

Notes: Online Fourier Transforms

M. Rule

July 27, 2011

1 The online Fourier transform

The online Fourier transform tracks a sliding window of the recent history of the time series. This window is updated by advancing one time step, and the FFT is used to get the power spectrum.

One can speed up computation by making updates recursive: on each time step, first subtract the contribution of the oldest timepoint, then add the contribution from the current timepoint. We can generalize this recursive approach by interpreting the online FT as a bank of band-pass filters.

2 A filter bank of damped, driven resonators

For each frequency bin, one constructs a pair of band-pass filters covering this bin, with their phases separated by 90° . The activation of each filter pair defines a vector, and the length of that vector is the amplitude for that frequency bin.

In this example, we consider a filter bank composed of first-order filters. This corresponds to a system of damped, driven oscillators with resonant frequencies equal to each frequency component. This is equivalent to estimating an online Fourier transform, where the signal is multiplied by an exponential window. Other window functions can be approximated by constructing higher-order filters with the desired impulse-response envelope.

This is fast, and easy to compute on e.g. microcontrollers with limited processing capability. Incrementally updating K filters costs $\mathcal{O}(K)$ time on each time-step. Contrast this with a windowed FFT, which would require $\mathcal{O}(K \log(K))$ computations per time step.

Let $X(t)$ be a real valued discrete time series with time step Δt . The component $z_\omega = r e^{i\phi}$, with magnitude r and phase ϕ at frequency $\omega = 2\pi f$, can be extracted by updating z on each timestep as

$$z_{\omega,t} = e^{i\omega\Delta t}(1 - \alpha)z_{\omega,t-1} + \alpha X(t)$$

The parameter $\alpha \in [0, 1]$, $\alpha = \Delta t/\tau$ determines the decay constant for the exponential envelope, and can be adjusted independently for each channel. Larger α allow the filter to respond faster, but give a worse estimate of the power. Smaller α integrate information over more time, but react more slowly to changes (Fig. 1, left).

3 State-space model interpretation

Inspired by the Kalman filter, we can view the components z_ω as a state space model for an underlying oscillatory process. We can use this model to predict future time points, and then compare our prediction to the observed value. It is then possible to update our state using the prediction error, rather than simply averaging in new state. This gives a more accurate estimate of the power spectrum (Fig 2, right).

The current state is predicted by advancing the phase of all frequency coefficients, and summing their real parts:

$$\hat{x}_t = \sum_{\omega} \Re(z_{\omega,t-1} e^{i\omega\Delta t})$$

Our error is a simple difference, and we attribute that error equally across all K coefficients

$$z_{\omega,t} = z_{\omega,t-1} e^{i\omega\Delta t} + (x - \hat{x})/K$$

This approach is best suited for stationary signals. The infinite memory means that it converges toward the true power spectrum, but also makes it respond slower to changes in the signal statistics.




Figure 1: **Figure 1: Frequency spectrum (magnitude) plots first-order filter bank and state-space online FT methods** Each plot show the estimated frequency spectrum from a signal composed of a sum of sine waves at 5, 10, 20, 40, 50, 100, 150, 200, 250, 300, and 350 Hz after two seconds of estimation. **Left:** estimated using the resonator model, with time constants equal to twenty times the wavelength for each frequency component. **Right:** estimated using the state space model. The state space model has converged to a much sharper estimate of the underlying frequency composition of the signal.