Stanford CS224W: Machine Learning with Graphs

CS224W: Machine Learning with Graphs
Jure Leskovec, Stanford University
http://cs224w.stanford.edu

Why Graphs?

Graphs are a general language for describing and analyzing entities with relations/interactions

Many Types of Data are Graphs (1)

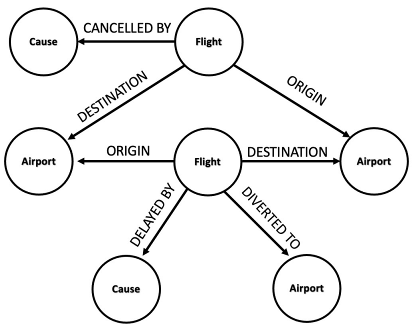

Image credit: SalientNetworks

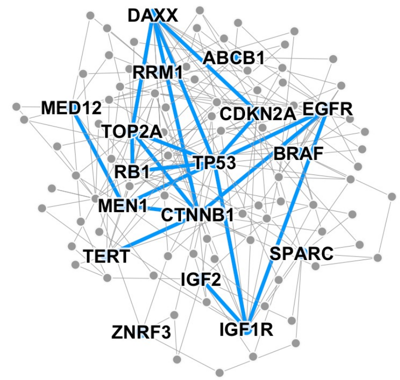

Event Graphs | **Computer Networks** | **Disease Pathways**

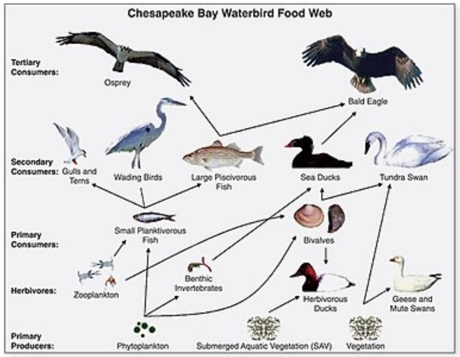

Image credit: Wikipedia

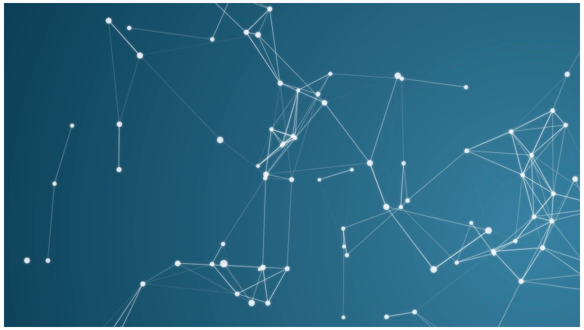

Image credit: Pinterest

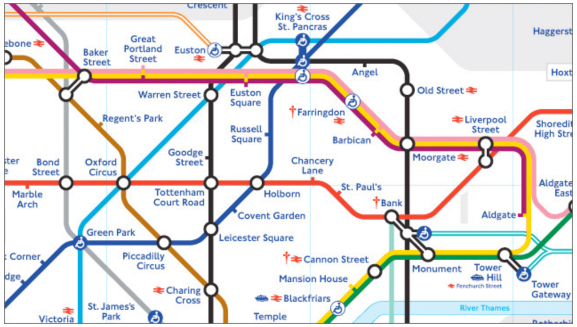

Image credit: visitlondon.com

Food Webs | **Particle Networks** | **Underground Networks**

Many Types of Data are Graphs (2)

Image credit: Medium

Social Networks

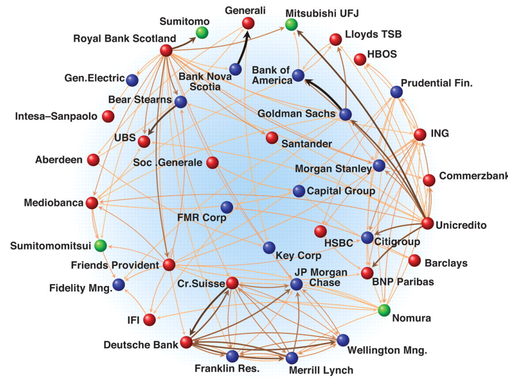

Image credit: Science

Economic Networks

Image credit: Lumen Learning

Communication Networks

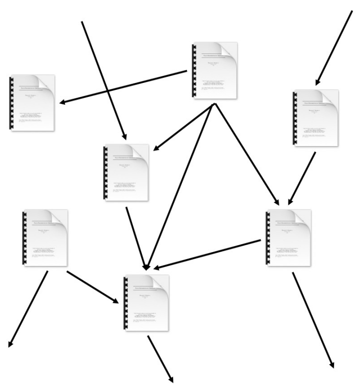

Citation Networks

Image credit: Missoula Current News

Internet

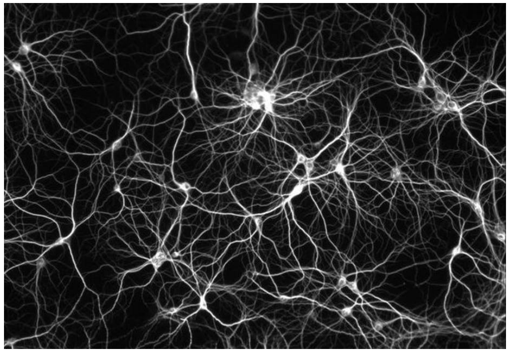

Image credit: The Conversation

Networks of Neurons

Many Types of Data are Graphs (3)

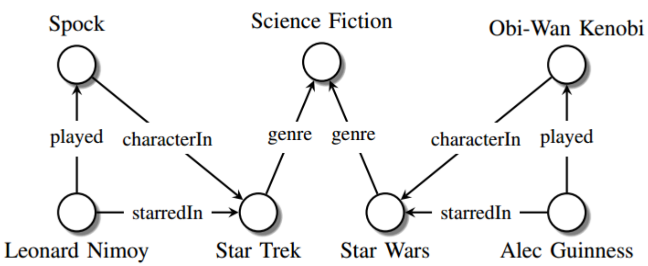

Image credit: Maximilian Nickel et al

Knowledge Graphs

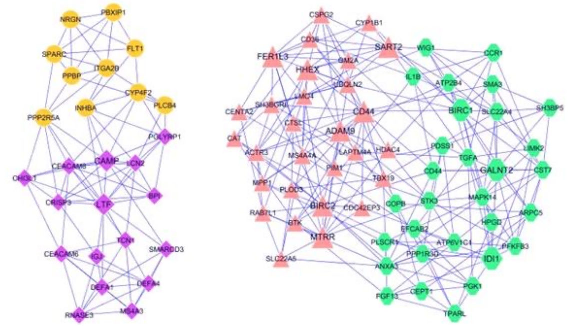

Image credit: ese.wustl.edu

Regulatory Networks

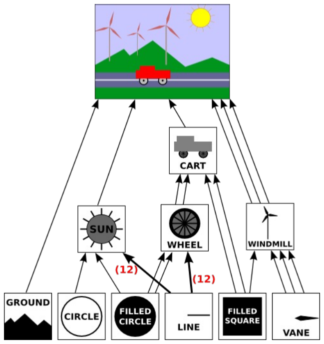

Image credit: math.hws.edu

Scene Graphs

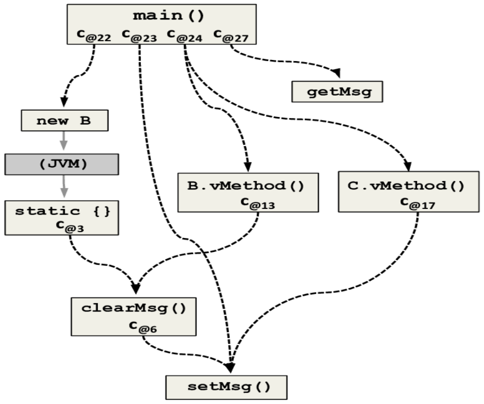

Image credit: ResearchGate

Code Graphs

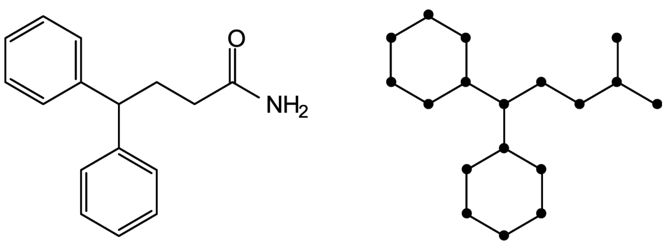

Image credit: MDPI

Molecules

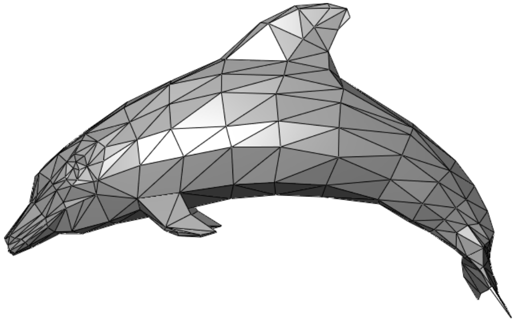

Image credit: Wikipedia

3D Shapes

Databases are Graphs!

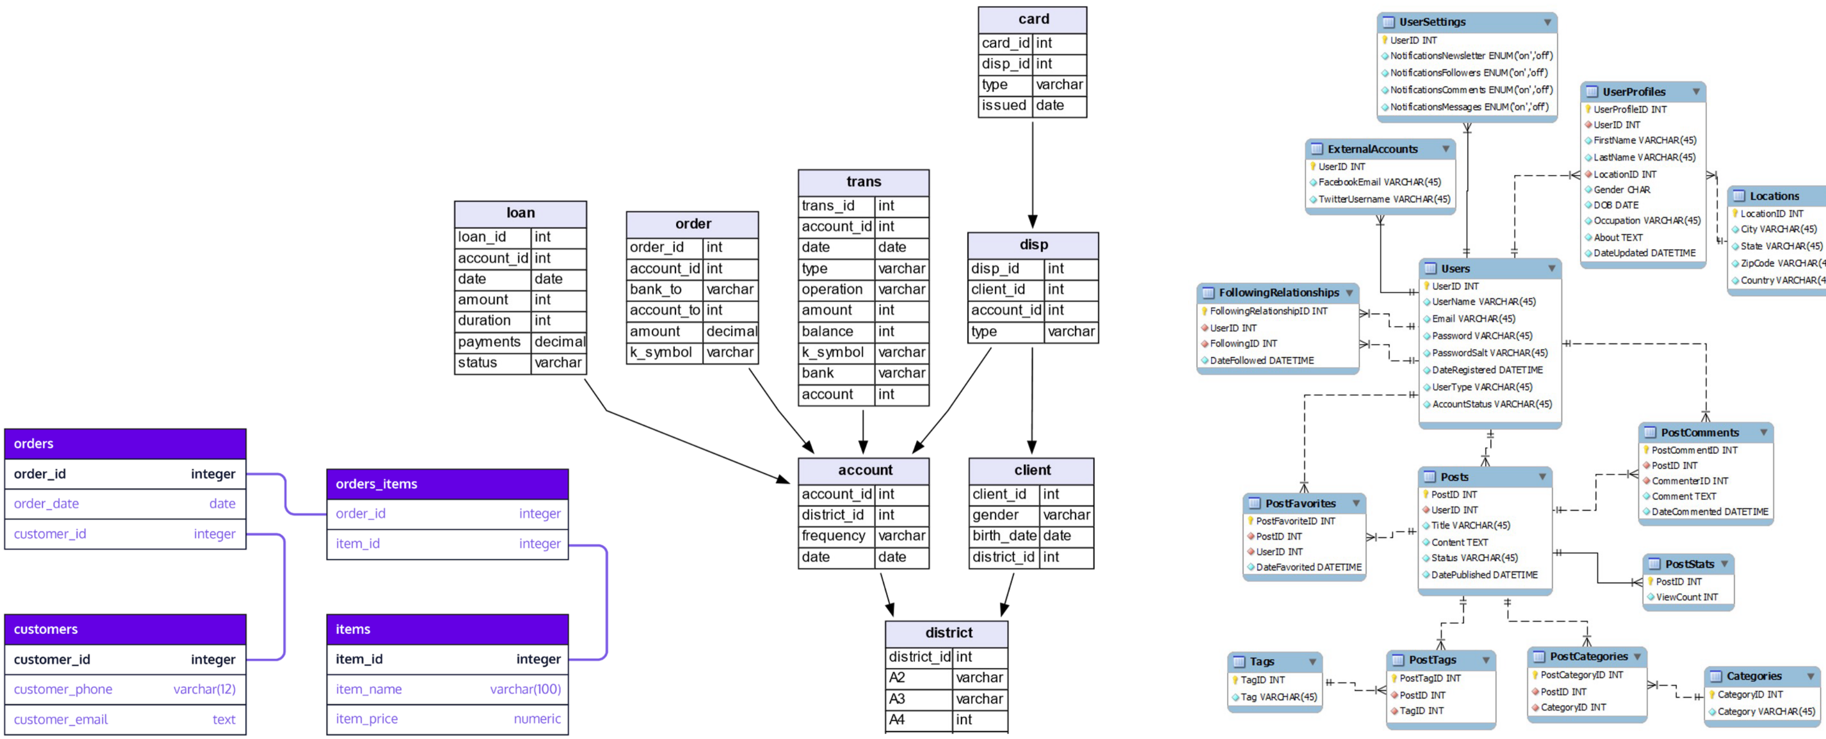

Commerce   Finance   Social Media

Relational Deep Learning

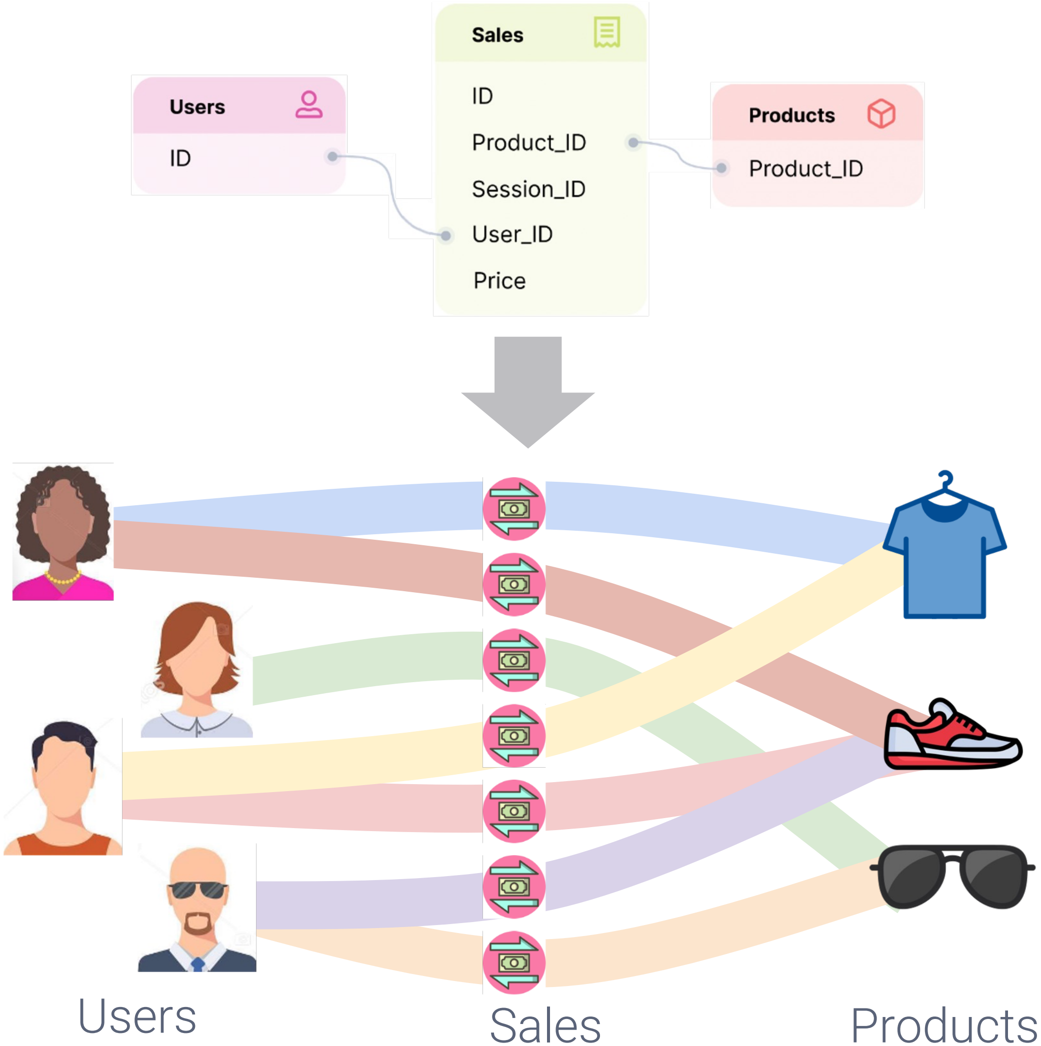

http://relbench.stanford.edu

Graphs: Machine Learning

Complex domains have a rich relational structure, which can be represented as a **relational graph**

By explicitly modeling relationships we achieve better performance!

Main question:

How do we take advantage of relational structure for better prediction?

Today: Modern ML Toolbox

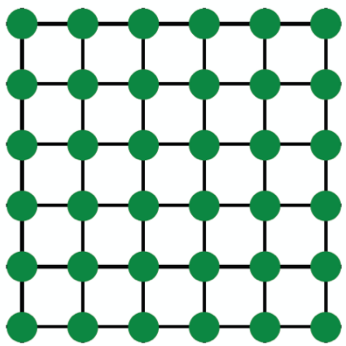

Images

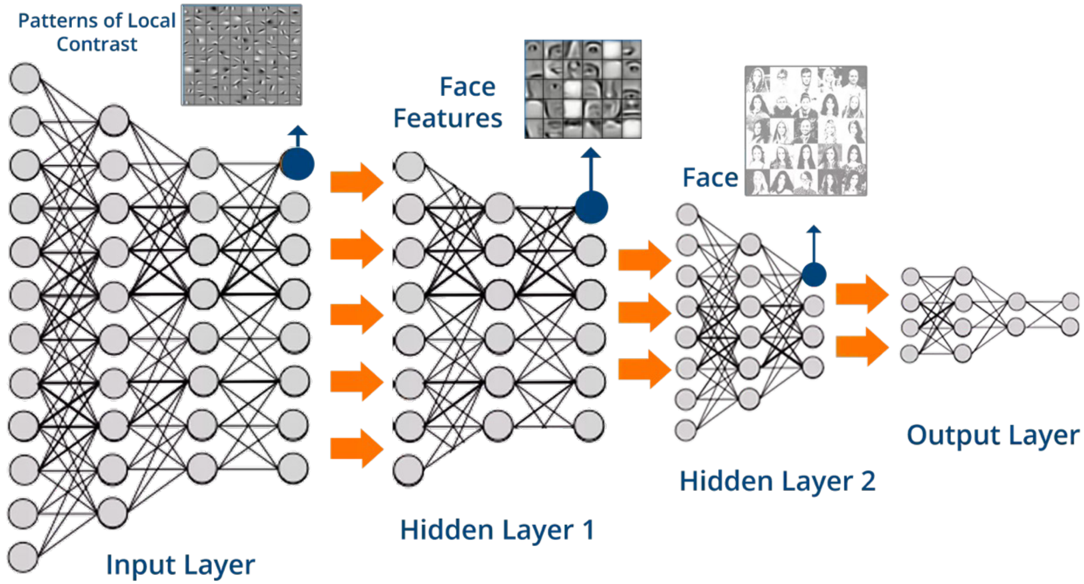

Text/Speech

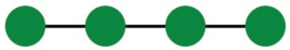

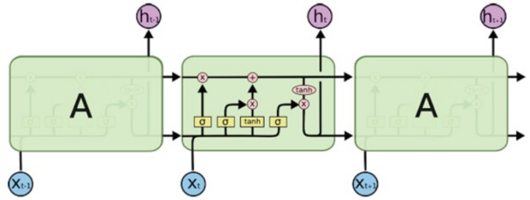

Modern deep learning toolbox is designed for simple sequences & grids

Doubt thou the stars are fire;
Doubt that the sun doth move;
Doubt truth to be a liar;
But never doubt I love...

Text

Audio signals

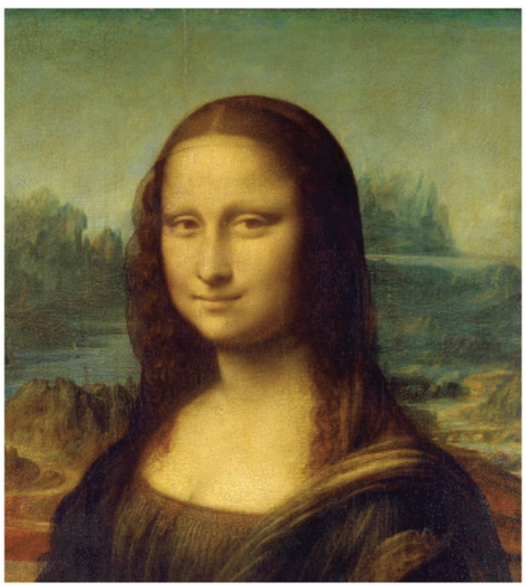

Images

Modern
deep learning toolbox
is designed for
sequences & grids

Why is Graph Deep Learning Hard?

Networks are complex.

- Arbitrary size and complex topological structure (*i.e.*, no spatial locality like grids)

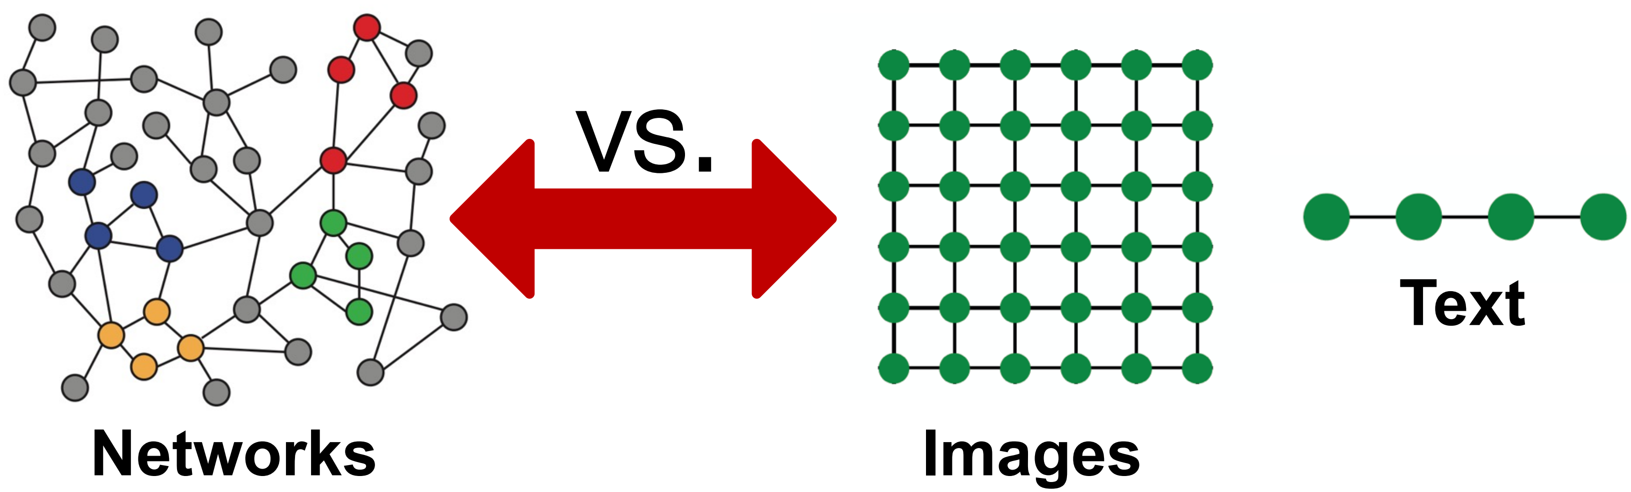

- No fixed node ordering or reference point
- Often dynamic and have multimodal features

This Course: CS224W

How can we develop neural networks that are much more broadly applicable?

Graphs are the new frontier of deep learning

Hot subfield in ML

ICLR 2023 keywords

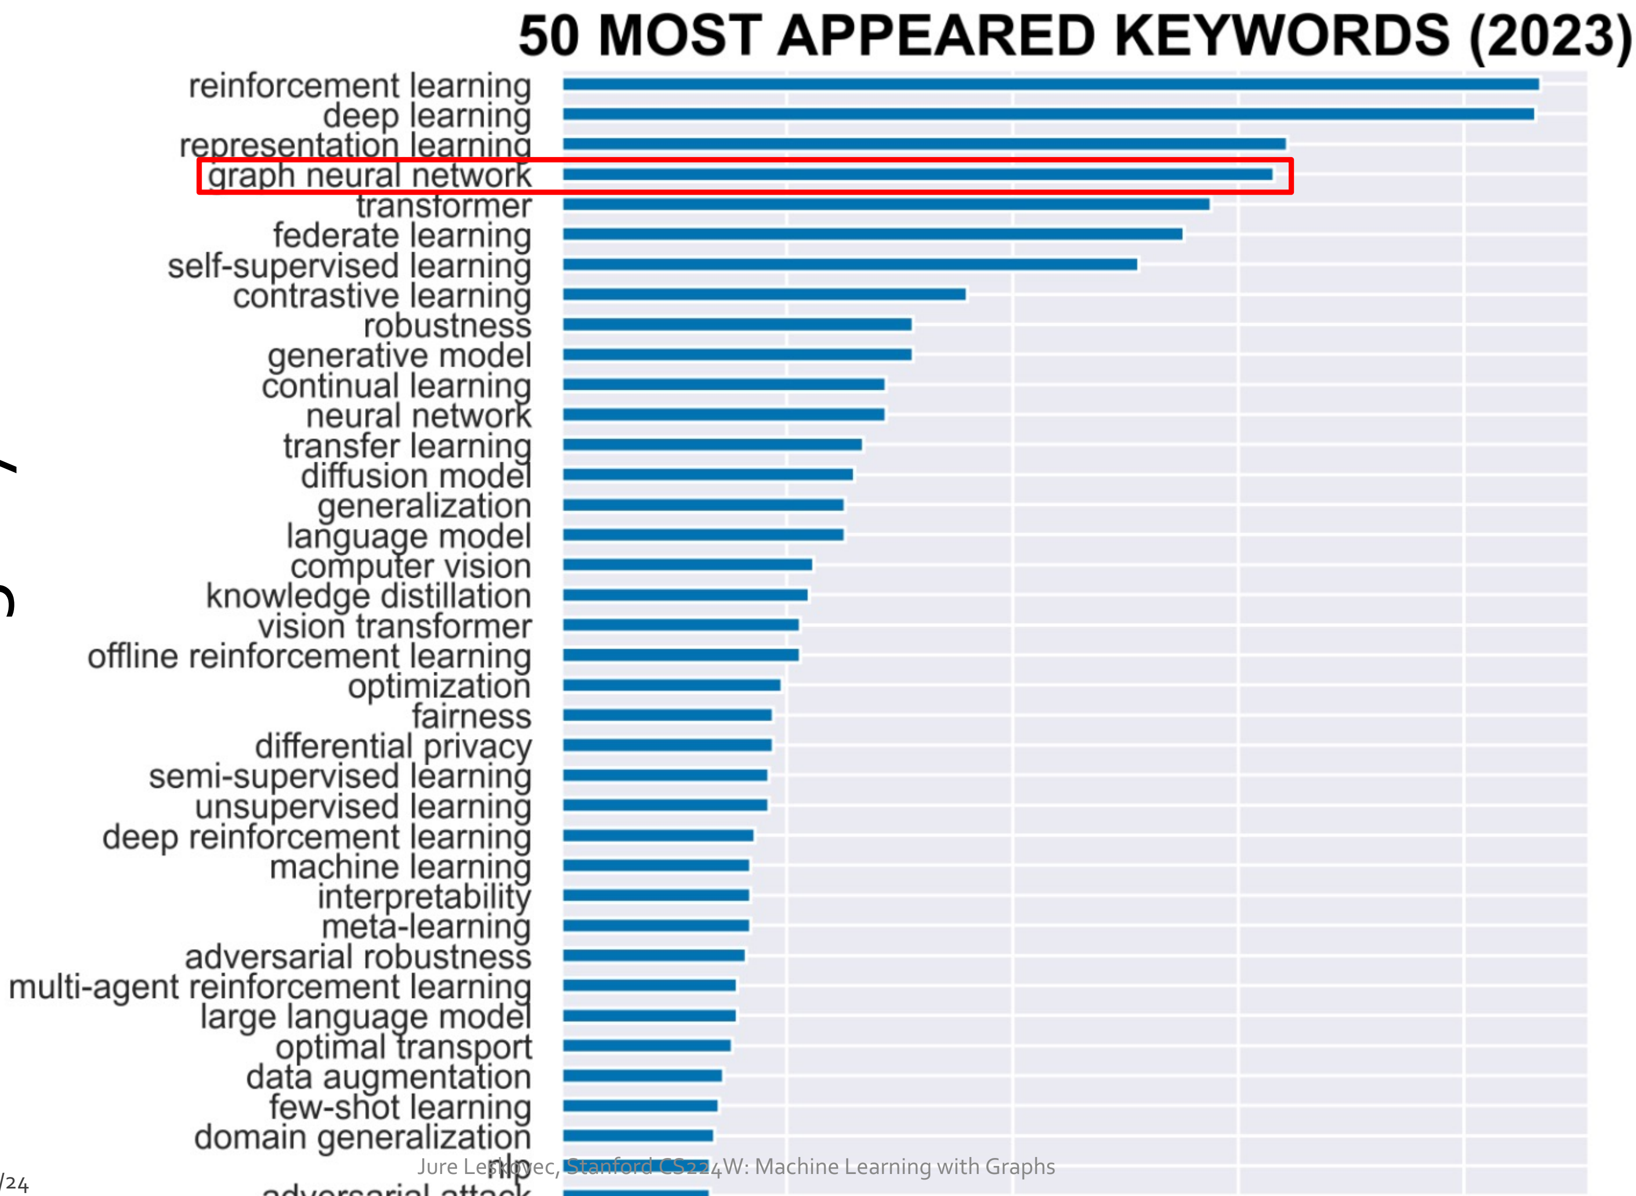

Stanford CS224W: Choice of Graph Representation

CS224W: Machine Learning with Graphs
Jure Leskovec, Stanford University
http://cs224w.stanford.edu

Graphs: A Common Language

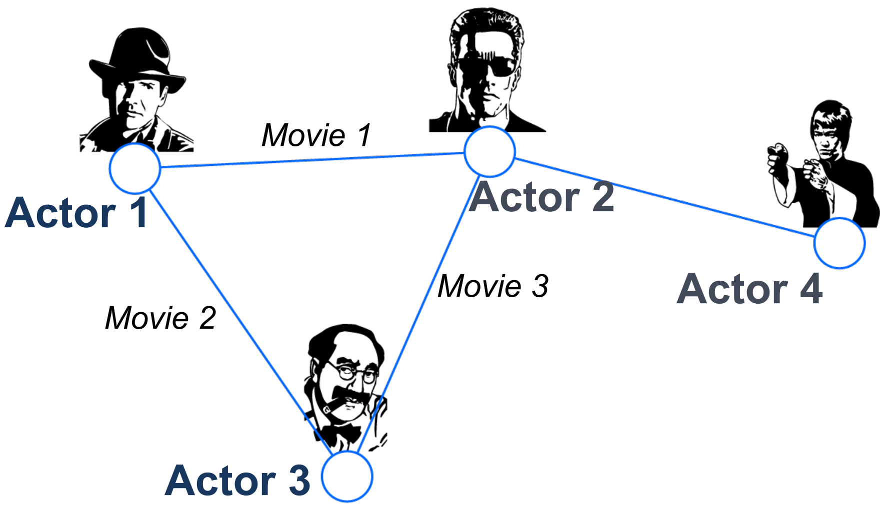

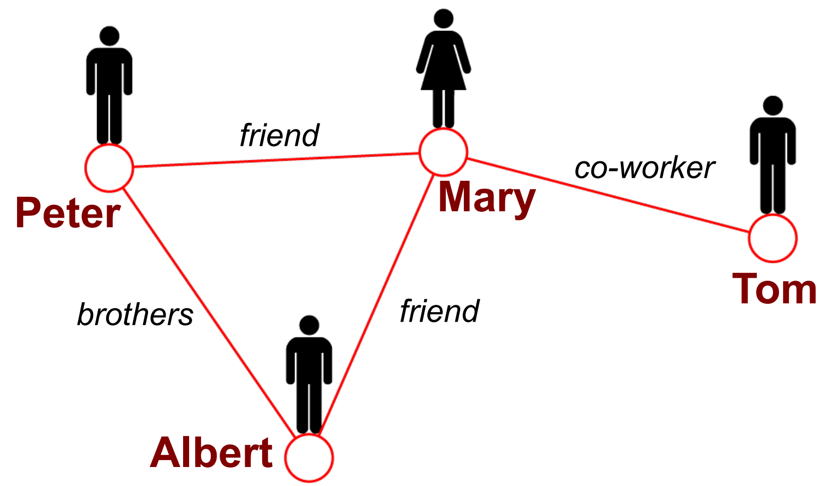

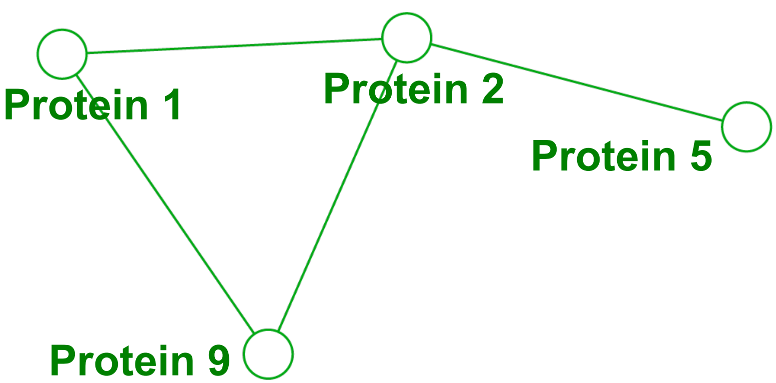

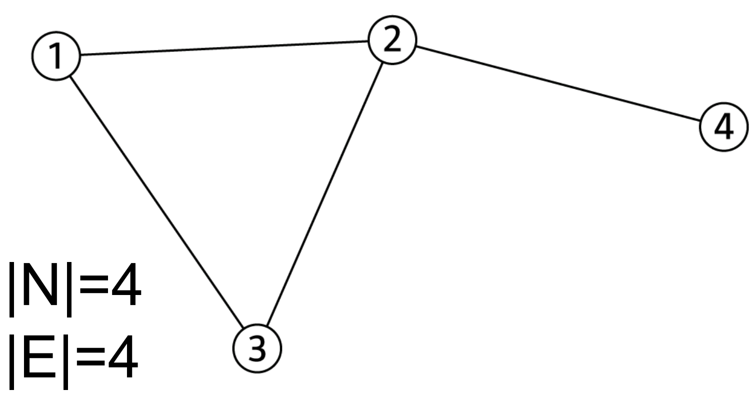

Heterogeneous Graphs

- **A heterogeneous graph is defined as**

$$\boldsymbol{G = (V, E, R, T)}$$

  - Nodes with node types $v_i \in V$
  - Edges with relation types $(v_i, r, v_j) \in E$
  - Node type $T(v_i)$
  - Relation type $r \in R$
  - Nodes and edges have **attributes/features**

Many Graphs are Heterogeneous

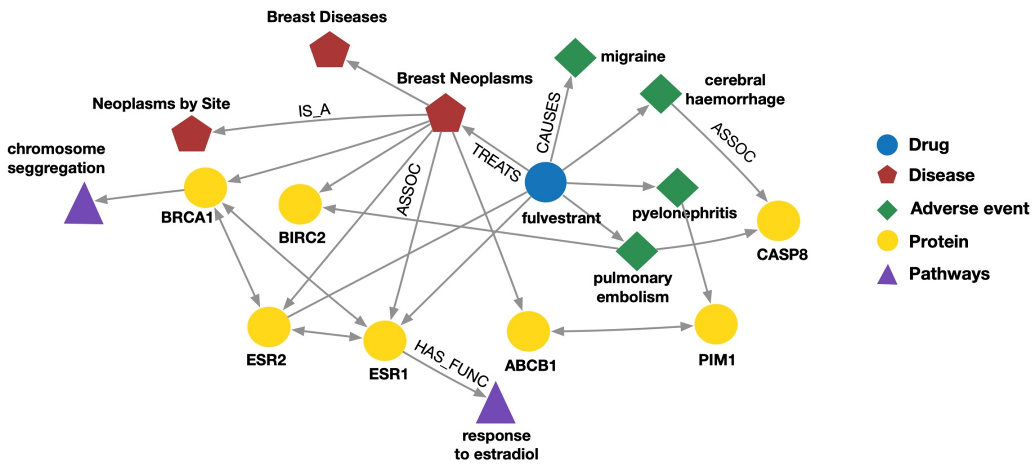

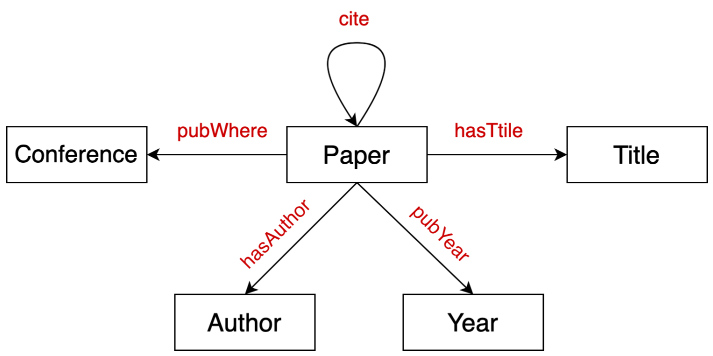

Biomedical Knowledge Graphs

Example node: Migraine
Example edge: (fulvestrant, Treats, Breast Neoplasms)
Example node type: Protein
Example edge type (relation): Causes

Academic Graphs

Example node: ICML
Example edge: (GraphSAGE, NeurIPS)
Example node type: Author
Example edge type (relation): pubYear

Choosing a Proper Representation

- **How to build a graph:**
  - **What are nodes?**
  - **What are edges?**
- **Choice of the proper network representation of a given domain/problem determines our ability to use networks successfully:**
  - In some cases, there is a unique, unambiguous representation
  - In other cases, the representation is by no means unique
  - The way you assign links will determine the nature of the question you can study

Directed vs. Undirected Graphs

Undirected

- **Links:** undirected (symmetrical, reciprocal)

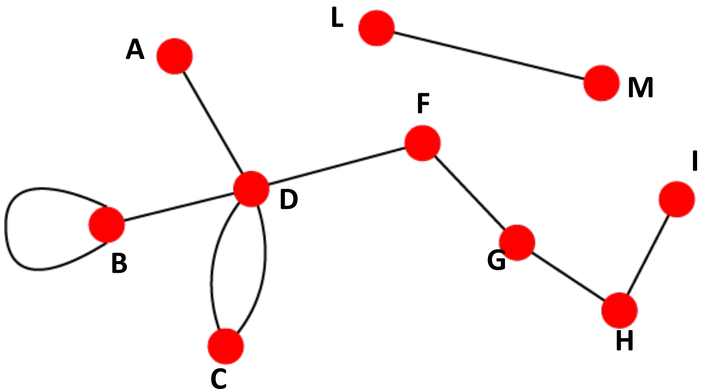

Directed

- **Links:** directed

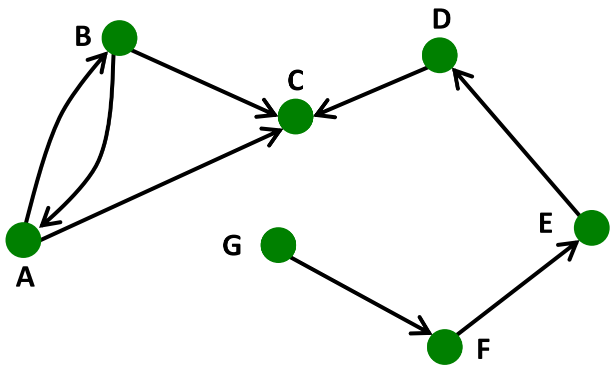

- **Other considerations:**
  - Weights
  - Properties
  - Types
  - Attributes

Bipartite Graph

- **Bipartite graph** is a graph whose nodes can be divided into two disjoint sets ***U*** and ***V*** such that every link connects a node in ***U*** to one in ***V***; that is, ***U*** and ***V*** are **independent sets**
- **Examples:**
  - Authors-to-Papers (they authored)
  - Actors-to-Movies (they appeared in)
  - Users-to-Movies (they rated)
  - Recipes-to-Ingredients (they contain)
- **"Folded" networks:**
  - Author collaboration networks
  - Movie co-rating networks

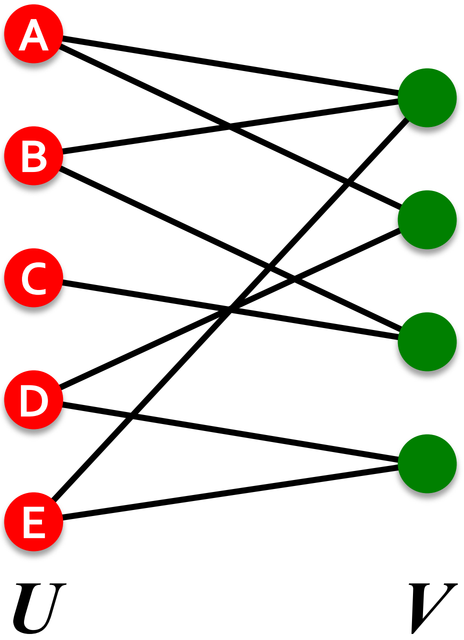

Stanford CS224W: Applications of Graph ML

CS224W: Machine Learning with Graphs
Jure Leskovec, Stanford University
http://cs224w.stanford.edu

Different Types of Tasks

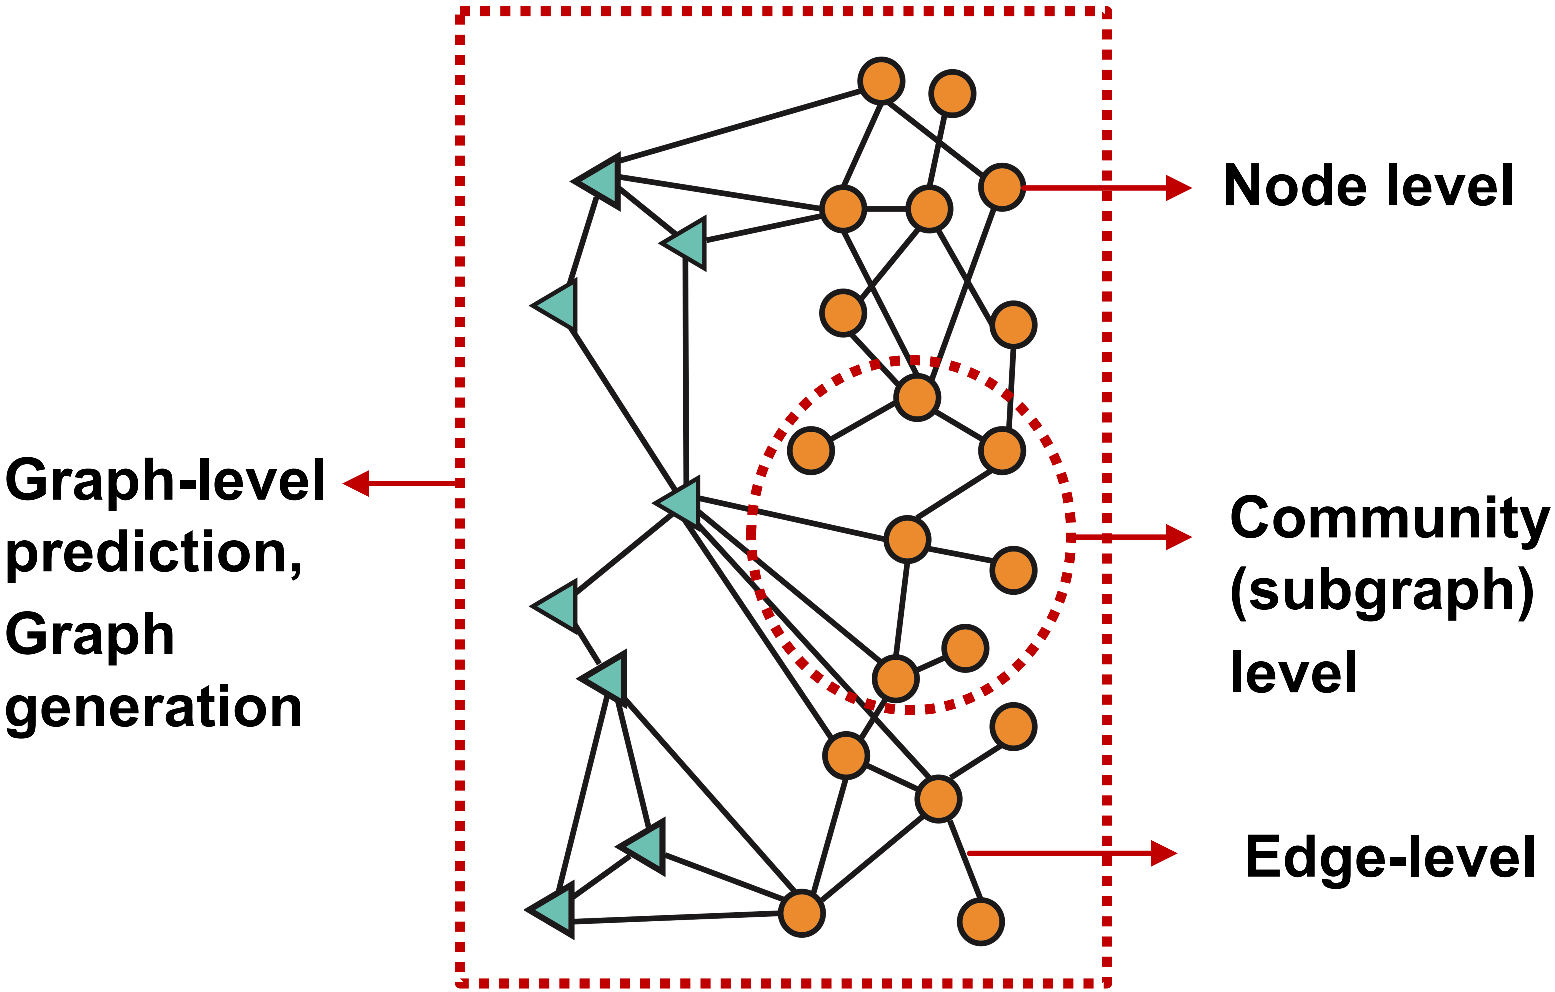

Node-Level Tasks

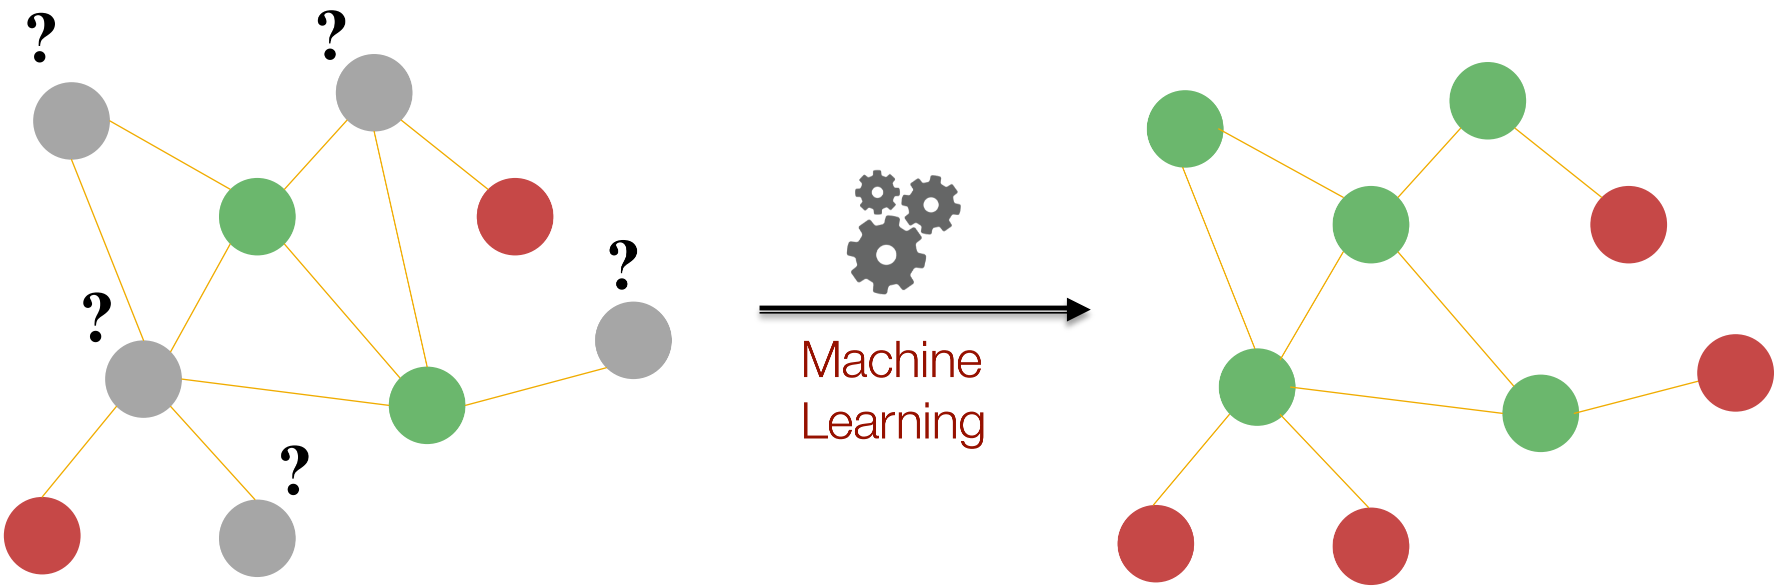

Node classification

Node-Level Network Structure

Goal: Characterize the structure and position of a node in the network:

- Node degree
- Node importance & position
  - E.g., Number of shortest paths passing through a node
  - E.g., Avg. shortest path length to other nodes
- Substructures around the node

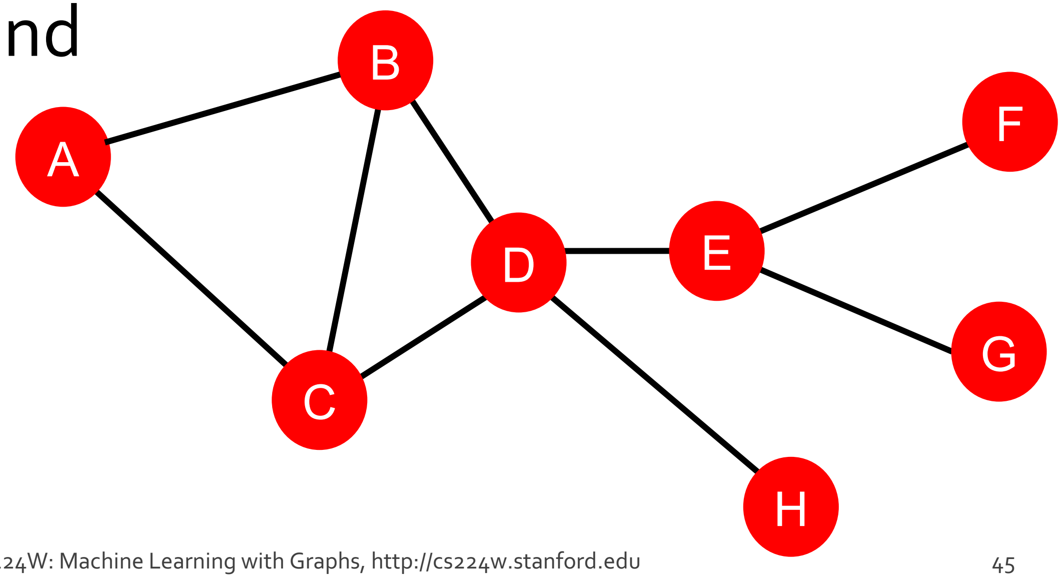

Node's Subgraphs: Graphlets

- **Graphlets:** A count vector of rooted subgraphs at a given node.
- **Example:**

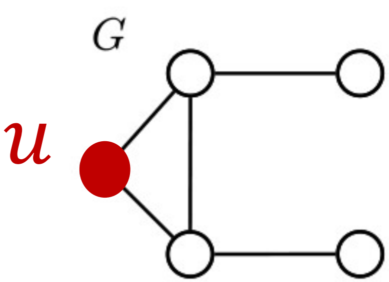

All possible graphlets on up to 3 nodes

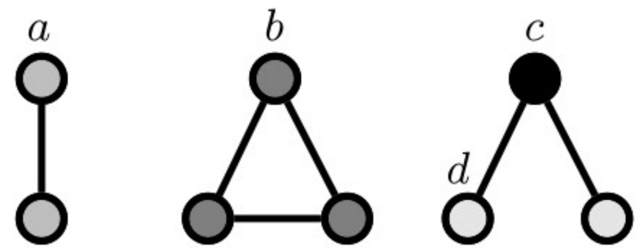

Graphlet instances of node u:

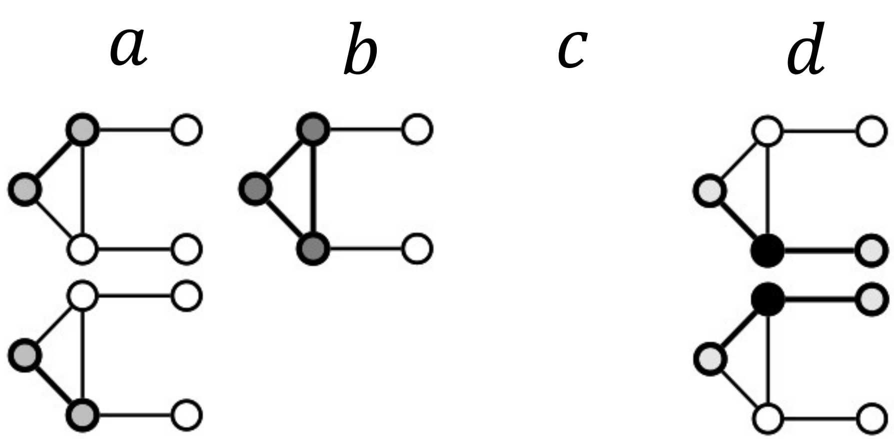

Graphlets of node u:
a, b, c, d
[2,1,0,2]

Discussion

Different ways to label nodes of the network:

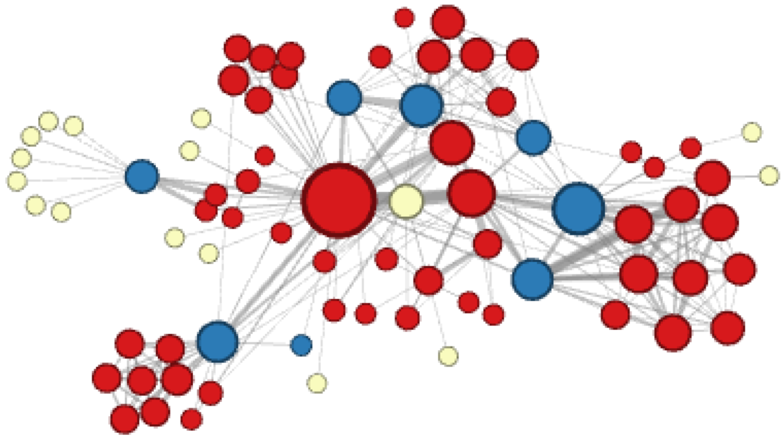

Node features defined so far would allow to distinguish nodes in the above example

However, the features defines so far would not allow for distinguishing the above node labelling

Example (1): Protein Folding

Computationally predict a protein's 3D structure based solely on its amino acid sequence:

For each node predict its 3D coordinates

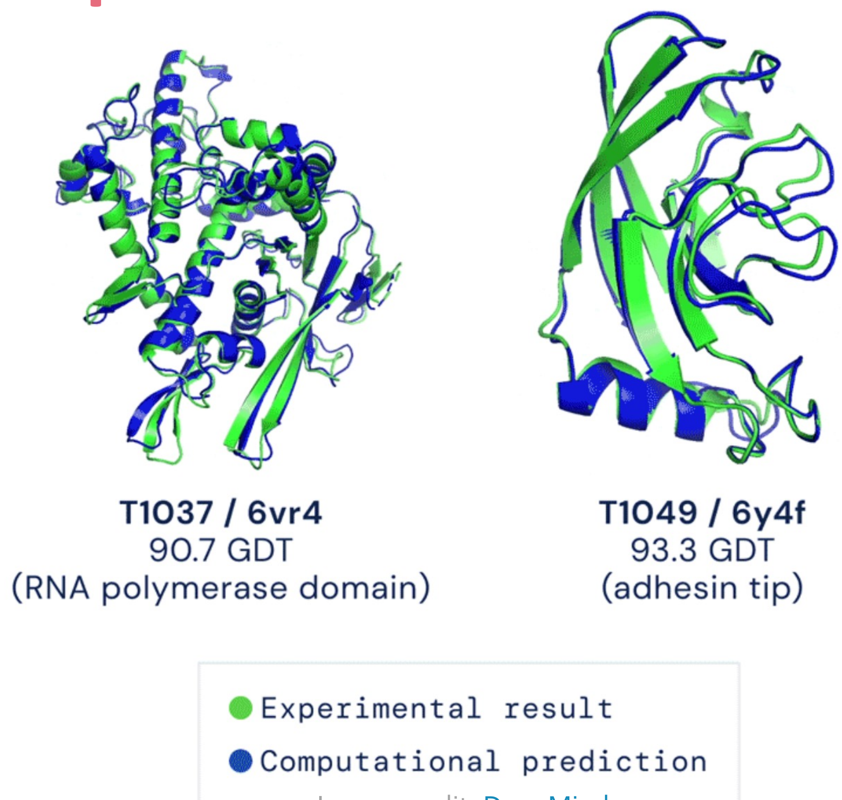

Image credit: [DeepMind]

AlphaFold: Impact

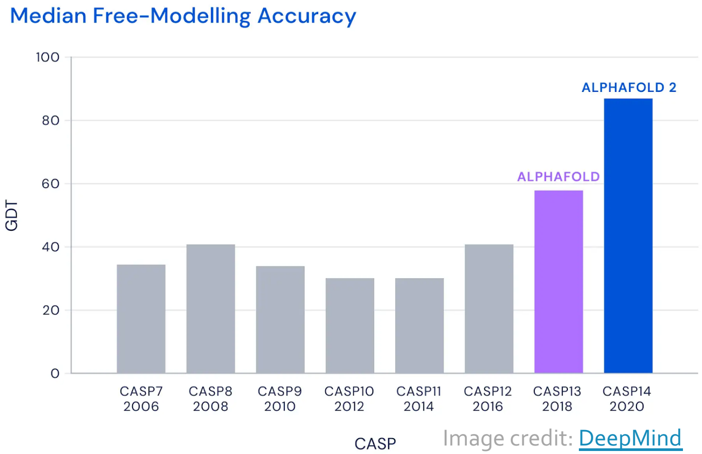

Image credit: [DeepMind]

Image credit: [SingularityHub]

AlphaFold's AI could change the world of biological science as we know it

DeepMind's latest AI breakthrough can accurately predict the way proteins fold

Has Artificial Intelligence 'Solved' Biology's Protein-Folding Problem?

12-14-20

DeepMind's latest AI breakthrough could turbocharge drug discovery

AlphaFold: Solving Protein Folding

- **Key idea:** "Spatial graph"
  - **Nodes:** Amino acids in a protein sequence
  - **Edges:** Proximity between amino acids (residues)

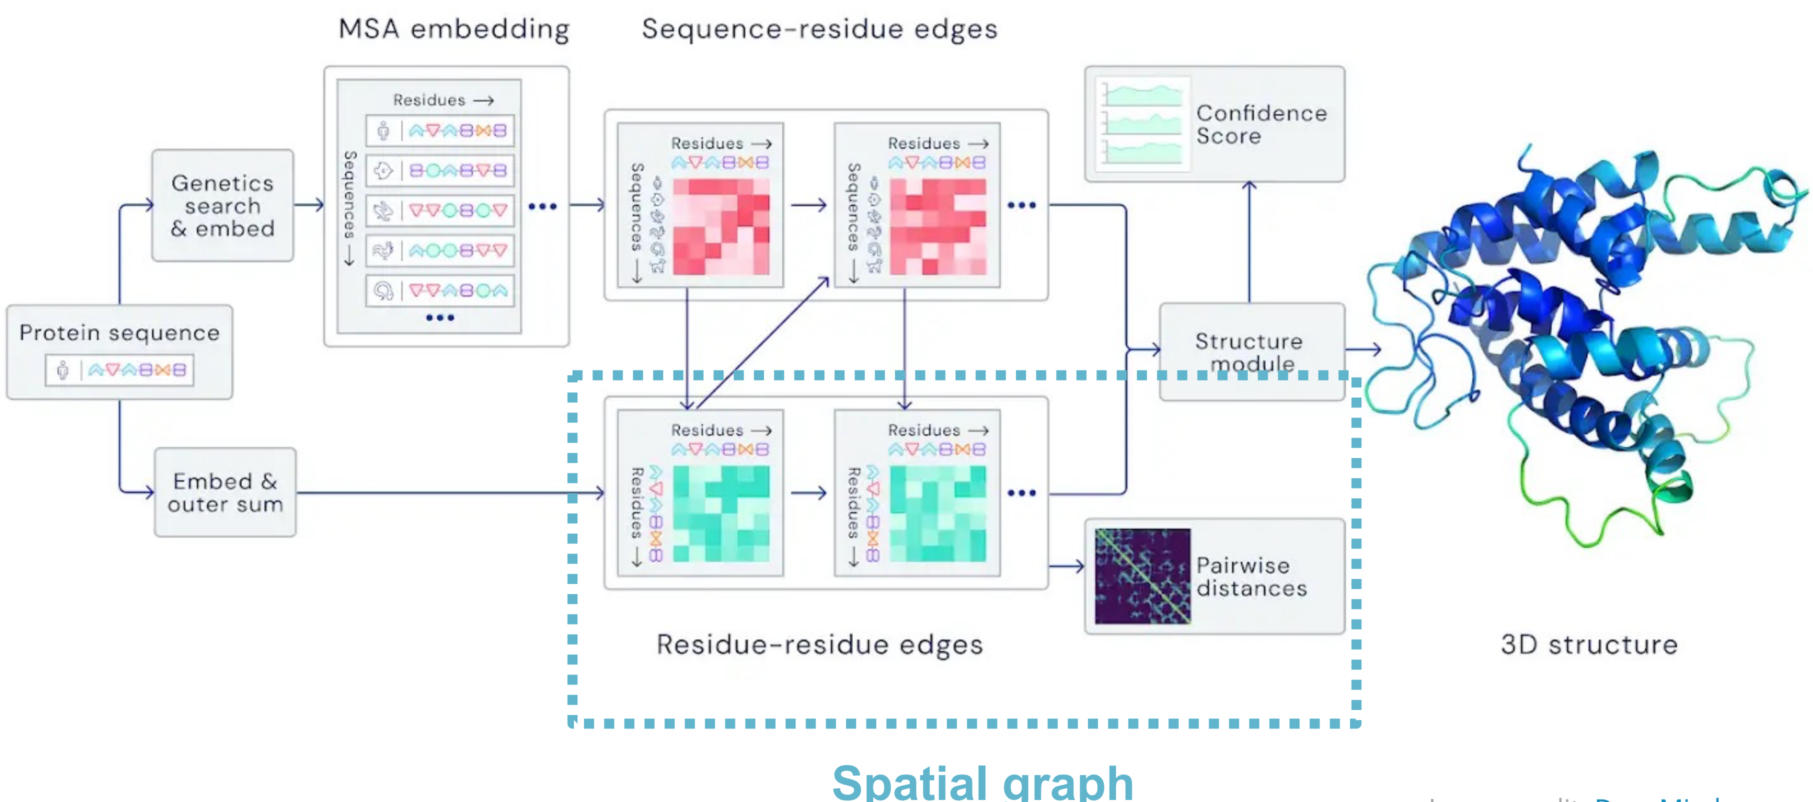

Image credit: [DeepMind]

Stanford CS224W: Link Prediction

CS224W: Machine Learning with Graphs

Jure Leskovec, Stanford University

http://cs224w.stanford.edu

Link-Level Prediction Task

- The task is to predict **new/missing/unknown links** based on the existing links.
- At test time, node pairs (with no existing links) are ranked, and top K node pairs are predicted.
- Task: Make a prediction for **a pair of nodes**.

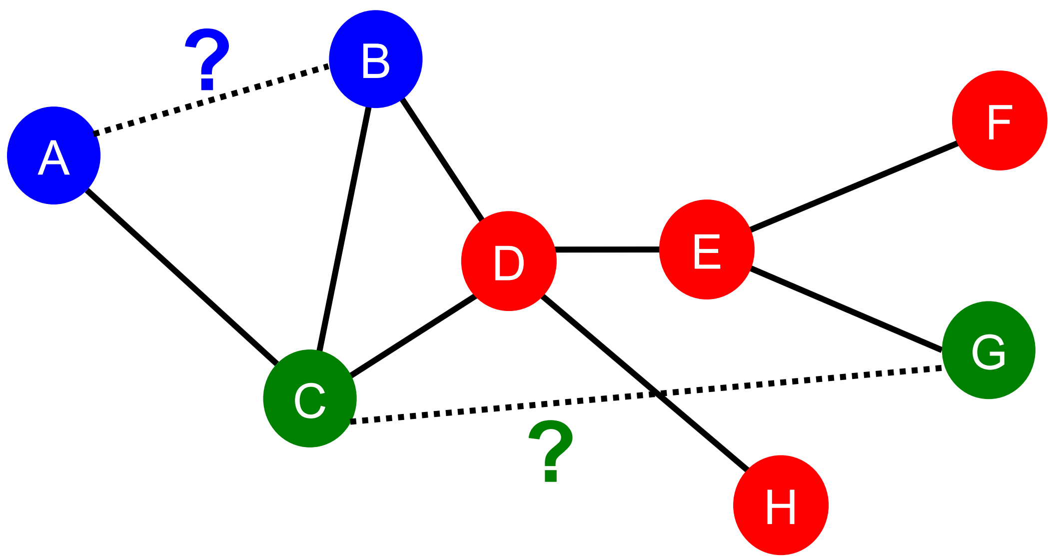

Example (1): Recommender Systems

- **Users interacts with items**
  - Watch movies, buy merchandise, listen to music
  - **Nodes:** Users and items
  - **Edges:** User-item interactions
- **Goal: Recommend items users might like**

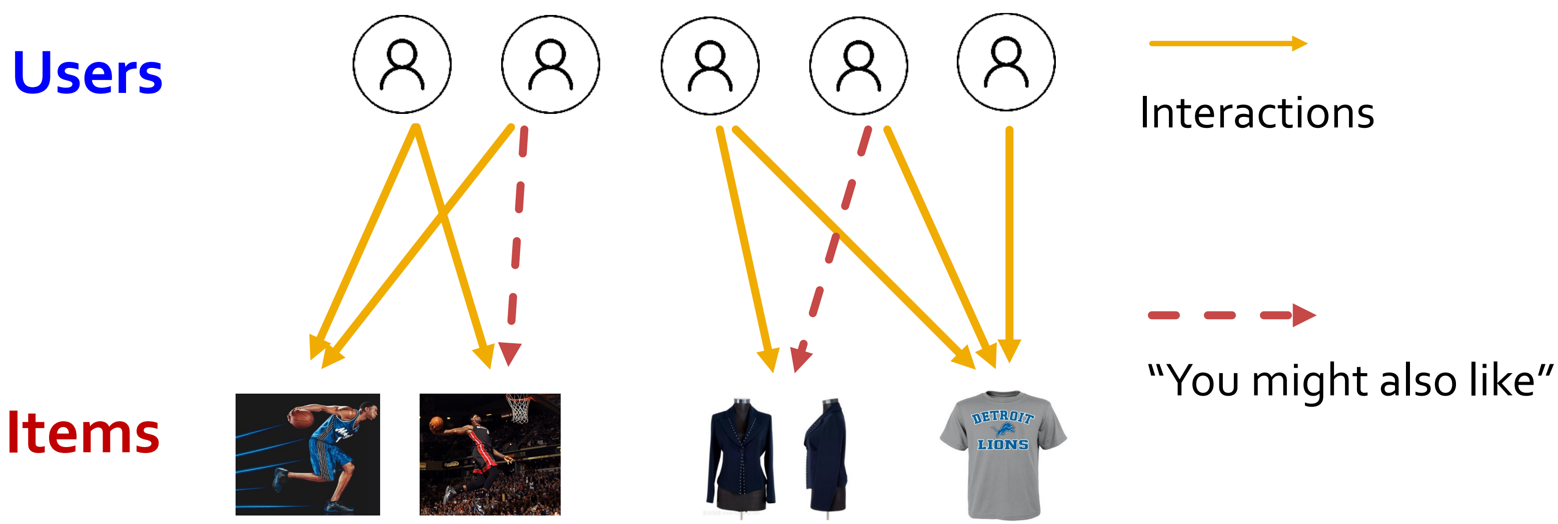

Ying et al., Graph Convolutional Neural Networks for Web-Scale Recommender Systems, KDD 2018

PinSage: Graph-based Recommender

Task: Recommend related pins to users

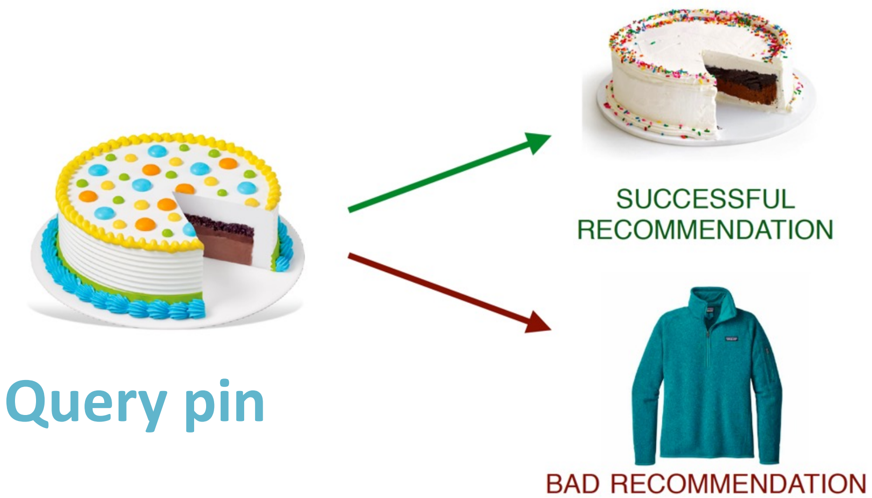

Task: Learn node embeddings z_i such that

$$d(z_{cake1}, z_{cake2}) < d(z_{cake1}, z_{sweater})$$

Predict whether two nodes in a graph are related

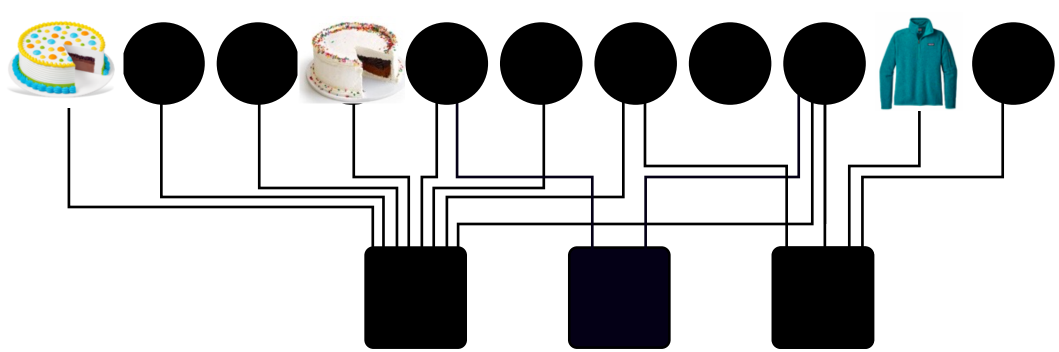

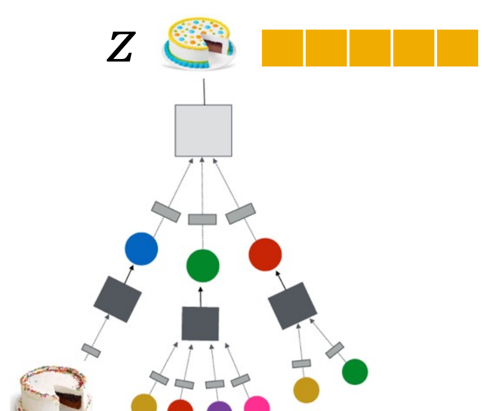

Stanford CS224W: Graph-Level Tasks

CS224W: Machine Learning with Graphs

Jure Leskovec, Stanford University

http://cs224w.stanford.edu

Graph-Level Prediction

- **Goal:** We want make a prediction for an entire graph or a subgraph of the graph.

- **For example:**

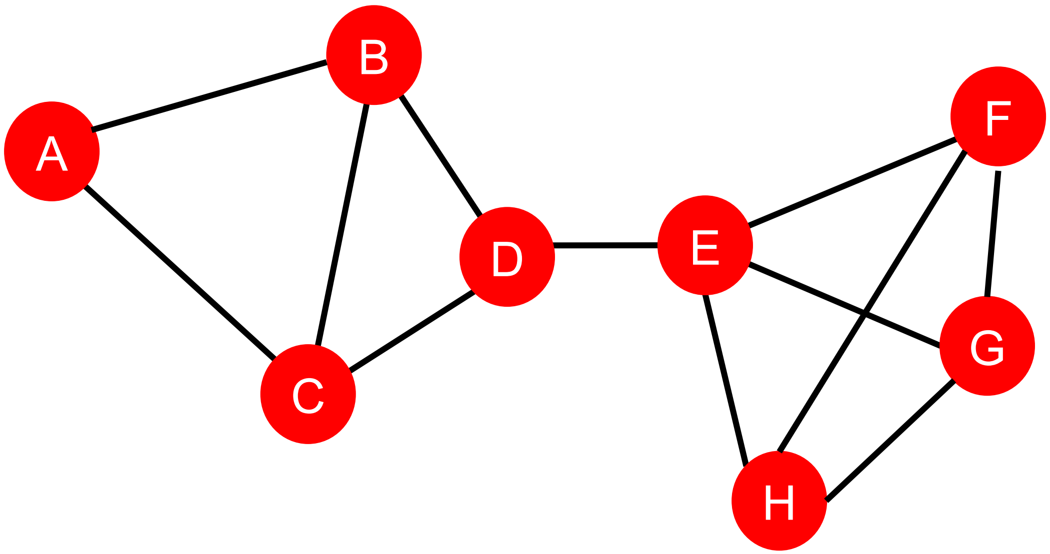

Example (1): Traffic Prediction

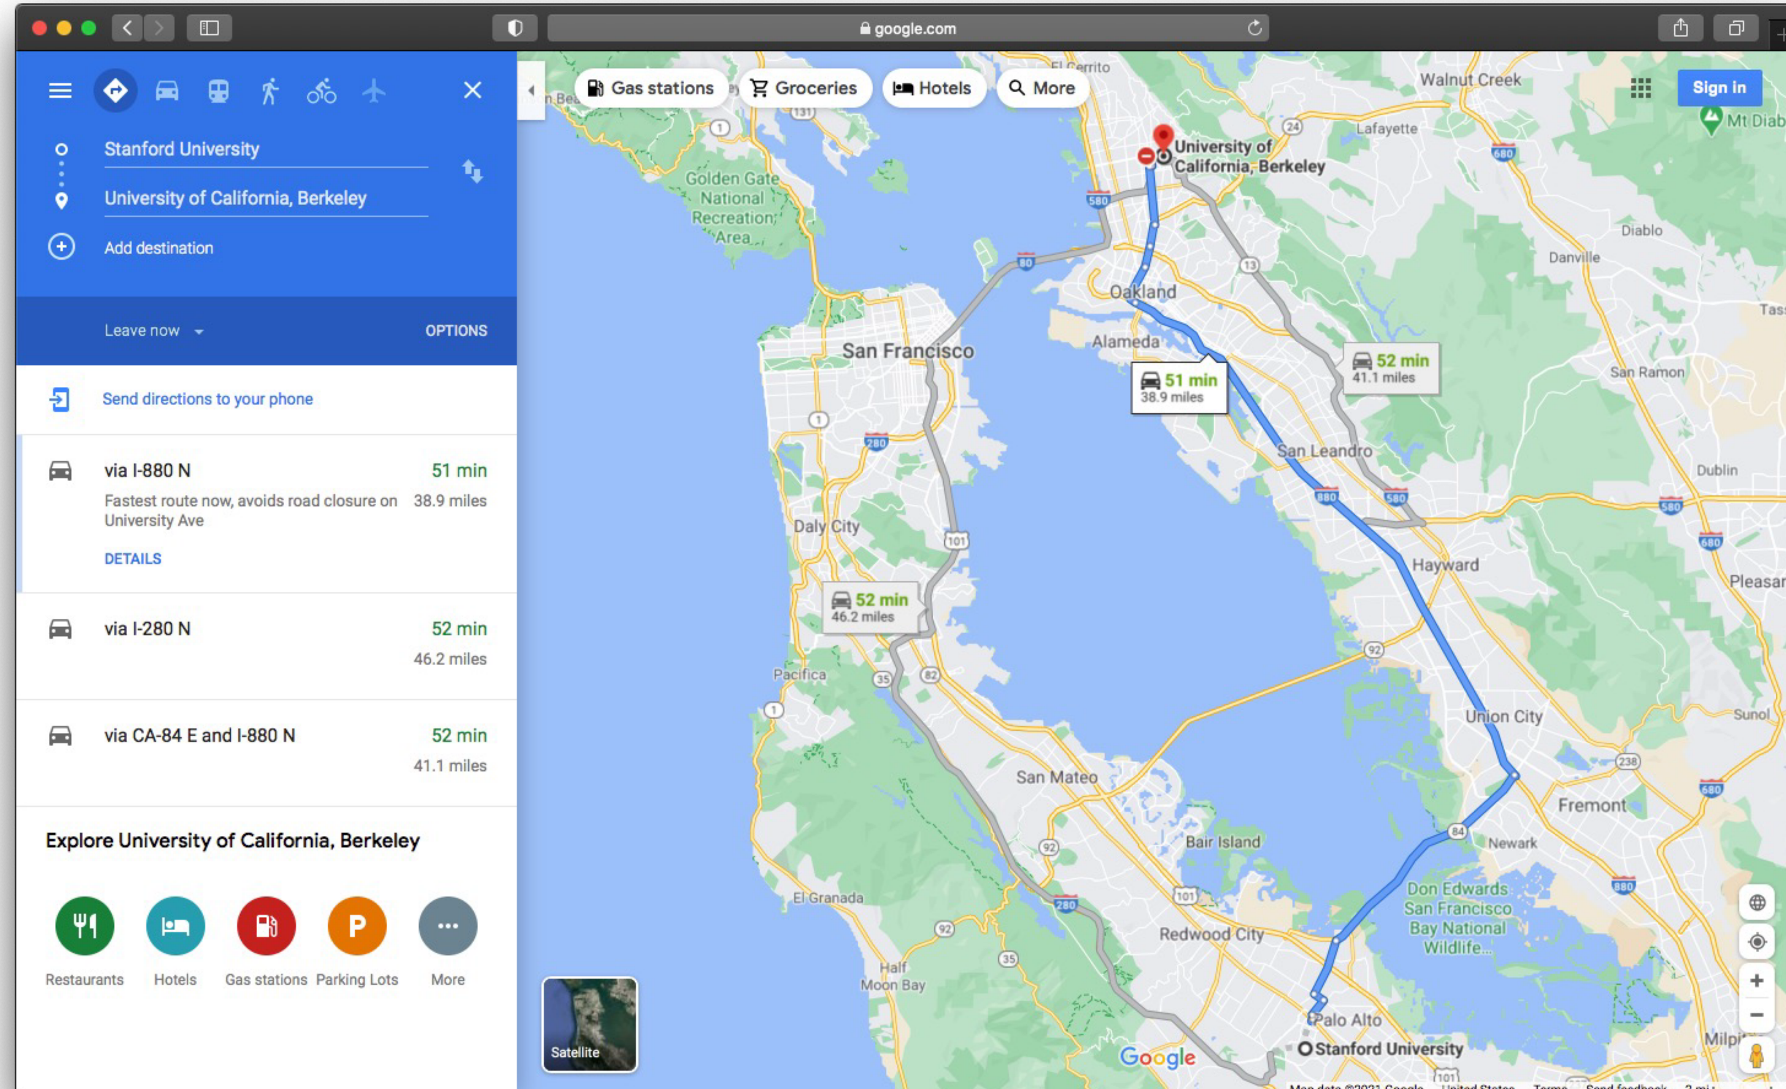

Road Network as a Graph

- **Nodes:** Road segments
- **Edges:** Connectivity between road segments
- **Prediction:** Time of Arrival (ETA)

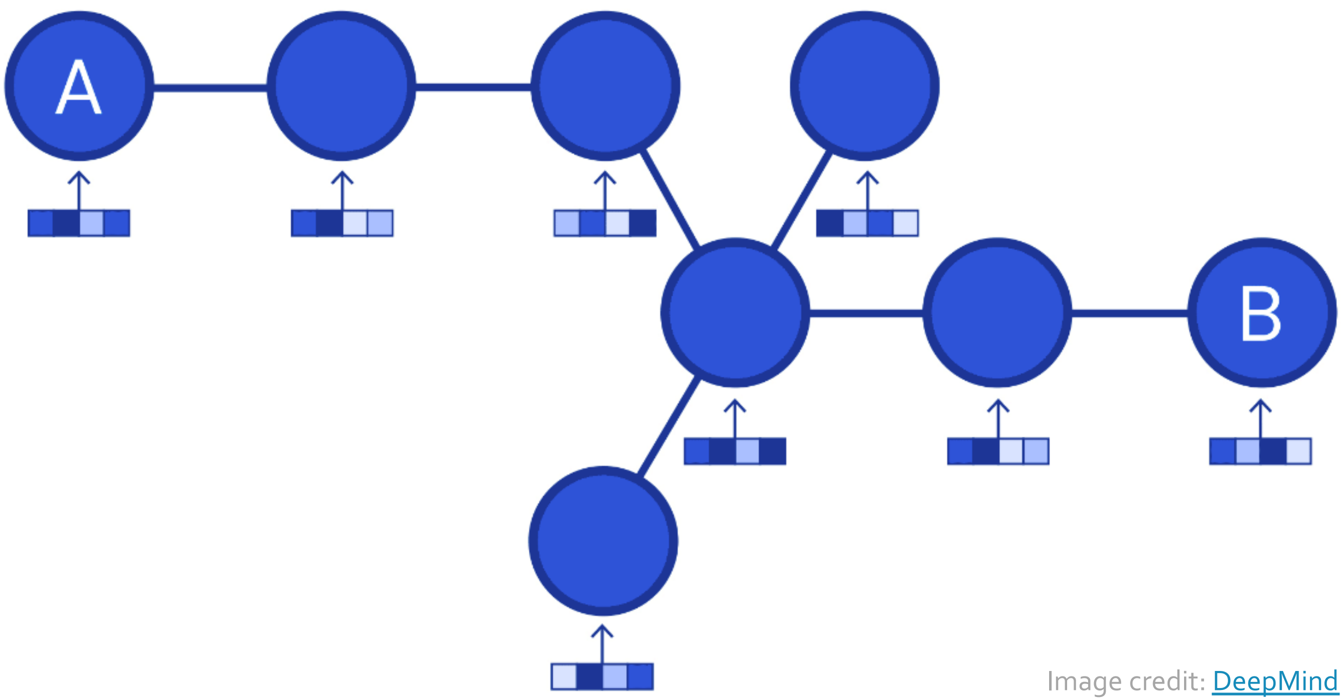

Image credit: DeepMind

Traffic Prediction via GNN

Predicting Time of Arrival with Graph Neural Networks

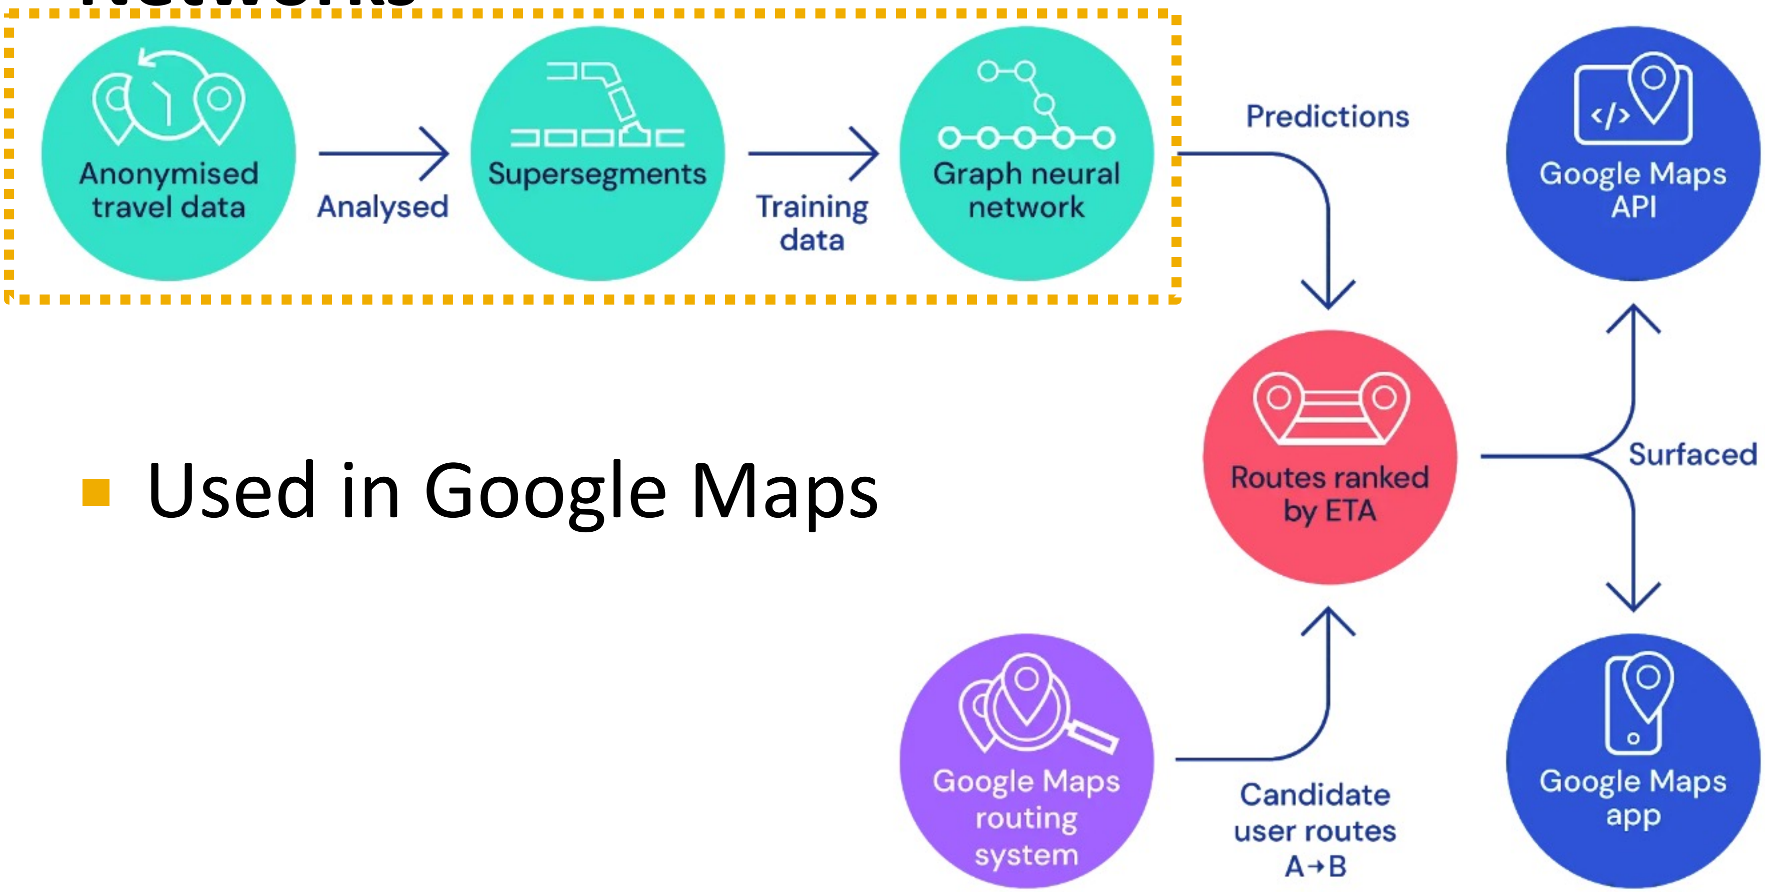

- Used in Google Maps

THE MODEL ARCHITECTURE FOR DETERMINING OPTIMAL ROUTES AND THEIR TRAVEL TIME.

Image credit: DeepMind

Example (2): Drug Discovery

- **Antibiotics are small molecular graphs**
  - **Nodes:** Atoms
  - **Edges:** Chemical bonds

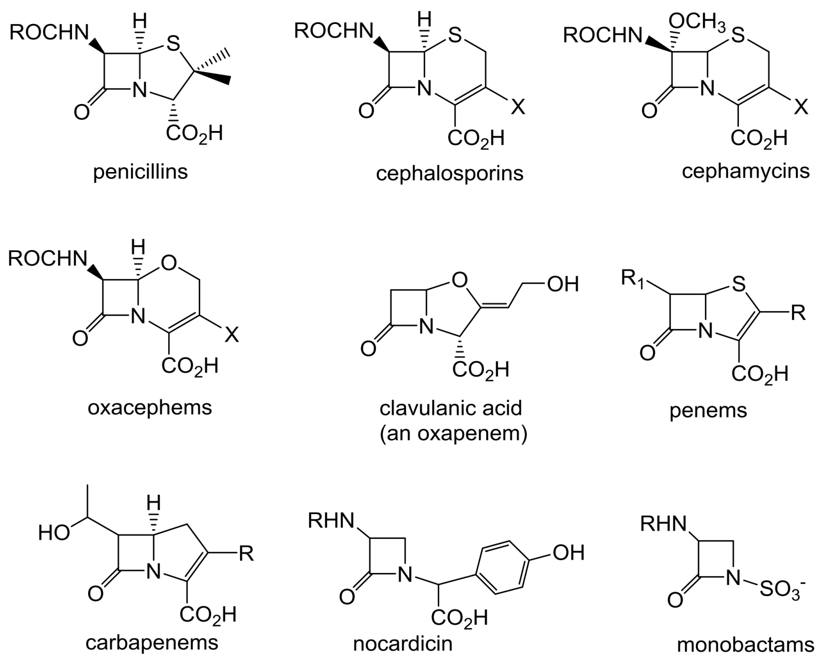

Konaklieva, Monika I. "Molecular targets of β-lactam-based antimicrobials: beyond the usual suspects." Antibiotics 3.2 (2014): 128-142.

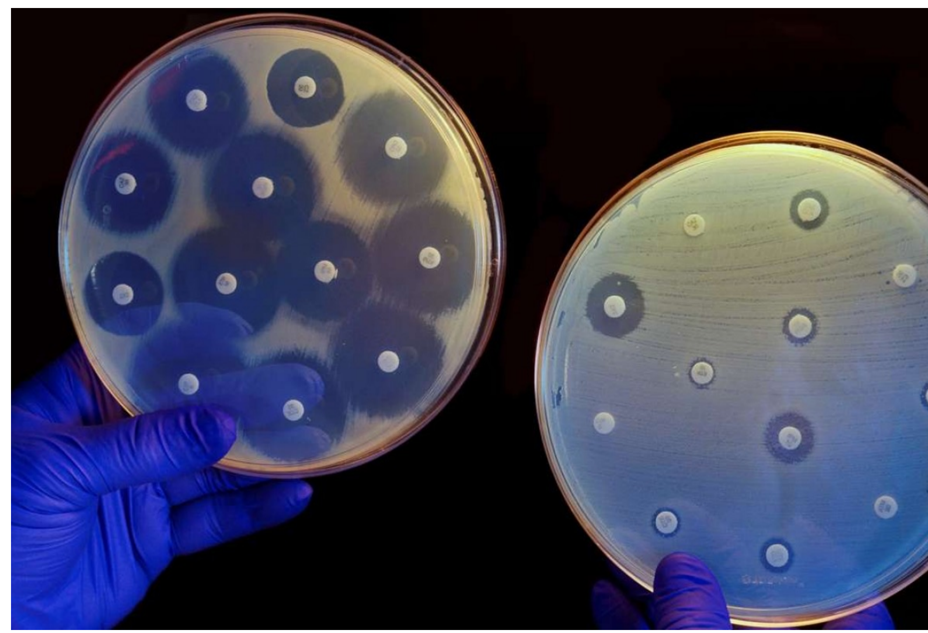

Image credit: [CNN]

Stokes et al., [A Deep Learning Approach to Antibiotic Discovery](), Cell 2020

Deep Learning for Antibiotic Discovery

- A Graph Neural Network **graph classification model**
- Predict promising molecules from a pool of candidates

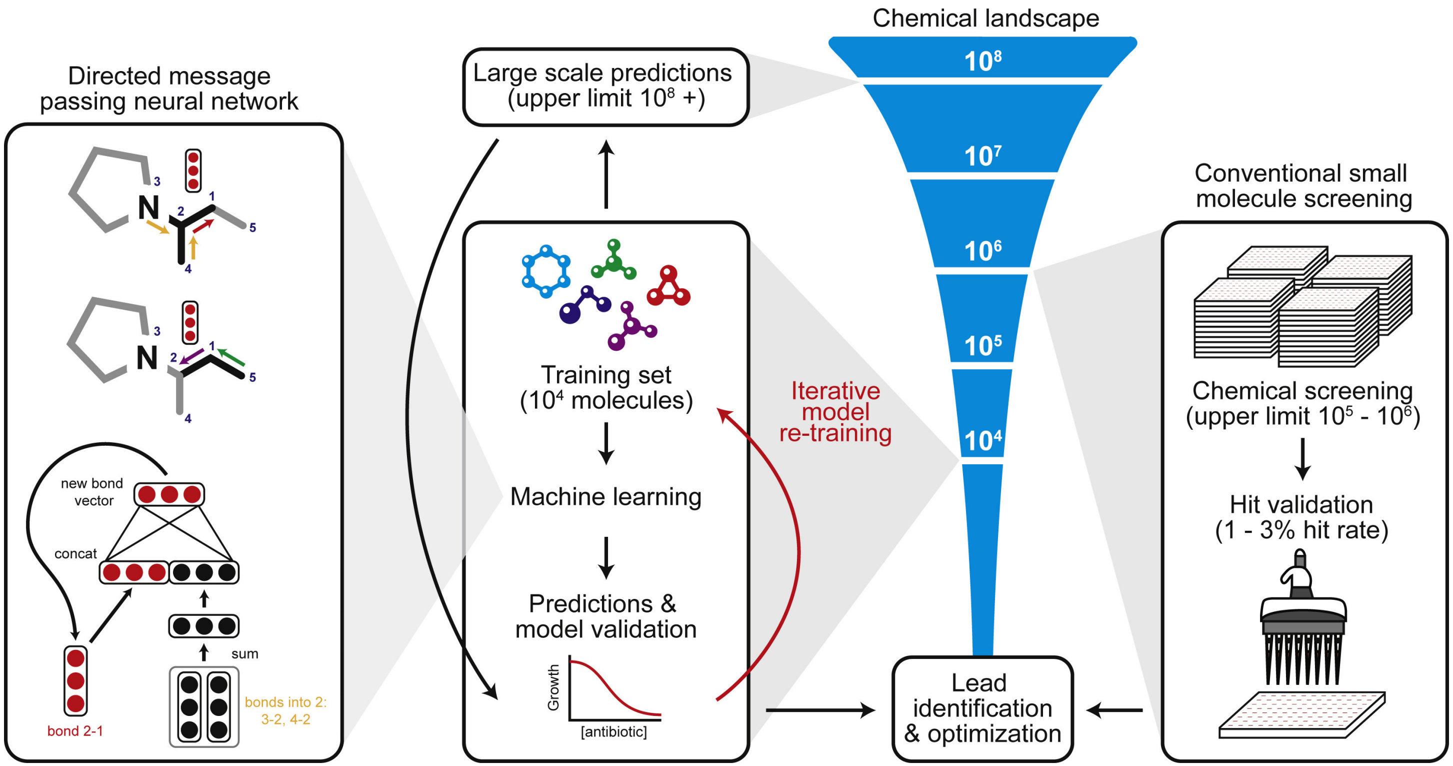

Stokes, Jonathan M., et al. "A deep learning approach to antibiotic discovery." Cell 180.4 (2020): 688-702.

Summary

ML in the language of graphs:

- Node-level:
  - Churn
  - Life-time value
  - Next best action
- Link-level:
  - Product affinity
  - Recommendations
- Graph-level:
  - Fraud, money laundering