




On the right track?

Batched trajectories and a semi-coherent likelihood for EMRI searches



Bert Depoorter (KU Leuven — Royal Observatory of Belgium)

Promotors: Henri Inchauspé (KU Leuven), Justin Janquart (Royal Observatory of Belgium)

Official promotor: Prof. Tjonnie G.F. Li (KU Leuven)







Extreme mass-ratio inspirals

$\sim 10^5$ orbits of a $\sim 10 M_\odot$ compact object into a $10^{5-7} M_\odot$ MBH

- ▶ Narrow, peaked posteriors: exceptional parameter constraints ($\Delta M/M \sim 10^{-5}$)
- ▶ Dynamics around MBH (galactic nuclei)
- ▶ Fundamental physics: precision tests of GR
- ▶ Environmental effects: discs, dark-matter, ...

Bottleneck: all of this science assumes the signal has been *found*, searching remains challenging.





Every advantage comes at a cost

Rich signals

10^5 cycles, multi-harmonic structure $\Rightarrow$
 10^{-5} fractional constraints

Prior-to-posterior volume ratio: coherent
matched filtering needs $\sim 10^{40}$ templates
(Gair et al. 2004)

$\Rightarrow$ brute-force coherent template bank is
impossible

Complicated waveforms

Hundreds of harmonics encode the
strong-field dynamics

FEW waveforms: ~ 0.1 s per evaluation
 $\Rightarrow$ any search strategy is limited by
waveform-throughput

Non-local degeneracies

Distant parameter degeneracies lead to
overlapping waveform harmonics

Non-local degeneracies in the intrinsic
space (Chua & Cutler 2022), samplers
get trapped

$\Rightarrow$ stochastic searches can't guarantee
coverage

The initial search must go beyond fully coherent matched filtering:

Which elements of the EMRI signal aid identification and can be leveraged across a large
part of the parameter space?



SNR localization




Figure: SNR² density of one Mojito-light EMRI (TDI A) in three bases

- ▶ **Time:** SNR slowly accumulates over the inspiral
- ▶ **Frequency:** smeared over neighbouring frequencies
- ▶ **Time–frequency:** each harmonic is a *thin, slowly-chirping track*, occupies $\sim 1/N_f$ of the pixels





Searching along time-frequency tracks

Tie the signal pixels in time frequency together along frequency tracks, using cheap semi-coherent methods.

Advantages of semi-coherent methods:

- ▶ **Dimensional reduction:** generating track requires (m_1, m_2, a, e_0, p_0) .
- ▶ **Cheap per system:** trajectory integration only, evaluate in batches
- ▶ **Robust to dephasing:** coherence demands phase accuracy to a fraction of a cycle over 10^5 – 10^6 cycles
- ▶ **Harmonics:** Generic Kerr orbits emit a combination set by $(f_\varphi, f_\theta, f_r)$.
- ▶ **Scaling:** statistics over the TF plane cost $\mathcal{O}(N_{\text{bins}} \times N_{\text{steps}})$





Figure: Track through SFT segments: predicted $f_{mn}(t)$ for multiple harmonics



EMRI dynamics

Adiabatic waveforms factorize into slowly evolving amplitudes and quickly evolving phases

$$h_+ - ih_\times = \frac{\mu}{d_L} \sum_{\ell m k n} A_{\ell m k n}(p(t), e(t), a) S_{\ell m k n}(\theta) e^{-i\Phi_{mkn}(t)}$$



- ▶ FEW solves flux-driven orbital elements: precomputed fluxes, interpolated.
- ▶ (p, e, Φ) are *smooth* functions of time: sparse solve + interpolant
- ▶ Tracks = phase derivatives, $f_{mn}(t) = \frac{1}{2\pi}(m\Omega_\varphi(t) + n\Omega_r(t))$





fewtrax: the FEW trajectory module in JAX

Reimplemented FEW trajectory module

- ▶ Same philosophy as FEW: Kerr eccentric equatorial OPA fluxes
- ▶ ODE solve in diffrax: end-to-end differentiable
- ▶ jit + vmap: fast and easy batch evaluation
- ▶ JAX advantage: CPU/GPU/TPU agnostic, easy to transfer.
- ▶ Forward & backward integration.

Just a tool, nothing new: what can we do with it?

fewtrax – FEW at FEW ODE nodes: $N = 200$ systems, $T \leq 2$ yr, $q \in [10^{-6}, 10^{-4}]$




Figure: fewtrax vs FEW at the FEW nodes.





Benchmarks: throughput, derivatives

Speed


- ▶ Single evaluation: comparable to FEW
- ▶ Batched (A100): $\sim 8 \mu\text{s}$ per trajectory ($\mathcal{O}(10^5)/\text{s}$)

End-to-end differentiable

- ▶ Exact derivatives through the ODE solve (diffrax adjoint)
- ▶ Hessians via jacfwd, batched with vmap.

Advantage: throughput, not single-call speed.

Trajectory	FEW (CPU)	fewtrax (CPU)	fewtrax (A100)
single call	50 ms	16 ms	200 ms
batched, per eval	—	3.2 ms	$8 \mu\text{s}$
throughput traj/s	20/s	300/s	$\mathcal{O}(10^5)/\text{s}$





A semi-coherent likelihood on track f_{20}

Per segment α : matched filter maximized over amplitude/phase [Speri et al., arXiv:2510.20891]

$$\mathcal{L}_\alpha = \frac{(\max_{\phi_\alpha} \langle \mathcal{D} | h_\alpha \rangle)^2}{\langle h_\alpha | h_\alpha \rangle}, \quad h_\alpha \sim A_\alpha e^{i\phi_\alpha + 2\pi i f_\alpha(t-t_\alpha) + \pi i \dot{f}_\alpha(t-t_\alpha)^2 + \dots}$$

Sum incoherently over N_s segments

$$\mathcal{L}(\vec{\theta}) = \sum_{\alpha=1}^{N_s} \mathcal{L}_\alpha, \quad \log p(\mathcal{D} | \vec{\theta}) = \frac{1}{2} \mathcal{L}(\vec{\theta}) + \text{const}$$

Search space: $\vec{\theta} = (m_1, m_2, a, e_f)$

- ▶ 4D, one extra dimension (inclination) for generic
- ▶ Backward integration: $(p_0, e_0) \mapsto (T_{\text{plunge}}, e_f)$

T_{plunge} : rolled over spectrogram, not searched.

Likelihood evaluation $\simeq$ trajectory solve.



Cost of likelihood evaluation

- ▶ Single jitted $\mathcal{L}(\vec{\theta})$: trajectory solve + Fresnel-kernel correlation over SFTs
- ▶ Batched with `vmap`: saturate the device

This likelihood is only marginally more expensive than the trajectory solve




Figure: Semi-coherent likelihood evaluation time vs batch size, CPU/GPU.





Trade-offs of semi-coherence

Cost: information

- ▶ Posterior broadening per parameter vs the coherent PE width
- ▶ SNR loss from dropping inter-segment phase coherence

Gain: runtime

- ▶ Semi-coherent $\mathcal{L}$ vs fully coherent matched-filter likelihood, per evaluation, CPU and GPU
- ▶ Batched on one GPU: orders of magnitude more parameter space per second

Tested on Mojito light dataset: mock data produced by DDPC.

Parameters: $m_1=2.7990e+06$, $m_2=15.394$, $a=0.9725$, $p_0=4.8621$, $e_0=0.1403$, $x_0=1.0$





Behaviour around Mojito-light EMRI 1




Figure: Corner-style grid scan of $\mathcal{L}$ over (m_1, m_2, a, e_f) .




Figure: Fully coherent PE posterior (MCMC) for intrinsic EMRI parameters.



Summary & next steps

Summary

- ▶ **Batched trajectories:** evaluates $\mathcal{O}(10^5)$ trajectories per second on one GPU (< 0.1 rad over 2 yr), fully differentiable
- ▶ Semi-coherent likelihood inherits trajectory throughput
- ▶ Enables large-scale exploration of EMRI parameter space.

Next steps

- ▶ Study semi-coherent likelihood in more detail: bias, inclusion of higher harmonics, scaling with T_{sft} .
- ▶ Inverse problem: fit a phenomenological track, then gradient-based global exploration / local optimization over the intrinsic space.
- ▶ Extend to full EMRI waveform in JAX.

Work in progress: Open source <https://github.com/BertDepoorter/fewtrax>

