Discussion: graph embeddings

- Pooling compresses graph information
- Could it hide a costly cross-GPU transfer?
- Which details must the embedding retain?

Discussion: generalisation and cost

- Zero-shot tests use related graph families
- Would it work on different architectures or hardware?
- 4-GPU NMT: 16.5% faster; 1.82× more evaluations
- When is the extra search cost worthwhile?