Taming Seismic Phase: From Statistical Characterization to Objective Phase Masking

Introduction

Modern land seismic acquisition increasingly operates in low signal-to-noise regimes. Dense sampling with single sensors and weaker sources, driven by operational, environmental, and permitting constraints, is now common. These conditions challenge processing assumptions that rely on stable trace-by-trace waveform attributes (Holt and Lubrano, 2020). Seismic data must pass through multiple intermediate stages long before stacking or migration, including velocity analysis, full-waveform inversion, statics estimation, and deconvolution, all of which depend on coherent signal attributes within local multichannel ensembles. Velocity models lack the resolution needed to explain small-scale heterogeneity and near-surface scattering, so model building itself operates on data already compromised by strong variability. At low signal-to-noise ratios, these algorithms become inefficient and unstable, often failing unevenly across the frequency band and making usable bandwidth difficult to diagnose and recover.

In such environments, waveform phase is no longer a deterministic trace-by-trace attribute (Rohatgi et al., 2025). It fluctuates strongly from trace to trace and varies with frequency, making broadband phase tracking fragile. Phase-sensitive methods do not fail outright, but break down selectively in noise-dominated frequency bands. For example, full-waveform inversion often becomes ineffective below ~2 Hz, where wavelengths increase and phase coherence is overwhelmed by scattering and residual noise. Phase must therefore be treated as a statistical variable defined over local multichannel ensembles. The key issue is not whether ensemble-based methods should be used, but how their impact on phase stability can be quantified and controlled frequency by frequency. Phase variability provides a sensitive diagnostic: where it can be reduced, coherency can be recovered; where it cannot, the data themselves identify unusable frequencies. This behavior follows a gradual, predictable progression with ensemble size, eliminating the notion of a "magic" aperture and enabling fit-for-purpose conditioning.

In this work, we adopt a data-driven framework for analyzing and stabilizing seismic phase using circular statistics (Rohatgi et al., 2025). Phase is treated as a circular variable, allowing both its dominant trend and variability to be quantified objectively within local ensembles. This separates phase from amplitude, avoids phase-wrapping ambiguities, and provides direct control over stabilization through ensemble size. Within this framework, phase masking becomes a statistically grounded tool for auditing and stabilizing pre-stack seismic data in low signal-to-noise environments.

Phase masking as a phase-only conditioning framework

Phase masking was introduced as a data-driven method for stabilizing seismic phase in low signal-to-noise environments using local trace ensembles (Bakulin et al., 2023). In its original formulation, it was implemented as a phase-only time–frequency masking operation, where a local stack provided a practical estimate of the dominant phase used to adjust individual trace phases while preserving amplitudes. This allowed anything from small phase corrections to full phase substitution and demonstrated that substantial coherency gains could be achieved without artifacts or amplitude distortion. Here, we generalize phase masking by operating directly on the raw trace phase and computing the reference phase as a circular mean within a moving window in space and frequency (Rohatgi et al., 2025). This removes the need for local stacking, treats phase consistently as a circular variable, and cleanly separates phase from amplitude. Amplitudes are preserved by construction, while phase variability is reduced in a statistically consistent manner. Circular variance provides an objective measure of phase stability and direct control of conditioning strength through ensemble size, making phase masking a predictable, fit-for-purpose operation rather than a heuristic one.

Synthetic example: Phase distortions caused by near-surface scattering

To illustrate these concepts, we explicitly simulate the impact of near-surface scattering using a frequency-dependent speckle-noise model (Figure 1). Starting from clean, laterally coherent traces with identical wavelets (Figure 1a), we apply random phase perturbations in the frequency domain following the speckle formulation of Bakulin et al. (2022), with perturbation magnitude increasing linearly with frequency. This captures the cumulative effect of small-scale heterogeneity and shallow scattering, producing scattering-like waveform distortions in the time domain (Figure 1b) while preserving the amplitude spectrum. Although the underlying signal remains perfectly coherent, the resulting data exhibit strong trace-to-trace phase variability, closely resembling low signal-to-noise field data affected by near-surface scattering (Bakulin et al., 2022; 2024).

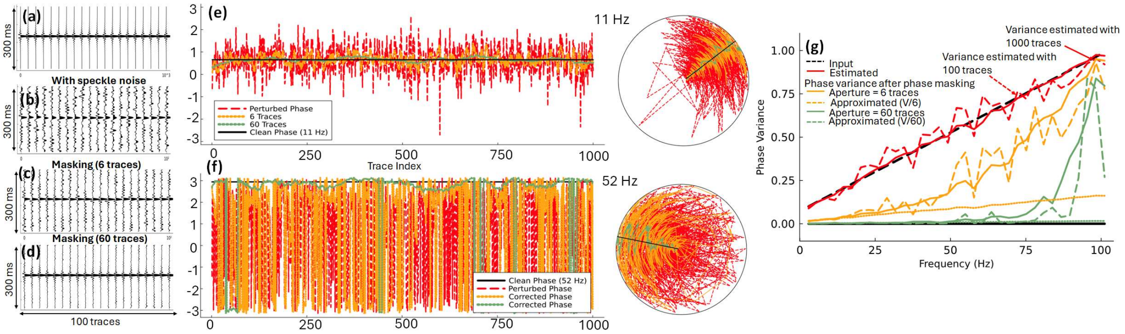

Figure 1. *Synthetic example illustrating frequency-dependent speckle phase perturbations and their systematic reduction by phase masking. (a) Clean synthetic traces with laterally consistent wavelets. (b) Traces contaminated by frequency-dependent speckle noise, producing scattering-like waveform distortions. (c,d) Traces after phase masking with apertures of 6 and 60 traces, showing mild and stronger phase stabilization. (e,f) Phase at 11 Hz (e) and 52 Hz (f) shown in two equivalent representations: wrapped linear phase versus trace index (left) and circular polar coordinates (right), for perturbed data (red), phase masked with 6 traces (orange), phase masked with 60 traces (green), and the clean reference phase (black). Linear representations at higher frequencies appear asymmetric and biased due to wrap-around at ±π, whereas circular coordinates remove artificial boundaries and reveal symmetric phase perturbations tightening around the signal phase with increasing aperture, without contradiction or bias. (g) Phase-variance spectra versus frequency for perturbed and masked data, showing systematic variance reduction with increasing ensemble size.*

Phase stabilization is then applied using circular-mean phase masking with different ensemble apertures. A small aperture of six traces produces a modest reduction in phase variability (Figure 1c), while a larger aperture of sixty traces yields substantially stronger stabilization (Figure 1d). In the linear wrapped-phase representation (left panels of Figures 1e and 1f), high-frequency phase appears asymmetric and biased due to artificial wrap-around at ±π. This impression is misleading. When the same samples are examined in circular coordinates (right panels of Figures 1e and 1f), the perturbations are symmetric about the signal phase regardless of whether the mean lies near 0 or ±π. The circular domain removes artificial boundaries, revealing a clear tightening of the phase distribution with increasing aperture. Across frequencies, phase masking reduces variability in a systematic and predictable manner without biasing the mean.

The remaining phase variability is quantified using the circular phase-variance metric (Figure 1g). Its apparent dependence on ensemble size is purely statistical rather than physical. Smaller ensembles (dashed lines) provide localized but noisier estimates, whereas larger ensembles (solid lines) yield smoother and more robust diagnostics. In all cases, phase variance decreases monotonically with increasing ensemble size, consistent with the expected $1/N$ scaling for statistical stabilization at low variance (Bakulin et al., 2024). The same underlying trend is observed at all levels of statistical resolution. Ensemble size therefore controls robustness, not behavior: large ensembles characterize global trends, while smaller ensembles resolve local variations and can be smoothed if needed.

Taken together, this example clarifies that two related but distinct choices are involved: the aperture used for phase masking and the ensemble size used to characterize phase variability. These parameters are complementary rather than conflicting. Both expose the same underlying statistical progression, differing only in resolution and robustness. Once this nested structure is recognized, aperture selection becomes deliberate rather than heuristic, and phase masking and phase-variance analysis form a consistent, non-ambiguous framework for conditioning and diagnosing phase stability.

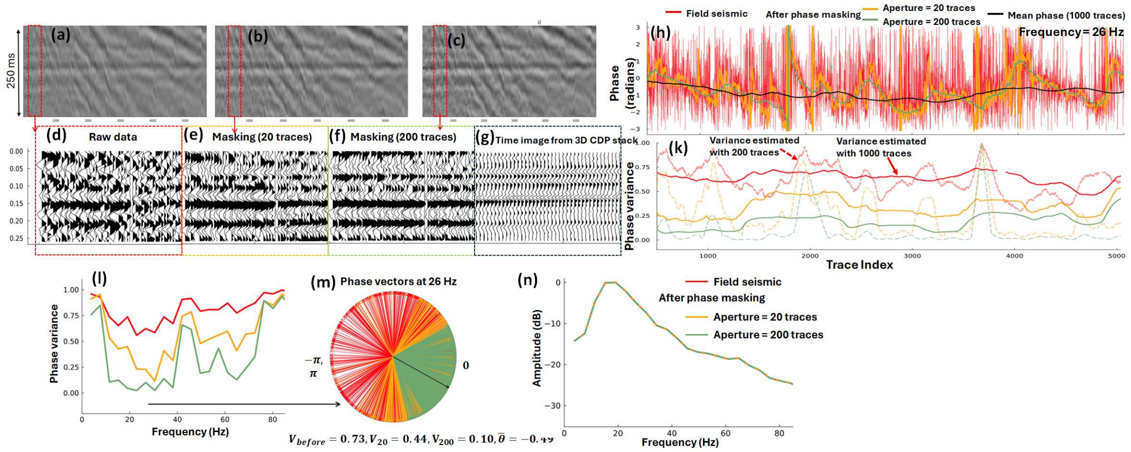

Figure 2. *Real-data example from a complex onshore U.S. land survey illustrating phase masking. (a,d) Prestack CDP super-gather before phase masking. (b,e) Same gather after phase masking with 20 traces, showing mild stabilization. (c,f) Phase masking with 200 traces, showing stronger stabilization. (g) Reference 3D CDP stack confirming that enhanced events are real and no artifacts are introduced. (h) Wrapped phase at 26 Hz along the gather for the original data (red) and after masking with 20 (orange) and 200 (green) traces. (k) Local phase variance along the gather. Phase masking reduces variance, while ensemble size controls emphasis on localized disruptions versus long-wavelength trends. (l) Average phase variance versus frequency for a single near-offset 2000-trace ensemble, defining usable bandwidth from phase coherence. (m) Polar phase vectors at 26 Hz showing progressive tightening of the phase distribution with increasing aperture. (n) Amplitude spectra before and after phase masking, confirming amplitudes are unchanged.*

Real data example: Phase stabilization in a complex near-surface environment

We apply the method to a complex onshore U.S. land dataset. A CDP super-gather after NMO correction is selected immediately after linear noise removal, where residual ground roll and other coherent noise remain and phase stability is still poor (Figures 2ad). Although reflection is visible, strong trace-to-trace phase fluctuations dominate, making the gather representative of low signal-to-noise conditions encountered for further processing. Phase masking is applied using two ensemble apertures: 20 traces for mild stabilization and 200 traces for stronger conditioning (Figures 2b,e and 2c,f). In both cases, phase variability is reduced and lateral coherency improves in a systematic and predictable manner, with increased stabilization for the larger ensemble. Importantly, this conditioning does not introduce artifacts. Comparison with the stacked section (Figure 2g) confirms that all enhanced events correspond to real reflectors present in the final image. Phase masking moves the data smoothly along a continuous statistical progression, without abrupt changes or the creation or suppression of events. Single-frequency phase analysis confirms the same behavior observed in the synthetic example. At 26 Hz (Figure 2h), the original data exhibit strong trace-to-trace variability. Masking with 20 traces already reduces this variability, while 200 traces produce a substantially smoother and more stable phase. As ensemble size increases, the estimated mean phase converges toward a smoothly varying reference obtained from a very large ensemble. This convergence reflects a statistical effect rather than a physical change: increasing ensemble size reduces variability without biasing the mean.

Once reflections are stabilized, a well-defined mean phase becomes evident. In polar coordinates (Figure 2m), phase vectors cluster symmetrically around the signal phase and tighten progressively with phase masking, providing a wrap-free representation that resolves ambiguities seen in linear phase plots. The mean phase is already stable in the unmasked data when estimated with a sufficiently large ensemble, indicating that the apparent disorder arises from variability rather than bias. The role of ensemble size in variance estimation is clarified in Figure 2k. When phase variance is computed using a large ensemble (1000 traces), it varies smoothly along offset, capturing the long-wavelength trend in phase stability common to both raw and masked data. Using a smaller ensemble (200 traces) reveals additional localized variability, including sharp spikes that highlight disrupted zones, likely associated with unsuppressed ground roll or other coherent noise. These local anomalies sit atop the same long-wavelength trend rather than contradicting it. After phase masking, the same relationship holds: local variability is reduced, while the underlying trend is preserved. Ensemble size therefore controls diagnostic resolution rather than physical behavior. A compact summary is provided by the average phase variance as a function of frequency (Figure 2l), computed using large local ensembles for statistical stability (near offsets). Phase variance decreases systematically across frequencies where signal is recoverable, quantifying improved temporal coherency. Where variance reduction stalls at the lowest and highest frequencies, the diagnostics identify unusable bands overwhelmed by noise. This distinction cannot be inferred from amplitude spectra, which conflate signal and noise and overestimate usable bandwidth. Figure 2n confirms that phase masking leaves amplitudes unchanged, whereas deconvolution typically boosts signal and noise alike. Phase-based diagnostics, supported by polar representations, provide a direct and meaningful indicator of where stabilization is effective.

Conclusions

We demonstrate that seismic phase variability can be analyzed and controlled in a fully data-driven manner using circular statistics. Phase coherency can be evaluated everywhere in the data, across all times, offsets, and traces, allowing every sample to be audited objectively. Circular mean and circular variance provide a direct diagnostic for seismic imaging and processing, independent of visual inspection or subjective judgment. Phase behavior is inherently statistical. Both phase masking and phase-variance analysis depend on ensemble size, but in complementary ways. The aperture used to stabilize the phase and the ensemble size used to quantify its variability are not competing parameters. They describe the same nested variability structure viewed at different levels of statistical resolution. Smaller ensembles provide localized but noisier estimates, larger ensembles provide smoother diagnostics, yet all reveal the same underlying trend. Once this structure is recognized, apparent contradictions disappear and aperture selection becomes deliberate rather than heuristic. This enables fit-for-purpose data conditioning. Phase masking can be applied only where needed and only to the degree required to support downstream tasks such as waveform inversion, velocity model building, statics estimation, deconvolution, or depth imaging. The objective is not to replace existing processing methods, but to compensate for what they cannot recover, namely the effects of unresolved small-scale heterogeneity and near-surface scattering. Stabilization is applied just enough to make subsequent algorithms reliable, without overconditioning or altering the underlying signal. Most importantly, phase stability becomes a measurable quantity rather than a judgment call. Parameter selection becomes objective. This logic can be encoded directly into processing and machine-learning workflows, eliminating gather-by-gather QC and enabling automated, scalable conditioning of large datasets. Each ensemble can be verified to be in the proper statistical state for the next processing step, improving reliability while significantly reducing turnaround time.

Acknowledgements

We thank TGS for providing the field data used in this study. The interpretations and conclusions are those of the authors.

References

Bakulin, A., Neklyudov, D., & Silvestrov, I. (2023). Seismic time-frequency masking for suppression of seismic speckle noise. Geophysics, **88**(5), V371–V385.

Bakulin, A., Neklyudov, D., & Silvestrov, I. (2024). The impact of receiver arrays on suppressing seismic speckle scattering noise caused by the meter-scale near-surface heterogeneity. Geophysics, **89**(6), V551–V561.

Bakulin, A., Neklyudov, D., and I. Silvestrov (2022). Multiplicative seismic noise caused by small-scale near-surface scattering and its transformation during stacking: Geophysics, **87**(5), V419-V435.

Holt, R., and Lubrano, A. (2020). Stabilizing the phase of onshore 3D seismic data. Geophysics, **85**(6), V473–V479.

Rohatgi, A., Bakulin, A., & Fomel, S. (2025). Data-driven analysis of seismic phase using circular statistics. The Leading Edge, **44**(9), 683–691.