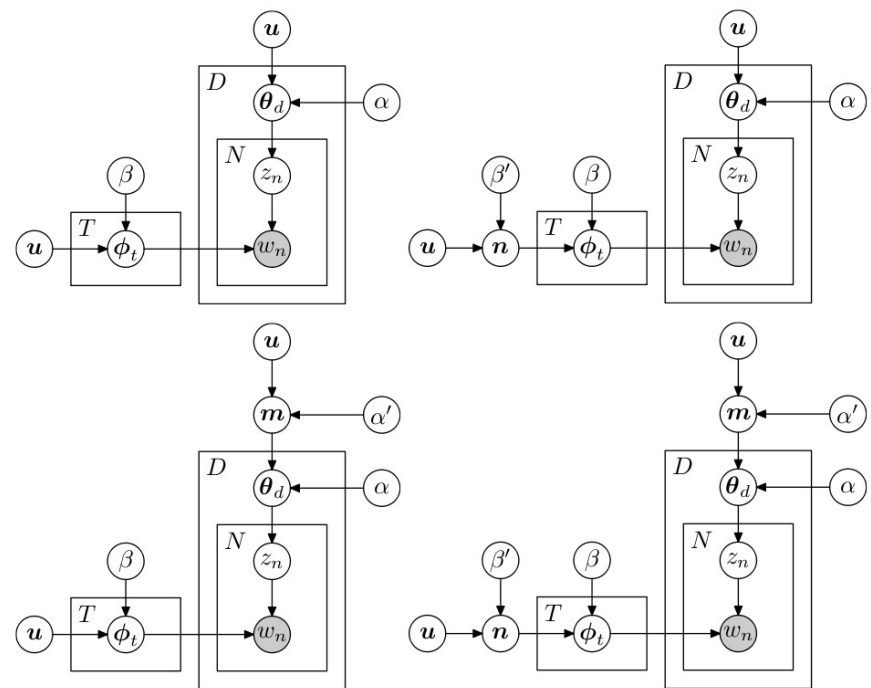


# Rethinking LDA: Why Priors Matter

Hanna M. Wallach, David Mimno, Andrew McCallum (UMass Amherst)

- Most work on LDA: symmetric Dirichlet priors
- This paper: compare several classes of structured priors
- Best choice: asymmetric prior for document-specific topic distributions, symmetric prior for topic-specific word distributions
- What you get: robustness to variation in # topics, better stopword handling...
- Efficient? OH YEAH!!!



This Friday!

## Applications for Topic Models: Text and Beyond

Special guests: Dave Blei, Mark Johnson,  
Tom Landauer, Fei-Fei Li, Eric Xing