We thank the referees for their time and the kind reviews. Brief responses follow.

Model selection
We explored different values of the parameter K. The value $K = 3$ achieves robust performance in both training and test data, and is interpretable biophysically, and so we focused our attention on $K = 3$ here. We found that smaller values of K led to worse performance, and higher values of K could lead to unstable learning. We plan to present further details of these analyses in an appendix to the final paper.

Real data
We are currently applying these methods to real voltage imaging data; preliminary results are encouraging. This work will most likely be described in a separate paper. We also emphasize that many experiments in our submission do use real voltage traces (corrupted with artificial noise, Figures 2+3) and neuron morphologies (with simulated voltage traces, Figure 5), allowing us to assess recovery of ground truth voltage via semi-synthetic data.

Quantitative comparison between the Kalman smoother and rSLDS
We have included this comparison in the right panel of Figure 3; we will clarify this point in the revised text.

Interpretation of $X^{(n)}$
These are auxiliary continuous latent variables that the rSLDS uses to model the voltage dynamics. Intuitively, the second dimension helps determine whether the voltage is rising or falling, which is an important signal for the discrete state transition probabilities in the rSLDS. We will clarify this in the revised text.

Other issues
We will upload the code (as suggested by R2) and fix typos; thank you for pointing these out. We will add material in the appendix clarifying the scalability of the method (as R1 suggested). We will introduce the biophysical meaning of the discrete latent variable earlier in the paper (as R1 suggested) and provide discussions about "slight spatial correlation errors" (as R3 suggested). Thanks again for these helpful suggestions.