Discussion of Alex Smola's Talk:

Remarks on parallelised MCMC
Questions to ponder

Yee Whye Teh
Gatsby Computational Neuroscience Unit, UCL

Parallelising MCMC

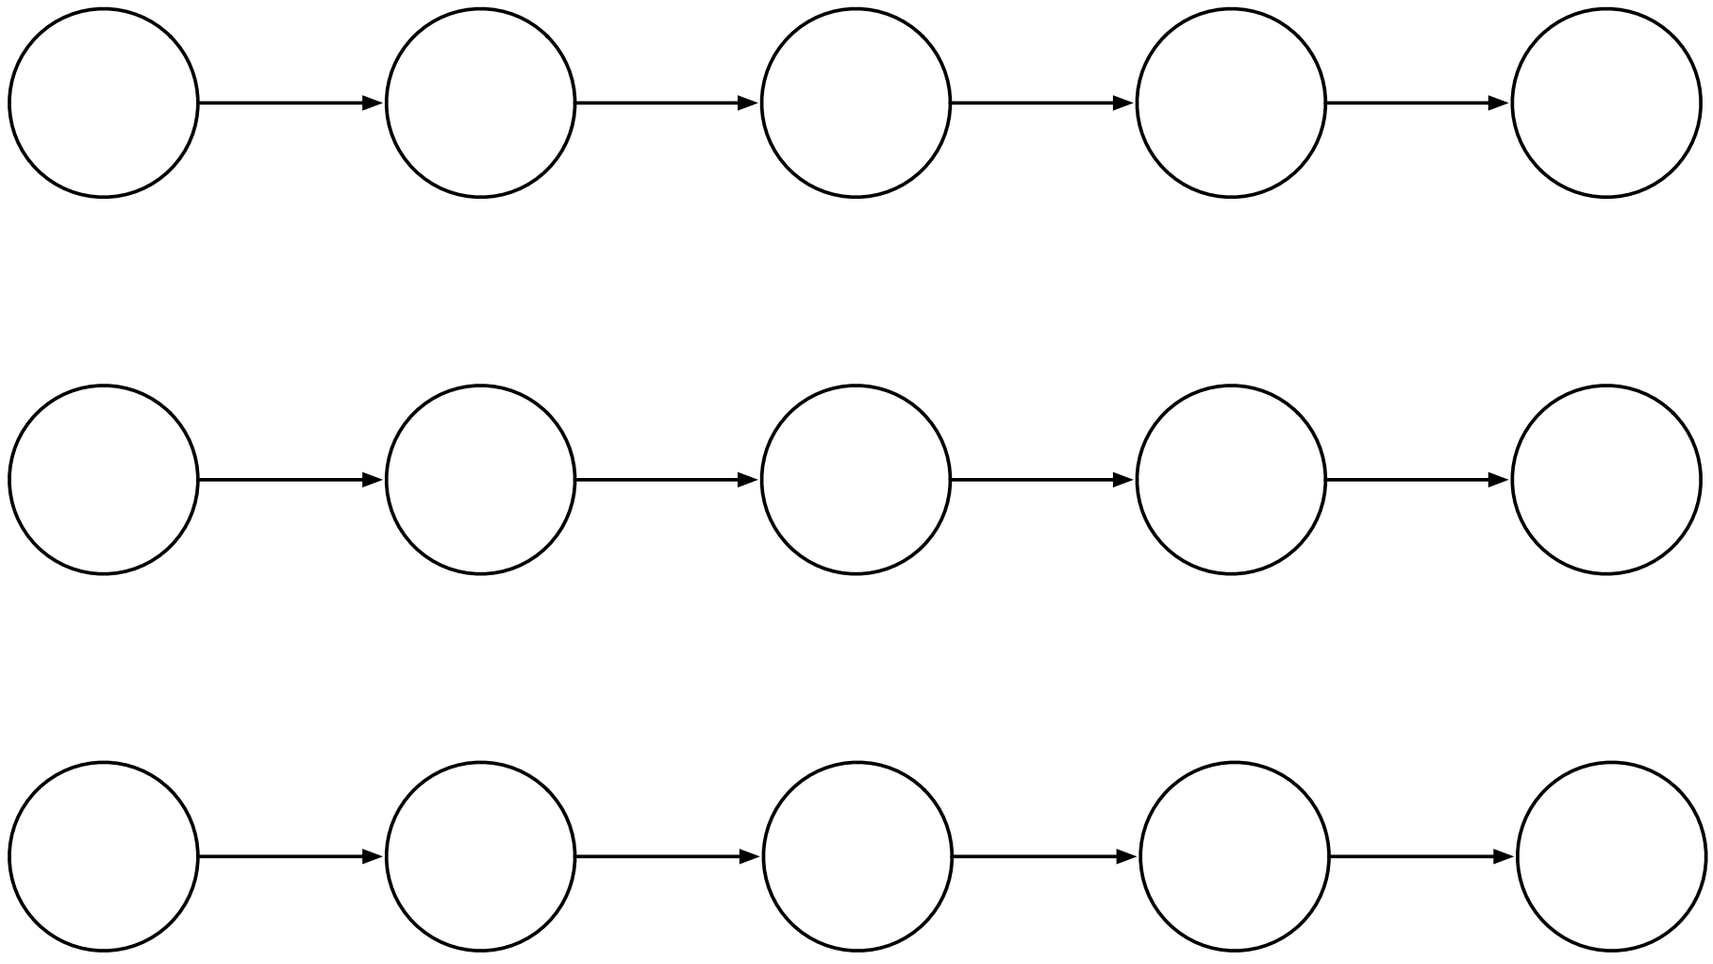

Parallelising MCMC

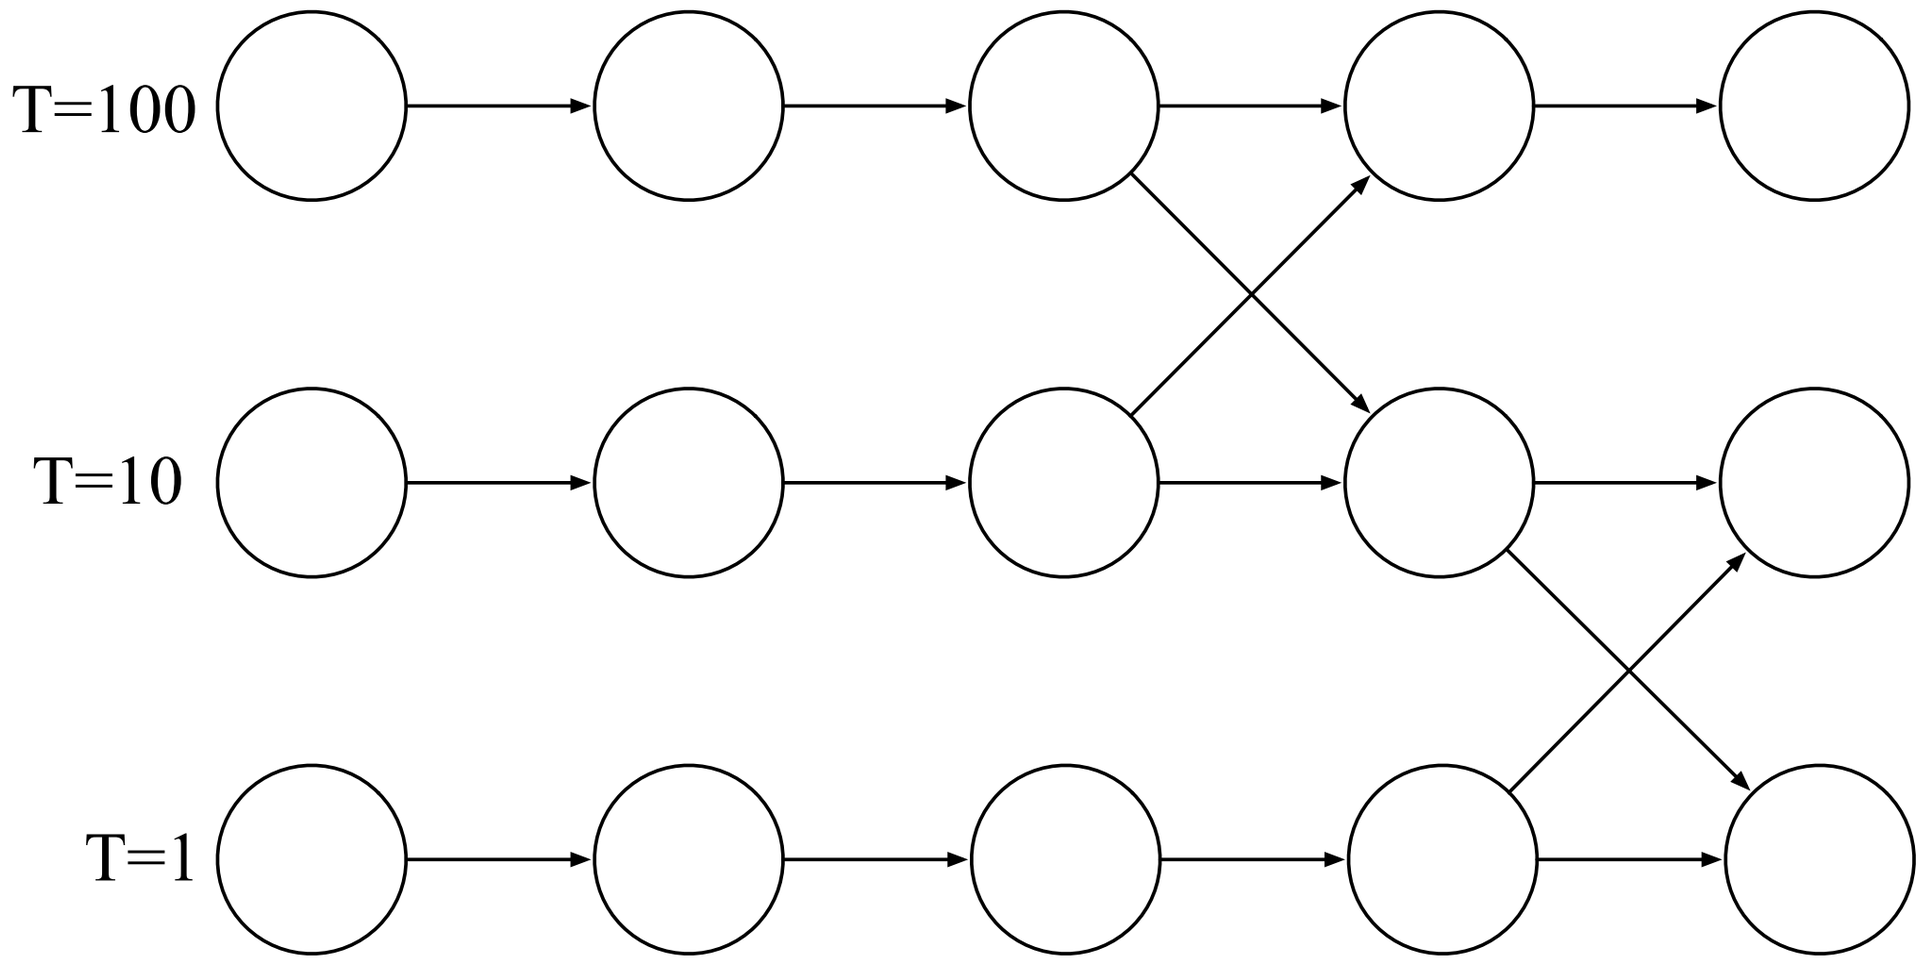

Parallelising MCMC

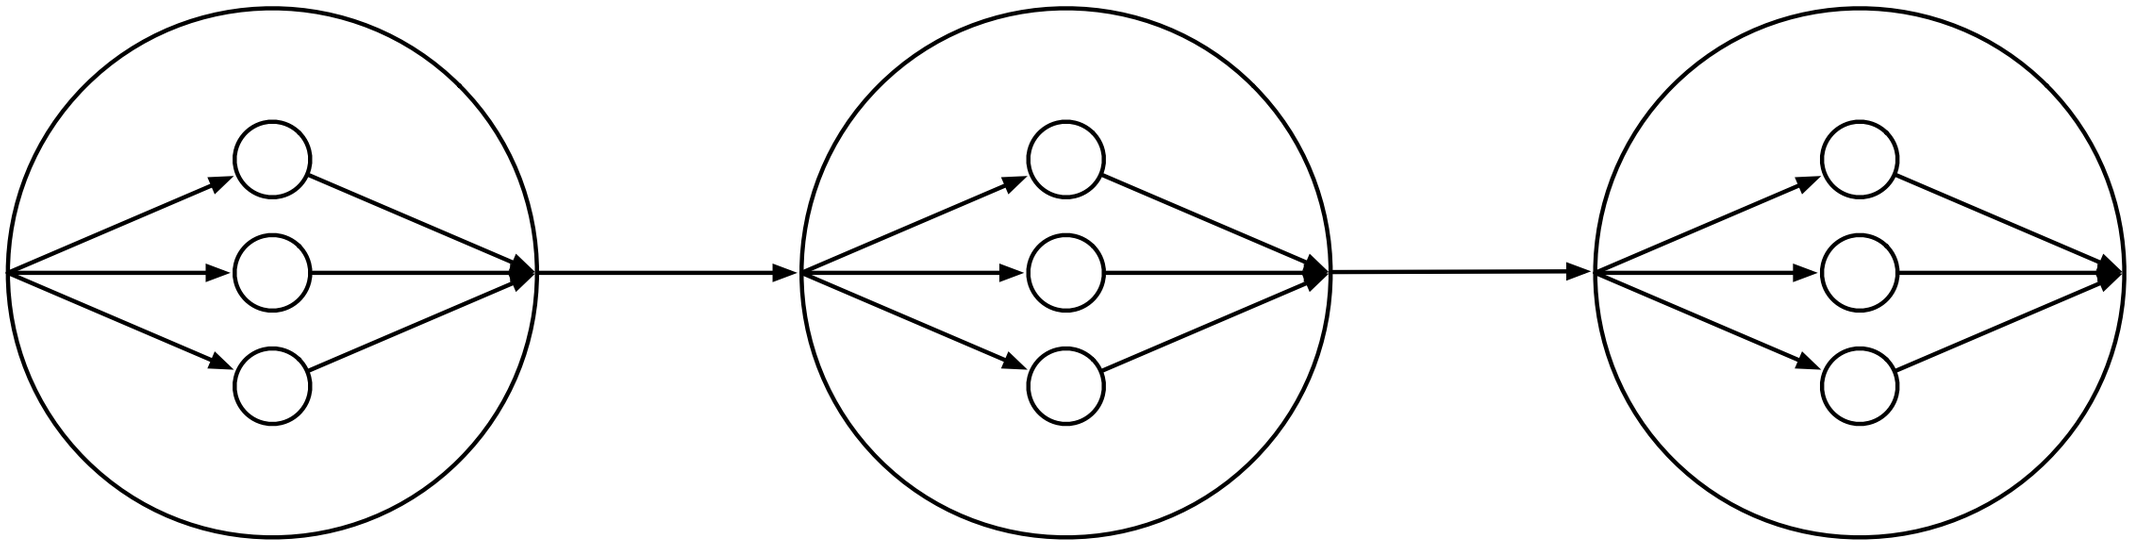

Circularly-coupled MCMC

Circularly-coupled MCMC

- $X_1 \sim P_0$
- $Z_1 \sim P_\infty$
- For i=2,3,...T
  - $U_i \sim$ Uniform[0,1]
  - $X_i = F(X_{i-1}, U_i)$
  - $Z_i = F(Z_{i-1}, U_i)$

Circularly-coupled MCMC

- $X_1 \sim P_0$
- $Z_1 \sim P_\infty$
- For i=2,3,...T
  - $U_i \sim$ Uniform[0,1]
  - $X_i = F(X_{i-1}, U_i)$
  - $Z_i = F(Z_{i-1}, U_i)$

- If $Z_T = X_T$, then Markov chain has "probably converged".

Circularly-coupled MCMC

- $X_1 \sim P_0$
- $Z_1 \sim P_\infty$
- For i=2,3,...T
  - $U_i \sim$ Uniform[0,1]
  - $X_i = F(X_{i-1}, U_i)$
  - $Z_i = F(Z_{i-1}, U_i)$

- If $Z_T = X_T$, then Markov chain has "probably converged".

- Use $Z_1 = X_T$ as a proxy for P_∞.

Circularly-coupled MCMC

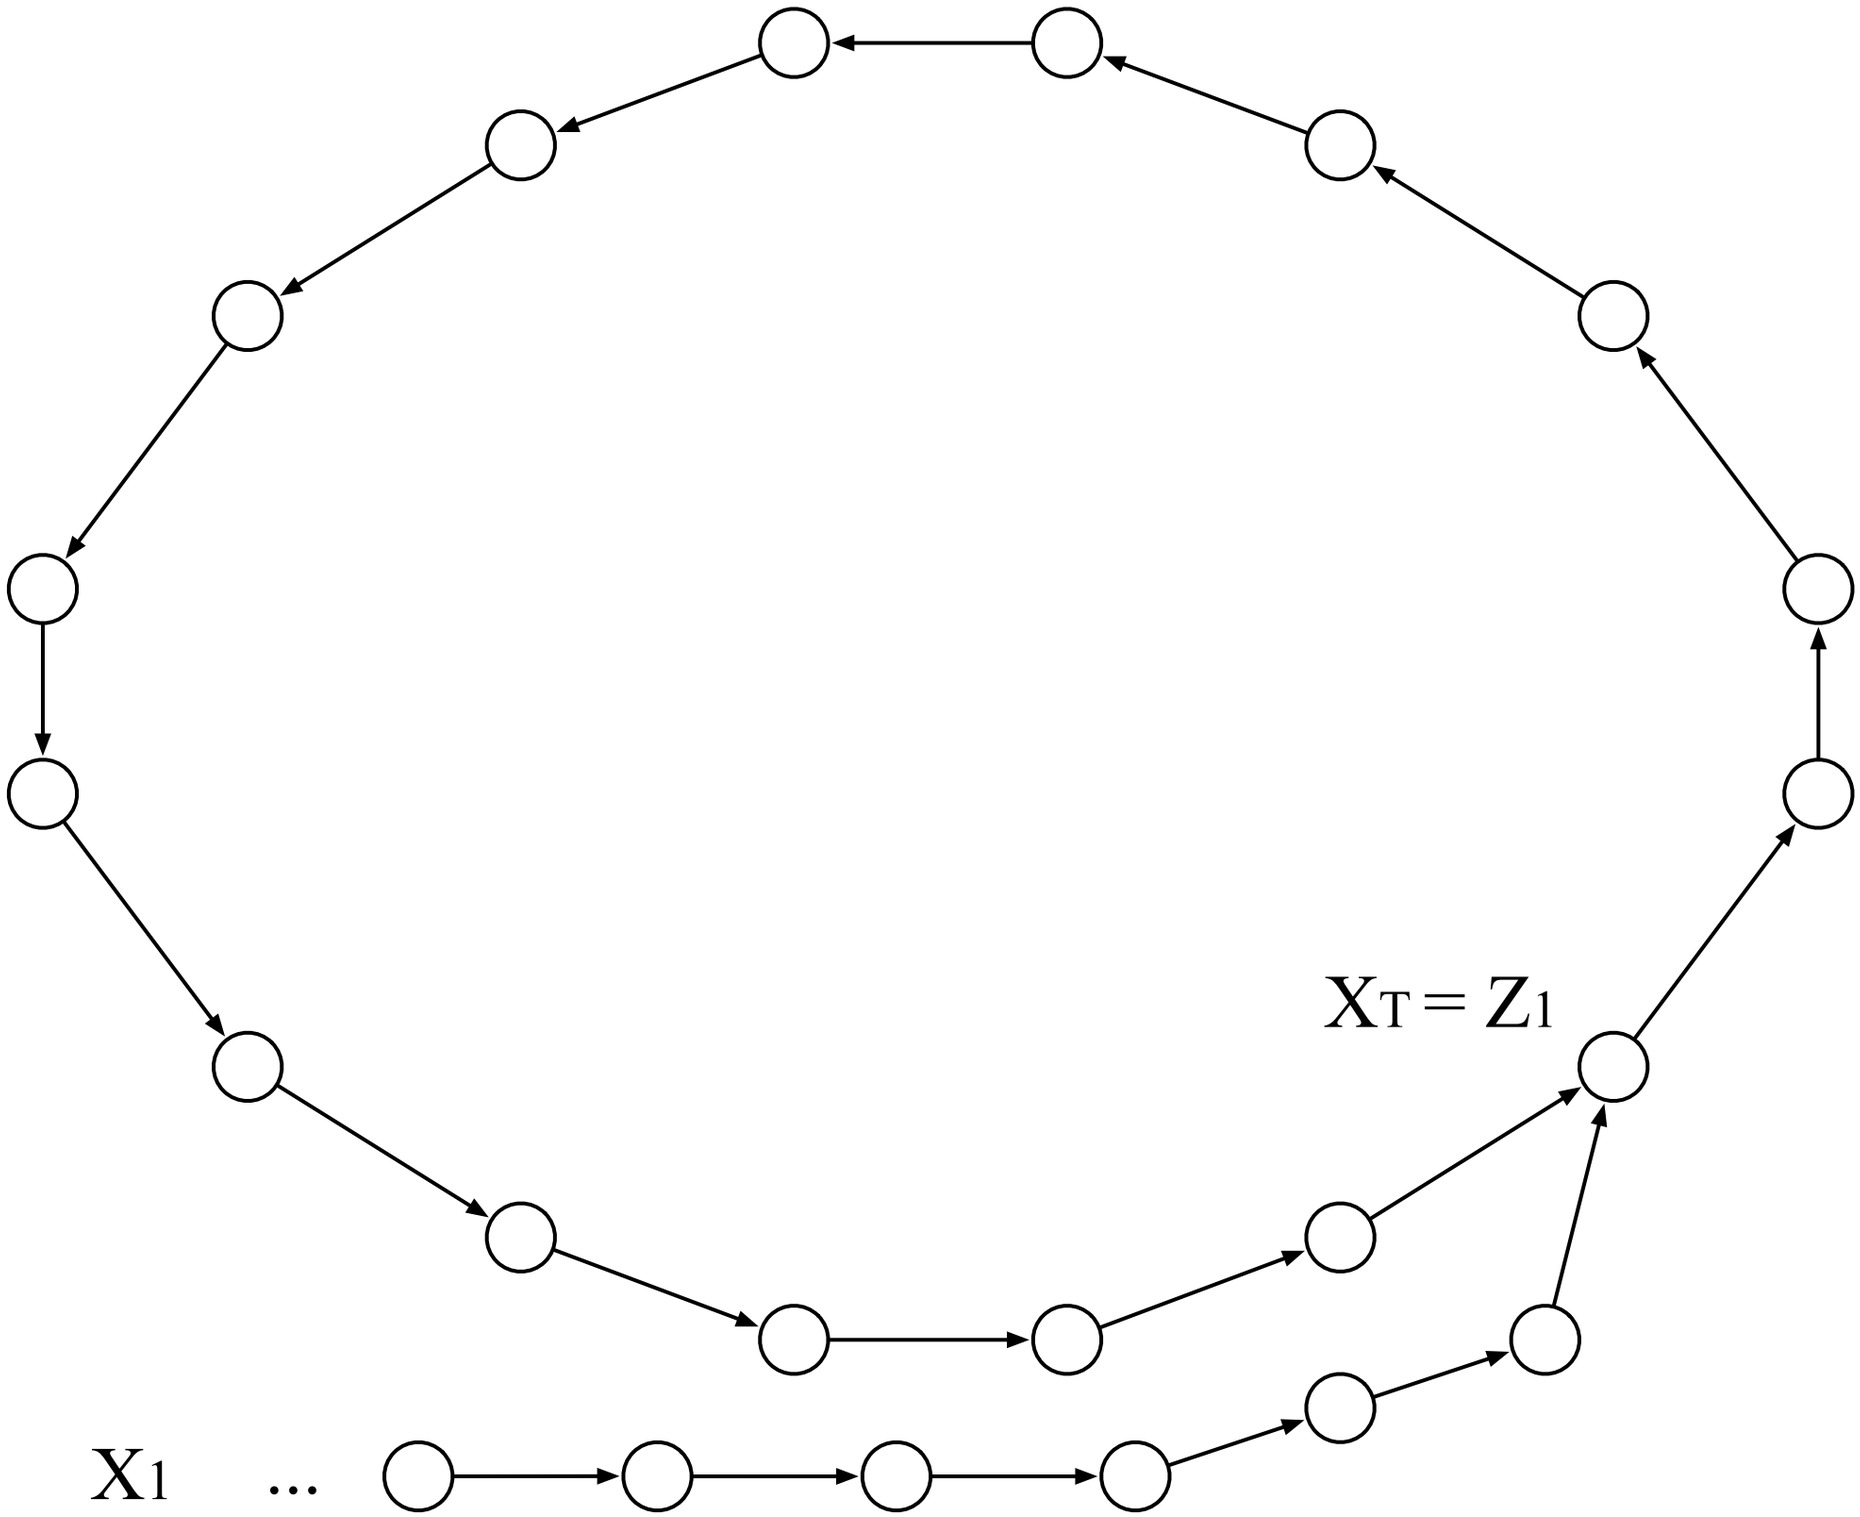

Circularly-coupled MCMC

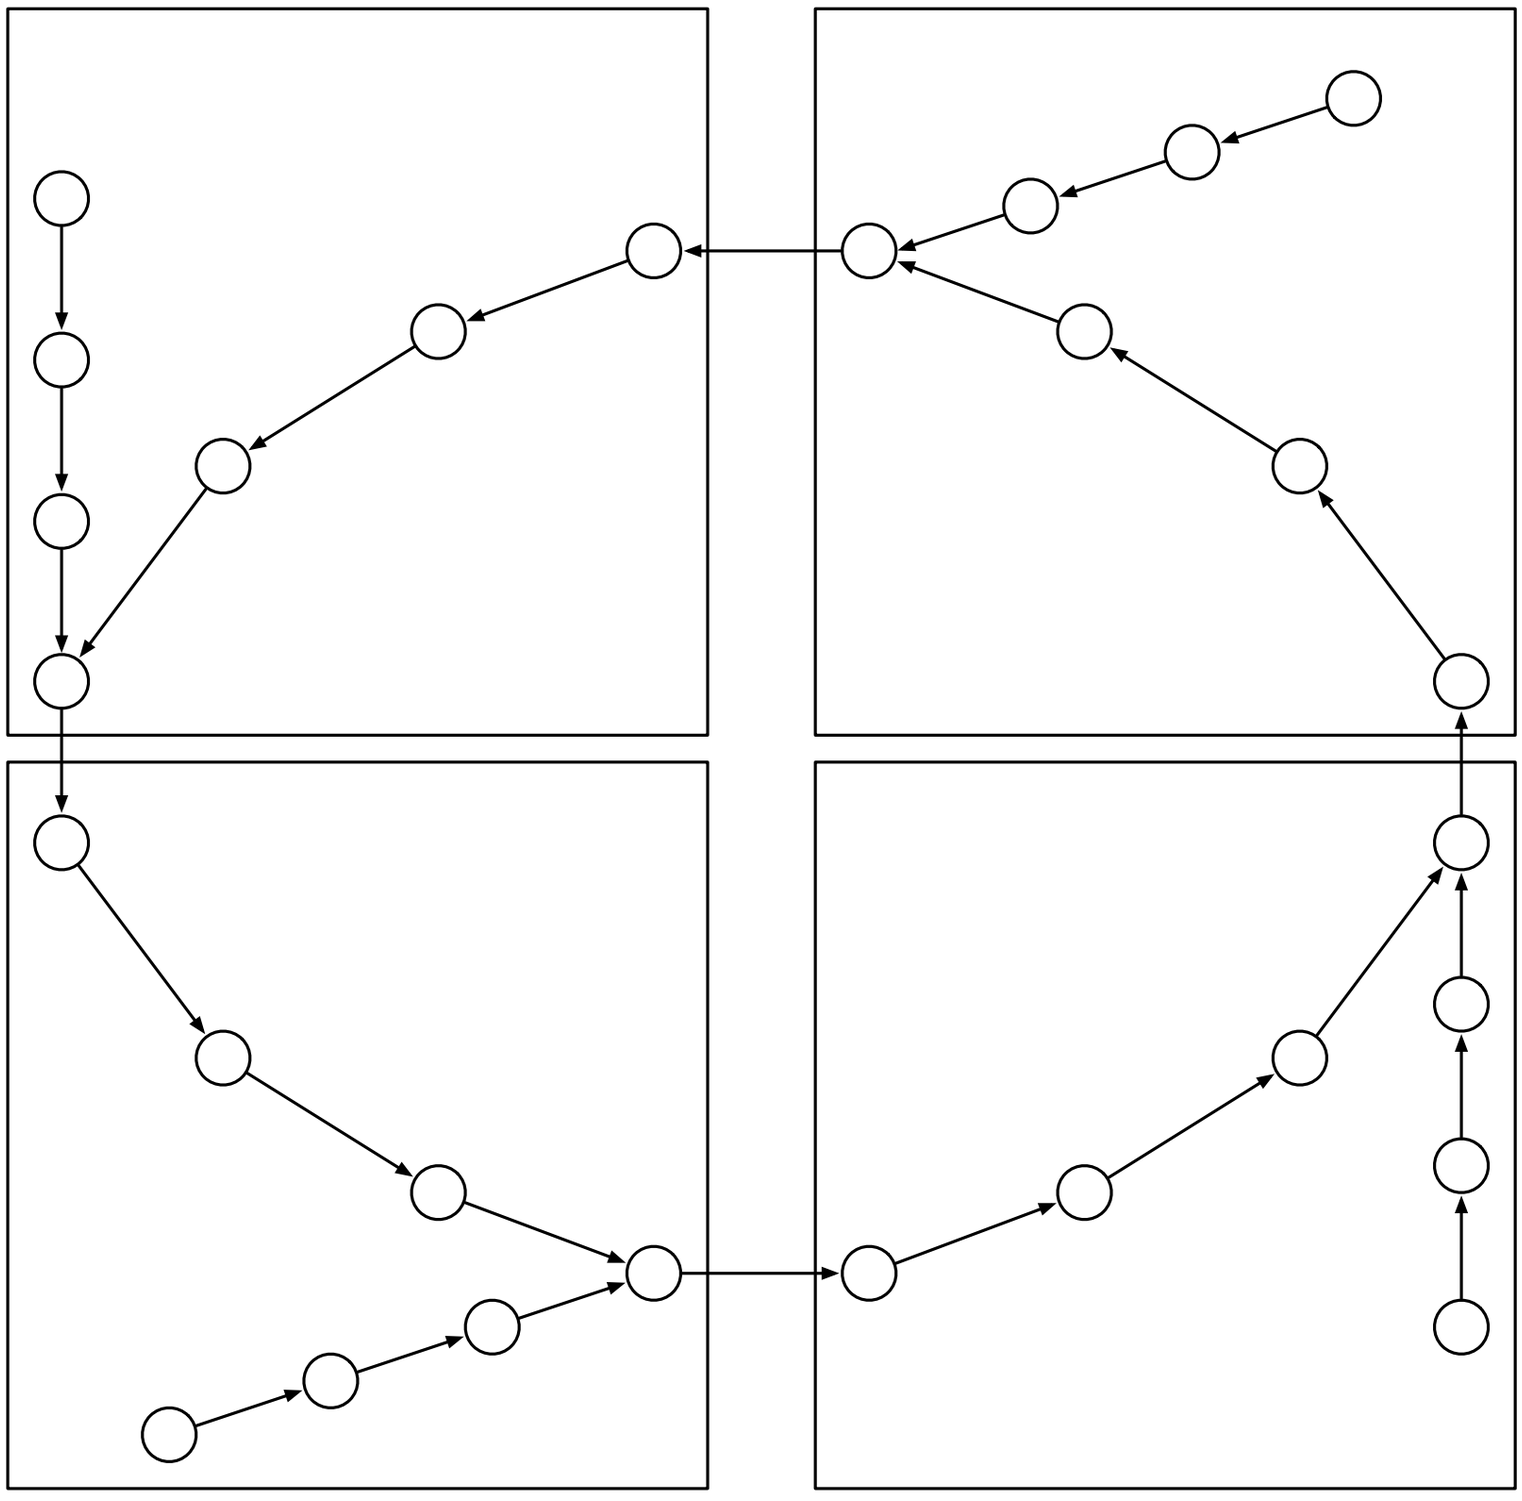

Circularly-coupled MCMC

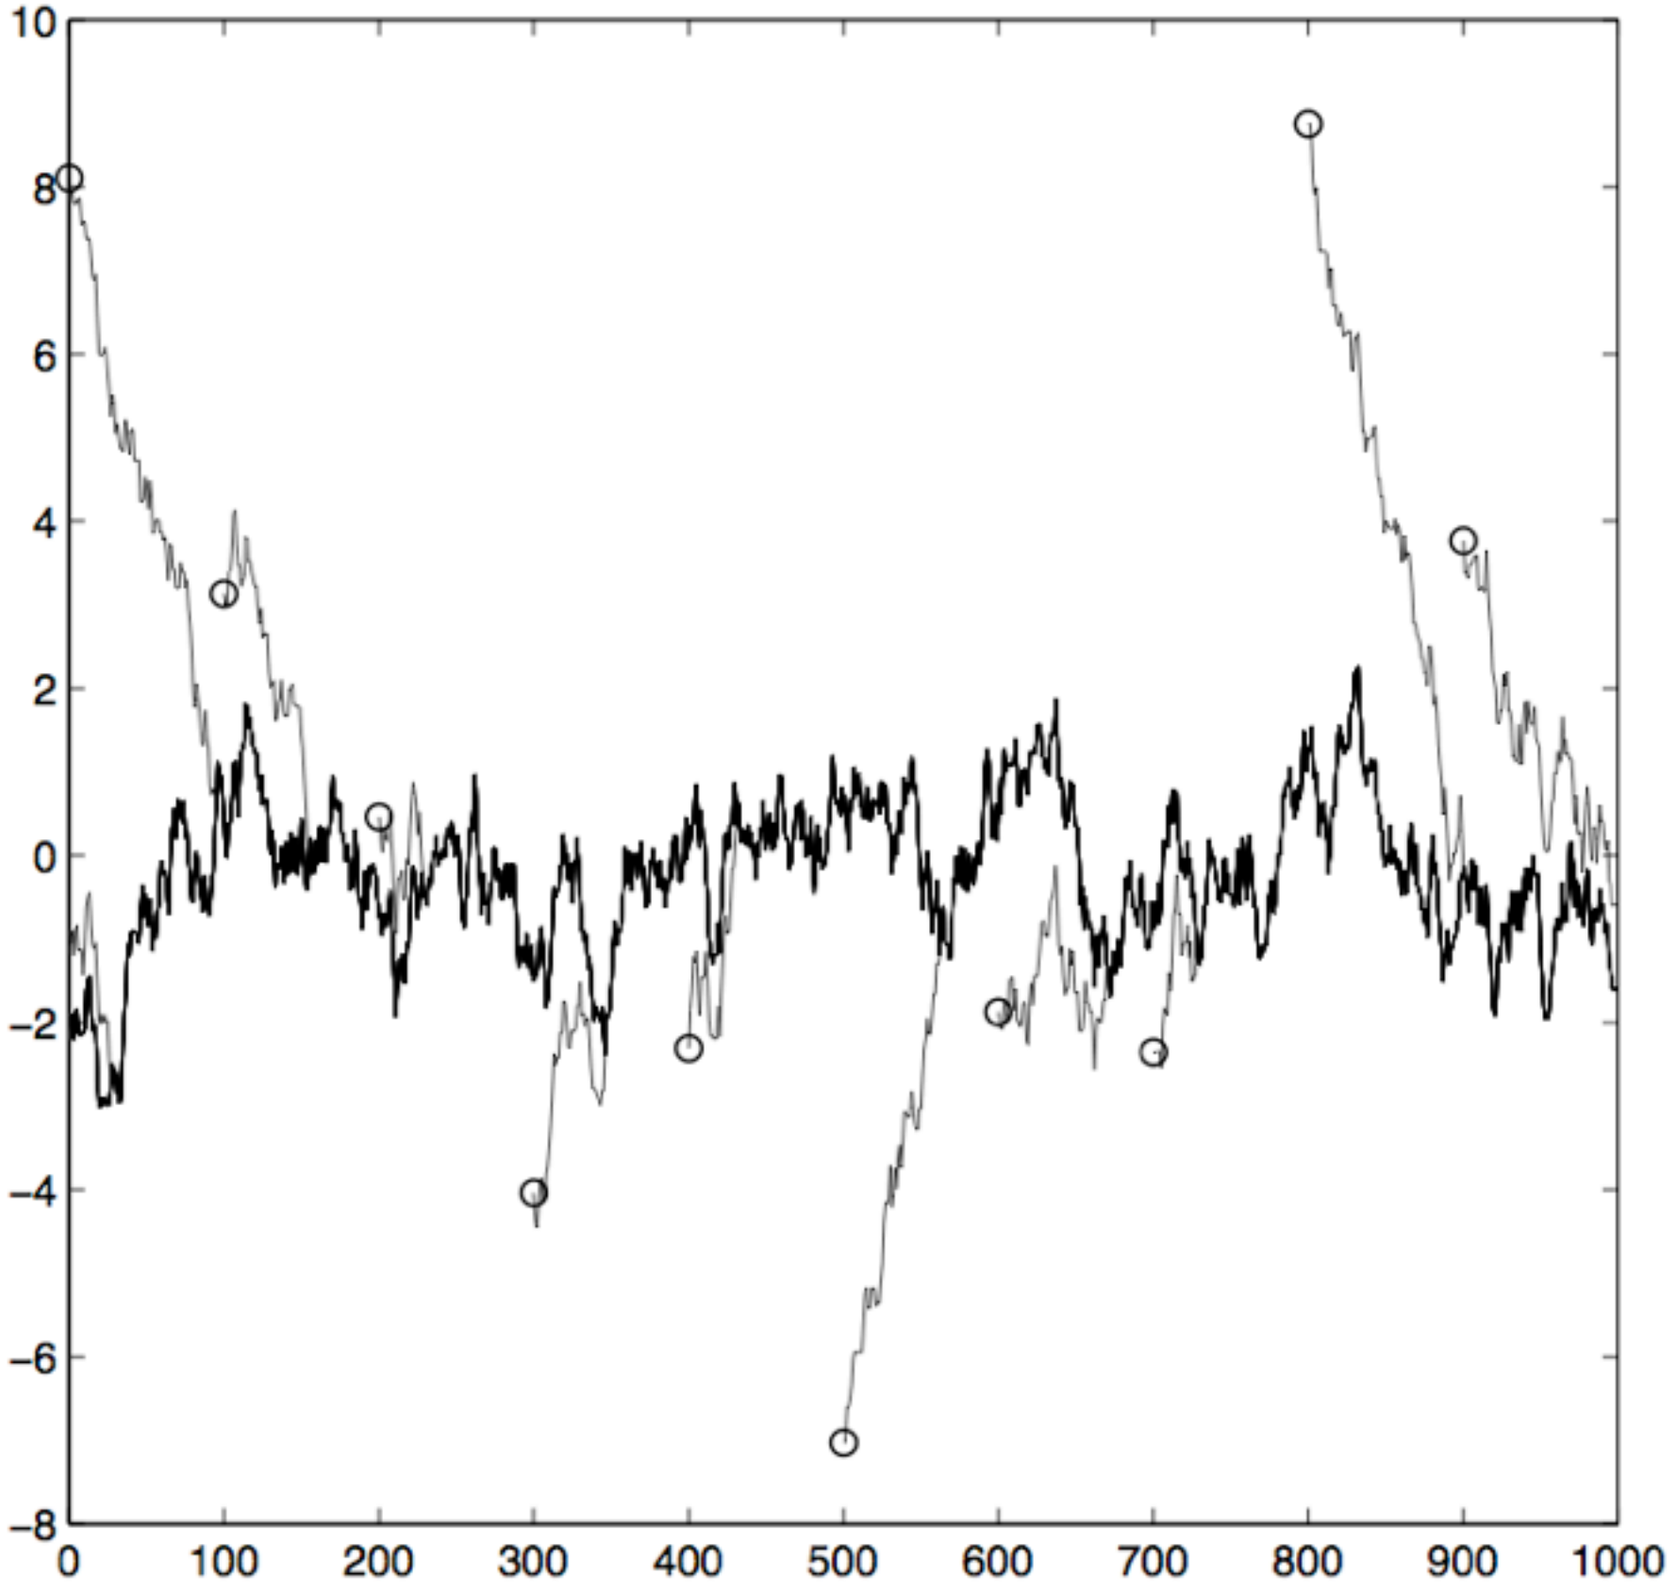

Questions to Ponder about MCMC:

- MCMC convergence:
  - Do we know what the parallelised sampler is converging to?
  - How do we know we have converged?
- Robustness against unreliable hardware
  - Any theory?

Questions to Ponder about Modelling:

- Large parametric or nonparametric?
  - Are these really nonparametric, or just large parametric?
  - DP ** model, or finite mixture of ** model?

- Complex models for a complex world:
  - But how to learn complex “patterns” from lots of data *efficiently*?

- Do we really need so much data?

- Is it feasible to engineer systems for more structured models?

Summary

- More consideration for computational concerns.
  - What good inference solutions do we have?

- A warning: avoid bloatware.