Kerr Linear Inversion in `bayRing`

Kerr-only model. For a Kerr ringdown without tail terms or quadratic amplitudes, and with the non-amplitude Kerr quantities fixed, the complex strain is linear in the QNM amplitudes:

$$h(t_j) = h_{\mathrm{fixed}}(t_j) + \sum_{k=1}^{N} C_k \, q_k(t_j), \qquad C_k = A_k e^{i\phi_k}. \tag{1}$$

Here $q_k(t)$ is the Kerr basis waveform for mode $k = (\ell, m, n)$, including the fixed Kerr frequency and damping time. The term h_{fixed} contains any constant offset or already-fixed mode amplitudes. In the simplest `QNM-modes = 220` example, the only unknown is $C_{220} = A_{220} e^{i\phi_{220}}$.

Because the model is linear in the complex amplitudes C_k, we do not need to do stochastic inference or non-linear minimization in this case. Once the non-amplitude quantities are fixed, the amplitudes follow from one deterministic linear-algebra solve.

Real weighted system. The code solves for real and imaginary amplitude parts,

$$x = (\mathrm{Re}\, C_1, \mathrm{Im}\, C_1, \ldots, \mathrm{Re}\, C_N, \mathrm{Im}\, C_N)^T. \tag{2}$$

After subtracting h_{fixed}, the real and imaginary data components are stacked and divided by their corresponding errors. The same weighting is applied to the basis columns, giving a real design matrix B and data vector d:

$$d \simeq Bx. \tag{3}$$

For each Kerr mode, B has two columns: the weighted waveform generated by $C_k = 1$ and by $C_k = i$, where i is the complex unit, $i^2 = -1$.

Closed-form solve. The weighted least-squares normal equations are

$$Fx = r, \qquad F = B^T B, \qquad r = B^T d. \tag{4}$$

The implementation symmetrizes F, diagonalizes it, and floors small eigenvalues using the configured tolerance ϵ:

$$F = V \Lambda V^T, \qquad \hat{x} = V \ \mathrm{diag}\left[\frac{1}{\max(\lambda_i, \epsilon)}\right] V^T r. \tag{5}$$

Each recovered complex amplitude $\hat{C}_k = \hat{x}_{2k-1} + i\hat{x}_{2k}$ is then written back in the `bayRing` parameterization,

$$\ln A_k = \log |\hat{C}_k|, \qquad \phi_k = \arg(\hat{C}_k) \bmod 2\pi. \tag{6}$$

The corresponding one-sigma error estimate is obtained from the same weighted Fisher matrix. In Cartesian amplitude coordinates the covariance is

$$\Sigma_x = V \ \mathrm{diag}\left[\frac{1}{\max(\lambda_i, \epsilon)}\right] V^T. \tag{7}$$

The code then propagates Σ_x to the $(\ln A_k, \phi_k)$ parameterization using the local Jacobian of the polar transformation.

Configuration. Use `method = Linear-inversion`. For the Kerr-only setup discussed here, leave each $\ln A$, ϕ amplitude pair free, and fix or omit all nonlinear parameters. The eigenvalue floor is set by the configured linear inversion tolerance. A minimal example is

```
config_files/config_SXS_0305_Kerr_220_
linear_inversion.ini
```