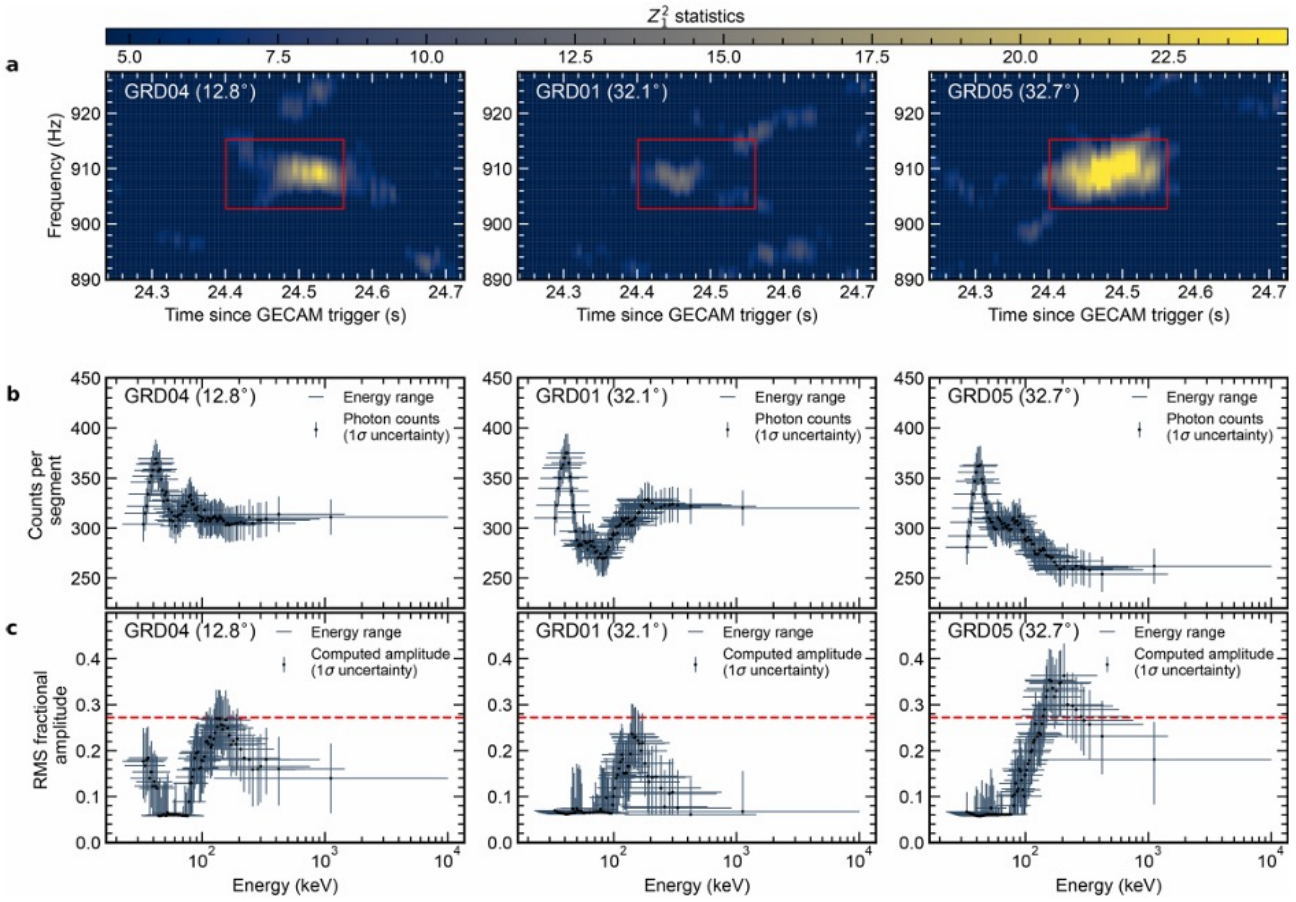

Consistent detection of the 909-Hz QPO across multiple detectors on GECAM-B.

Panel a: This panel displays the dynamical power spectra from individual detectors. The angle in the parentheses is the incident angle of each detector relative to the GRB source direction. The red box marks the core time interval ([24.401, 24.561] s) and frequency (909 Hz) where the signal appears.

In the dynamical power spectra of all three detectors, the yellow bright spots representing the signal appear simultaneously within the red box. This indicates that three independent electronic recording channels registered this signal at the exact same time and frequency.

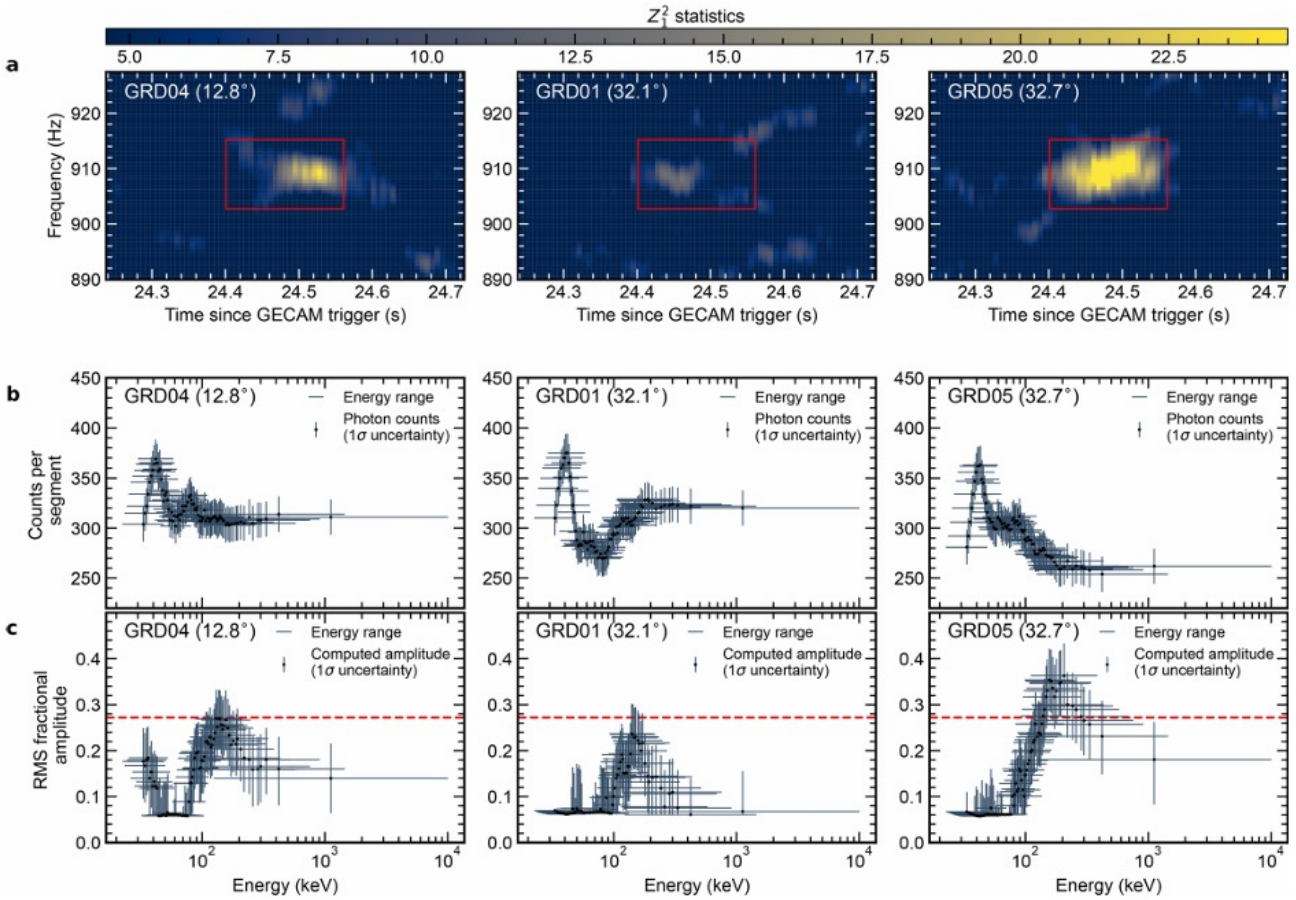

Panel b: Shows the raw photon count rates in different energy bins for the three detectors during the signal's appearance. It demonstrates that due to different incident angle, GRD04 received significantly more photons across all energy bands than the other two detectors.

Panel c (Root Mean Square (RMS) Fractional Amplitude at 909 Hz): The red dashed line marks an amplitude level of approximately 27.2%, which is the peak amplitude calculated from the combined data.

Although the total number of photons received by the three detectors differs, the trend of their measured signal strength as a function of energy is highly consistent. In all cases, the strength is weaker at lower energies, peaks in the mid-energy range (approx. 100-300 keV), and then decreases. The peak signal strengths measured by the three detectors are consistent with the 27.2% level marked by the red dashed line, within the error margins.

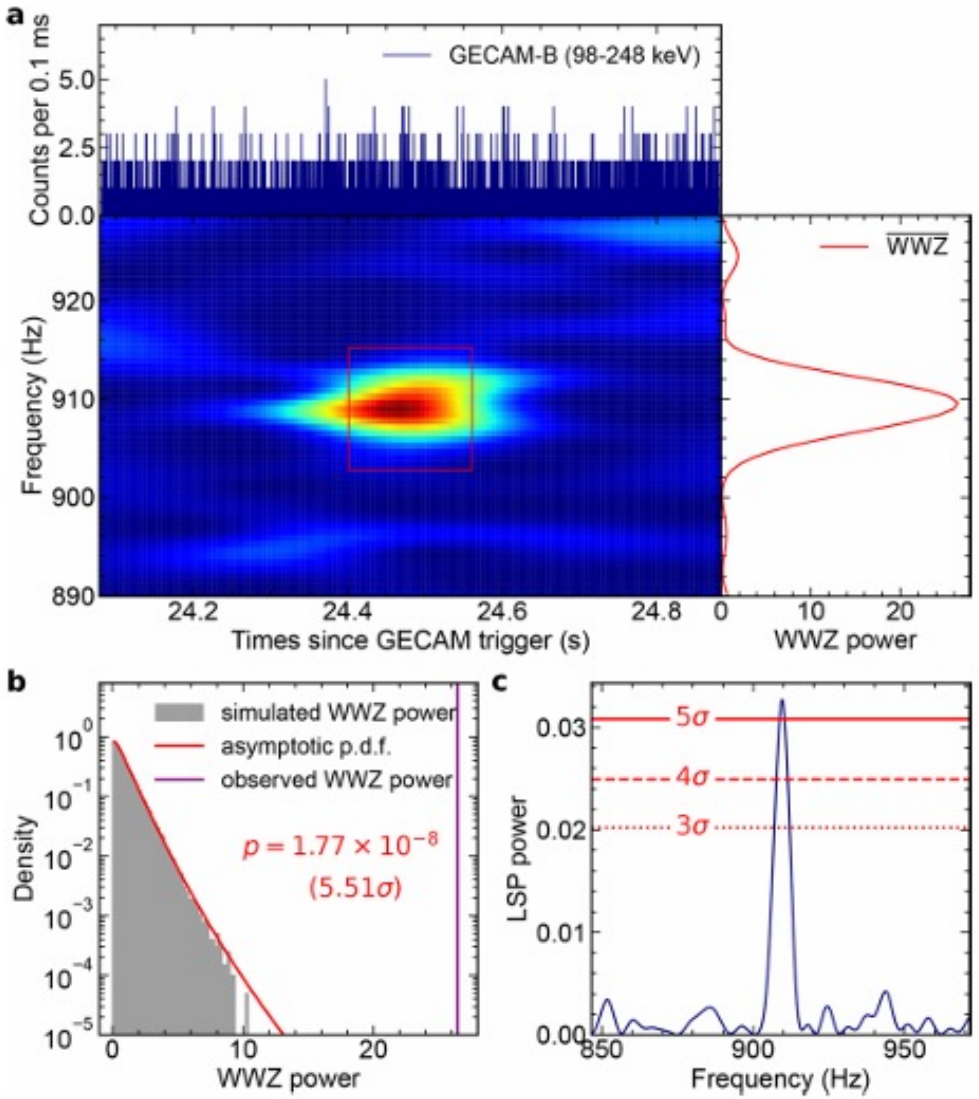

Monte Carlo simulations to verify the 909-Hz signal in GRB 230307A.

Panel a, Middle Panel: The WWZ dynamical power spectrum. The red box again marks the core region where the signal appears (in time, 24.4-24.56 s; in frequency, 909 Hz).

Right Panel: The power spectra from the region within the red box are averaged over time. A sharp peak near 909 Hz is clearly visible.

Panel b (Monte Carlo Simulation and Probability Calculation):

Gray Histogram (simulated WWZ power): This represents the results of a Monte Carlo simulation. Based on the noise characteristics of the observed light curve, the research team generated "pseudo light curves" that contained no real periodic signals. They then analyzed this pure noise data using the exact same methods as the real data to obtain the "maximum power values generated by noise."

Red Curve: This is the theoretical probability density function (an F-distribution) fitted to the gray histogram.

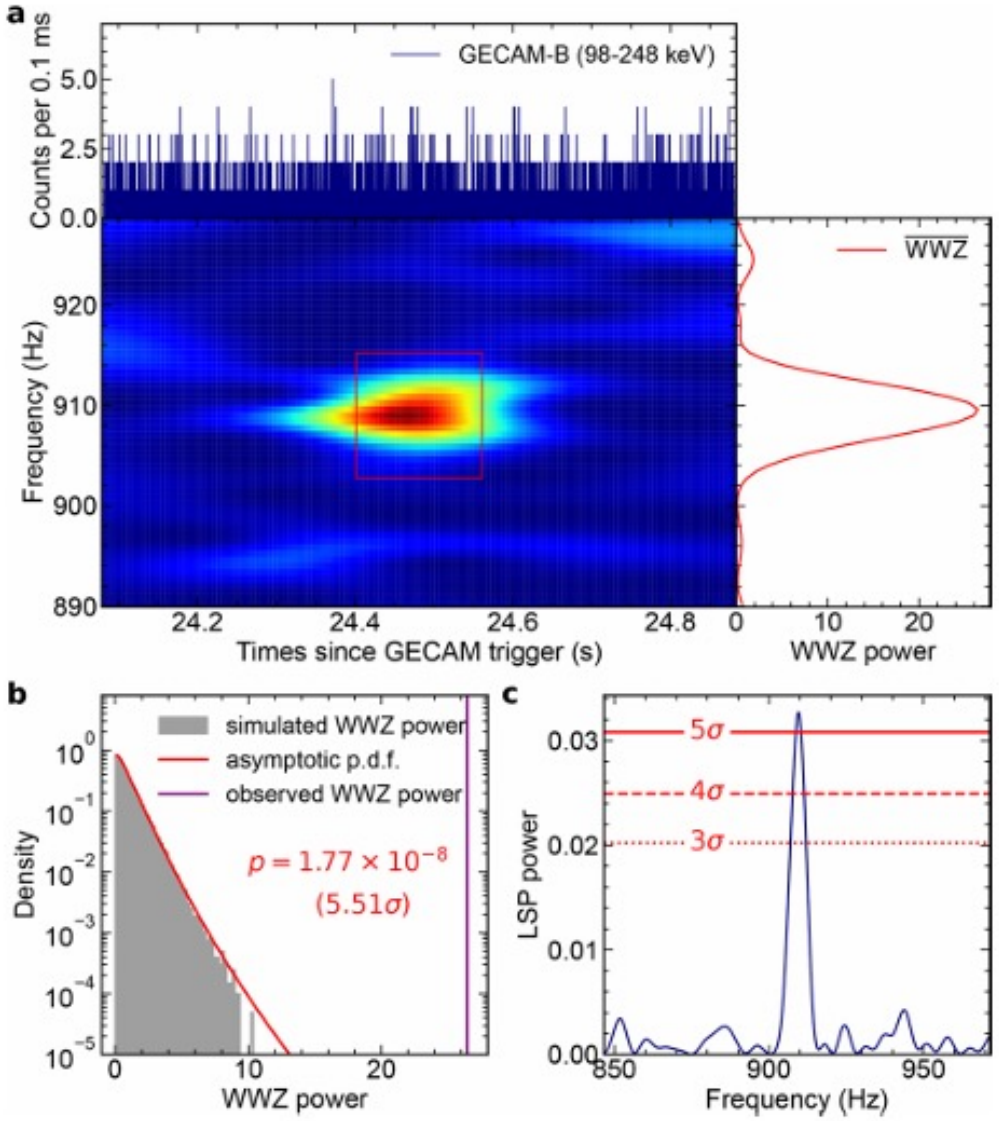

The purple vertical line, representing the real observation, falls far into the extreme right tail of the gray noise distribution, even beyond the range of the histogram. This implies that in 100,000 simulations, not a single instance of random noise produced a signal as strong as the one observed. Using the red theoretical curve, the team calculated the probability of noise producing a power value greater than or equal to the observed value (i.e., the area under the curve to the right of the purple line) to be $p = 1.77 \times 10^{-8}$. This probability corresponds to a significance of 5.51σ.

Panel c (Lomb-Scargle Periodogram): Red Horizontal Lines (3σ, 4σ, 5σ) represent the false alarm probability thresholds for different significance levels. The result shows a very sharp peak at 909 Hz, and the height of this peak exceeds the 5σ threshold line.

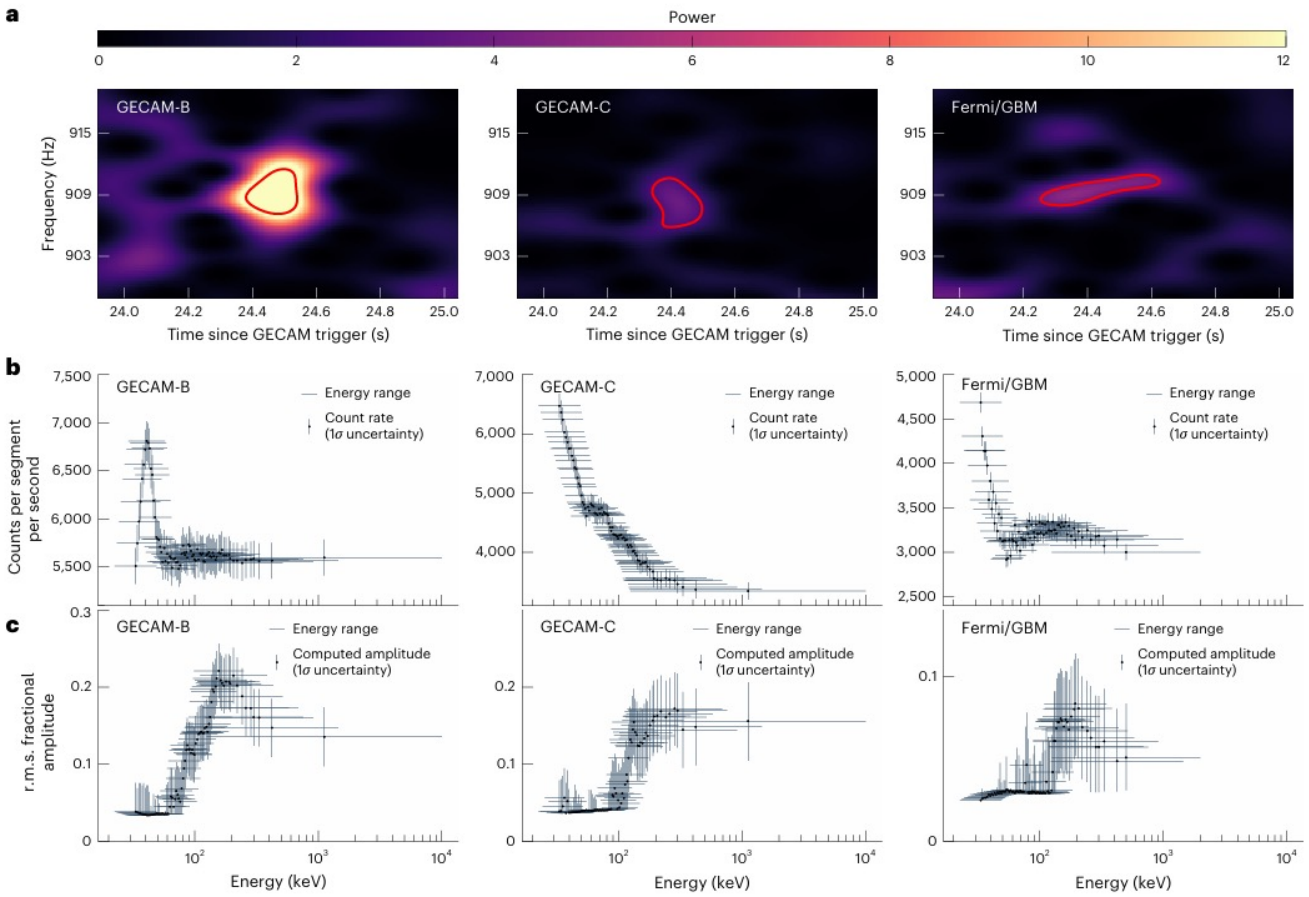

Consistent QPO detections in GRB 230307A across multiple instruments.

Panel a: This panel displays the dynamical power spectra recorded by three different instruments.

The time centers of the three red contour regions are almost perfectly aligned at approximately 24.48 seconds. The frequency centers of the three regions are also identical, all locked at 909 Hz.

Panel b: This panel shows the number of photons collected by each instrument in different energy bins. It is evident that GECAM-B's photon count in the key energy range is the highest, which explains why it detected the strongest signal.

Panel c: This panel illustrates the trend of RMS Fractional Amplitude as a function of energy. The three plots show a highly consistent trend: the signal strength is not constant but is weaker at lower energies, peaks in the mid-energy range of approximately 100-300 keV, and then decreases at higher energies.

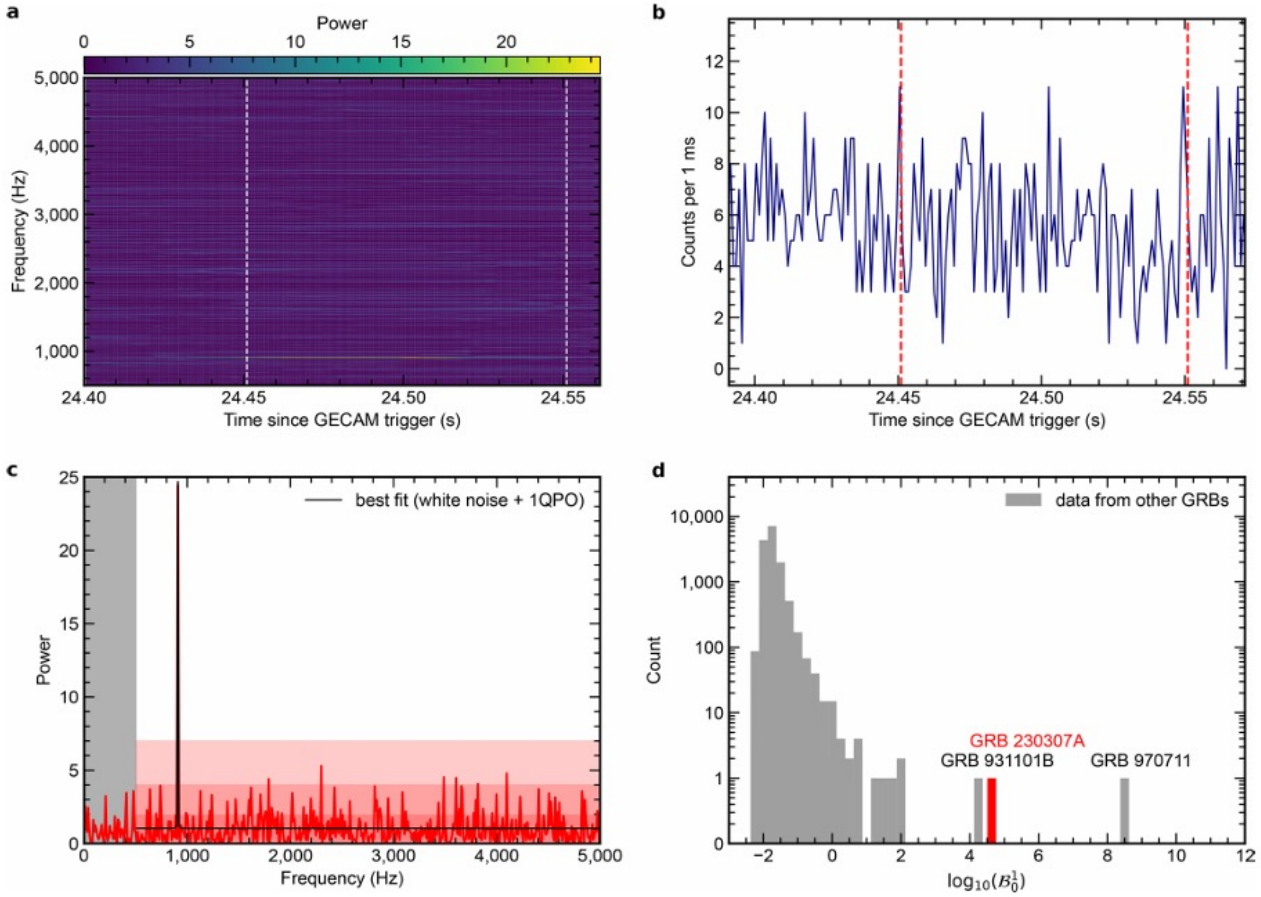

Bayesian model fitting for the 909-Hz QPO in GRB 230307A.

Panel a: A dynamical power spectrum based on Groth normalized power spectra. The white vertical lines mark a 100-millisecond time window where the power is maximal.

Panel b: The raw light curve. The red dashed lines indicate the time window ([24.451, 24.551] s) selected from Panel a.

Panel c, Gray Shaded Area: Marks the "red noise" region between 10-500 Hz. It is clear that this region has no significant excess power, indicating that the noise is primarily white noise (Poisson noise). Red Bands: These represent the confidence intervals predicted by a "pure white noise" model. The sharp peak at 909 Hz can be seen to far exceed the 3σ red band, once again demonstrating that it cannot be white noise. Black Curve (best fit): This is the best-fit result obtained by fitting the data with a "white noise + one signal (Lorentzian model)". The black curve can be seen to pass perfectly through the 909 Hz peak.

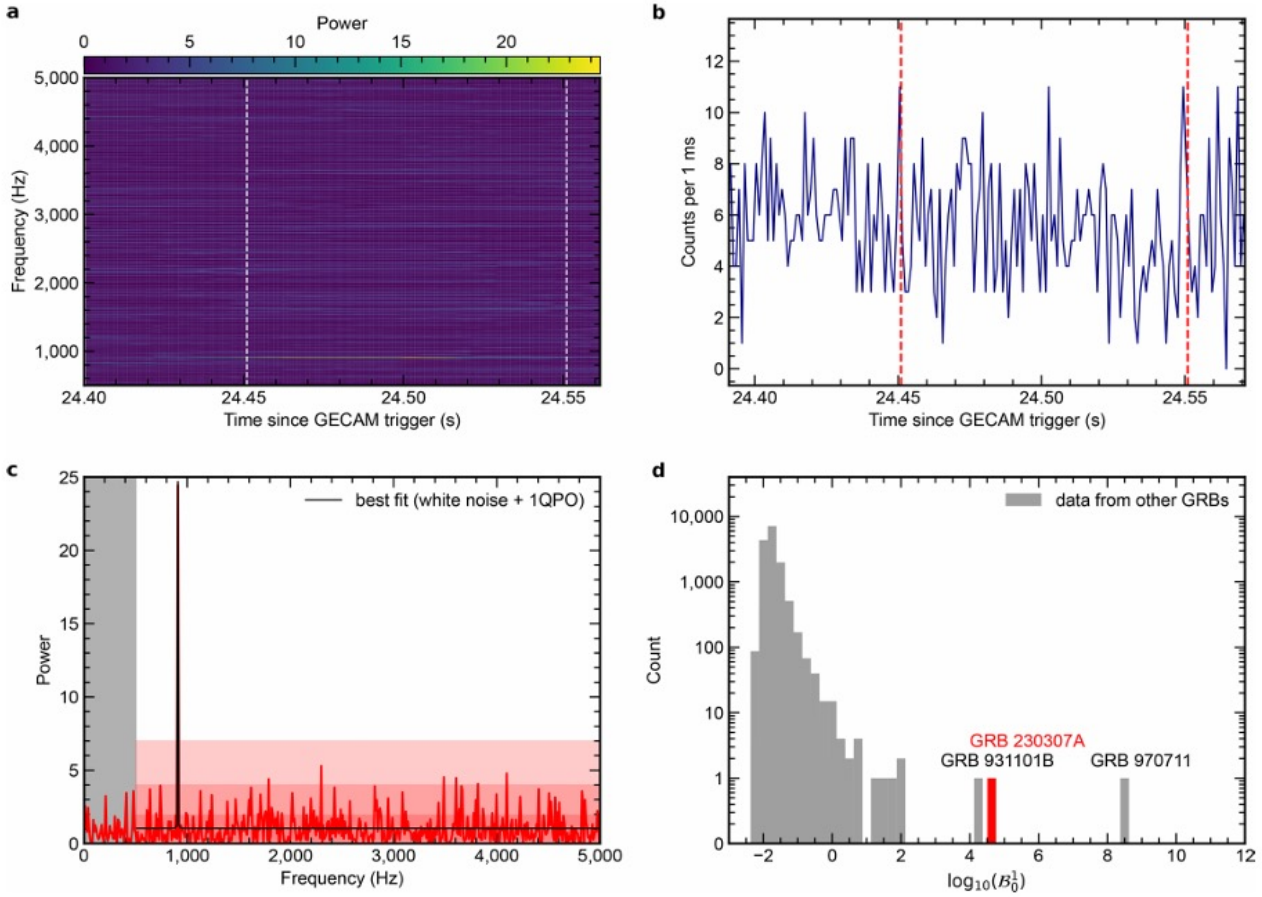

From the fitting results, the quality factor of this signal can be calculated as Q = 476. The Q factor (central frequency / peak width) is a measure of the signal's "purity." A higher Q value means a narrower frequency range for the signal, making it more like a single-frequency "beat." A Q value of 476 is extremely high, far exceeding that of typical quasi-periodic oscillations (QPOs), which strongly suggests that this is a true periodic signal rather than a broad oscillation.

Panel d: The Bayes factor is an indicator used to compare the goodness-of-fit of two models. A B value greater than 100 (i.e., log(B) > 2) is generally considered "decisive evidence" that model 2 is far superior to model 1.

Gray Histogram (data from other GRBs): It is evident that most GRBs have small Bayes factors, suggesting their Power Density Spectra (PDS) can be well explained by a pure noise model.

Red Histogram: GRB 230307A, its Bayes factor is extremely high. This indicates that the "noise + signal" model is better than the "pure noise" model.

GRB 931101B & GRB 970711: These are two other short GRBs from the literature where QPOs have been previously reported. Their Bayes factors are also very high.

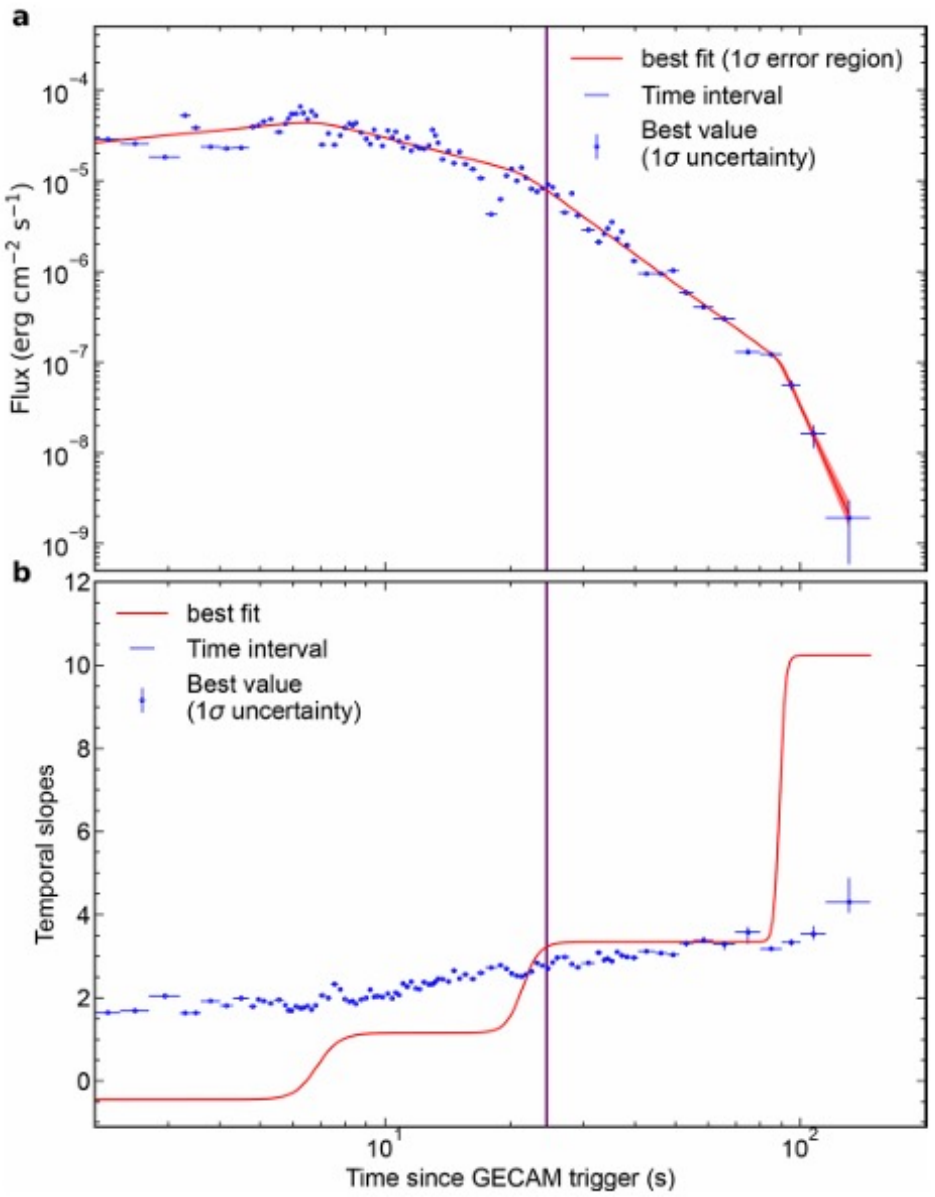

Coincidence between the 909-Hz signal's detection interval and the transition to high-latitude emission.

Panel a: This panel shows the evolution of the GRB's energy flux over time. Red Curve (best fit): This is the best fit to the blue data points using a "smoothly broken power-law" function.

Panel b, Red Curve: This shows the evolution of the slope over time, calculated from the red best-fit curve in Panel a. Blue Data Points (curvature effect): These represent the predicted evolution of the slope based on the "high-latitude emission" or "curvature effect" theoretical model.

It is evident that starting from approximately 20 seconds, the red curve (representing the actual observation) begins to align closely with the blue data points (representing the "curvature effect" theory). This strongly suggests that from this moment on, the dominant component of the observed GRB radiation transitioned from the **"prompt emission"** of the main burst to the **"afterglow" of high-latitude emission**. The purple vertical band once again falls precisely on the starting point of this transition. In other words, the appearance of the 909 Hz signal perfectly marks the exact moment when the GRB's radiation mechanism shifted from one phase (prompt emission) to another (high-latitude emission / curvature effect).

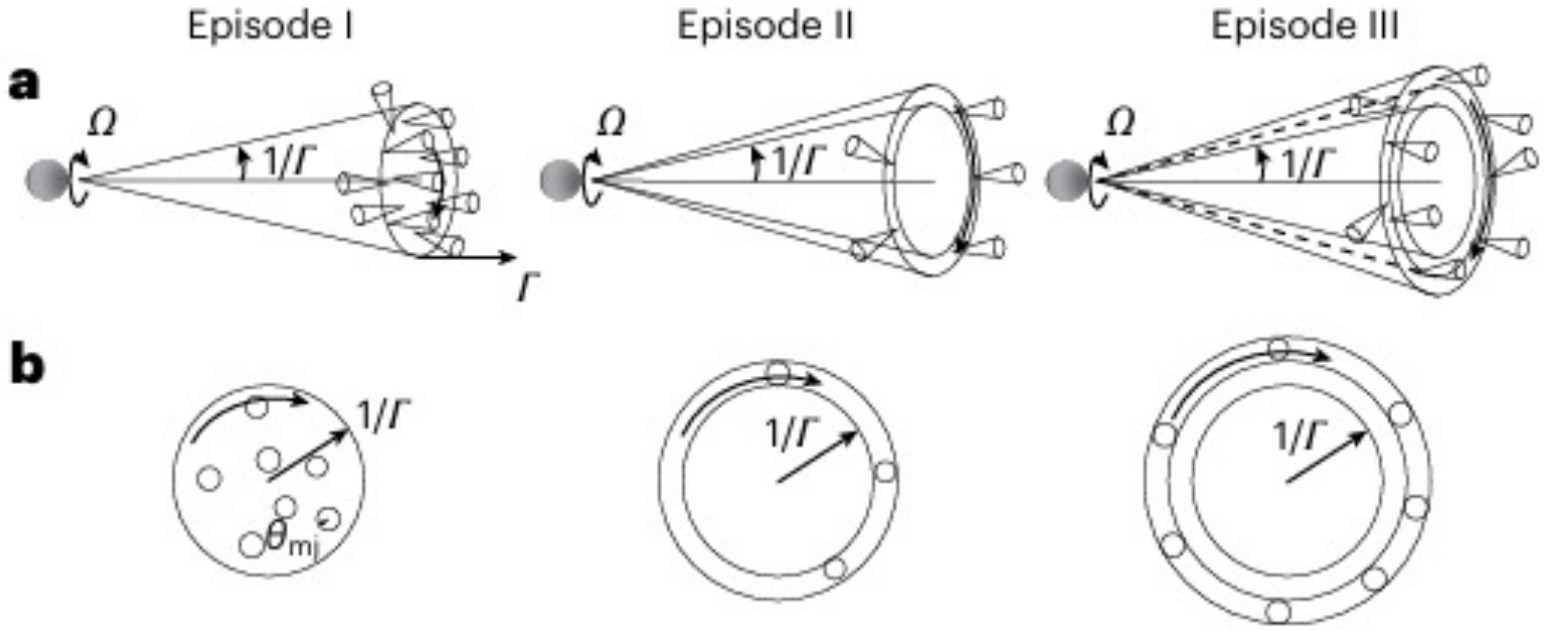

A physical scenario to interpret the 909-Hz signal from GRB 230307A.

Episode I: Quasi-symmetrically distributed mini-jets contribute to the observed gamma-ray emission without detectable periodicity.

Episode II: As the transition from the prompt emission to the HLE occurs, a small number of high-latitude mini-jets introduces asymmetry in their spatial distribution, enabling the detection of periodicity.

Episode III: At later times, the number of observed mini-jets increases, causing the HLE to become quasi-symmetrically distributed, rendering the periodicity undetectable again.

This physical picture interprets the existence of the periodic signal during a brief time interval.

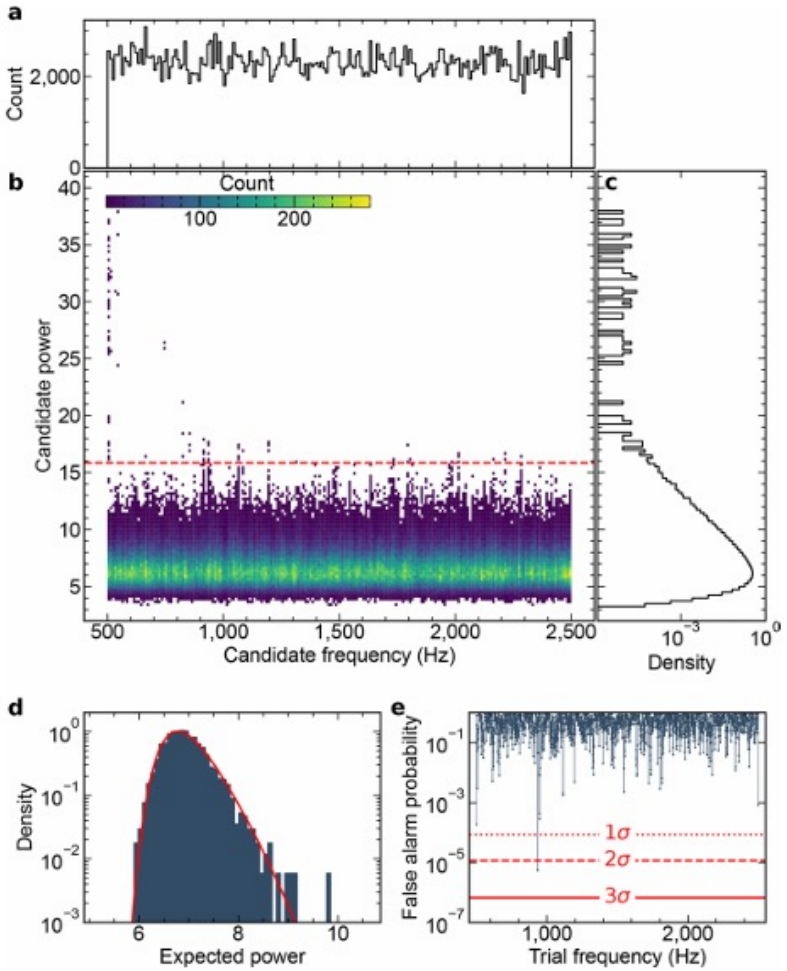

Results of the blind search for periodicity in GRB 211211A.

Another very similar astrophysical event—GRB 211211A—was analyzed using the exact same methods as those applied to GRB 230307A.

At approximately 546 Hz, although the power was highest, its false alarm probability was as high as 31.09%, which means there is nearly a one-in-three chance that it is just noise. This is statistically insignificant.

At approximately 935 Hz (a frequency of interest because it is close to the 909 Hz of GRB 230307A), although there is a small "dip," its false alarm probability is still as high as 1.09%, failing to even meet the 2σ standard. In a blind search where a large number of trials must be considered (the "look-elsewhere effect"), this is also statistically insignificant.

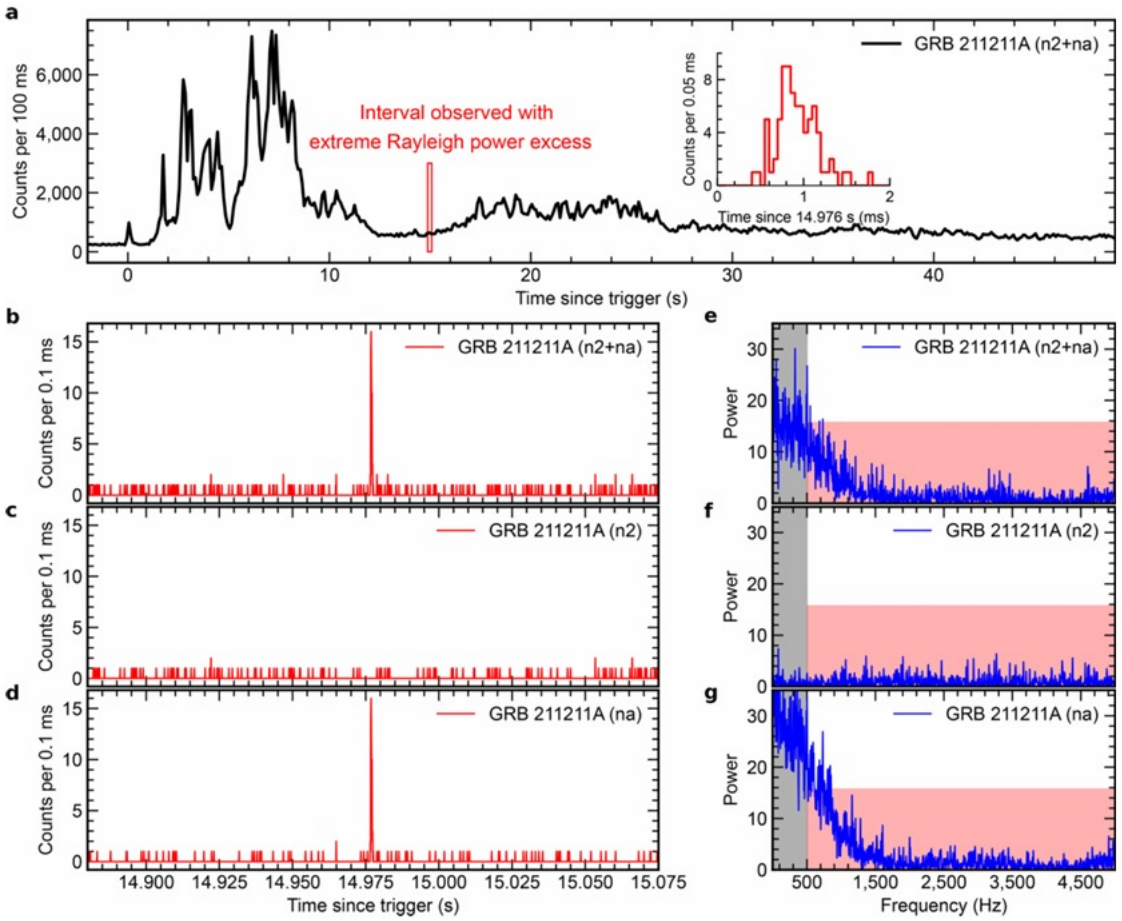

Very strong red noise during GRB 211211A.

Panel a: The overall light curve of GRB 211211A.

Red Box: Marks the time interval of the strongest candidate signal.

Inset: This is a high-time-resolution zoom-in of the red-boxed region. We can clearly see that the source of the so-called "signal" is not a continuous, periodic oscillation, but an extremely bright, extremely brief (duration of about 1 ms), isolated spike.

Panels b, c, d: These panels show the light curves recorded by different detector combinations or by individual detectors. The signal is clearly recorded only in detector na, while being completely absent in detector n2. This makes it highly probable that it is an instrumental noise, a cosmic ray event, or an electronic malfunction specific to detector na.

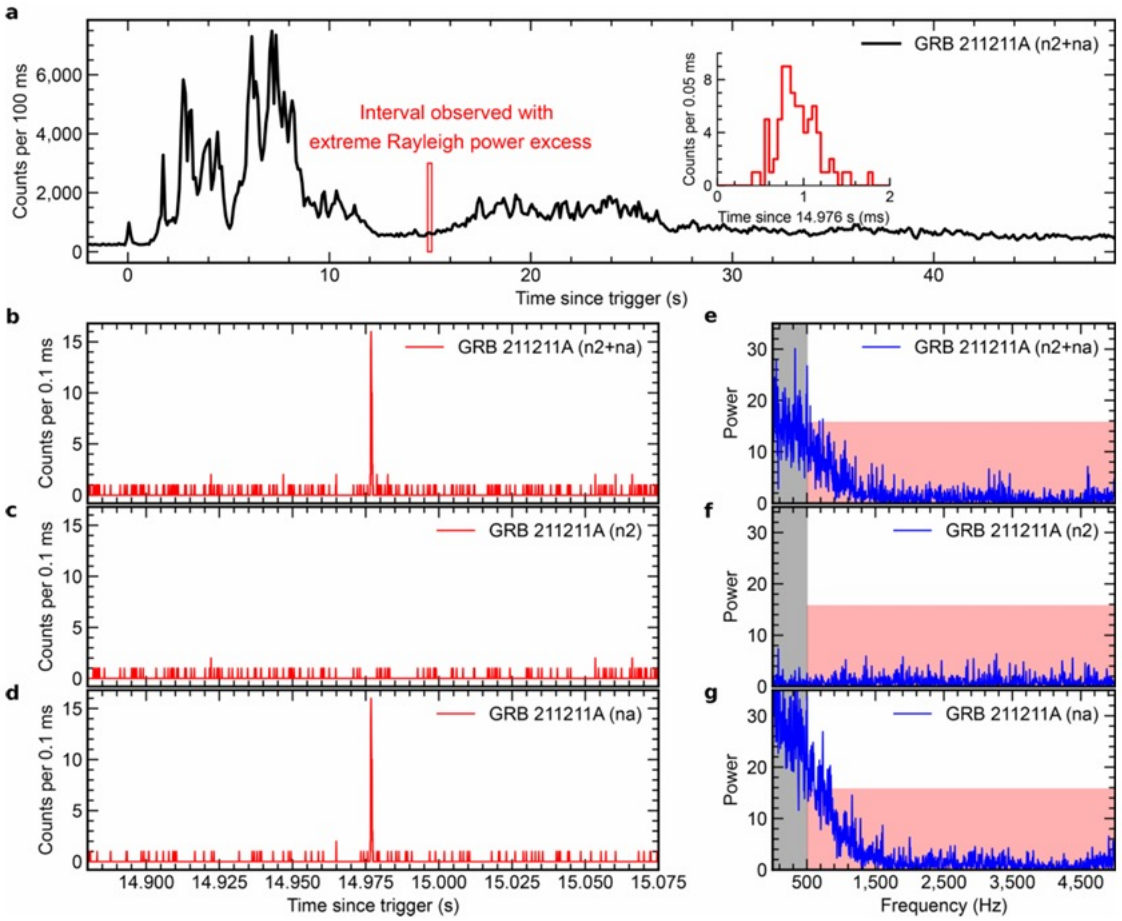

Panels e, f, g: These panels show the Groth-normalized power spectra corresponding to Panels b, c, and d.

Gray shaded area: low-frequency region from 0-500 Hz.

Red shaded area: The region from 500-5000 Hz and the threshold power from Extended Data Fig. 5.

In Panel g (detector na), we can see that in the low-frequency region (gray area), the power is very high and extends towards higher frequencies, forming a long "tail." This type of noise, which is strong at low frequencies and weak at high frequencies, is known as "red noise." The "tail" of this red noise extends into the signal search region above 500 Hz (red area) and exceeds the noise threshold. It is this noise "tail" that was incorrectly identified as a "signal" in the 500-549 Hz band.

In Panel f (detector n2), because the spike is absent, the power spectrum is very flat, consisting entirely of white noise, with no part exceeding the threshold.

The power spectrum in Panel e (combined data) is a superposition of those in f and g.

Summary

- A 909 Hz periodic signal with a 4.2σ significance was detected in gamma-ray burst GRB 230307A, providing strong evidence for a millisecond magnetar as the merger remnant.
- The signal was brief, lasting only 160 ms, and coincided precisely with the critical transition from the prompt emission to the high-latitude emission phase.
- Joint detection by multiple independent instruments (GECAM and Fermi) confirmed the signal's astrophysical origin and ruled out instrumental artifacts.
- Bayesian analysis verified it as a genuine periodic signal (not a QPO), indicating its origin from a stable rotator.
- The signal's transient nature is explained by a model based on asymmetric "mini-jets." During the transition phase, the observational symmetry of the jet was temporarily broken, briefly revealing the spin period of the central engine.