Collapsed Gibbs Sampling for LDA

- **Key idea:** construct a well-behaved Markov chain such that
  1. the states of the chain represent an assignment of $\boldsymbol{Z}$
  2. state transitions occur between states that differ in only one $z_{j,t}$
  3. transition probabilities are based on the **full conditional**:

     $$P\big(z_{j,t} = k \mid \boldsymbol{Z}_{\neg j,t}, \boldsymbol{W}, \alpha, \beta \big) = \frac{P(z_{j,t}=k, \boldsymbol{Z}_{\neg j,t}, \boldsymbol{W} | \alpha, \beta)}{P(\boldsymbol{Z}_{\neg j,t}, \boldsymbol{W} | \alpha, \beta)} \propto P\big(z_{j,t} = k, Z_{\neg j,t}, \boldsymbol{W} \mid \alpha, \beta \big)$$

     where $\boldsymbol{Z}_{\neg j,t}$ is the set of assignments for $\boldsymbol{Z}$ without position (j, t)

- **Guaranteed** to produce **true samples** from the posterior!
  IF you run it for long enough…