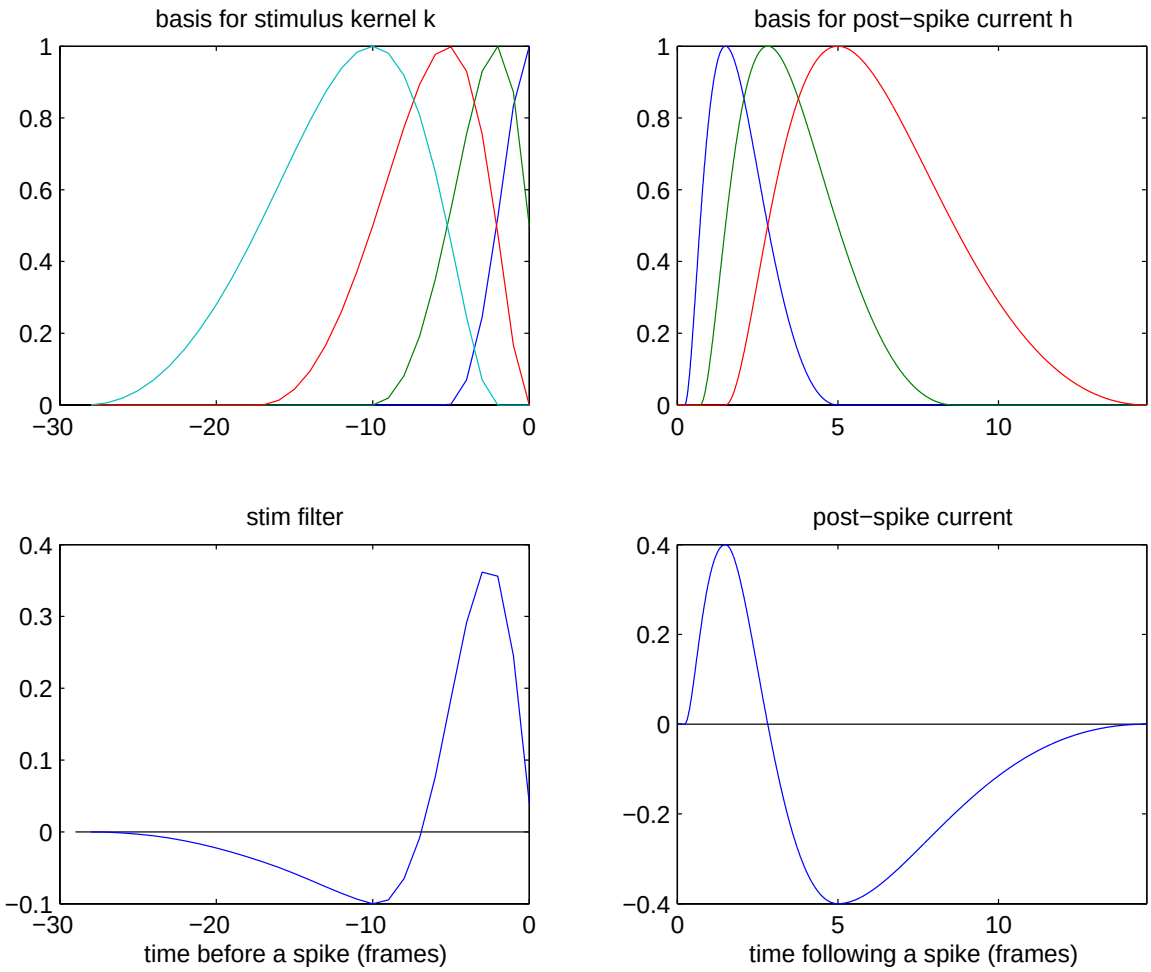

Figure 1: Parameters used for simulating a generalized integrate-and-fire (gif) model. Top row shows the bases used to describe the stimulus kernel k (left) and post-spike kernel h (right), both consisting of raised cosine "bumps" with a log-scaling of the time axis (allowing the representation of fine temporal structure near the time of a spike, and smoother structure further away). The bi-phasic h induces burstiness in the model's response. The full model is here specified by 10 parameters: k (4 params), h (3 params), σ_{noise}, V_{leak}, and g_{leak} (parametrized in the code as $\tau = 1/g_{leak}$, the membrane time constant).

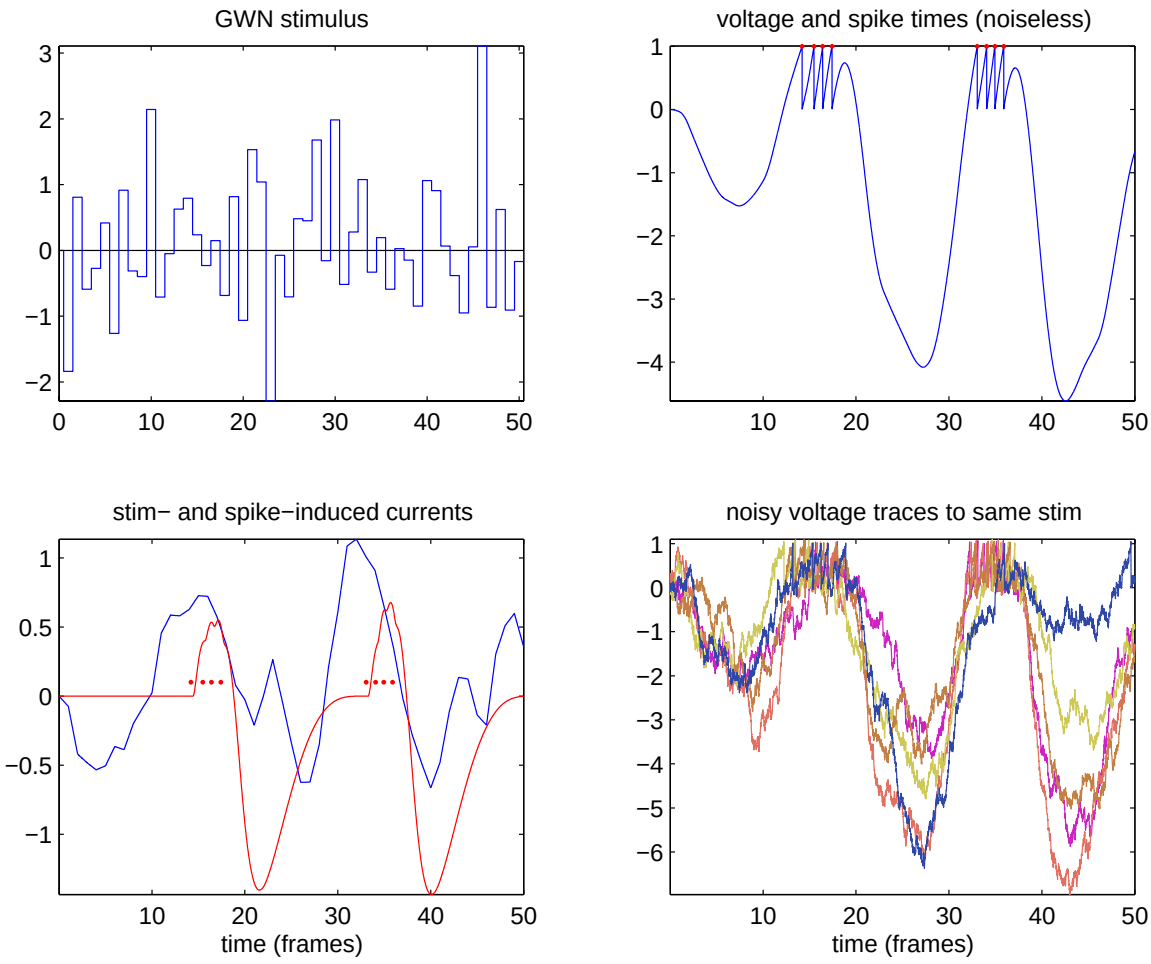

Figure 2: GIF model simulation. Upper-left: a 50-frame temporal Gaussian white noise stimulus. Lower-left: Stimulus-induced current (blue) and spike history induced current (red); dots indicate spike times. Total current = stimulus current + spike-history current + noise current. Upper-right: voltage trace and spike times (red dots) when simulated with noise amplitude set to zero. Lower-right: five noisy voltage traces drawn using the same stimulus current, with $\sigma_{noise} = 0.5$.

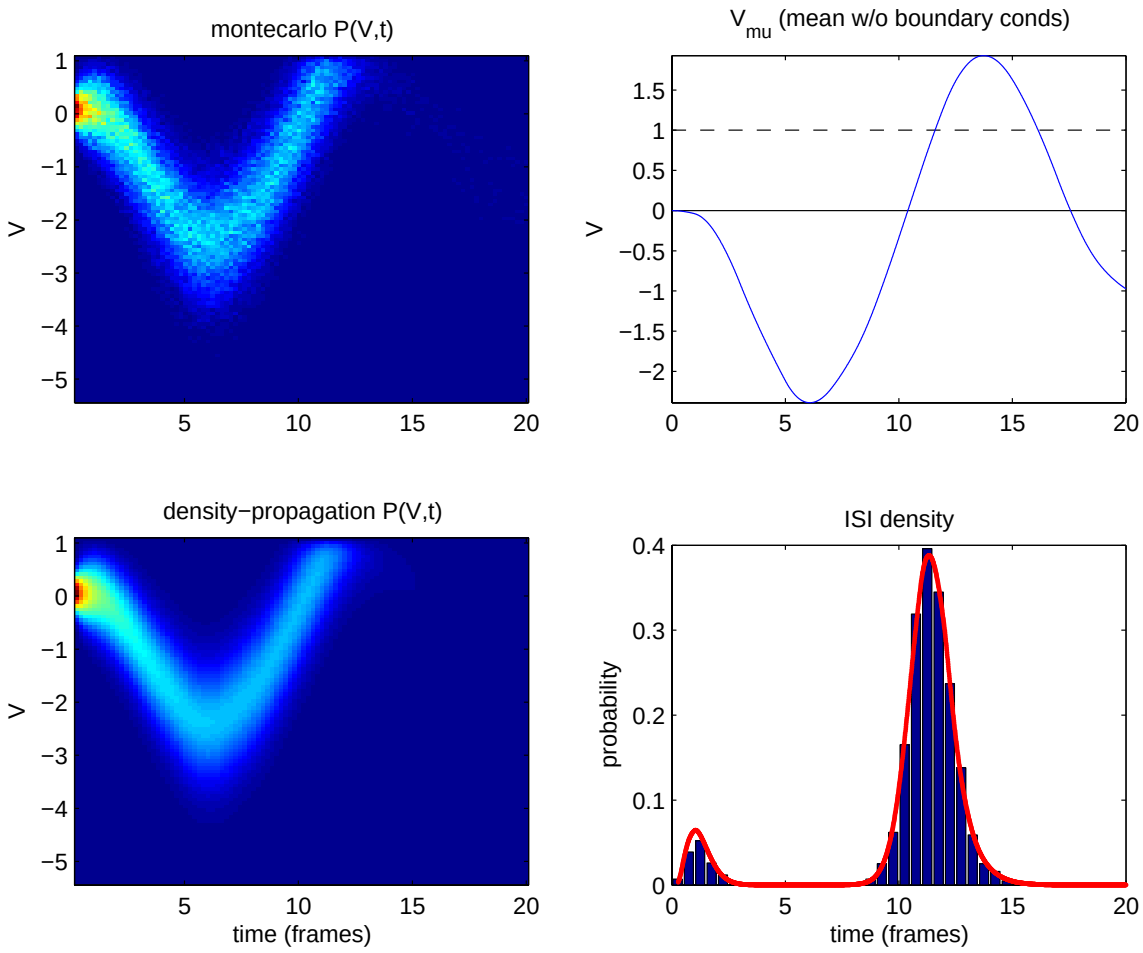

Figure 3: Likelihood validation. Upper-left: Probability density $P(V,t)$ over voltage as a function of time, computed by sampling noisy voltage trajectories from the model and then taking a histogram for every moment in time. Each vertical slice gives the distribution $P(V)$ at a fixed time t since the previous spike, representing our uncertainty about V given the injected current, the previous spike time t_0, and the fact that no spike has occured on $(t_0, t]$. Upper-right: noiseless voltage trajectory $V_\mu(t)$ for this same stimulus, which is also gives the mean of the Gaussian density on sub-threshold V the obtains in the absence of boundary conditions. Lower-right: interspike-interval distribution for the first spike after time 0, computed with monte-carlo samples (blue) and density-propagation (red). Also known as the density of first-passage times.

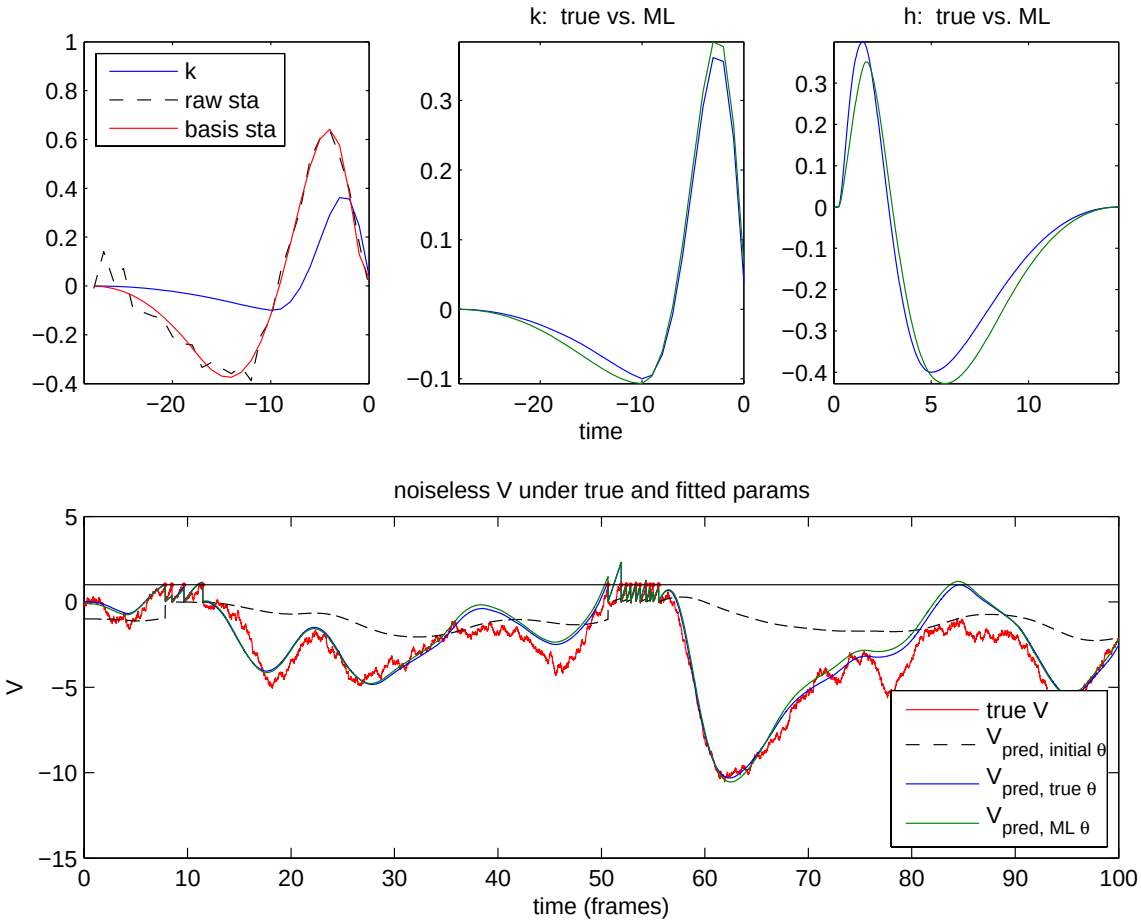

Figure 4: Maximum-likelihood estimation of the GIF parameters, using a 2000-sample stimulus (eliciting 380 spikes), which takes ≈ 5 min. A scaled version of the STA (upper-left) was used as an initial guess of k, while h was set to zero and other parameters were initialized with arbitrary values. The ML estimates (green) roughly match the true parameters (blue) for both k and h (upper middle and upper right, respectively), though accuracy improves with more data. Bottom: true voltage response to the first 100 frames of the stimulus (red trace; dots indicate spike times), and noiseless voltage traces generated by different settings of the model parameters: initialization for fitting (dashed), true parameters (blue) and ML estimates (green).