Beyond the DP mixture model setting, split-merge MCMC moves are not well studied. Both Meeds et al. (2006) and Mørup, Schmidt and Hansen (2011) mention adapting an RG procedure for relational models with latent features based on the beta process. However, neither work provides details on constructing proposals, and both lack experimental validation that split-merge moves improve inference.

7.2. *Split-merge MCMC for the BP-AR-HMM.* In standard mixture models, such as considered by Jain and Neal (2004), a given data item i is associated with a single cluster k_i, so selecting two anchors i and j is equivalent to selecting two cluster indices k_i, k_j. However, in feature-based models such as the BP-AR-HMM, each data item i possesses a *collection* of features indicated by $\mathbf{f}_i$. Therefore, our split-merge requires a mechanism not only for selecting anchors, but also for choosing candidate features to split or merge from $\mathbf{f}_i, \mathbf{f}_j$. After proposing modified feature vectors, the associated state sequences must also be updated. Following the motivations for our data-driven birth–death proposals, our split-merge proposals create new feature matrices $\mathbf{F}^*$ and state sequences $\mathbf{z}^*$, collapsing away HMM parameters $\boldsymbol{\theta}, \boldsymbol{\eta}$. Figure 4 illustrates $\mathbf{F}$ and $\mathbf{z}$ before and after a split proposal. Motivated by the efficiencies of sequential allocation [Dahl (2005)], we adopt a sequential approach. Although a RG approach that samples all variables $(\mathbf{F}, \mathbf{z}, \boldsymbol{\theta}, \boldsymbol{\eta})$ is also possible and relatively straightforward, our experiments [Supplement I of Fox et al. (2014)] show that our sequential collapsed proposals are vastly preferred. Intuitively, constructing high acceptance rate proposals for $\boldsymbol{\theta}, \boldsymbol{\eta}$ can be very difficult since each behavior-specific parameter is high dimensional.

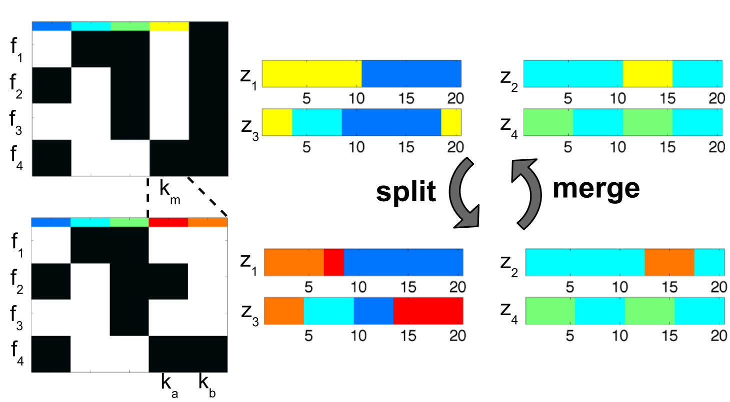

FIG. 4. *Illustration of split-merge moves for the BP-AR-HMM, which alter binary feature matrix* $\mathbf{F}$ (*white indicates present feature*) *and state sequences* $\mathbf{z}$. *We show* $\mathbf{F}, \mathbf{z}$ *before* (top) *and after* (bottom) *feature* k_m (yellow) *is split into* k_a, k_b (red, orange). *An item possessing feature* k_m *can have either* k_a, k_b, *or both after the split, and its new* $\mathbf{z}$ *sequence is entirely resampled using any features available in* $\mathbf{f}_i$. *An item without* k_m *cannot possess* k_a, k_b, *and its* $\mathbf{z}$ *does not change. Note that a split move can always be reversed by a merge.*