

Relative Performance guarantees for Approximate Inference in Latent Dirichlet Allocation

Indraneel Mukherjee, David M. Blei

- Inference: $\text{Prob}(\text{latent variables} \mid \text{data})$. **Intractable**
- [TehNeWe06] Collapse latent space for better approximation.
- **Q:** Is the improvement worth the extra computation?
- Theorem: *For LDA, when documents are long, computation not worth it.*
- Our rigorous mathematical analysis is also verified empirically:

