Collapsed Gibbs Sampling: (Non-exhaustive) Pros/Cons

- Pros:
  - Ease of implementation
  - Fast iterations (no transcendental functions), fast convergence (at least relative to a full Gibbs sampler)
  - Low memory usage (only require $O(MN)$ storage for the current values of $z_{j,t}$)

- Cons:
  - No obvious parallelization strategy (each iteration depends on previous)
  - Can be difficult to assess convergence
    - "Variational inference is that thing you implement while waiting for your Gibbs sampler to converge." – David Blei