Graph Representation Learning

William L. Hamilton and Jian Tang
McGill University, HEC, and Mila

Quebec Institute for Learning Algorithms (Mila)

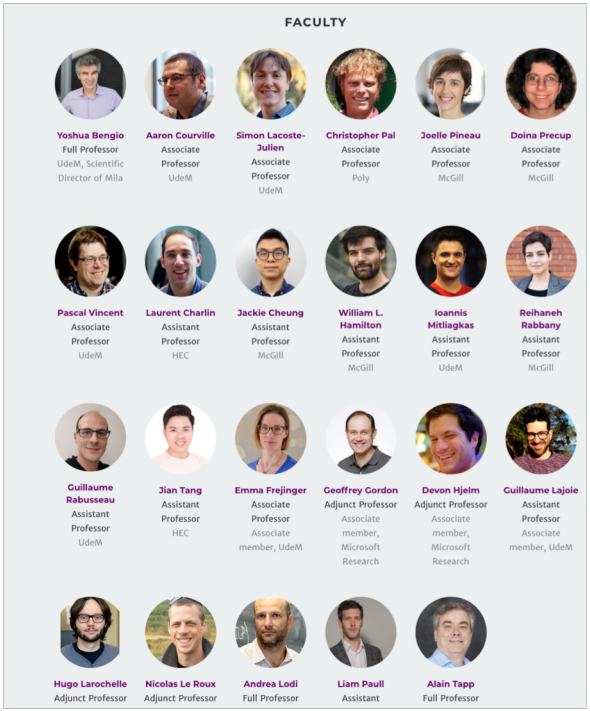

- Lead by the Deep Learning pioneer: Yoshua Bengio
- The largest academic lab on deep learning and reinforcement learning
- >30 professors (14 core member), ~ 300 students
- Multiple Postdoc, Ph.D., Master, and Interns positions are available.

Why graphs?

Graphs are a general language for describing and modeling complex systems

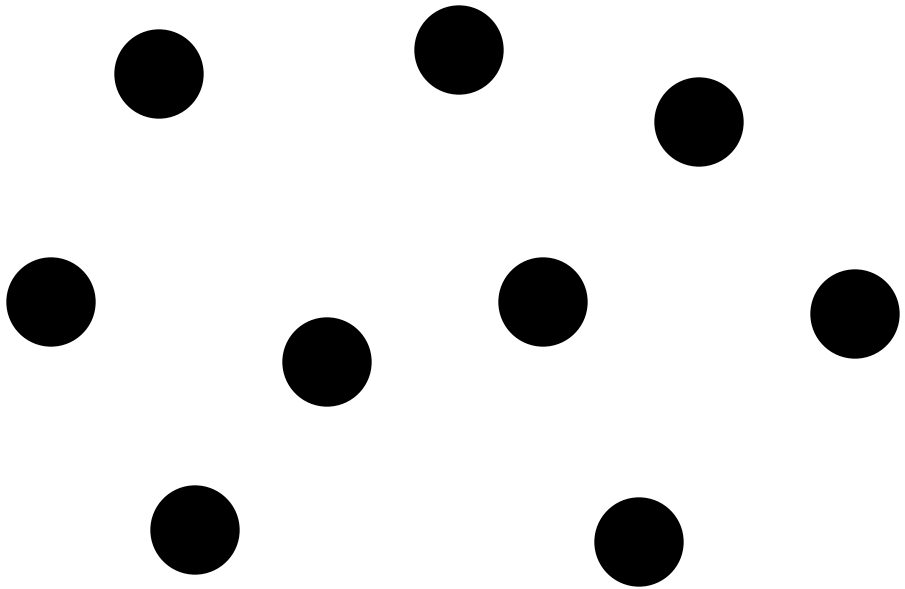

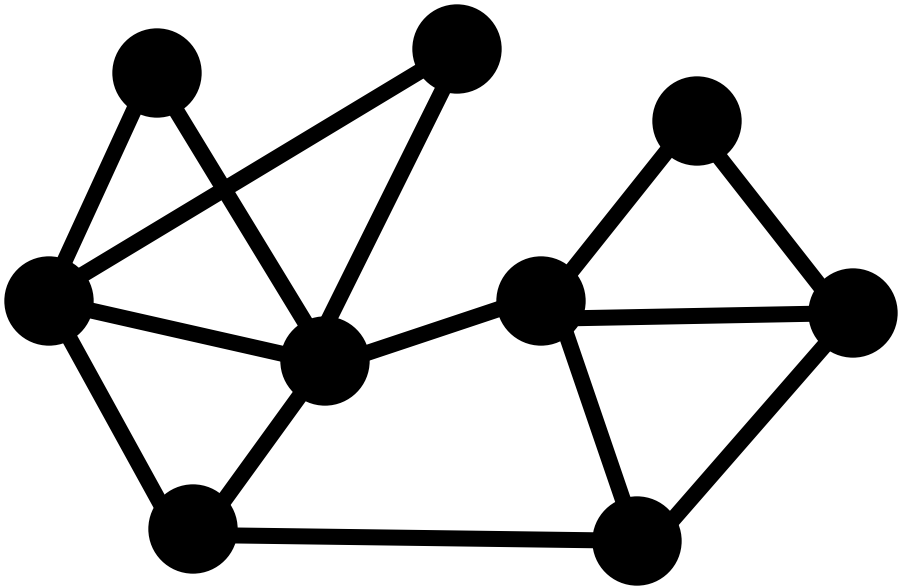

Graph!

Many Data are Graphs

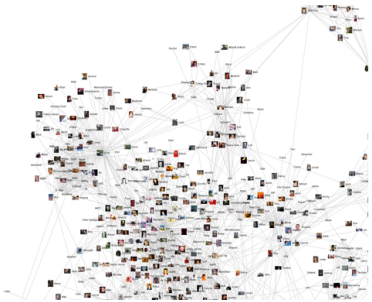

Social networks

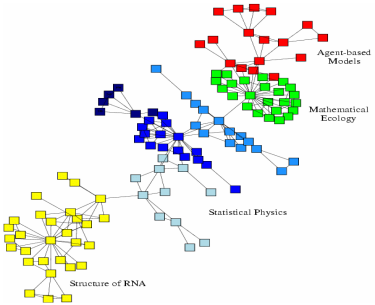

Economic networks

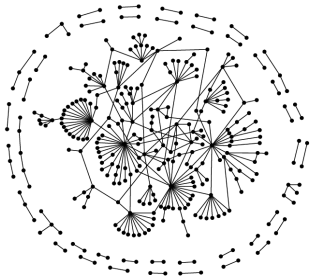

Biomedical networks

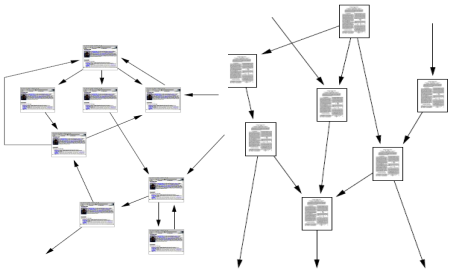

Information networks: Web & citations

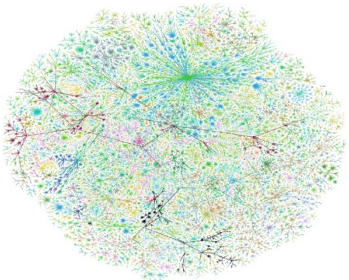

Internet

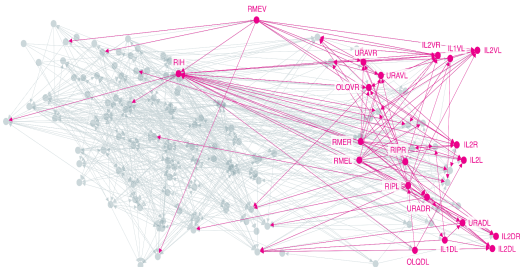

Networks of neurons

Why Graphs? Why Now?

- Universal language for describing complex data
  - Networks/graphs from science, nature, and technology are more similar than one would expect
- Shared vocabulary between fields
  - Computer Science, Social science, Physics, Economics, Statistics, Biology
- Data availability (+computational challenges)
  - Web/mobile, bio, health, and medical
- Impact!
  - Social networking, Social media, Drug design

Machine Learning with Graphs

Classical ML tasks in graphs:

- Node classification
  - Predict a type of a given node
- Link prediction
  - Predict whether two nodes are linked
- Community detection
  - Identify densely linked clusters of nodes
- Network similarity
  - How similar are two (sub)networks

Example: Node Classification

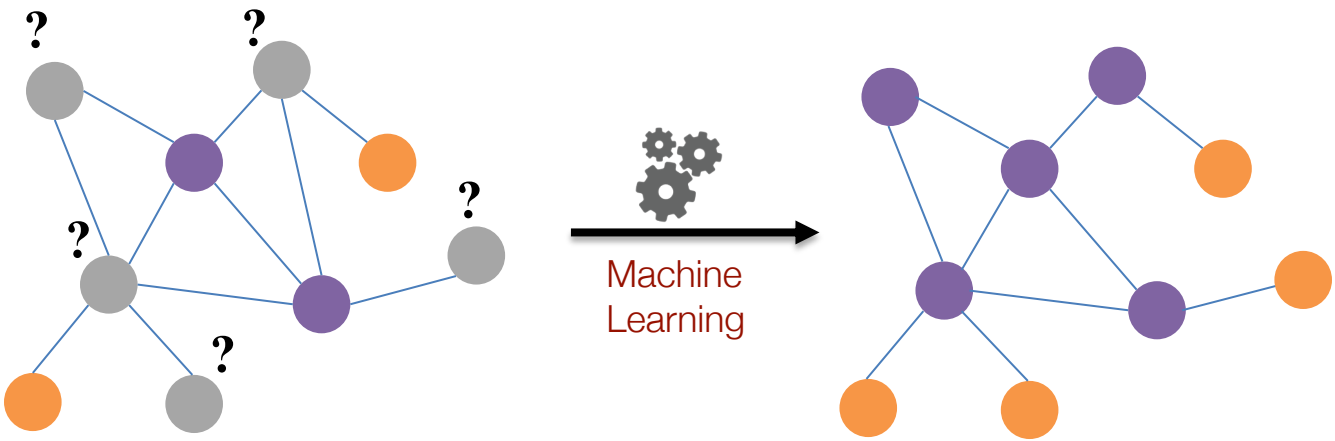

Example: Node Classification

Classifying the function of proteins in the interactome!

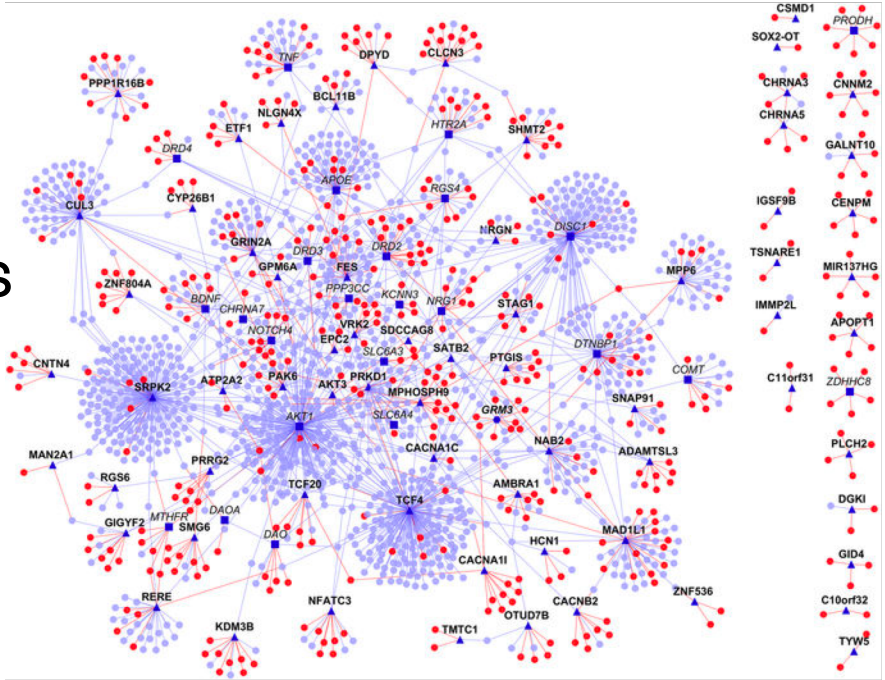

Image from: Ganapathiraju et al. 2016. Schizophrenia interactome with 504 novel protein–protein interactions. *Nature*.

Example: Link Prediction

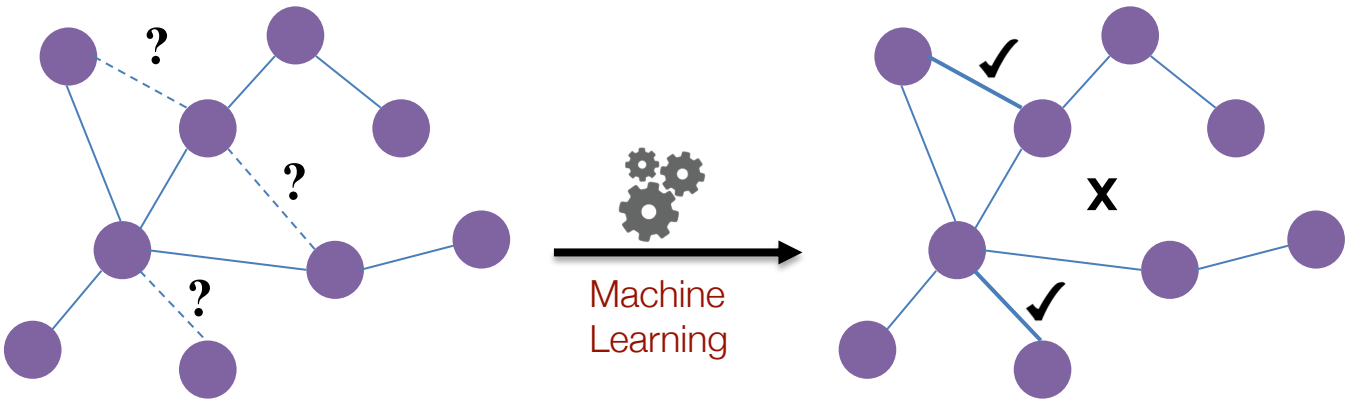

Example: Link Prediction

Content recommendation is link prediction!

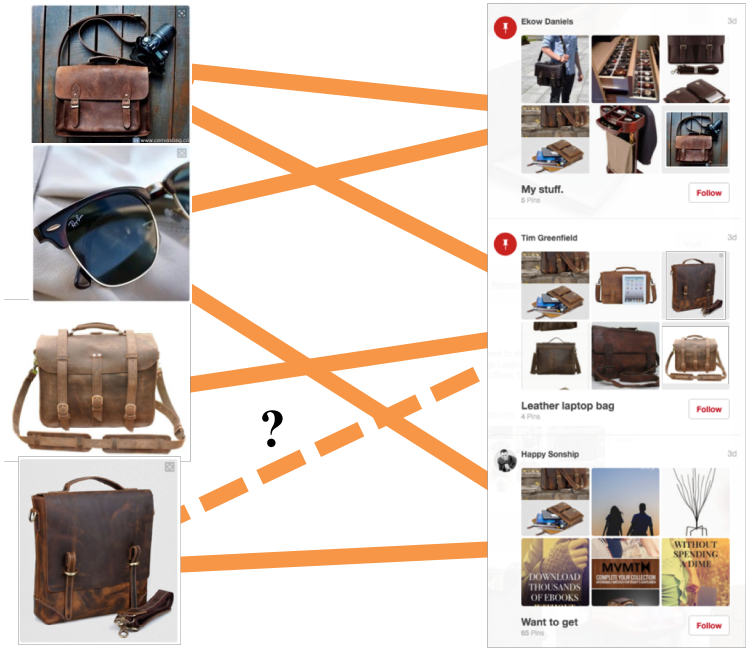

Machine Learning Lifecycle

- (Supervised) Machine Learning Lifecycle: This feature, that feature. Every single time!

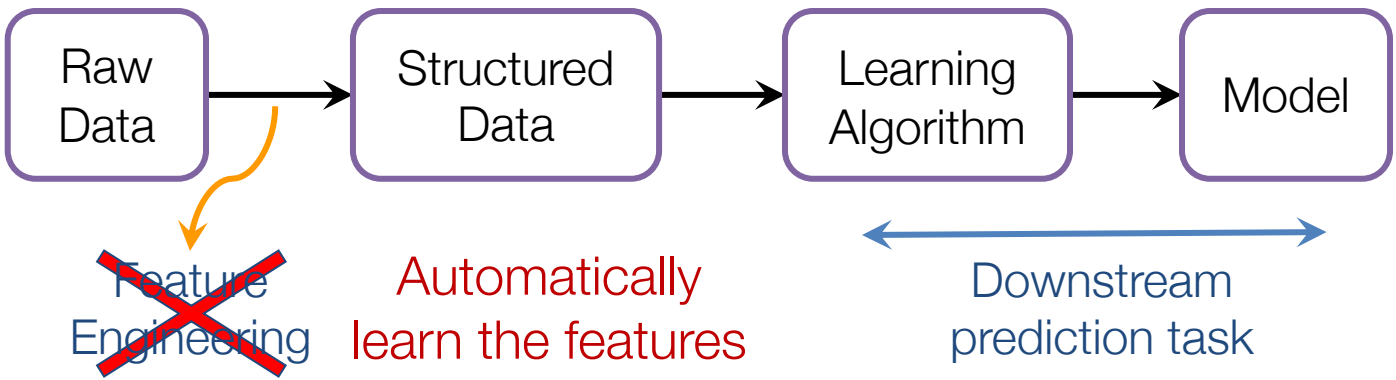

Feature Learning in Graphs

Goal: Efficient task-independent feature learning for machine learning in graphs!

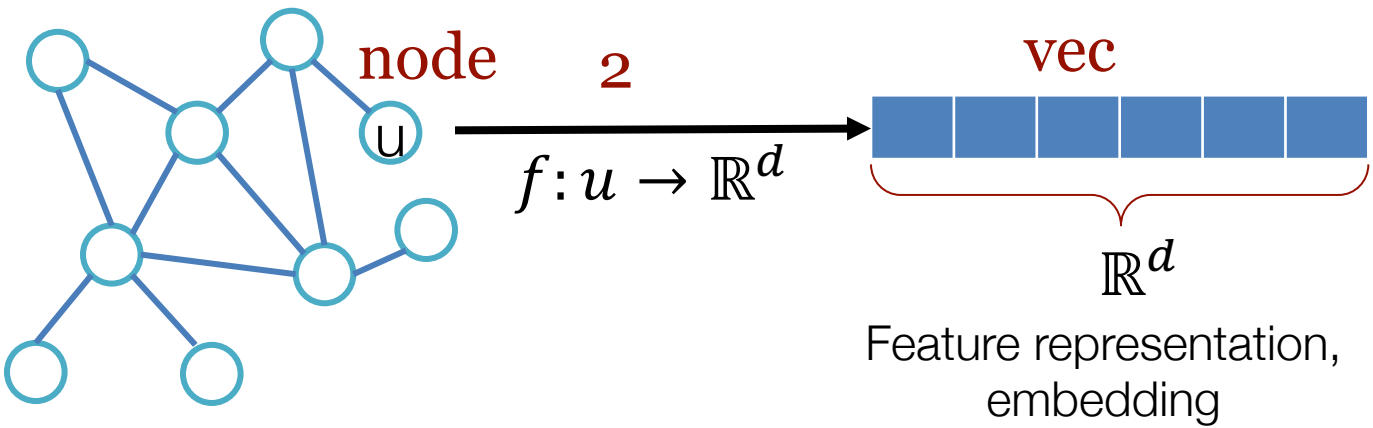

Example

- Zachary's Karate Club Network:

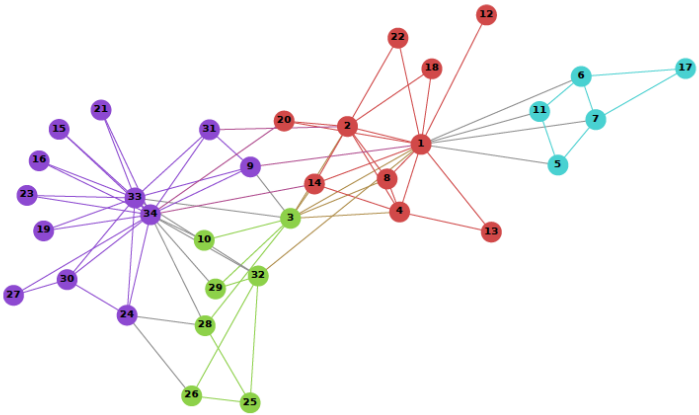

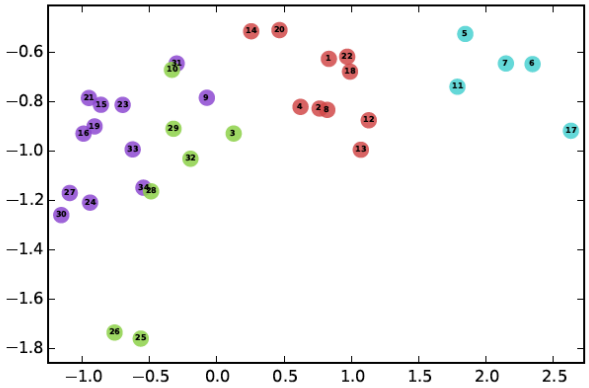

Input Output

Image from: Perozzi et al. 2014. DeepWalk: Online Learning of Social Representations. *KDD.*

Why Is It Hard?

- Modern deep learning toolbox is designed for simple sequences or grids.
  - CNNs for fixed-size images/grids….


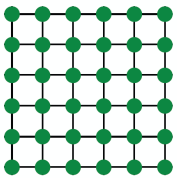
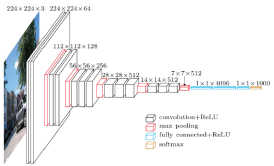

  - RNNs or word2vec for text/sequences…



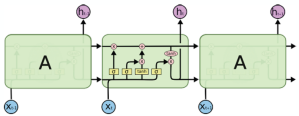

Why Is It Hard?

- But graphs are far more complex!
  - Complex topographical structure (i.e., no spatial locality like grids)

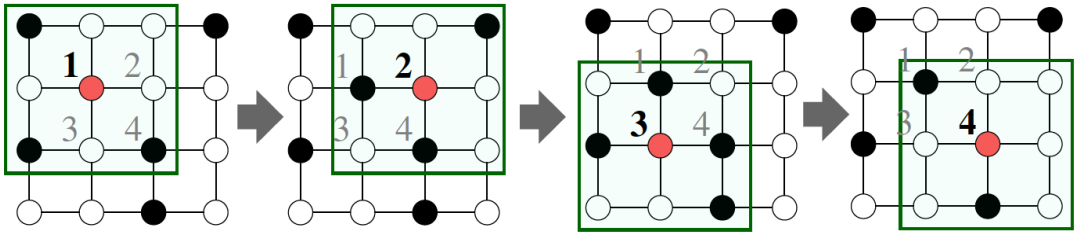

  - No fixed node ordering or reference point (i.e., the isomorphism problem)
  - Often dynamic and have multimodal features.

This talk

- 1) Node embeddings
  - Map nodes to low-dimensional embeddings.

~30min break

- 2) Graph neural networks
  - Deep learning architectures for graph-structured data
- 3) Generative graph models
  - Learning to generate realistic graph data.