A. Supplemental Material

A.1. Factorial hidden Markov model

An example inference result from the factorial hidden Markov model is shown in Figure 7. The algorithm successfully recovers differing interpretations for the same recorded energy usage data.

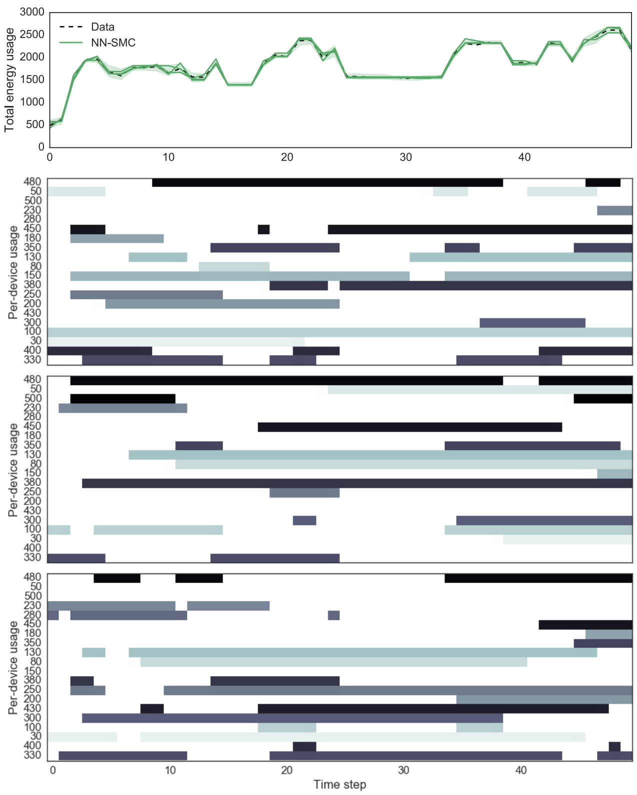

Figure 7. Scenario analysis for a synthetic additive factorial HMM example, from a single SMC sweep with 500 particles. In all plots, the horizontal axis denotes time in the state space. The top plot shows the reconstructed signal $y_{1:T}$, with two standard deviations around the mean shown as a light-green band, and three individual scenarios called out as individual lines. These scenarios are shown in detail in the separate plots below: each row represents the energy usage of a particular device, with darker colors showing higher mean energy usage; white indicates the device is off. The green lines in the top plot are recovered by summing vertically across the rows of each individual scenario. The very different recovered device activities $x_{1:T}^{1:D}$ yield output signals $y_{1:T}$ which are indistinguishable up to noise.

A.2. Training the neural network

The training procedure for each epoch, using synthetic training and validation data, proceeds as follows:

1. Sample a synthetic dataset $\{\mathbf{x}_\ell, \mathbf{y}_\ell\}_{\ell=1}^{N_{train}}$ and a validation set $\{\mathbf{x}_\ell, \mathbf{y}_\ell\}_{\ell=1}^{N_{validate}}$
2. Compute initial validation error, and loop:
   (a) Perform a mini-batch gradient update on η, from the synthetic dataset
   (b) Compute a new validation error on the sampled validation set
   (c) Continue until validation error increases, or until a set maximum number of steps is reached.