# TibGM: A Transferable and Information-Based Graphical Model Approach for Reinforcement Learning

Tameem Adel and Adrian Weller
Presented by: Tameem Adel

Machine Learning Group, University of Cambridge, UK.

ICML 2019

Most standard RL algorithms do not address generalisation and transferability.

- Graphical models (GMs) grant the possibility of:
  - beginning from precise objectives,
  - expressing them in the interpretable form of a graph, and
  - potentially achieving the objectives via efficient inference

Most standard RL algorithms do not address generalisation and transferability.

- Graphical models (GMs) grant the possibility of:
  - beginning from precise objectives,
  - expressing them in the interpretable form of a graph, and
  - potentially achieving the objectives via efficient inference
- We take full advantage of this cycle by:
  - proposing a novel information theoretic objective aiming at:
    - maximising ‘local’ reward, and
    - facilitating transfer learning (transferability) and exploration, ultimately leading to improved ‘global’ reward maximisation.
  - proposing a latent space disentangled into local and global components, as parts of a variational inference procedure.

# Contributions

1. A graphical model based on which an introduced information theoretic objective leads to the solution of RL problems.
2. The derivation of a correspondence between the mutual information-based objective and a two-fold RL objective targeting both reward maximisation and generalisation & transferability.
3. The latent space is disentangled into 'local' (reward maximisation) and 'global' (generalisation & transferability) components.
4. An information theoretic, unsupervised pretraining procedure further focusing on exploration, in cases with sparse, deceptive or very delayed extrinsic rewards.
5. State-of-the-art results on 16 benchmark tasks.

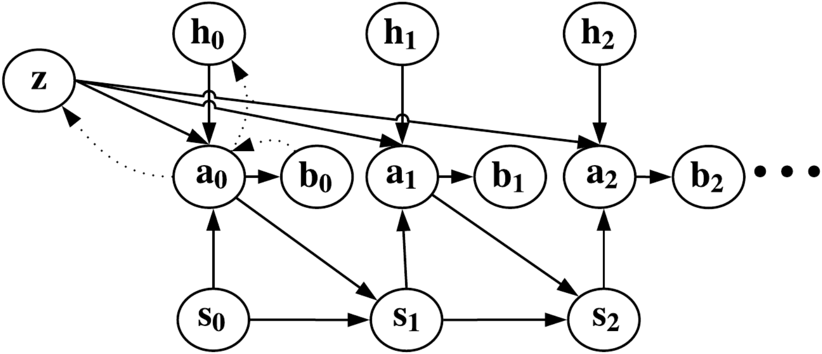


$$\max_{\theta,\phi} \quad \mathbf{I}(\mathbf{a_t}, \mathbf{b_t}) - \beta\, \mathbf{I}(\mathbf{z}, \mathbf{b_t} | \mathbf{s_t}), \tag{1}$$