Exercise: EM for for HMMs with tied mixtures

In many applications, it is common that the observations are high-dimensional vectors (e.g., in speech recognition, $\mathbf{x}_t$ is often a vector of cepstral coefficients and their derivatives, so $\mathbf{x}_t \in \mathbb{R}^{39}$), so estimating a full covariance matrix for KM values (where M is the number of mixture components per hidden state), as in Exercise **??**, requires a lot of data. An alternative is to use just M Gaussians, rather than MK Gaussians, and to let the state influence the mixing weights but not the means and covariances. This is called a **semi-continuous HMM** or **tied-mixture HMM**.

- Draw the corresponding graphical model.
- Derive the E step.
- Derive the M step.