


Divide and Infer

Semi-coherent methods for robust gravitational-wave analysis

 /perturber/SPLIT · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

Shubham Kejriwal
National University of Singapore




SPLIT

Global Fit II, NUS

Motivation: Precision Inference

- Extreme- and intermediate-mass-ratio inspirals will last in the LISA band for months to years.



Motivation: Precision Inference

- Extreme- and intermediate-mass-ratio inspirals will last in the LISA band for months to years.
- Waveform accuracy requirements for vacuum-GR signals of ~ 1 rad over the entire inspiral.

$$\Phi(t) = \frac{1}{\epsilon} \Phi^{(0)}(t) + \Phi^{(1)}(t) + \mathcal{O}(\epsilon)$$

Diagram illustrating the components of the waveform phase $\Phi(t)$ and their relative magnitudes:

- $\Phi^{(0)}(t)$ is associated with $\sim 10^5$ rads (indicated by an upward arrow).
- $\Phi^{(1)}(t)$ is associated with ~ 1 rad (indicated by a downward arrow).
- $\mathcal{O}(\epsilon)$ is associated with negligible (indicated by an upward arrow).

Motivation: Precision Inference

- Extreme- and intermediate-mass-ratio inspirals will last in the LISA band for months to years.
- Waveform accuracy requirements for vacuum-GR signals of ~ 1 rad over the entire inspiral.
- (Unmitigated) glitches and other non-stationarities in the data can induce systematic biases in EMRI/IMRI inference.


Mitigation Threshold	Canonical (Prograde)		Retrograde		Strong-field	
	$\max \beta^i $	Range of σ^i	$\max \beta^i $	Range of σ^i	$\max \beta^i $	Range of σ^i
SNR > 8	0.007	≈ 1	0.1	(1.0, 3.0)	0.01	≈ 1
SNR > 90	0.04	(1.05, 1.24)	0.6	(1.0, 13.0)	0.07	(1.23, 1.37)
SNR > 400	0.3	(1.9, 3.0)	0.4	(1.6, 44.0)	0.7	(3.0, 5.0)
No mitigation	1	(3.6, 8.5)	5	(4.1, 73.0)	3	(8.7, 12.6)

Biases and precision degradation in EMRI inference due to mitigated and unmitigated glitches in the LISA data (Boumerdassi et al. 2026)



Motivation: Precision Inference


- Extreme- and intermediate-mass-ratio inspirals will last in the LISA band for months to years.
- Waveform accuracy requirements for vacuum-GR signals of ~ 1 rad over the entire inspiral.
- (Unmitigated) glitches and other non-stationarities in the data can induce systematic biases in EMRI/IMRI inference.
- Beyond-vacuum-GR effects secularly and perturbatively modify the vacuum-GR evolution, inducing potentially large systematic biases during inference.



Impact of an unmodelled planetary migration effect on the measured vacuum-GR parameters (Kejriwal et al. 2024)

The Semi-coherent Approach

- Segment the data into shorter, contiguous blocks, each analyzed coherently and combined *incoherently*.




SPLIT

Semi-coherent Posteriors for
Long-Inspiral Trajectories

 [/perturber/SPLIT](https://github.com/perturber/SPLIT) · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

The Semi-coherent Approach

- Segment the data into shorter, contiguous blocks, each analyzed coherently and combined *incoherently*.
- Decompose the parameter space into “static” (m_1, m_2, a) and “evolving” (p, e, ι, Φ_a) parameters.




SPLIT

Semi-coherent Posteriors for
Long-Inspiral Trajectories

 [/perturber/SPLIT](https://github.com/perturber/SPLIT) · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

The Semi-coherent Approach

- Segment the data into shorter, contiguous blocks, each analyzed coherently and combined *incoherently*.
- Decompose the parameter space into “static” (m_1, m_2, a) and “evolving” (p, e, ι, Φ_a) parameters.
- Waveform accuracy requirements need only be satisfied over the much smaller block length.




SPLIT

Semi-coherent Posteriors for
Long-Inspiral Trajectories

 [/perturber/SPLIT](https://github.com/perturber/SPLIT) · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

The Semi-coherent Approach

- Segment the data into shorter, contiguous blocks, each analyzed coherently and combined *incoherently*.
- Decompose the parameter space into “static” (m_1, m_2, a) and “evolving” (p, e, ι, Φ_a) parameters.
- Waveform accuracy requirements need only be satisfied over the much smaller block length.
- Glitches and other instrumental artifacts are limited to individual blocks, do not bias the inference of other blocks.




SPLIT

Semi-coherent Posteriors for
Long-Inspiral Trajectories

 [/perturber/SPLIT](https://github.com/perturber/SPLIT) · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

The Semi-coherent Approach

- Segment the data into shorter, contiguous blocks, each analyzed coherently and combined *incoherently*.
- Decompose the parameter space into “static” (m_1, m_2, a) and “evolving” (p, e, ι, Φ_a) parameters.
- Waveform accuracy requirements need only be satisfied over the much smaller block length.
- Glitches and other instrumental artifacts are limited to individual blocks, do not bias the inference of other blocks.
- Secular perturbations from beyond-vacuum-GR effects do not build up significantly over individual blocks, allowing consistency of the block-level posteriors with the truth.



SPLIT

Semi-coherent Posteriors for
Long-Inspiral Trajectories

 [/perturber/SPLIT](https://github.com/perturber/SPLIT) · [arXiv:2606.13728](https://arxiv.org/abs/2606.13728)

Some novel techniques

- Heavy-tailed Student-t distribution-based likelihood for maximizing overlap in the presence of systematic biases in individual blocks.

The diagram illustrates the components of the likelihood function. A blue arrow labeled "semi-coherent likelihood" points to the overall expression. A blue arrow labeled "#blocks" points to the product index $i=0$ to $N_{\text{blocks}}-1$. Two blue arrows labeled "evolving set in i th block" and "non-evolving set in i th block" point to the parameters $\vec{\varphi}_i$ and $\vec{\vartheta}$ respectively within the block likelihood $p(\vec{d}_i | \vec{\varphi}_i, \vec{\vartheta})$.

$$p(\{\vec{d}_i\} | \{\vec{\varphi}_i\}, \vec{\vartheta}) = \prod_{i=0}^{N_{\text{blocks}}-1} p(\vec{d}_i | \vec{\varphi}_i, \vec{\vartheta})$$

$$p(\vec{d}_i | \vec{\varphi}_i, \vec{\vartheta}) \propto \exp \left(\frac{\nu_i + 2}{2} \sum_{f_k} \log \left[1 + \frac{\langle \vec{d}_i - \vec{h} | \vec{d}_i - \vec{h} \rangle_k}{\nu_i} \right] \right)$$

Motivated by C. Rover "A Student-t based filter for robust signal detection"(2011)

Some novel techniques

- Heavy-tailed Student-t distribution-based likelihood for maximizing overlap in the presence of systematic biases in individual blocks.
- Markovian Student-t prior for weakly-coherent trajectory priors across blocks (trade-off between degeneracies and robustness).

Markovian initial condition

$$p(\{\vec{\varphi}_i\}|\vec{\vartheta}) = \pi(\vec{\varphi}_0|\vec{\vartheta}) \prod_{i=1}^{N_{\text{blocks}}-1} \pi(\vec{\varphi}_i|\vec{\varphi}_{i-1}, \vec{\vartheta})$$

Transition probability priors are also Student-t

The diagram illustrates the components of the Markovian Student-t prior. A blue arrow points from the text 'Markovian initial condition' to the term $\pi(\vec{\varphi}_0|\vec{\vartheta})$ in the equation. Another blue arrow points from the text 'Transition probability priors are also Student-t' to the product term $\prod_{i=1}^{N_{\text{blocks}}-1} \pi(\vec{\varphi}_i|\vec{\varphi}_{i-1}, \vec{\vartheta})$.

Some novel techniques

- Heavy-tailed Student-t distribution-based likelihood for maximizing overlap in the presence of systematic biases in individual blocks.
- Markovian Student-t prior for weakly-coherent trajectory priors across blocks (trade-off between degeneracies and robustness).
- Sequential Gibbs sampling for “feasible” Bayesian inference: lowered dimensionality for healthy convergence.

$$D_{\text{infer}} = N_{\text{blocks}} \times D_{\text{evol}} + D_{\text{static}}$$

$$\vec{\varphi}_0 \rightarrow \vec{\varphi}_1 \rightarrow \dots \rightarrow \vec{\varphi}_{N_{\text{blocks}}-1} \rightarrow \vec{\vartheta}$$

Sequential Gibbs sampling framework in SPLIT: instead of inferring D_{infer} globally, we sequentially sample each block, conditioned on the current state in all other blocks.

Results: vacuum-GR (consistency)


- Year-long plunging, equatorial eccentric EMRI injection (PN5AAK) with SNR = 50.0 with $N_{blocks} = 5$, $D_{static} = 3$, and $D_{evol} = 4$.

Parameter (units)	Symbol	Value
Intrinsic static parameters		
Primary mass ($M_{\odot}$)	m_1	10^6
Secondary mass ($M_{\odot}$)	m_2	10
Primary spin (-)	a	0.8
Intrinsic evolving parameters (at $t = 0.0$)		
Semi-latus rectum (m_1)	p_0	7.447
Eccentricity (-)	e_0	0.2
Cosine of inclination (-)	x_{I0}	1.0
Initial azimuthal phase (rad)	$\Phi_{\varphi,0}$	1.0
Initial polar phase (rad)	$\Phi_{\theta,0}$	0.0
Initial radial phase (rad)	$\Phi_{r,0}$	1.0
Extrinsic static parameters		
Luminosity distance (Gpc)	d_L	3.251
Source polar angle (rad)	θ_S	$\pi/4$
Source azimuthal angle (rad)	ϕ_S	$\pi/3$
Spin polar angle (rad)	θ_K	$\pi/6$
Spin azimuthal angle (rad)	ϕ_K	$\pi/8$

T = 1.0 year vacuum-GR EMRI injection in Example I,
Kejriwal 2026

Results: vacuum-GR (consistency)


- Year-long plunging, equatorial eccentric EMRI injection (PN5AAK) with SNR = 50.0 with $N_{\text{blocks}} = 5$, $D_{\text{static}} = 3$, and $D_{\text{evol}} = 4$.
- 1D residuals consistent with the truth. Fractional loss of optimal SNR due to segmentation (but also windowing) leads to broader posteriors v/s FC.



1D residuals (vacuum-GR injection) SC v/s FC,
Kejriwal 2026

Results: vacuum-GR (consistency)

- Year-long plunging, equatorial eccentric EMRI injection (PN5AAK) with SNR = 50.0 with $N_{\text{blocks}} = 5$, $D_{\text{static}} = 3$, and $D_{\text{evol}} = 4$.
- 1D residuals consistent with the truth. Fractional loss of optimal SNR due to segmentation (but also windowing) leads to broader posteriors v/s FC.
- Full-dim posteriors show retention of correlations between evolving and static parameter sets; Phases take the biggest hit, as expected.



Full-dim posteriors (vacuum-GR injection) SC v/s FC,
Kejriwal 2026

Results: accretion (robustness)


- 0.5 year-long plunging, equatorial eccentric IMRI injection with SNR = 50.0 with $N_{blocks} = 5$, $D_{static} = 2$, and $D_{evol} = 4$.

Parameter (units)	Symbol	Value
Intrinsic static parameters		
Primary mass ($M_{\odot}$)	m_1	10^6
Secondary mass ($M_{\odot}$)	m_2	2×10^3
Primary spin (--)	a	0.0
Surface density (g/cm^2)	Σ_0	5.25×10^5
Aspect ratio (--)	h_0	0.025
Disk profile ()	Σ_p	-3/2
Intrinsic evolving parameters (at $t = 0.0$)		
Semi-latus rectum (m_1)	p_0	24.442
Eccentricity (--)	e_0	0.04
Cosine of inclination (· ·)	x_{I0}	1.0
Initial azimuthal phase (rad)	$\Phi_{\varphi,0}$	1.0
Initial polar phase (rad)	$\Phi_{\theta,0}$	0.0
Initial radial phase (rad)	$\Phi_{r,0}$	1.0
Extrinsic static parameters		
Luminosity distance (Gpc)	d_L	16.344
Source polar angle (rad)	θ_S	$\pi/4$
Source azimuthal angle (rad)	ϕ_S	$\pi/3$
Spin polar angle (rad)	θ_K	$\pi/6$
Spin azimuthal angle (rad)	ϕ_K	$\pi/8$

T = 0.5 year accretion IMRI injection in Example III,
Kejriwal 2026

Results: accretion (robustness)

- 0.5 year-long plunging, equatorial eccentric IMRI injection with $\text{SNR} = 50.0$ with $N_{\text{blocks}} = 5$, $D_{\text{static}} = 2$, and $D_{\text{evol}} = 4$.




Deviation of the environment-rich injection from a vacuum-GR template at the true parameters (based on Duque et al. 2025)

Results: accretion (robustness)

- 0.5 year-long plunging, equatorial eccentric IMRI injection with SNR = 50.0 with $N_{blocks} = 5$, $D_{static} = 2$, and $D_{evol} = 4$.
- Posterior coverage show (much) better consistency with the truth v/s the fully-coherent method.

Parameters	SC ($i = 0$)	SC ($i = 1$)	SC ($i = 2$)	SC ($i = 3$)	SC ($i = 4$)	FC
Mahalanobis distance						
(Full-dim)	0.646	0.693	0.784	0.858	0.639	1207.835
1D z-scores						
m_1	0.017	-	-	-	-	0.154
m_2	0.031	-	-	-	-	0.075
p_i	0.324	0.246	0.176	0.085	0.027	0.689
e_i	0.469	0.401	0.355	0.222	0.065	4.802
$\Phi_{\varphi,i}$	0.041	0.036	0.004	0.120	0.029	0.213
$\Phi_{r,i}$	0.071	0.015	0.019	0.032	0.039	0.116



Outlook

- Cost of SC inference ~order-of-magnitude higher than FC. SBI may help.
- Current pipeline fixes equal-length blocks. Not necessary in principle.
- Impact of instrumental artifacts currently unexplored. To change in the near future.
- Application beyond LISA possible. 3G will face many similar issues.
- Best applied as a step between search and final posterior inference(?) Ultimately, FC has the constraining power for high-precision astronomy promised by LISA EMRIs.



python `≥3.9` license `MIT` status `alpha`
DOI `10.5281/zenodo.20290209` docs `stable`


Dependencies

FastEMRIWaveforms (Chapman-Bird et al. 2026)
Eryn (Karnesis et al. 2023)



Optional Dependencies

fastlisaresponse (Katz et al. 2022)
LISAAalysisTools (Katz et al. 2025)

Bonus Slides



Results: stationary noise



Decorrelation timescale

- The semi-coherent likelihood assumes independence of noise between blocks.
- Noise is typically correlated across the entire observation window. However, it can be shown to be approx. uncorrelated given a decorrelation threshold.
- Checked for LWA and TDI-1st Gen A-channel PSDs. Found that the decorrelation timescales are small enough for any reasonable size of N_{blocks} .

$$p(\{\vec{d}_i\}|\{\vec{\varphi}_i\},\vec{\vartheta}) = \prod_{i=0}^{N_{blocks}-1} p(\vec{d}_i|\vec{\varphi}_i,\vec{\vartheta})$$

