CS6208 : Advanced Topics in Artificial Intelligence
Graph Machine Learning

Lecture 1 : Introduction to Graph Machine Learning

Semester 2 2022/23

Xavier Bresson

https://twitter.com/xbresson

Department of Computer Science
National University of Singapore (NUS)

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

AI : Birth, winter and renaissance

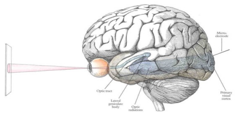

Visual primary cortex
Hubel-Wiesel
1959

1958
Perceptron
Rosenblatt

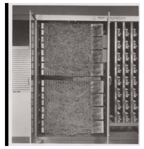

1962
Birth of
Data Science
Split from Statistics
Tukey

1986
Backprop
Hinton-Rumelhart

1987
Neocognitron
Fukushima

First
NeurIPS

First
KDD
1989

1995
SVM/Kernel techniques
Vapnik

RNN/LSTM
Schmidhuber
1997

CNN
LeCun-Bengio
1998

1999
First NVIDIA
GPU

First Amazon
Cloud Center
2006

Big Data
Volume 2x/1.5 year

Hardware
GPU speed 2x/1.5 year

Kaggle
Platform
2010

2012
Deep Learning
Breakthrough
Hinton-Krizhevsky

GAN, Goodfellow
ResNet, He
Deep RL, Mnih
Attention, Bengio
2014
Facebook Center
OpenAI Center

Google AI
TensorFlow
Facebook AI
PyTorch
2015
Data scientist
1st Job in US

4th industrial revolution?
Digital intelligence
revolution
or new AI bubble?

AI Birth

AI Winter [1960's-2012]

Kernel techniques
Graphical models
Wavelets, numerical PDEs
Handcrafted Features

AI Renaissance

Supervised big data

- Google's Law of Data : Volume of data double every 1.5 years.
  - Eric Schmidt : "Every 2 days we create as much information as we did up to 2003."

- Data generation/collection is easy but data labeling is challenging.
  - It is time and money consuming.
  - It requires humans (machine ability bounded by human intelligence).

- Why do we need labeled/annotated data? Because supervised deep learning is designed to predict labels.

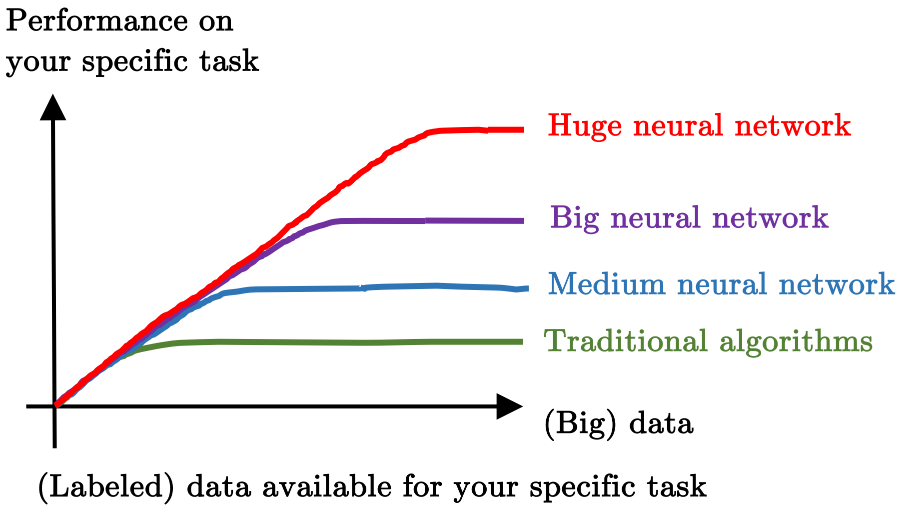

Self-supervised big data

- How much labeled data to achieve super-human performances?
  - ImageNet for image recognition has 1.4M images (machine error is 3% and human error is 5%).
  - AlphaGo Zero played 4.9 million games.
  - Self-driving cars : Tesla's self-driving cars collected 2 billion miles in 2019 (and no autonomous cars so far).

- The next AI revolution will not be supervised !
  - It will be self-supervised with e.g. LLMs/ChatGPT.

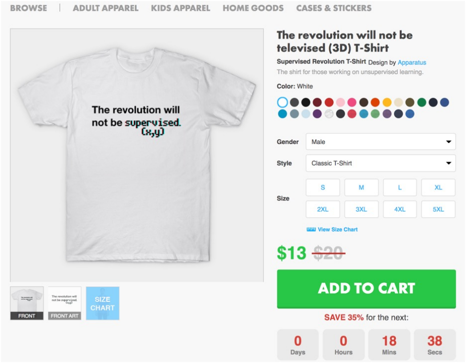

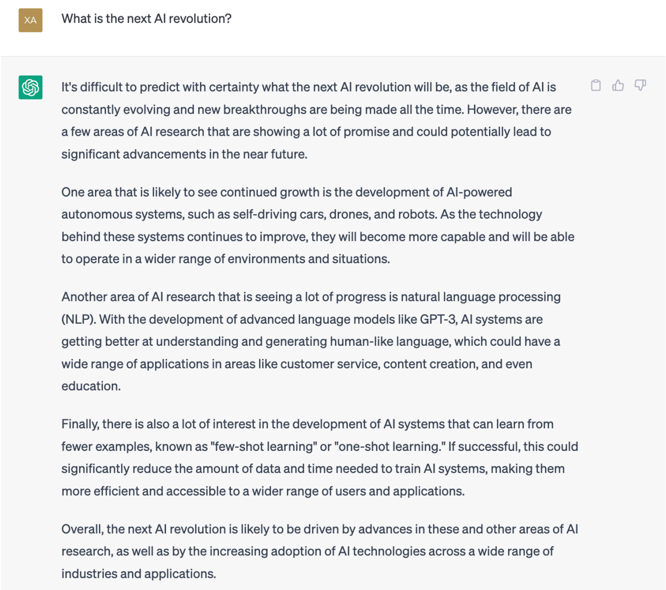

XA What is the next AI revolution?

It's difficult to predict with certainty what the next AI revolution will be, as the field of AI is constantly evolving and new breakthroughs are being made all the time. However, there are a few areas of AI research that are showing a lot of promise and could potentially lead to significant advancements in the near future.

One area that is likely to see continued growth is the development of AI-powered autonomous systems, such as self-driving cars, drones, and robots. As the technology behind these systems continues to improve, they will become more capable and will be able to operate in a wider range of environments and situations.

Another area of AI research that is seeing a lot of progress is natural language processing (NLP). With the development of advanced language models like GPT-3, AI systems are getting better at understanding and generating human-like language, which could have a wide range of applications in areas like customer service, content creation, and even education.

Finally, there is also a lot of interest in the development of AI systems that can learn from fewer examples, known as "few-shot learning" or "one-shot learning." If successful, this could significantly reduce the amount of data and time needed to train AI systems, making them more efficient and accessible to a wider range of users and applications.

Overall, the next AI revolution is likely to be driven by advances in these and other areas of AI research, as well as by the increasing adoption of AI technologies across a wide range of industries and applications.

Computational power

- 1975 Moore's Law of Computers :
  - 100x faster every decade
  - CPU, GPU/TPU, cloud computing
  - Brain : 10^{14} synapses/10^3 Tflops = (standard) computers in 2030 (100 times faster than today)

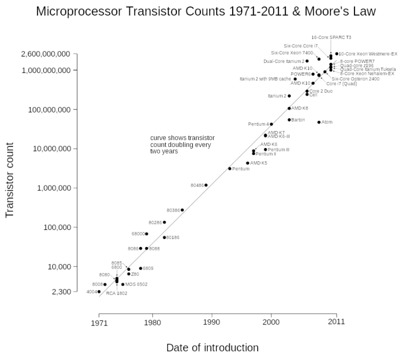

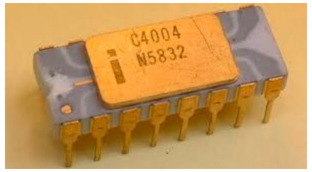

1st Intel CPU 1971

1st Nvidia GPU in 1999

1st Amazon Cloud in 2006

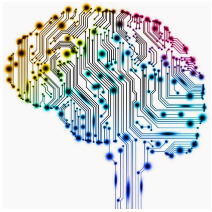

2030
=

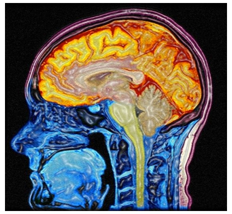

Neural network architectures

- Most research works focus on deep learning architectures (along with training and optimization)
  - Industry labs leverage academic ideas to large-scale/impressive experiments (e.g. generative models)
- Still, no architecture can solve simultaneously many tasks like the human brain but progress has been made e.g. Transformers.

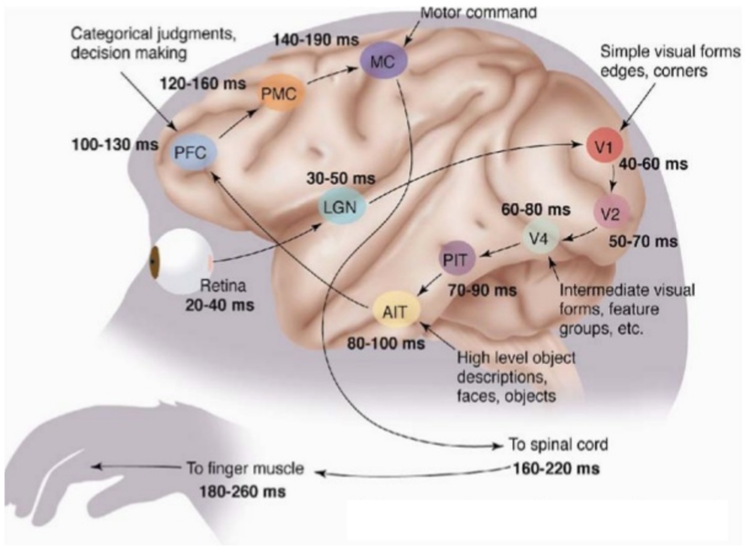

Biological Neural Network

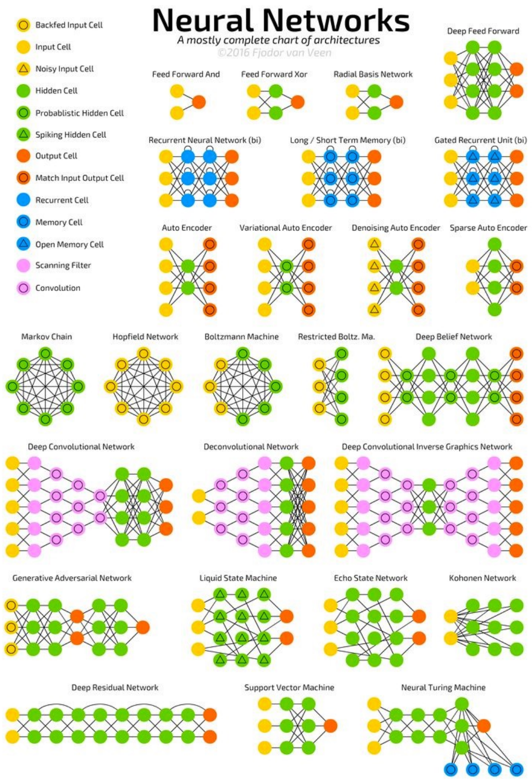

Artificial Neural Networks

The deep learning revolution

- Breakthrough in Computer Vision (CNNs)

  LeCun, Bottou, Bengio, Haffner 1998
  Krizhevsky, Sutskever, Hinton, 2012

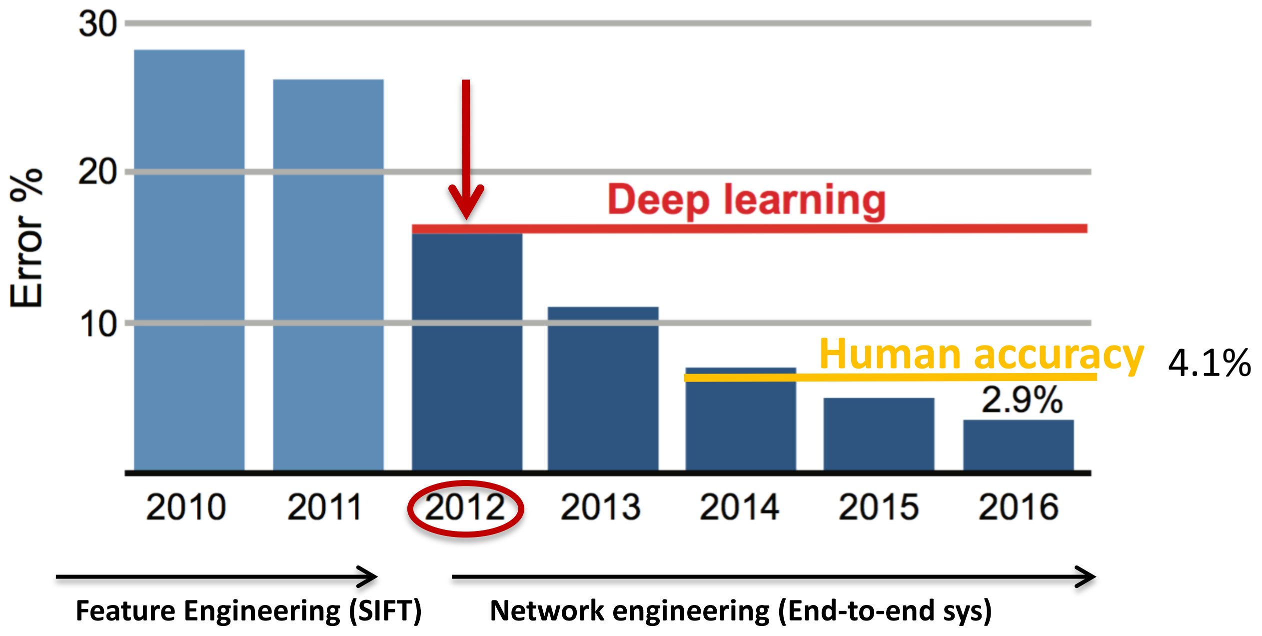

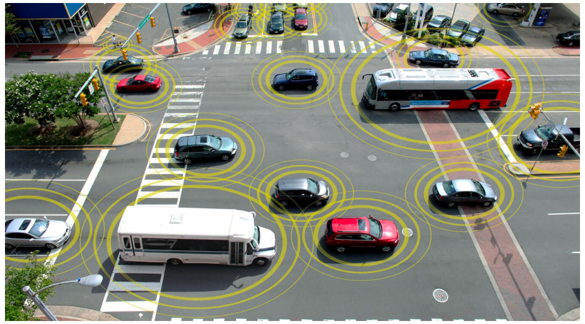

- Breakthrough in Speech and Natural Language Processing (RNNs/Transformers)

  Hochreiter, Schmidhuber, 1997 and Vaswani et-al, 2017

The power of deep learning

- CNNs/RNNs/Transformers are powerful architectures to solve high-dimensional learning problems.

- Curse of dimensionality :

  dim(image) = dim(1000 x 1000) = 10^6

  For N=10 samples/dim ⇒ $10^{1,000,000}$ points

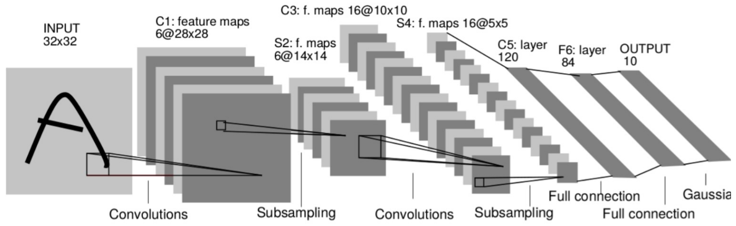

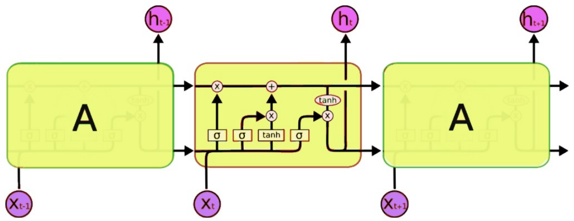

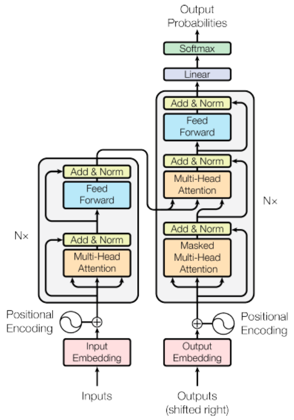

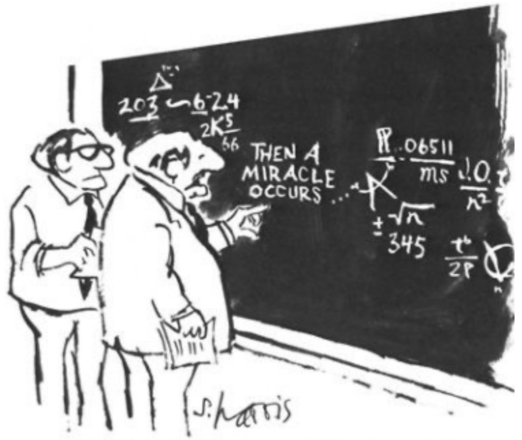

"I think you should be more explicit here in step two."

Applications

- Computer Vision : Autonomous driving, Face recognition
- Natural Language Processing : Machine Translation, Text generation, Chatbot (ChatGPT)
- Speech Recognition : Virtual assistants (Alexa/Siri/Google/Cortana)

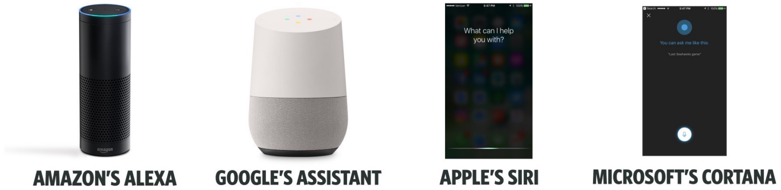

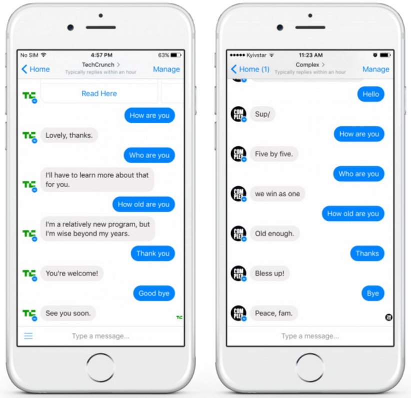

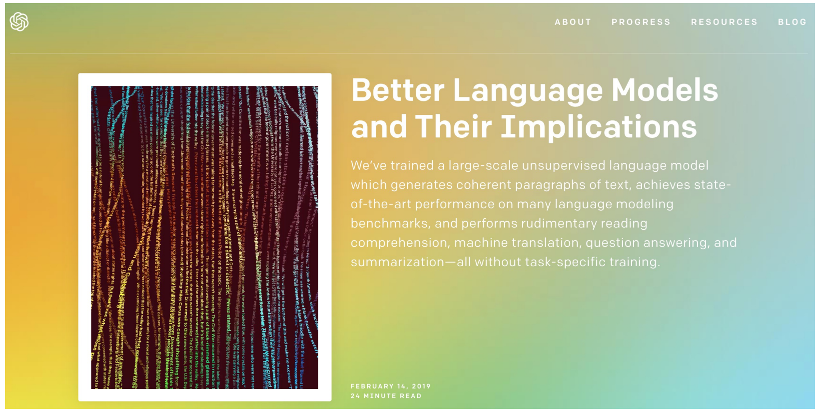

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

Limitations of deep learning

- Standard DL such as:
  - CNNs, RNNs, Transformers
  - Modern DL toolboxes (TensorFlow, PyTorch)
  - Successful products for CV/NLP/SR (Google, Meta, etc.)

  but limited to work with grid/sequence data.

- Do we have data NOT based on grid/sequence?

  Yes, quite a lot! (see next slide)

- How to generalize NNs beyond CV/NLP/SR?

- How to design universal and broadly applicable architectures?
  - A solution is Graph Neural Networks (GNNs)
  - With universal architecture (unified all NNs) – New NN generation

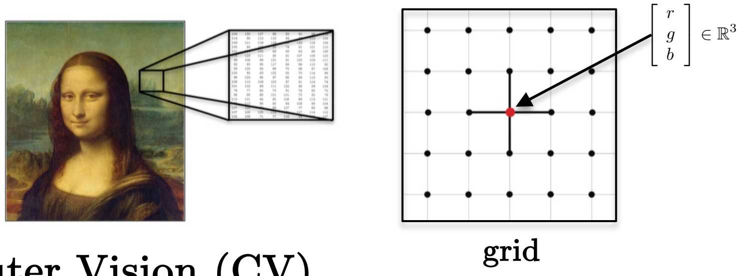

Computer Vision (CV)

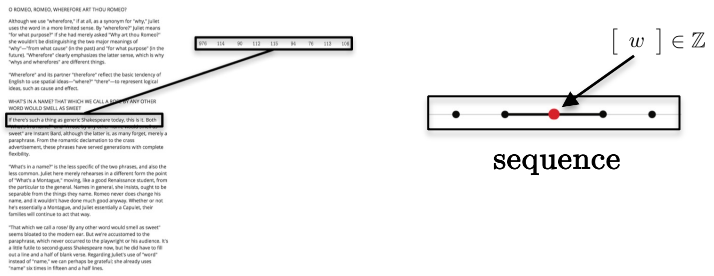

Natural Language Processing (NLP)

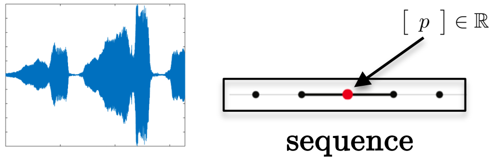

Speech Recognition (SR)

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- **Graph-Structured data**
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

Graph-structured data

Social networks (Advertisement/ recommendation)

User$_i$ $\begin{bmatrix} \text{messages} \\ \text{images} \\ \text{videos} \end{bmatrix}_i \in \mathbb{R}^d$

User connection$_{ij}$ $A_{ij} = \begin{cases} 1 & \text{if ij friends} \\ 0 & \text{otherwise} \end{cases}$

User$_j$ $\begin{bmatrix} \text{messages} \\ \text{images} \\ \text{videos} \end{bmatrix}_j \in \mathbb{R}^d$

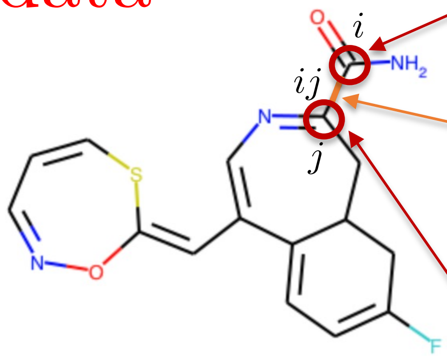

Quantum Chemistry (novel molecules for drugs and materials)

Atom$_i$ $\begin{bmatrix} \text{type} \\ \text{coordinates} \\ \text{charge} \end{bmatrix}_i \in \mathbb{R}^{d_v}$

Bond$_{ij}$ $A_{ij} = \begin{cases} 1 & \text{if ij bond} \\ 0 & \text{otherwise} \end{cases}$

$\begin{bmatrix} \text{type} \\ \text{energy} \end{bmatrix}_{ij} \in \mathbb{R}^{d_e}$

Atom$_j$ $\begin{bmatrix} \text{type} \\ \text{coordinates} \\ \text{charge} \end{bmatrix}_j \in \mathbb{R}^{d_v}$

Brain connectivity (sMRI)

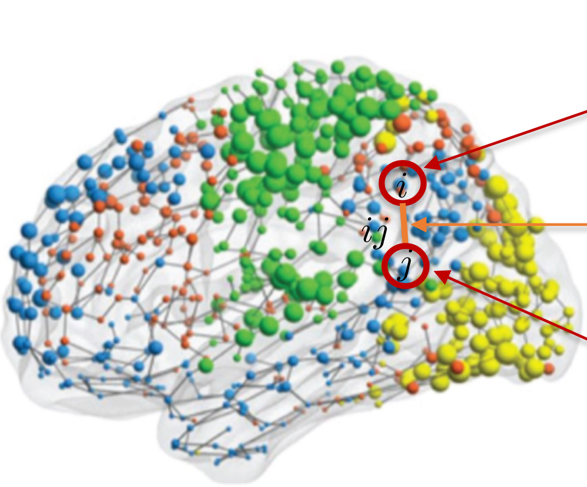

Brain analysis (Neuroscience/neuro-diseases)

ROI$_i$ $\begin{bmatrix} a_1 \\ \vdots \\ a_T \end{bmatrix}_i \in \mathbb{R}^T$

Cerebral connection$_{ij}$ $A_{ij} \in \mathbb{R}_+$

ROI$_j$ $\begin{bmatrix} a_1 \\ \vdots \\ a_T \end{bmatrix}_j \in \mathbb{R}^T$

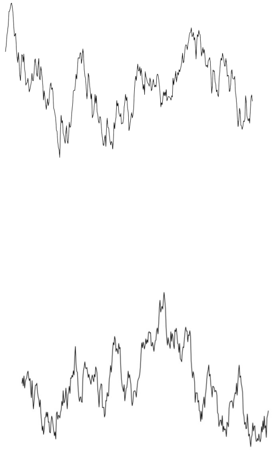

Functional activations (fMRI)

Graph-structured data

Social networks (Advertisement)

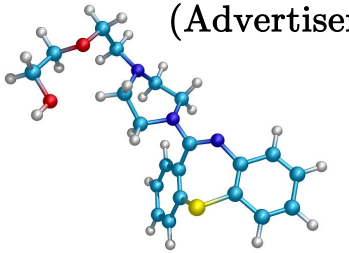

Drug/Material molecules (Chemistry)

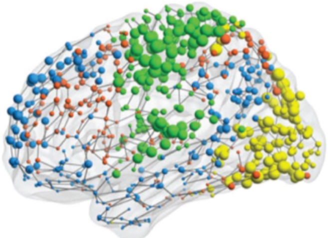

Brain connectivity (Neuroscience)

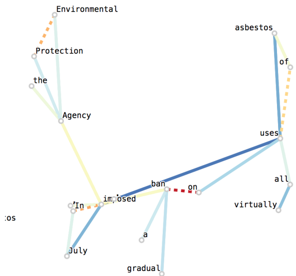

Word relationships (NLP)

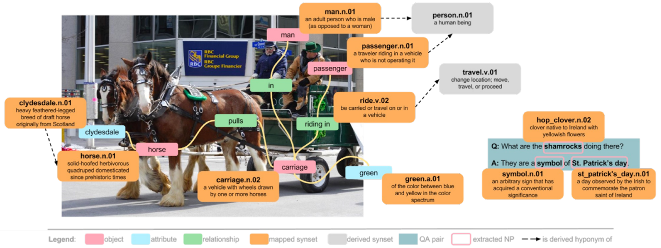

Scene understanding

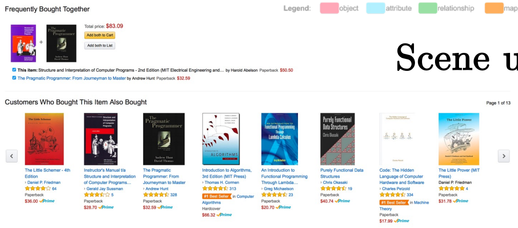

Recommender systems (Amazon, Netflix)

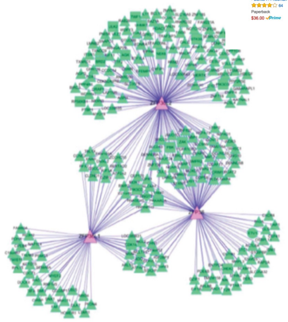

Gene Regulatory Network

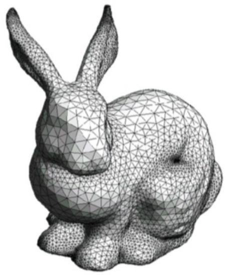

3D Meshes (Computer Graphics)

Transportation networks

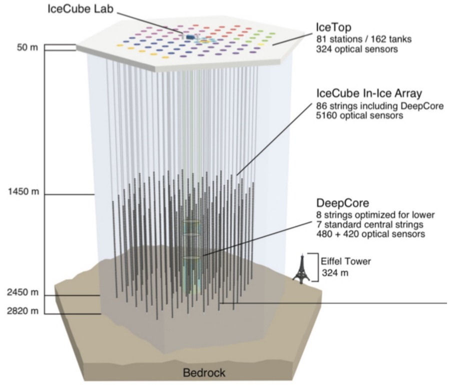

Neutrino detection (High-energy Physics)

=

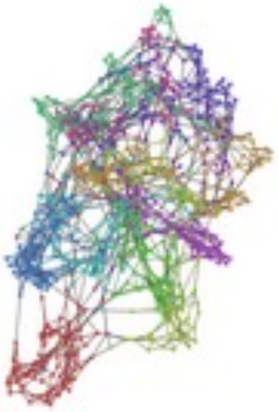

Graph

Graph-structured data

- A graph G is defined by :
  - Vertices V
  - Edges E
  - Adjacency matrix A

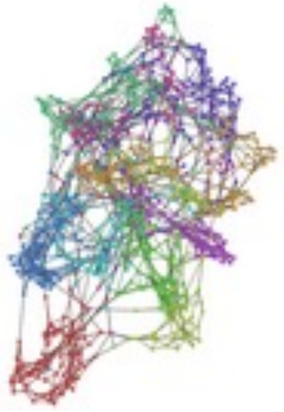

=

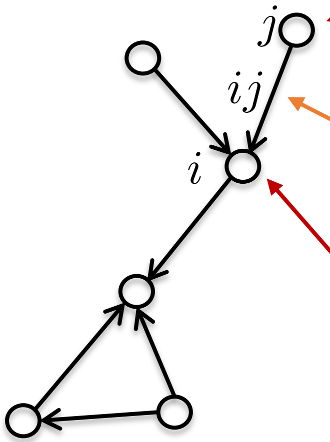

$$h_j = \begin{bmatrix} \vdots \end{bmatrix} \in \mathbb{R}^{d_v}$$

$$e_{ij} = \begin{bmatrix} \vdots \end{bmatrix} \in \mathbb{R}^{d_e}$$

$$h_i = \begin{bmatrix} \vdots \end{bmatrix} \in \mathbb{R}^{d_v}$$

$$\mathcal{G} = (V, E, A)$$

$$V = \{1, .., n\}$$

$$g = \begin{bmatrix} \vdots \end{bmatrix} \in \mathbb{R}^{d_g}$$

- Graph features :
  - Node features : h_i, h_j (atom type)
  - Edge features : e_{ij} (bond type)
  - Graph features : g (molecule energy)

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

Graph Neural Networks (GNNs)

Why generalizing Deep Learning to graphs is hard ?

- Irregular grid/topology: the size of graphs varies and the number of neighbors changes
- No node ordering: no coordinates/directions on graphs like up, down, right or left
- Lots of symmetries:
  - Permutation equivariance/invariance
  - Isomorphic nodes, i.e. nodes with the same neighborhood structures
- Dynamic topology: node/edge can appear and be discarded

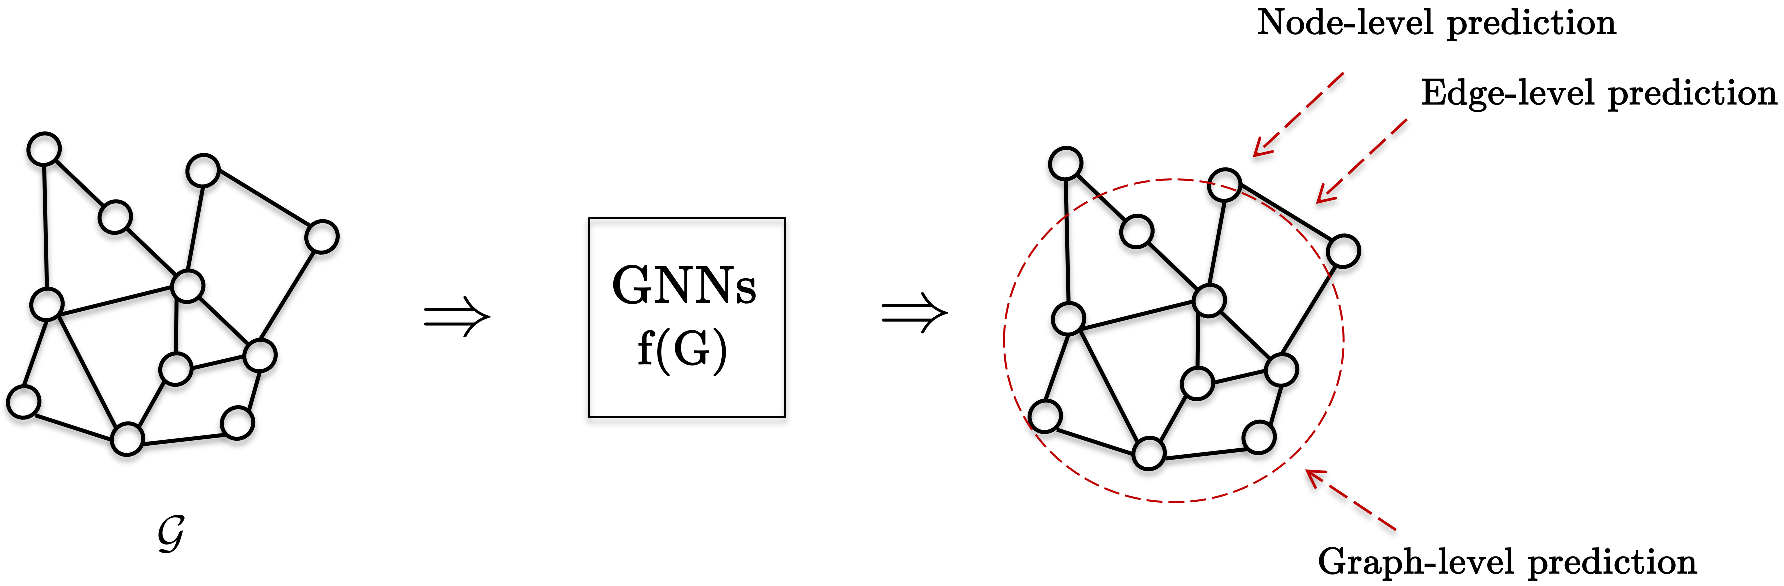

GNNs in research

- One of the hottest machine learning topics since 2021 !
- Standard toolkit for analyzing and learning from data on graphs
- Flexible to adapt to complex data structure and combine distinct modalities

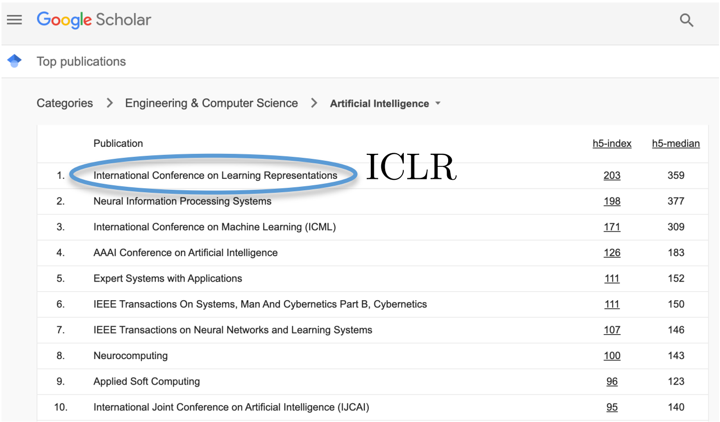

	Publication	h5-index	h5-median
1.	International Conference on Learning Representations	203	359
2.	Neural Information Processing Systems	198	377
3.	International Conference on Machine Learning (ICML)	171	309
4.	AAAI Conference on Artificial Intelligence	126	183
5.	Expert Systems with Applications	111	152
6.	IEEE Transactions On Systems, Man And Cybernetics Part B, Cybernetics	111	150
7.	IEEE Transactions on Neural Networks and Learning Systems	107	146
8.	Neurocomputing	100	143
9.	Applied Soft Computing	96	123
10.	International Joint Conference on Artificial Intelligence (IJCAI)	95	140

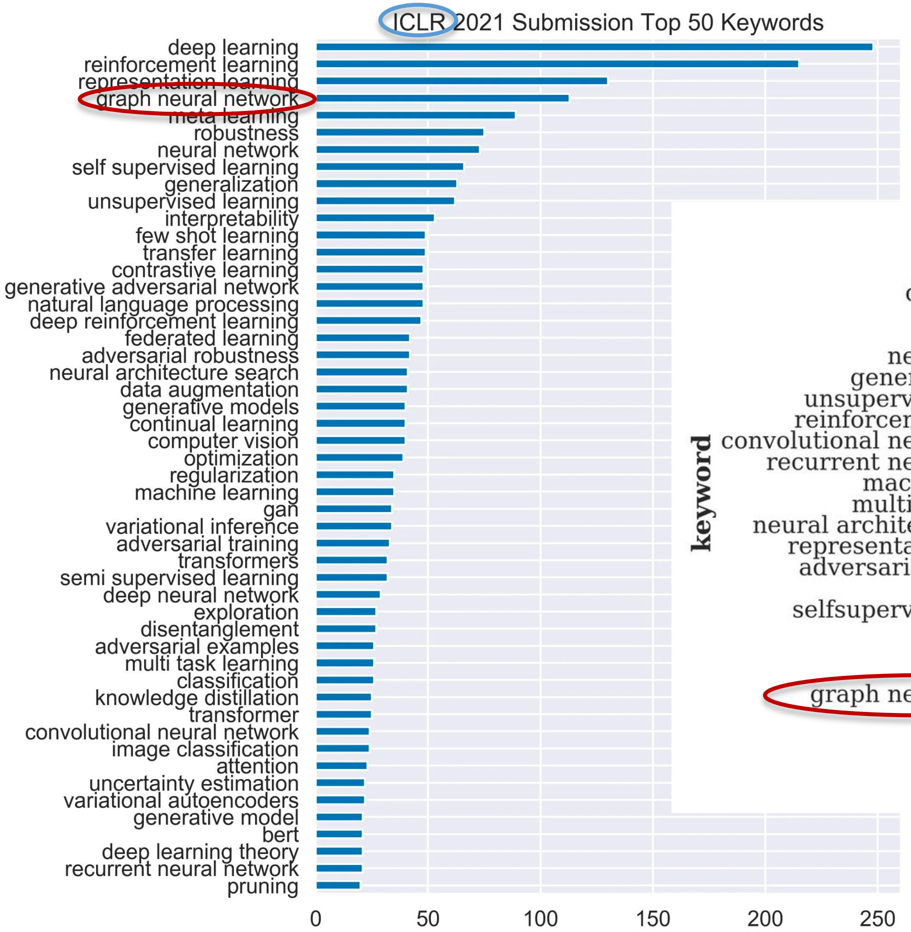

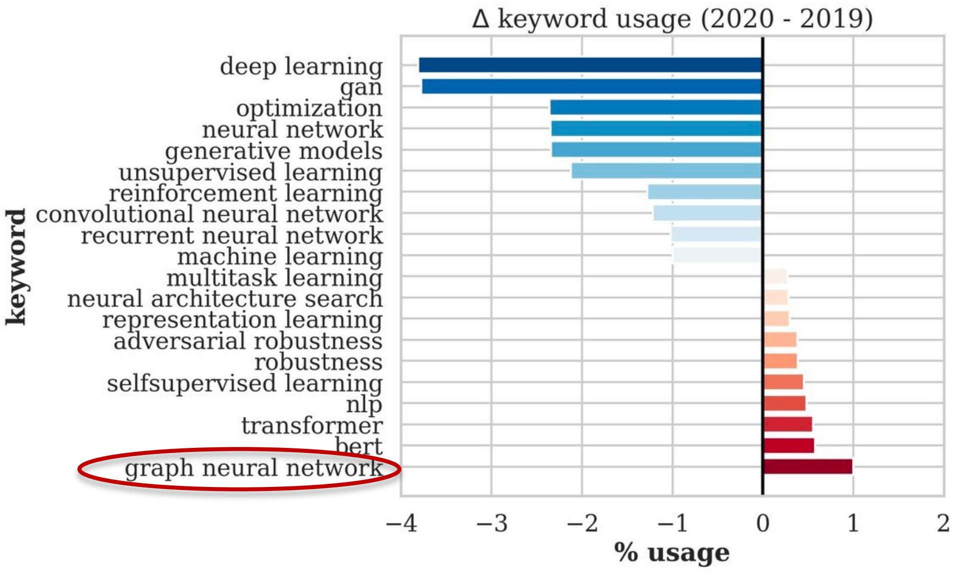

My research

- Design minimalist and mathematically sound GNNs for a wide range of applications
- Co-inventor of Graph Convolutional Networks in 2016 (10^{th} most cited paper at NeurIPS)
- Popularized GNNs w/ Michael Bronstein (Oxford), Yann LeCun (NYU/Meta/Turing) with tutorials at NeurIPS'17, CVPR'17, SIAM AM'18
- Developed with Yoshua Bengio'20'21 (MILA/Turing) and Yann LeCun'22'23 new classes of GNNs
- Invited speaker at KDD'23, NeurIPS'22, AAAI'21, KDD'21'23, ICLR'20, ICML'20 (AI conferences)
- Conference/workshop organizer at LoG'23'22, UCLA'23'21'19'18, ICLR'22, MLSys'21
- Awarded the US$2M NRF Fellowship in 2017 (largest individual grant in Singapore)
- 16,000+ citations, 60+ articles, 70,000+ YouTube views, 8,000+ Twitter followers

NeurIPS'17
1,000-2,000
participants

CVPR'17
500-1,000
participants

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

GNN case studies

- Chip design (Google)
- Resource management
- Scene reasoning (Meta)
- Recommendation (UberEats/ Pinterest/ Alibaba)
- Fake news detection
- Finance
- Natural Language Processing
- Knowledge graphs (Amazon)
- Transportation (Google Map ETA)
- Autonomous driving (NVIDIA)
- Protein & drug design (Google DeepMind/ Microsoft/ AstraZeneca)
- Energy physics & simulations (Google DeepMind)
- Code bug detection
- Genomics

GNNs for chip design

Google AI Blog

The latest news from Google AI

https://ai.googleblog.com/2020/04/chip-design-with-deep-reinforcement.html

Chip Design with Deep Reinforcement Learning

Thursday, April 23, 2020

Posted by Anna Goldie, Senior Software Engineer and Azalia Mirhoseini, Senior Research Scientist, Google Research, Brain Team

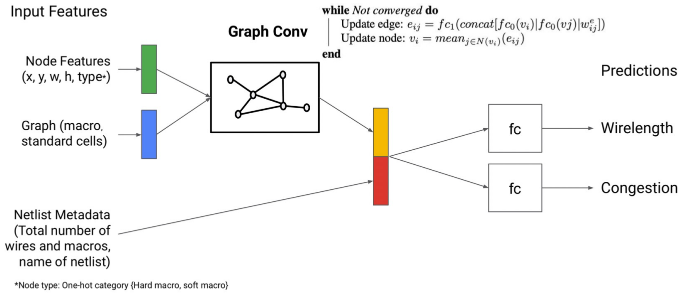

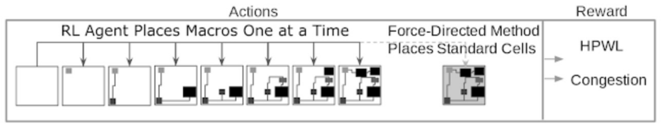

During each training iteration, the macros are placed by the policy one at a time and the standard cell clusters are placed by a force-directed method. The reward is calculated from the weighted combination of approximate wirelength and congestion.

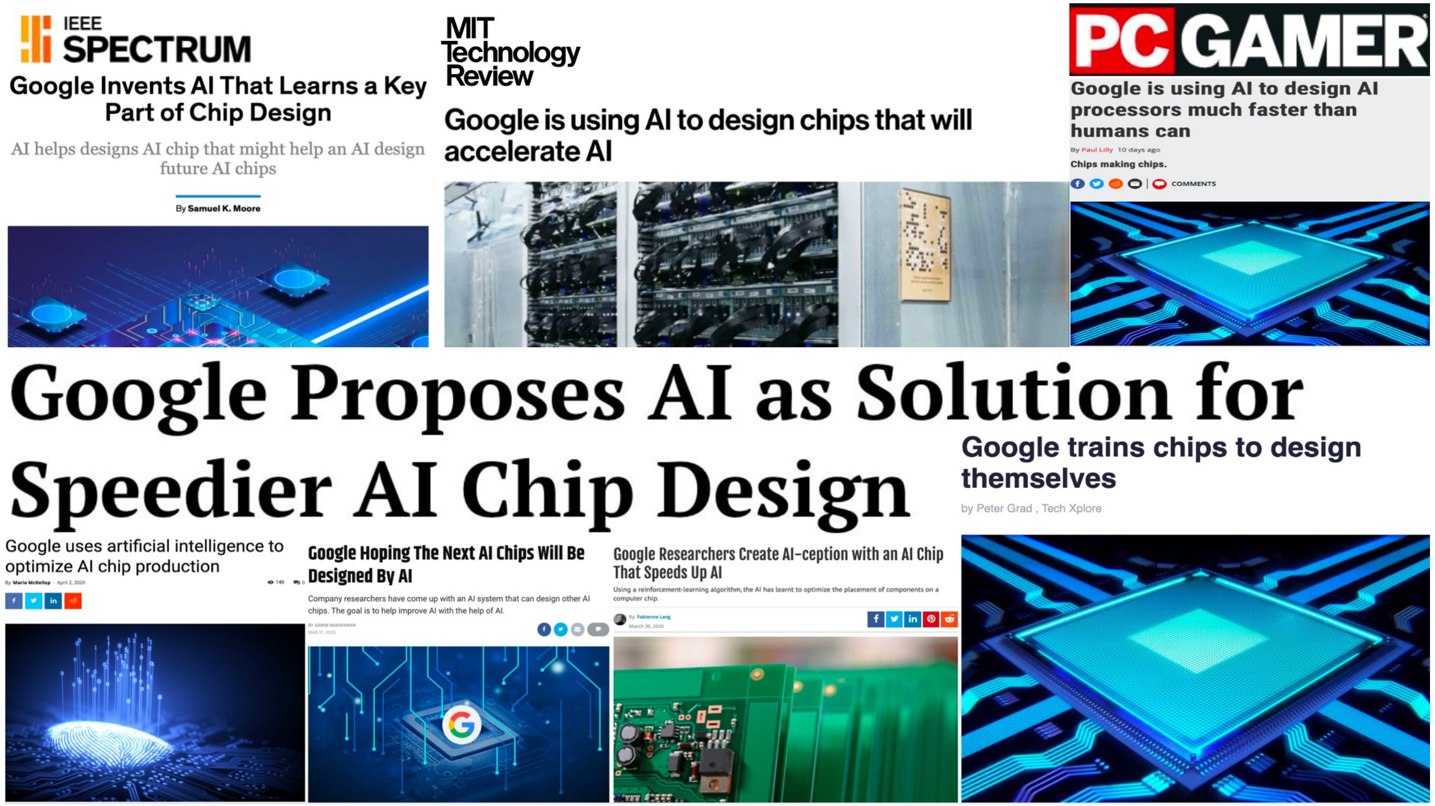

https://www.theverge.com/2021/6/10/22527476/google-machine-learning-chip-design-tpu-floorplanning

GNNs for resource allocation

- Operations Research/Combinatorial Optimization s.a. assignment, routing, planning, supply chain, scheduling are used every day in revenue management, transportation, manufacturing, supply chain, public policy, hardware design, etc.

- Most OR problems are NP-hard.

- Traditional OR solvers are hand-crafted algorithms with years of research work and significant specialized knowledge.

- DL can learn universal high-quality algorithms to solve OR problems with GNNs and RL.

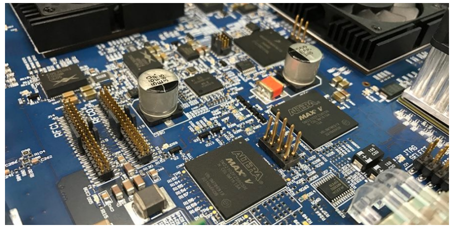

Hardware pieces placement

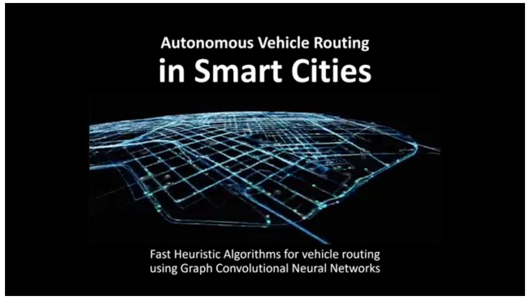

Uber taxis allocation

Can Transformers Solve This 90-Year-Old Classic Computer Science Problem Better Than Human Algorithms?

Deep learning can't beat human solutions — yet

Andre Ye Mar 19 · 6 min read ★

The Travelling Salesman Problem was formulated in 1930, and is a classical computer science problem for optimization. It's a simple problem:

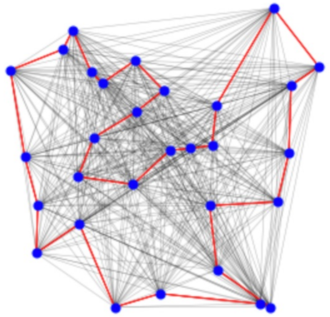

https://towardsdatascience.com/can-transformers-solve-this-90-year-old-classic-computer-science-problem-better-than-human-

Amazon warehouse management

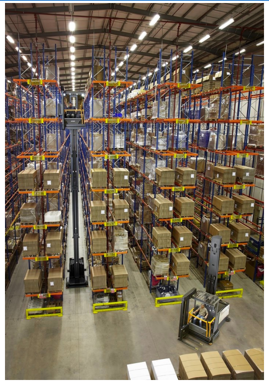

Solving the TSP for Warehouses

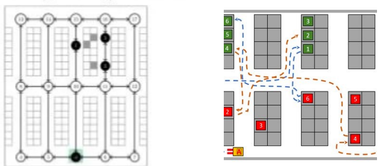

GNNs for scene reasoning

A neural approach to relational reasoning

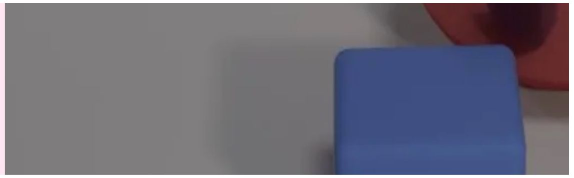

SHARE

AUTHORS

AS Adam Santoro

DR David Raposo

NW Nick Watters

Consider the reader who pieces together the evidence in an Agatha Christie novel to predict the culprit of the crime, a child who runs ahead of her ball to prevent it rolling into a stream or even a shopper who compares the relative merits of buying kiwis or mangos at the market.

We carve our world into relations between things. And we understand how the world works through our capacity to draw logical conclusions about how these different things – such as physical objects, sentences, or even abstract ideas – are related to one another. This ability is called relational reasoning and is central to human intelligence.

https://deepmind.com/blog/article/neural-approach-relational-reasoning

Questions in CLEVR test various aspects of visual reasoning including **attribute identification**, **counting**, **comparison**, **spatial relationships**, and **logical operations**.

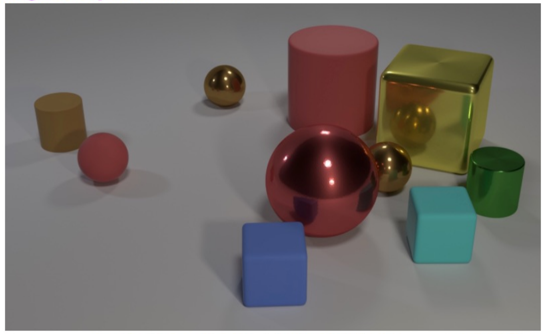

Q: Are there an **equal number** of **large things** and **metal spheres**?
Q: **What size** is the **cylinder that is left of** the **brown metal** thing **that is left of** the **big sphere**?
Q: There is a **sphere** with the **same size as** the **metal cube**; is it **made of the same material as** the **small red sphere**?
Q: **How many** objects are **either small cylinders** or **red** things?

Graph R-CNN for Scene Graph Generation

Jianwei Yang[1*], Jiasen Lu[1*], Stefan Lee[1], Dhruv Batra[1,2], and Devi Parikh[1,2]

[1]Georgia Institute of Technology [2]Facebook AI Research
{jw2yang, jiasenlu, steflee, dbatra, parikh}@gatech.edu

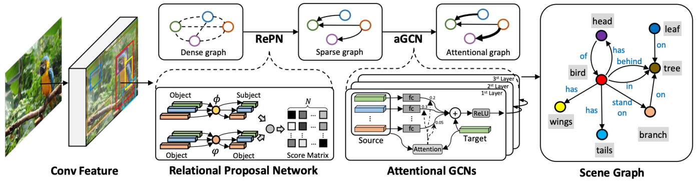

Fig. 2. The pipeline of our proposed Graph R-CNN framework. Given an image, our model first uses RPN to propose object regions, and then prunes the connections between object regions through our relation proposal network (RePN). Attentional GCN is then applied to integrate contextual information from neighboring nodes in the graph. Finally, the scene graph is obtained on the right side.

GNNs for autonomous driving

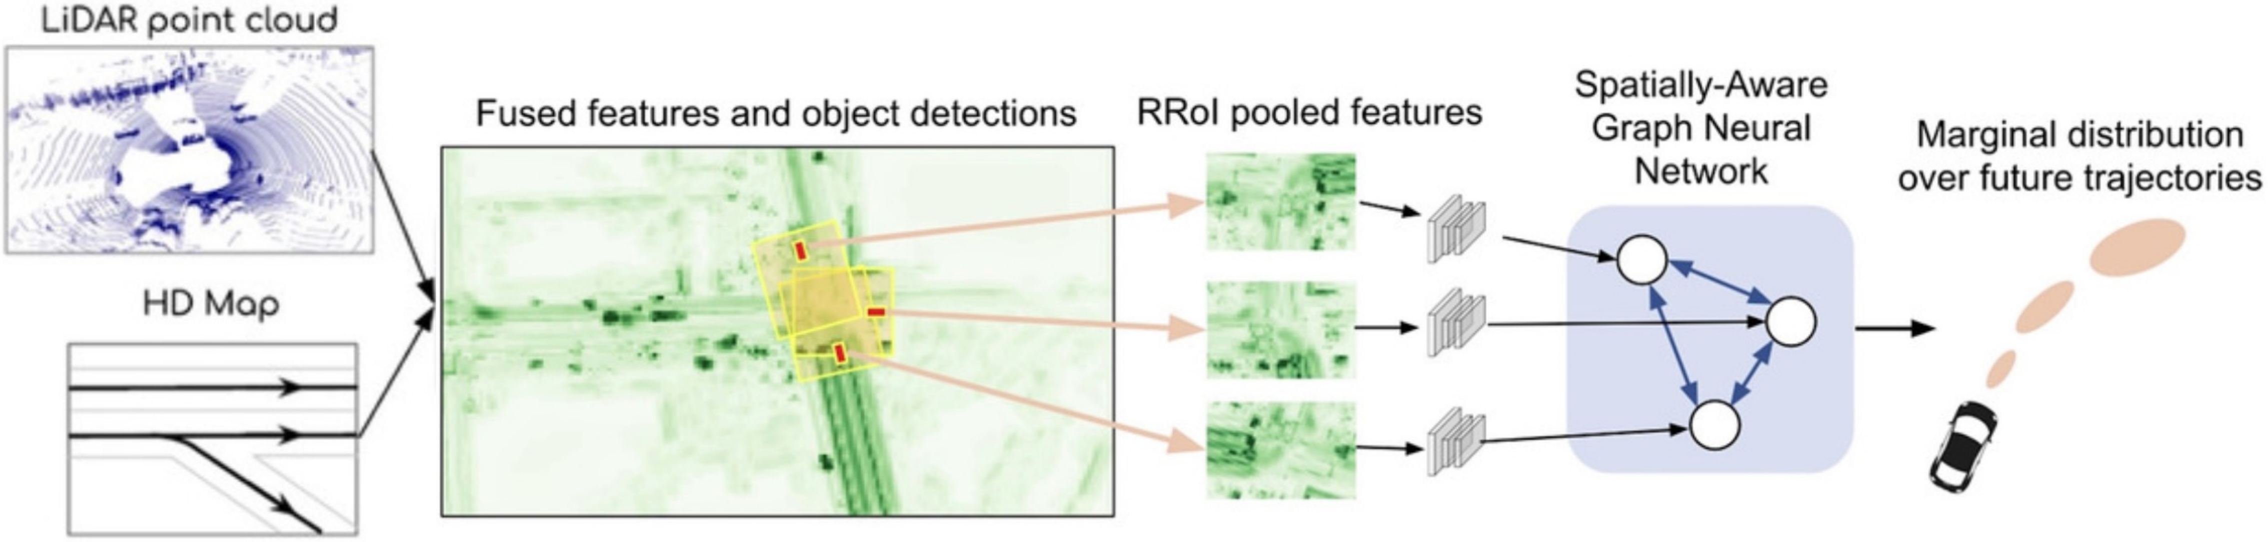

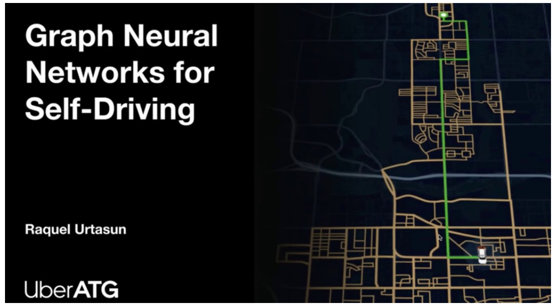

https://slideslive.com/38930570/graph-neural-networks-for-selfdriving

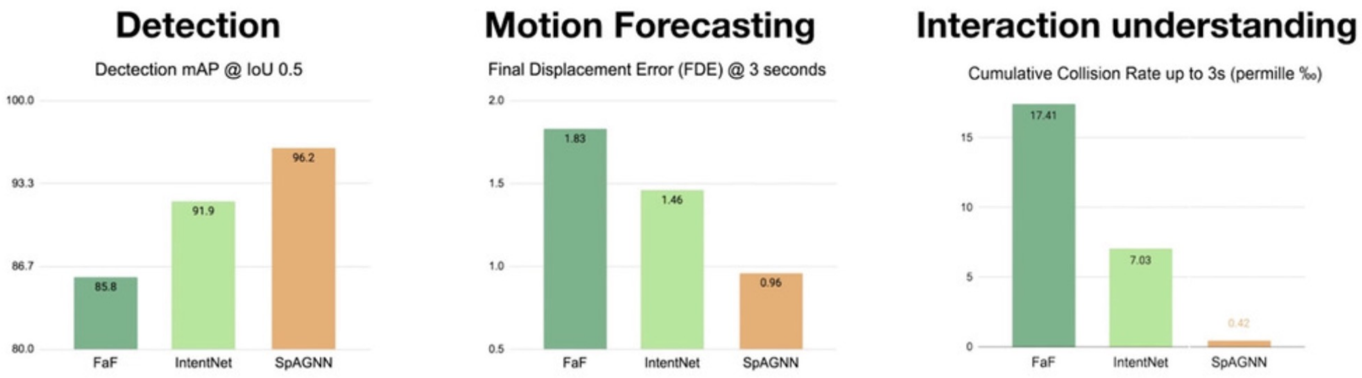

GNNs for autonomous driving

- Computer Vision with visual common sense
  - Reasoning on scene graph

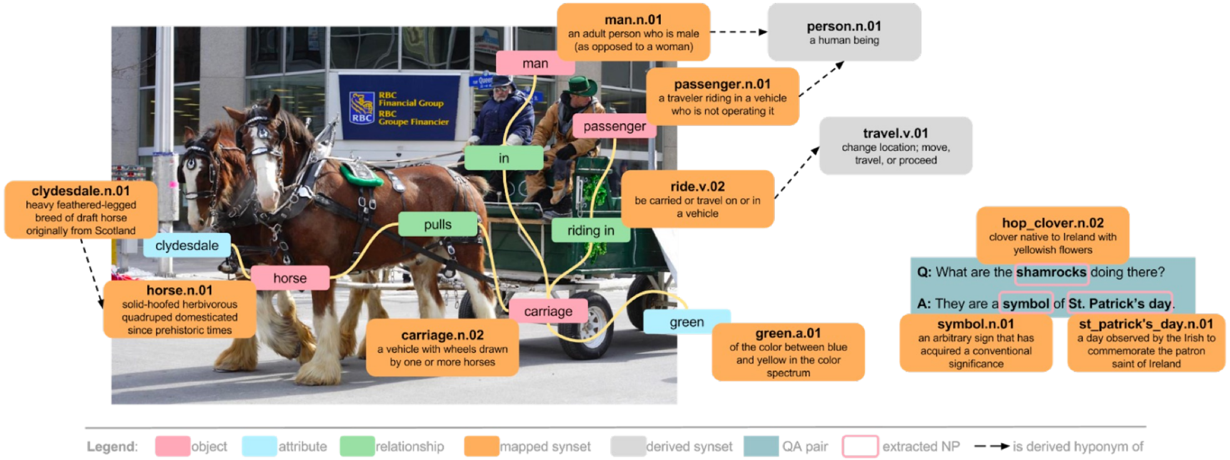

Scene graph understanding

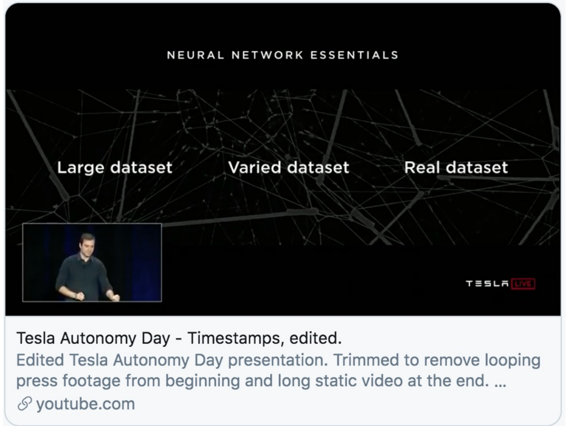

A. Karpathy (Tesla AI Autopilot Director)
https://youtu.be/2PpNmSdFP7Q?t=3369

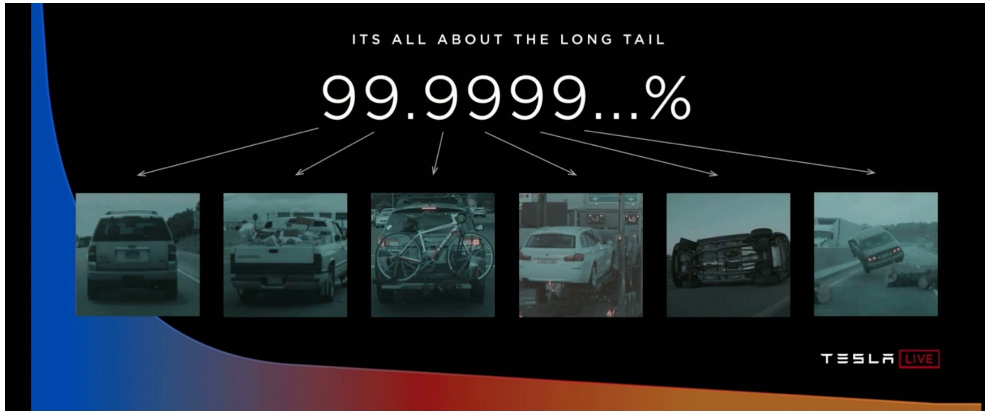

GNNs for recommendation

Uber Engineering

AI | General Engineering

Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations

Ankit Jain, Isaac Liu, Ankur Sarda, and Piero Molino — December 4, 2019

0

Graph learning for dish and restaurant recommendation at Uber

There are several recommendation surfaces within the Uber Eats app, depicted in Figure 3, below:

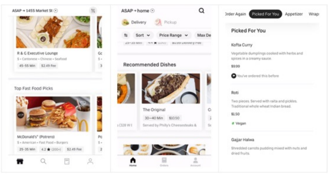

Figure 3: The Uber Eats UI surfaces a wealth of options for hungry users informed by past orders and previously specified user preferences.

139 | 3 | 142 SHARES

The Uber Eats app serves as a portal to more than 320,000 restaurant-partners in over 500 cities globally across 36 countries. In order to make the user experience more seamless and easy-to-navigate, we show users the dishes, restaurants, and cuisines they might like up front. To this end, we previously developed ML models to better understand queries and for multi-objective optimization in Uber Eats search and recommender system in Uber Eats searches and surfaced food options.

https://eng.uber.com/uber-eats-graph-learning

Cornell University

arXiv.org > cs > arXiv:1902.08730

Computer Science > Distributed, Parallel, and Cluster Computing

[Submitted on 23 Feb 2019]

AliGraph: A Comprehensive Graph Neural Network Platform

Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, Jingren Zhou

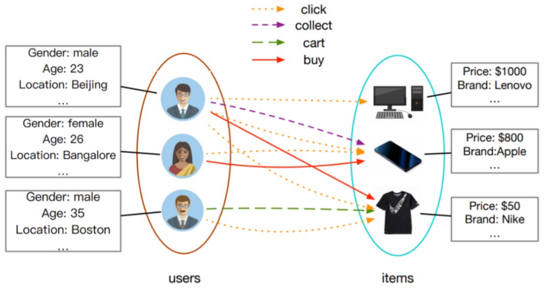

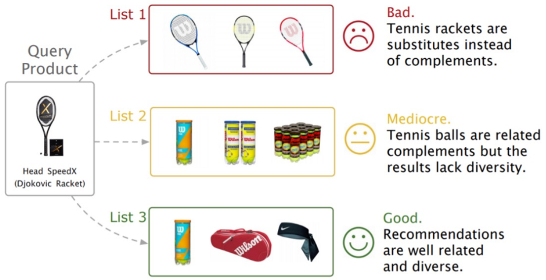

Pinterest Engineering Blog | MACHINE LEARNING | INFRASTRUCTURE | OPEN SOURCE | DATA SCIENCE | MOBILE

PinSage: A new graph convolutional neural network for web-scale recommender systems

Pinterest Engineering · Follow

Aug 16, 2018 · 8 min read

Ruining He | Pinterest engineer, Pinterest Labs

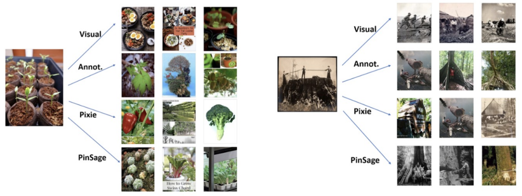

Figure 3: Examples of pins recommended by different algorithms. The image to the left is the query pin. Recommended items to the right are computed using Visual embeddings, Annotation embeddings, Pixie (purely graph-based method), and PinSage.

https://medium.com/pinterest-engineering/pinsage-a-new-graph-convolutional-neural-network-for-web-scale-recommender-systems-88795a107f48

GNNs for fake news detection

- Social networks :
  - Fake News (2016 US Presidential Election) /(human) adversarial attacks

Fake News Detection on Social Media using Geometric Deep Learning

Federico Monti[1,2] **Fabrizio Frasca**[1,2] **Davide Eynard**[1,2] **Damon Mannion**[1,2]

Michael M. Bronstein[1,2,3]

[1]Fabula AI, United Kingdom [2]USI Lugano, Switzerland [3]Imperial College, United Kingdom

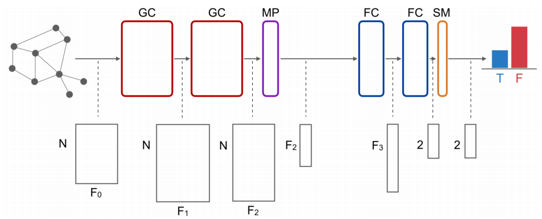

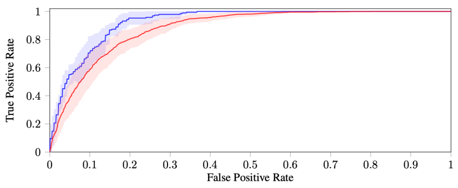

Figure 6: Performance of URL-wise (blue) and cascade-wise (red) fake news detection using 24hr-long diffusion time. Shown are ROC curves averaged on five folds (the shaded areas represent the standard deviations). ROC AUC is 92.70 ± 1.80% for URL-wise classification and 88.30 ± 2.74% for cascade-wise classification, respectively. Only cascades with at least 6 tweets were considered for cascade-wise classification.

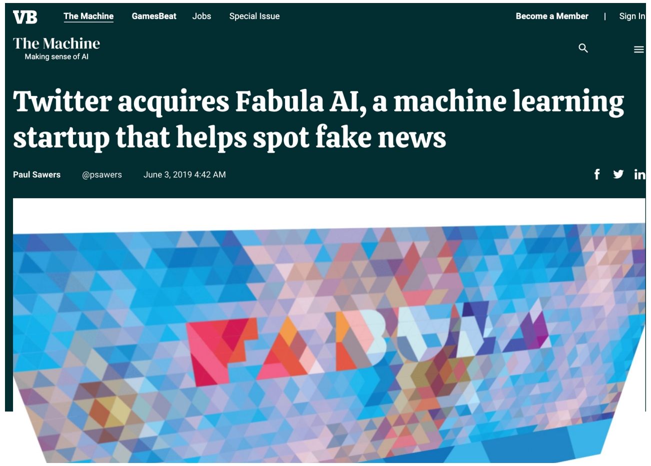

Twitter network

Facebook network

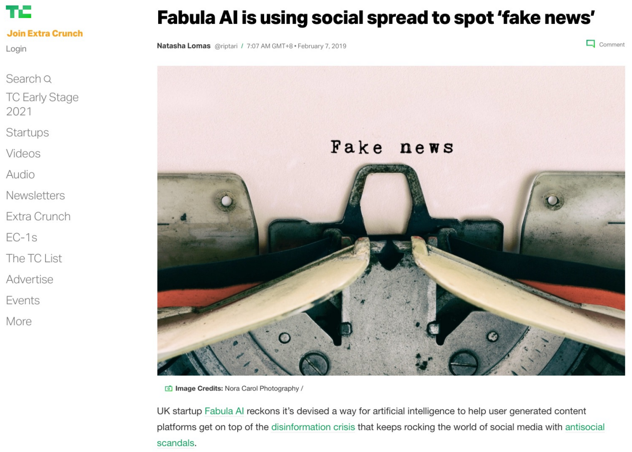

GNNs for finance

- Financial networks model financial entities and their relationships.
- Applications :

International banking
- *Nodes: Countries*
- *Edges: Capital flows*

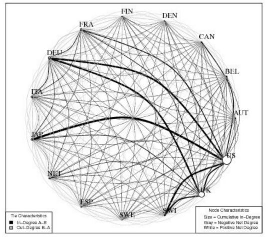

Image credit: The Political Economy of Global Finance: A Network Model

Bitcoin transactions
- *Nodes: BTC wallets*
- *Edges: Transactions*

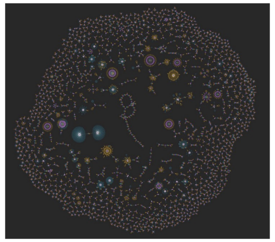

Image credit: https://dailyblockchain.github.io/

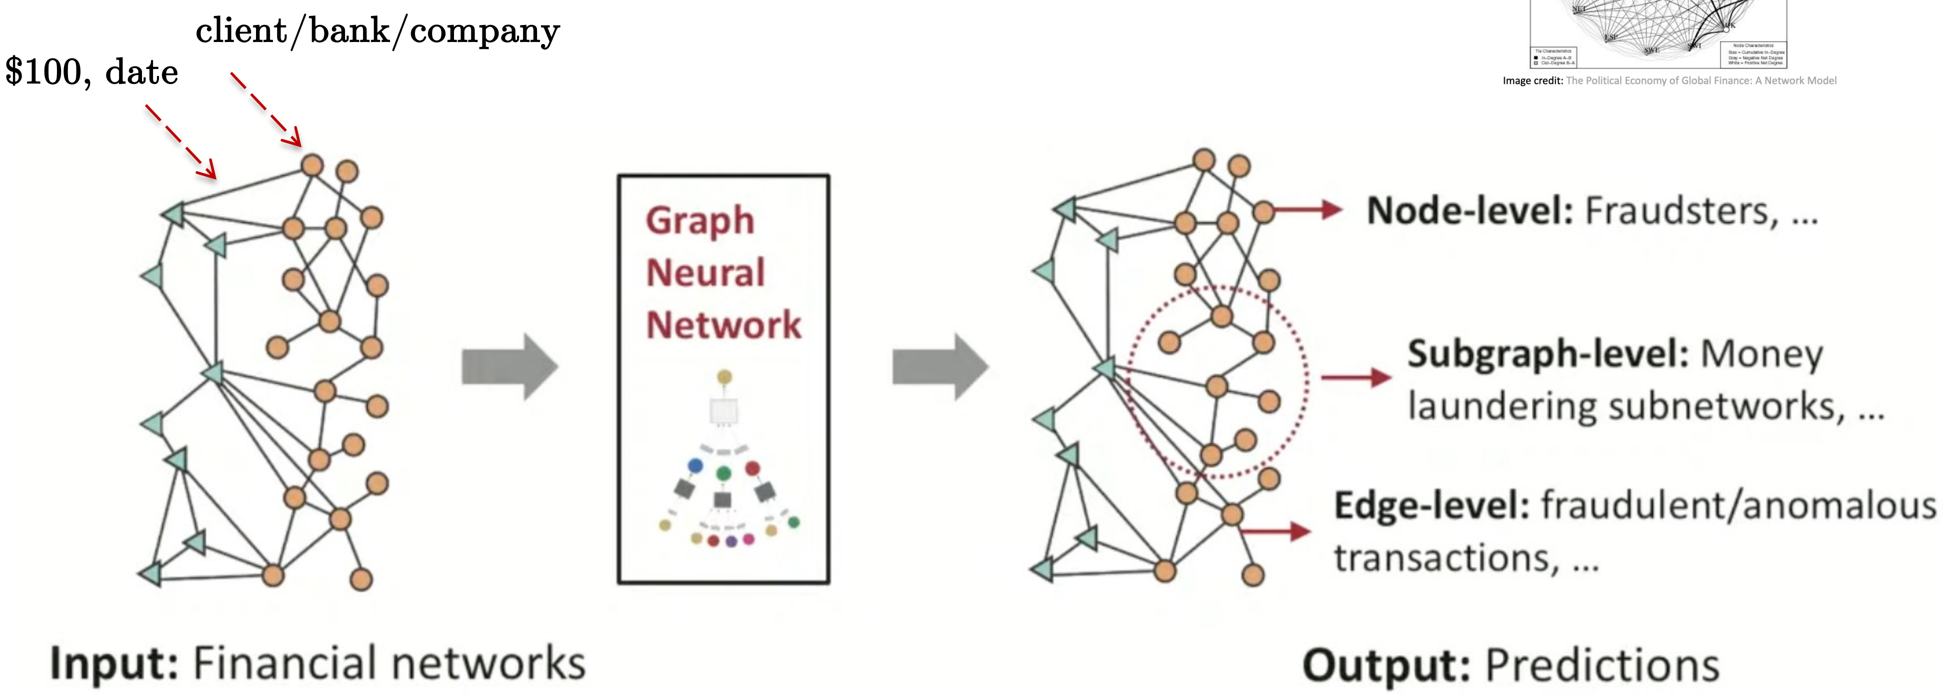

Credit: You et-al, Arxiv 2021

GNNs for knowledge graphs

amazon | science

INFORMATION AND KNOWLEDGE MANAGEMENT

Combining knowledge graphs, quickly and accurately

Novel cross-graph-attention and self-attention mechanisms enable state-of-the-art performance.

By Hao Wei

March 19, 2020

Knowledge graphs are a way of representing information that can capture complex relationships more easily than conventional databases. At Amazon, we use knowledge graphs to represent the hierarchical relationships between product types on amazon.com; the relationships between creators and content on Amazon Music and Prime Video; and general information for Alexa's question-answering service — among other things.

https://www.amazon.science/blog/combining-knowledge-graphs-quickly-and-accurately

- Knowledge graphs represent large-scale data information in the form of triples.
  - Triples are two entities and their type of relationship.

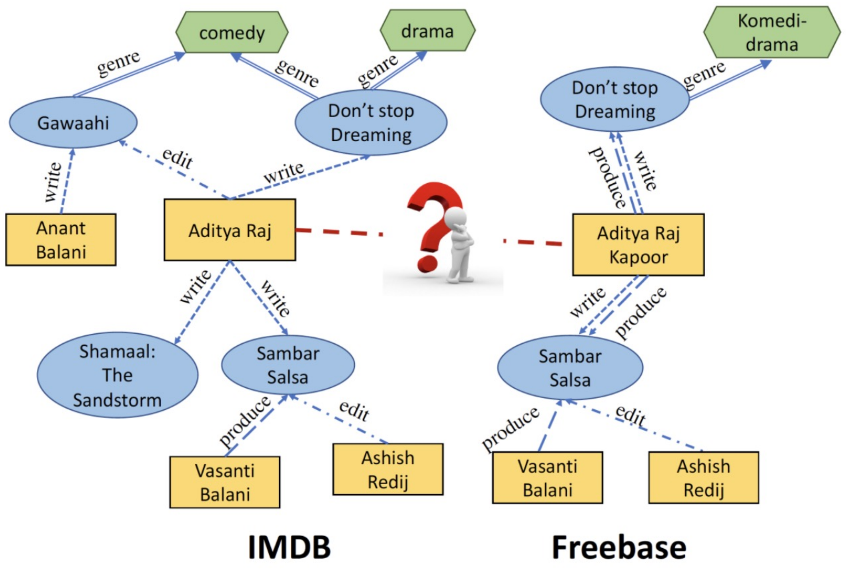

This example illustrates the challenge of entity alignment. IMDB lists the writer of the movie *Don't Stop Dreaming* as Aditya Raj, but the (now defunct) Freebase database lists him as Aditya Raj Kapoor. Are they the same person?

GNNs for knowledge graphs

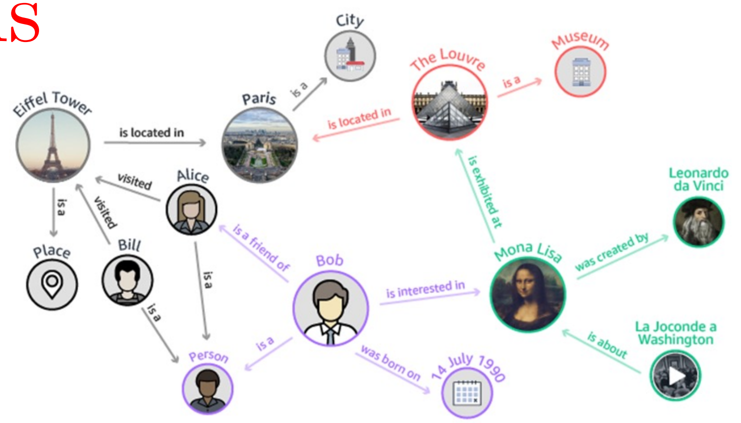

Credit: Amazon Neptune ML

- Reasoning on KGs can provide critical new information.
  - Predict drugs that target proteins associated to Covid-19.
  - Where did Canadian Turing Award winners graduate?
- Limitations of non-ML reasoning on KG :
  - KGs are incomplete, noisy.
  - Answer query on large graphs can become intractable with standard graph search algorithms.
- GNNs can represent complex queries in a continuous embedding space.
  - Answer is given by the closest entity (w.r.t. the Euclidean distance).

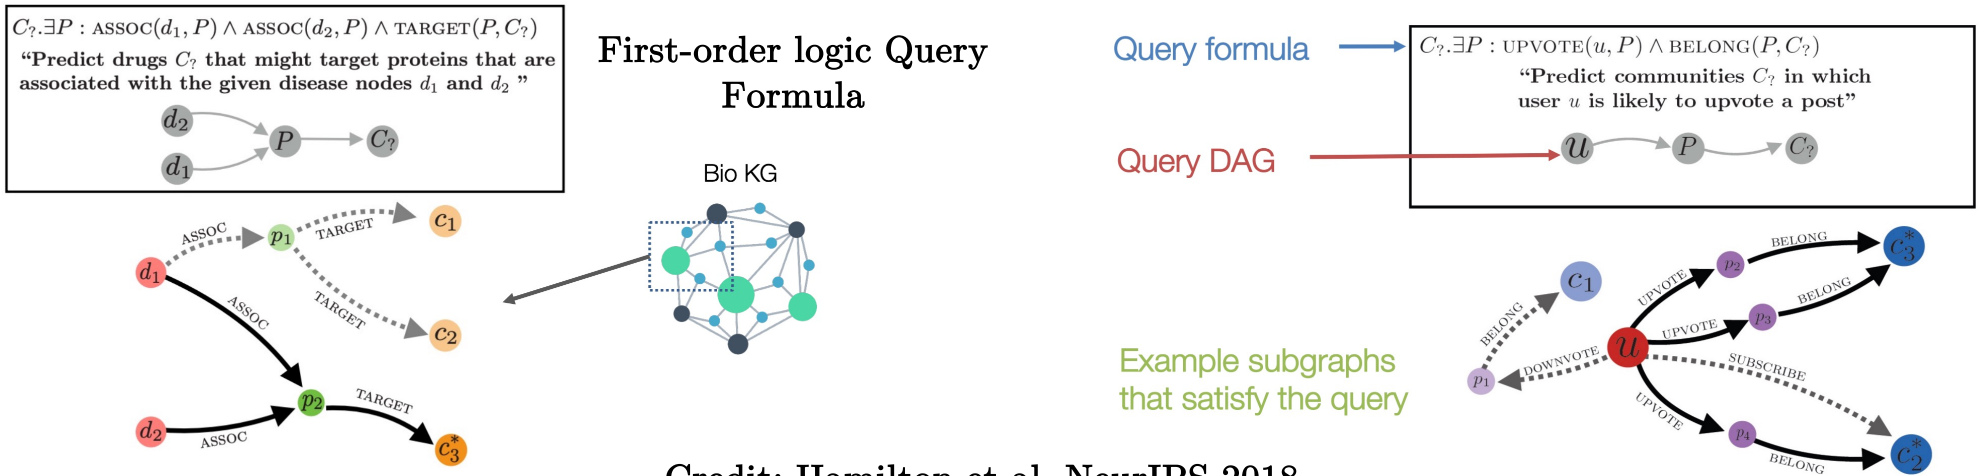

Credit: Hamilton et al, NeurIPS 2018

GNNs for NLP

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

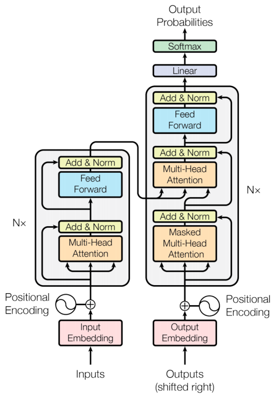

Figure 1: The Transformer - model architecture.

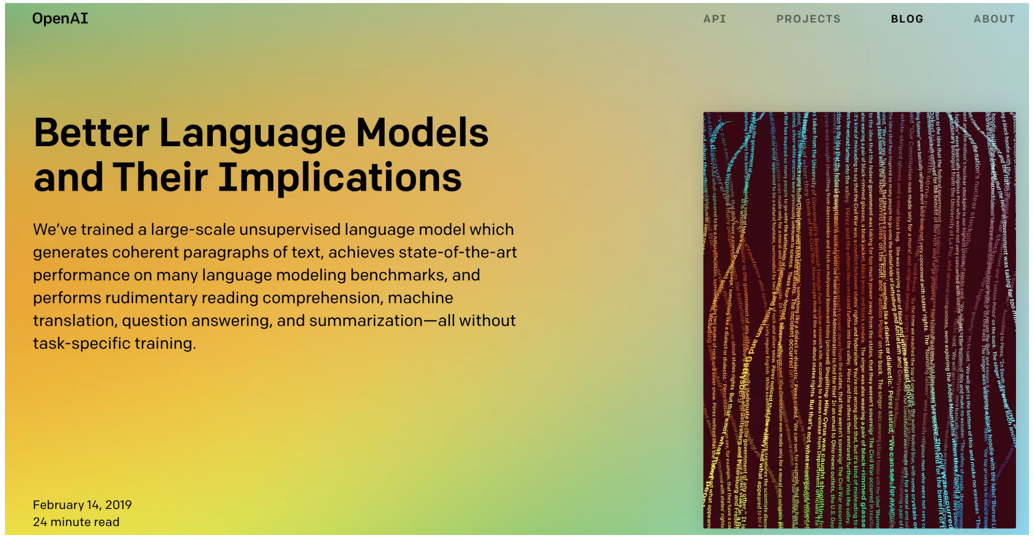

https://openai.com/blog/better-language-models

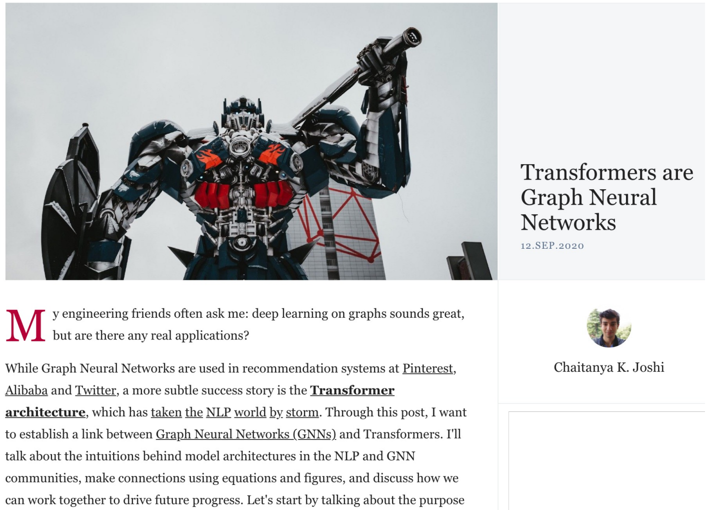

Transformers are Graph Neural Networks

12.SEP.2020

Chaitanya K. Joshi

My engineering friends often ask me: deep learning on graphs sounds great, but are there any real applications?

While Graph Neural Networks are used in recommendation systems at Pinterest, Alibaba and Twitter, a more subtle success story is the **Transformer architecture**, which has taken the NLP world by storm. Through this post, I want to establish a link between Graph Neural Networks (GNNs) and Transformers. I'll talk about the intuitions behind model architectures in the NLP and GNN communities, make connections using equations and figures, and discuss how we can work together to drive future progress. Let's start by talking about the purpose

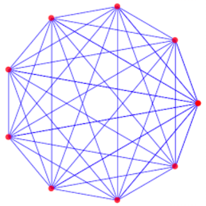

https://thegradient.pub/transformers-are-graph-neural-networks

Transformers are (fully-connected) Graph Neural Networks.

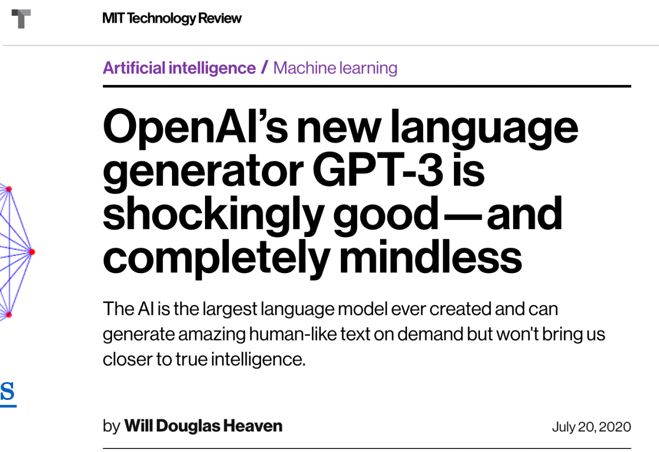

MIT Technology Review

Artificial intelligence / Machine learning

OpenAI's new language generator GPT-3 is shockingly good—and completely mindless

The AI is the largest language model ever created and can generate amazing human-like text on demand but won't bring us closer to true intelligence.

by **Will Douglas Heaven** July 20, 2020

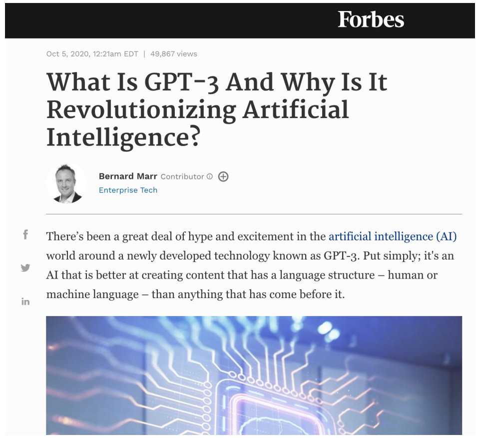

Forbes

Oct 5, 2020, 12:21am EDT | 49,867 views

What Is GPT-3 And Why Is It Revolutionizing Artificial Intelligence?

Bernard Marr Contributor
Enterprise Tech

There's been a great deal of hype and excitement in the artificial intelligence (AI) world around a newly developed technology known as GPT-3. Put simply; it's an AI that is better at creating content that has a language structure – human or machine language – than anything that has come before it.

Limitation of Transformers

- Impressive results at large-scale
  - 2020 GPT-3, 175B parameters, 285,000 CPUs, 10,000 GPUs, 400Gb/sec network connectivity, 500 billion tokens, US$12 Million to train
  - Transformers capture dynamic word representation depending on the context.
  - Example : The vase broke. The news broke. Sandy broke the world record. Sandy broke the law. We broke even. The burglar broke into the house. Etc.
  - DL has not yet reached human performance (no common sense).

xkcd: Machine Learing

- What is missing to get to human's level?
  - More data? Yes but not sufficient.
  - Reasoning, but with what inductive bias?
    - Knowledge graphs

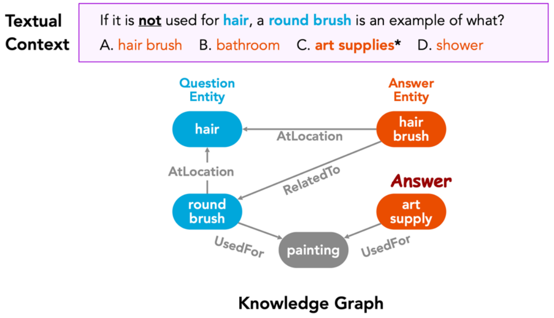

Credit: Antoine Bosselut, EPFL

GNNs for NLP

- Transformers vs GNNs with knowledge graphs

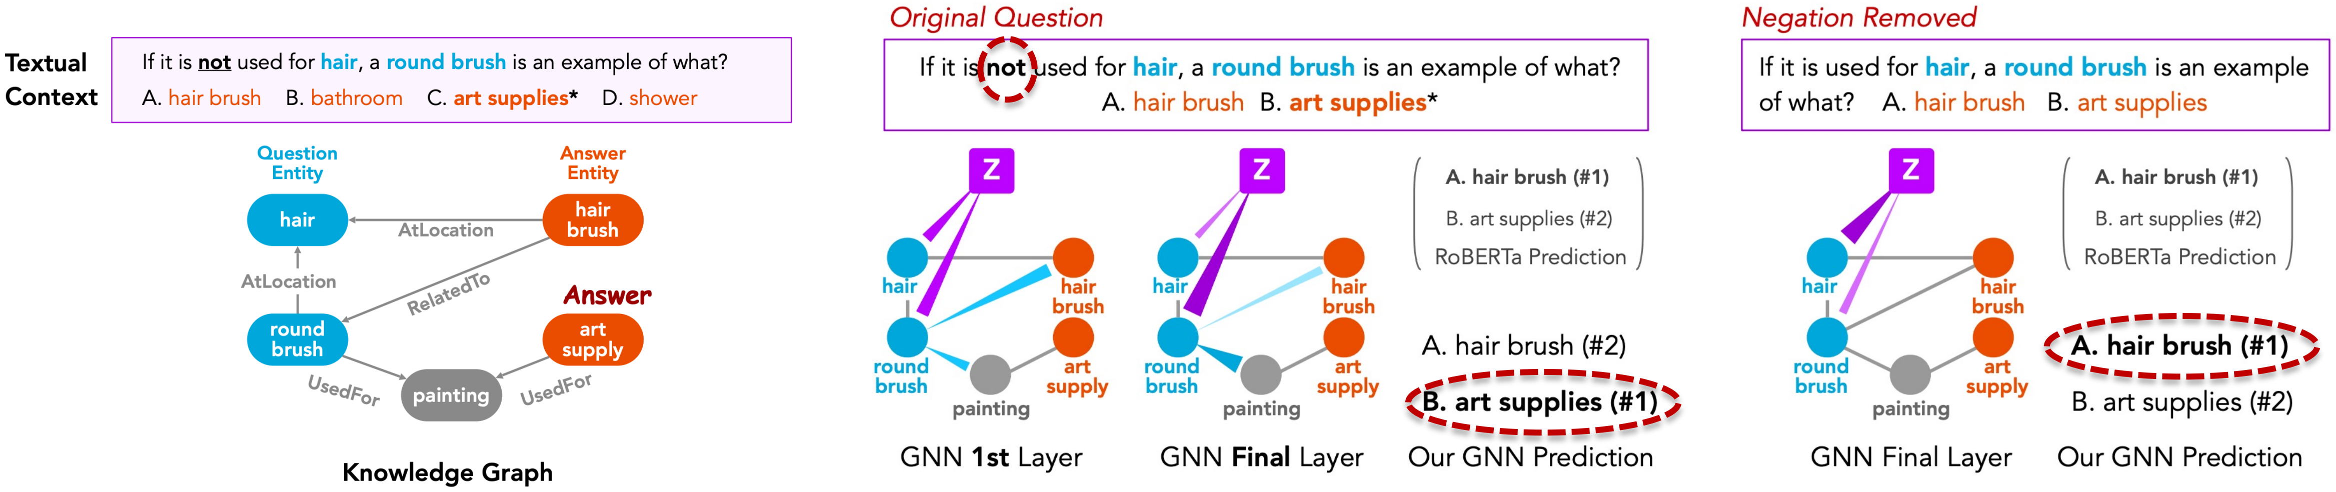

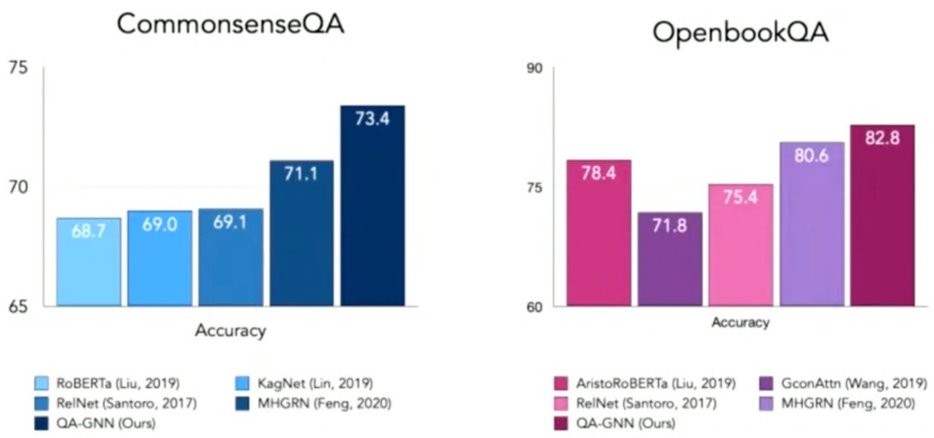

Credit: Yasunaga et al, NAACL, 2021

GNNs for transportation


BLOG POST
RESEARCH

03 SEP 2020

Traffic prediction with advanced Graph Neural Networks

SHARE

AUTHORS

OL Oliver Lange*

LP Luis Perez

By partnering with Google, DeepMind is able to bring the benefits of AI to billions of people all over the world. From reuniting a speech-impaired user with his original voice, to helping users discover personalised apps, we can apply breakthrough research to immediate real-world problems at a Google scale. Today we're delighted to share the results of our latest partnership, delivering a truly global impact for the more than one billion people that use Google Maps.

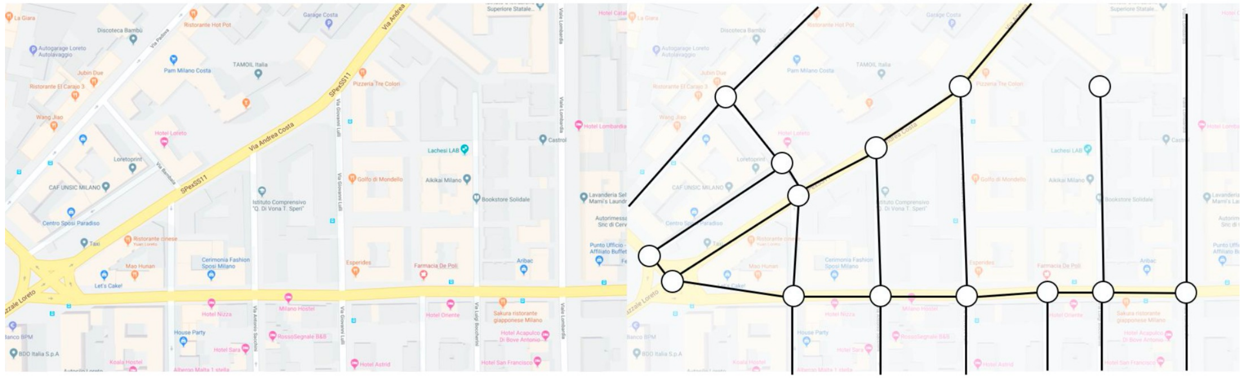

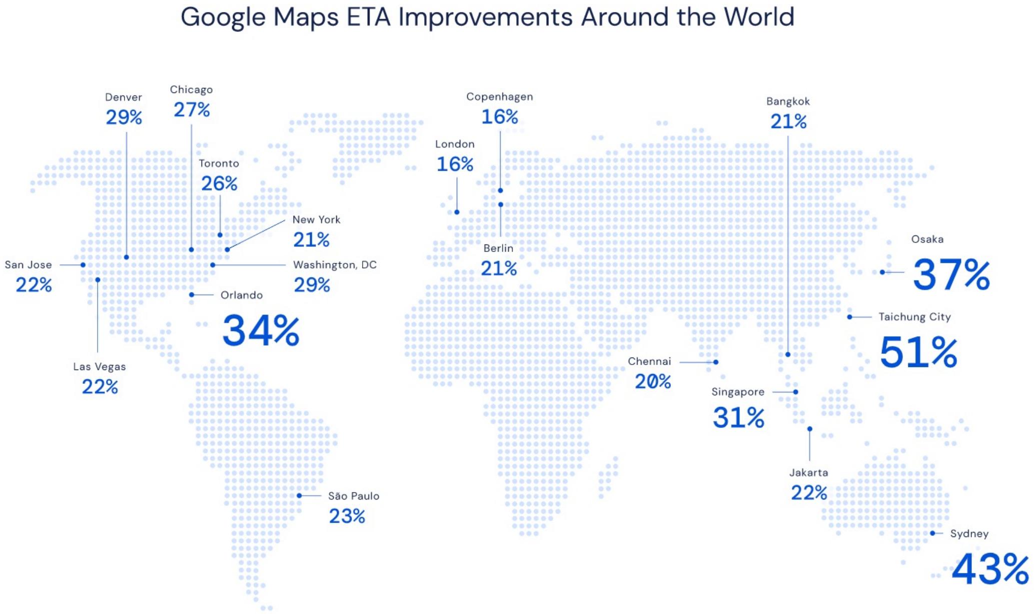

https://deepmind.com/blog/article/traffic-prediction-with-advanced-graph-neural-networks

GNNs for protein folding


AlphaFold: a solution to a 50-year-old grand challenge in biology

SHARE

AUTHORS

TAt The AlphaFold team

Proteins are essential to life, supporting practically all its functions. They are large complex molecules, made up of chains of amino acids, and what a protein does largely depends on its unique 3D structure. Figuring out what shapes proteins fold into is known as the "protein folding problem", and has stood as a grand challenge in biology for the past 50 years. In a major scientific advance, the latest version of our AI system AlphaFold has been recognised as a solution to this grand challenge by the organisers of the biennial Critical Assessment of protein Structure Prediction (CASP). This breakthrough demonstrates the impact AI can have on scientific discovery and its potential to dramatically accelerate progress in some of the most fundamental fields that explain and shape our world.

https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology

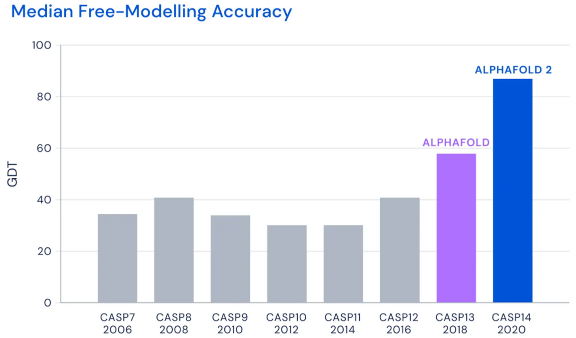

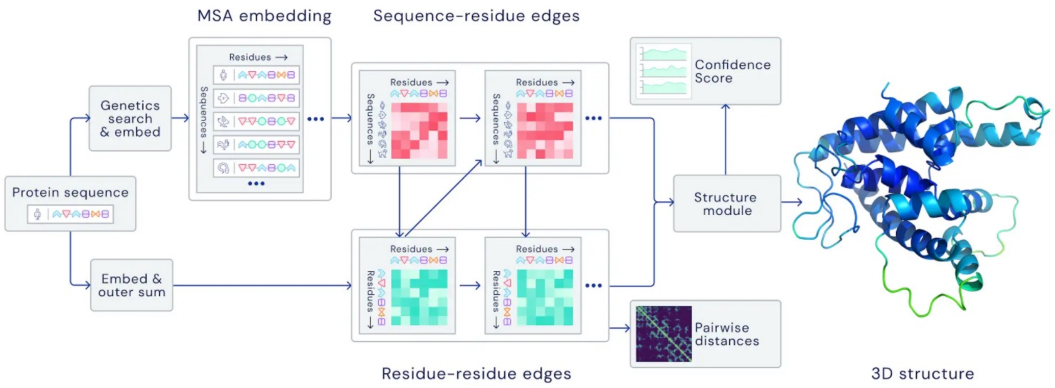

Graph Transformers

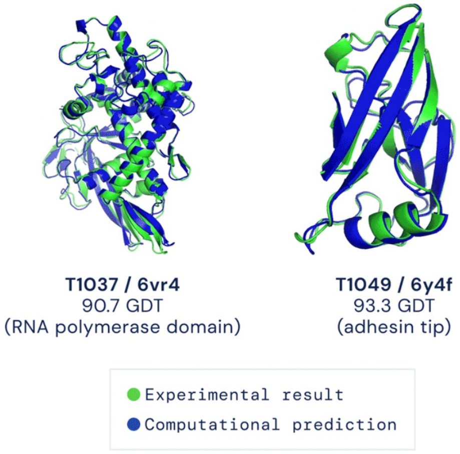

GNNs for protein function & interaction

nature > nature communications > articles > article

Article | Open Access | Published: 26 May 2021

Structure-based protein function prediction using graph convolutional networks

Vladimir Gligorijević, P. Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C. Taylor, Ian M. Fisk, Hera Vlamakis, Ramnik J. Xavier, Rob Knight, Kyunghyun Cho & Richard Bonneau

Nature Communications **12**, Article number: 3168 (2021) | Cite this article

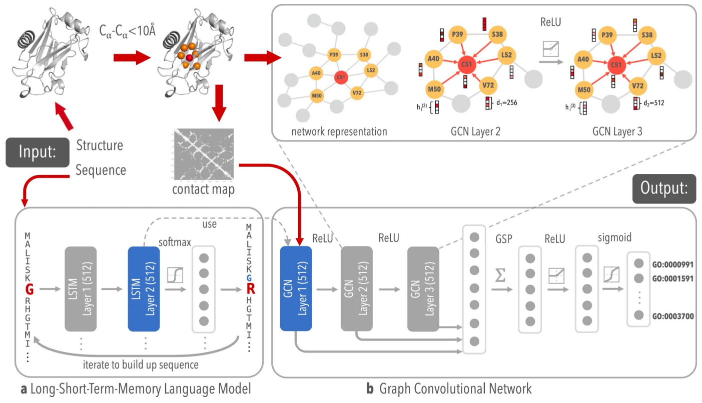

nature > scientific reports > articles > article

Article | Open Access | Published: 03 December 2020

SkipGNN: predicting molecular interactions with skip-graph networks

Kexin Huang, Cao Xiao, Lucas M. Glass, Marinka Zitnik & Jimeng Sun

Scientific Reports **10**, Article number: 21092 (2020) | Cite this article

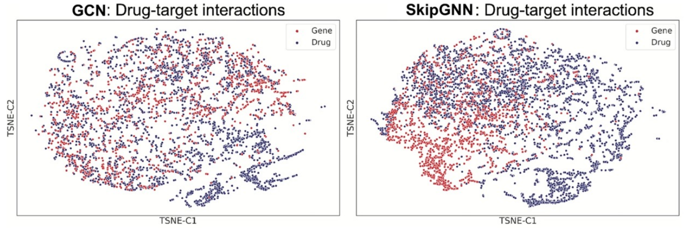

DeepRank-GNN: A Graph Neural Network Framework to Learn Patterns in Protein-Protein Interfaces

M. Réau[1,#], N. Renaud[2,#], L. C. Xue[3], A. M. J. J. Bonvin[1,*]

[1]Computational Structural Biology Group, Department of Chemistry, Bijvoet Centre, Faculty of Science, Utrecht University, Utrecht, 3584CH, The Netherlands., [2]Netherlands eScience Center, Science Park 140, 1098 XG,

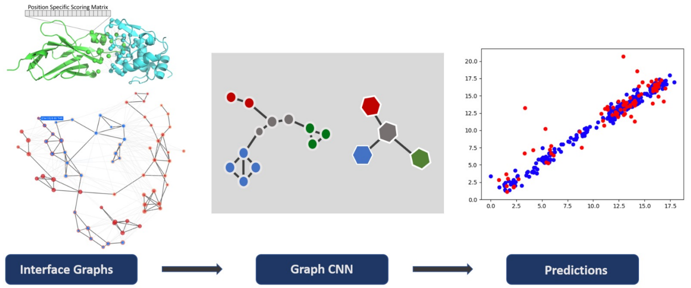

arXiv > q-bio > arXiv:2202.05146

Quantitative Biology > Biomolecules

[Submitted on 7 Feb 2022 (v1), last revised 4 Jun 2022 (this version, v4)]

EquiBind: Geometric Deep Learning for Drug Binding Structure Prediction

Hannes Stärk, Octavian-Eugen Ganea, Lagnajit Pattanaik, Regina Barzilay, Tommi Jaakkola

Predicting how a drug-like molecule binds to a specific protein target is a core problem in drug discovery. An extremely fast computational binding method would enable key applications such as fast virtual screening or drug engineering. Existing methods are computationally expensive as they rely on heavy candidate sampling coupled with scoring, ranking, and fine-tuning steps. We challenge this paradigm with EquiBind, an SE(3)-equivariant geometric deep learning model performing direct-shot prediction of both i) the receptor binding location (blind docking) and ii) the ligand's bound pose and orientation. EquiBind achieves significant speed-ups and better quality compared to traditional and recent baselines. Further, we show extra improvements when coupling it with existing fine-tuning techniques at the cost of increased running time. Finally, we propose a novel and fast fine-tuning model that adjusts torsion angles of a ligand's rotatable bonds based on closed-form global minima of the von Mises angular distance to a given input atomic point cloud, avoiding previous expensive differential evolution strategies for energy minimization.

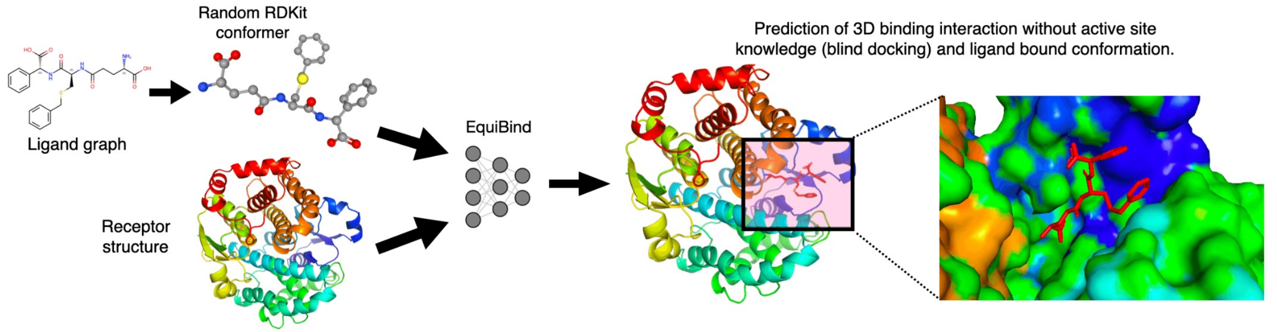

GNNs for drug design

Cell

Volume 180, Issue 4, 20 February 2020, Pages 688-702.e13

CellPress

Article

A Deep Learning Approach to Antibiotic Discovery

Jonathan M. Stokes [1, 2, 3], Kevin Yang [3, 4, 10], Kyle Swanson [3, 4, 10], Wengong Jin [3, 4], Andres Cubillos-Ruiz [1, 2, 5], Nina M. Donghia [1, 5], Craig R. MacNair [6], Shawn French [6], Lindsey A. Carfrae [6], Zohar Bloom-Ackermann [2, 7], Victoria M. Tran [2], Anush Chiappino-Pepe [5, 7], Ahmed H. Badran [2], Ian W. Andrews [1, 2, 5], Emma J. Chory [1, 2], George M. Church [5, 7, 8], Eric D. Brown [6], Tommi S. Jaakkola [3, 4] … James J. Collins [1, 2, 5, 8, 9, 11]

Highlights

- A deep learning model is trained to predict antibiotics based on structure
- Halicin is predicted as an antibacterial molecule from the Drug Repurposing Hub
- Halicin shows broad-spectrum antibiotic activities in mice
- More antibiotics with distinct structures are predicted from the ZINC15 database

https://www.sciencedirect.com/science/article/pii/S0092867420301021

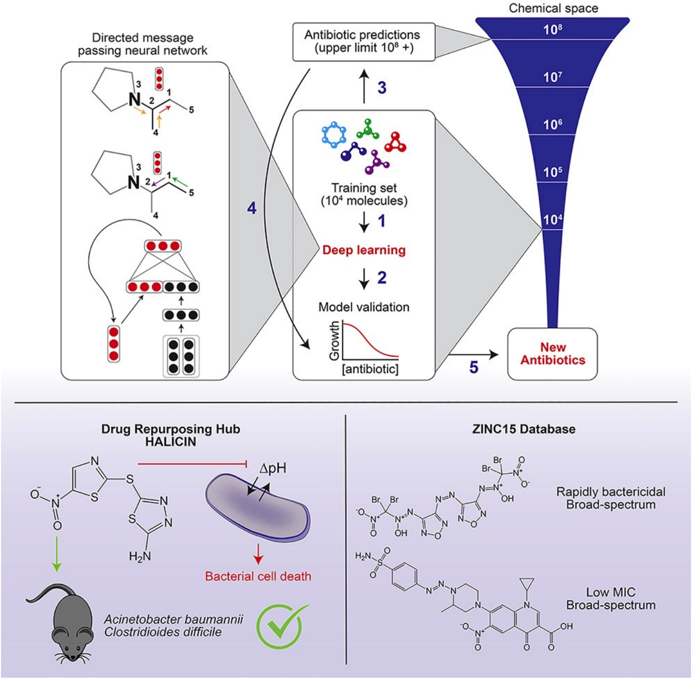

nature

NEWS · 20 FEBRUARY 2020

Powerful antibiotics discovered using AI

Machine learning spots molecules that work even against 'untreatable' strains of bacteria.

FINANCIAL TIMES

Artificial intelligence

AI discovers antibiotics to treat drug-resistant diseases

Machine learning uncovers potent new drug able to kill 35 powerful bacteria

Quanta magazine

ARTIFICIAL INTELLIGENCE

Machine Learning Takes On Antibiotic Resistance

To combat resistant bacteria and refill the trickling antibiotic pipeline, scientists are getting help from deep learning networks.

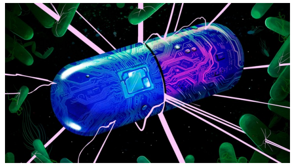

MIT Technology Review

77 Mass Ave

AI vs. bacteria

GNNs for energy physics


Fermilab
Home | About | Science | Jobs | Contact | Phone Book
Newsroom | DUNE at LBNF | Come visit us | Resources for
News
Newsroom
News and features
Press releases
Fermilab in the news

The next big thing: the use of graph neural networks to discover particles

September 24, 2020 | Zack Savitsky

Share Tweet Email

> "What was a week ago just an object of research is now a widely usable tool that could transform our ability to analyze data from particle physics experiments." – Lindsey Gray

"You can still apply all of the same things we're learning about graph neural networks in the HGCal to other detectors in other experiments," Gray said. "The rate at which we're adopting machine learning in high-energy physics is not even close to saturated yet. People will keep finding more and more ways to apply it."

https://news.fnal.gov/2020/09/the-next-big-thing-the-use-of-graph-neural-networks-to-discover-particles

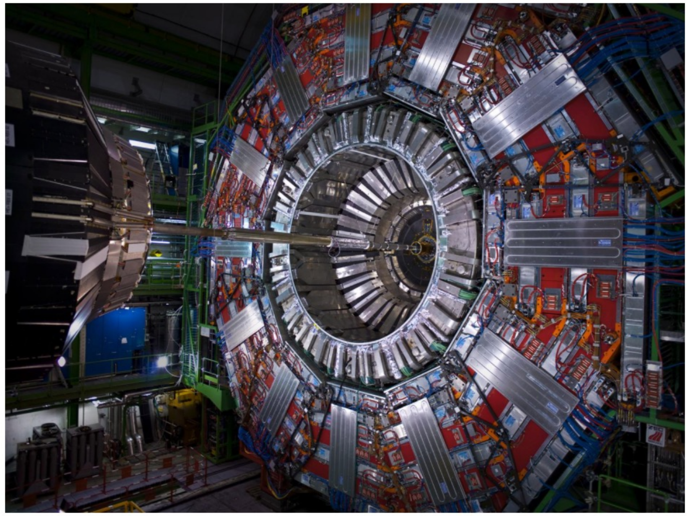

The CMS detector at the Large Hadron Collider takes billions of images of high-energy collisions every second to search for evidence of new particles. Graph neural networks expeditiously decide which of these data to keep for further analysis. Photo: CERN

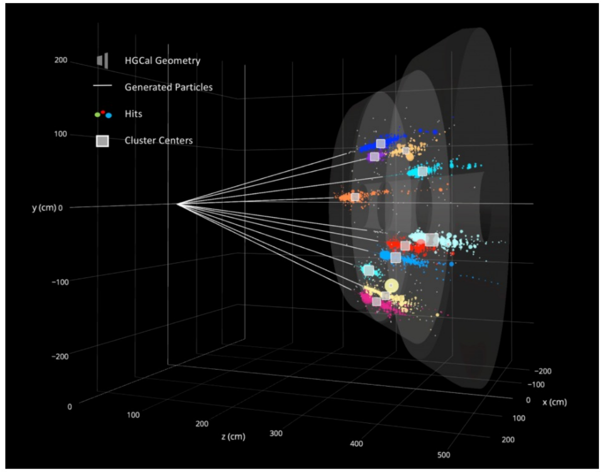

The upgraded high-granularity calorimeter — a component of the CMS detector at the Large Hadron Collider — produces complicated images of particles generated from collisions. Researchers are working to implement graph neural networks to optimize the analysis of this data to better identify and characterize particle interactions of interest. Image courtesy of Ziheng Chen, Northwestern University

GNNs for energy physics

- High-energy physics with neutrino detection (hard to detect because they have a very small chance of interacting with regular matter).

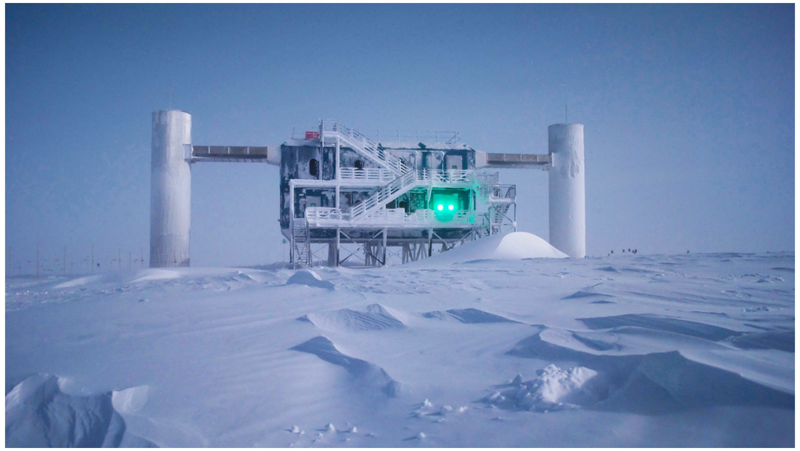

IceCube Neutrino Observatory

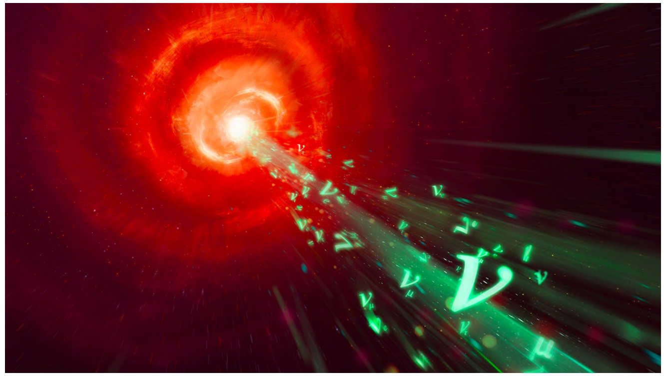

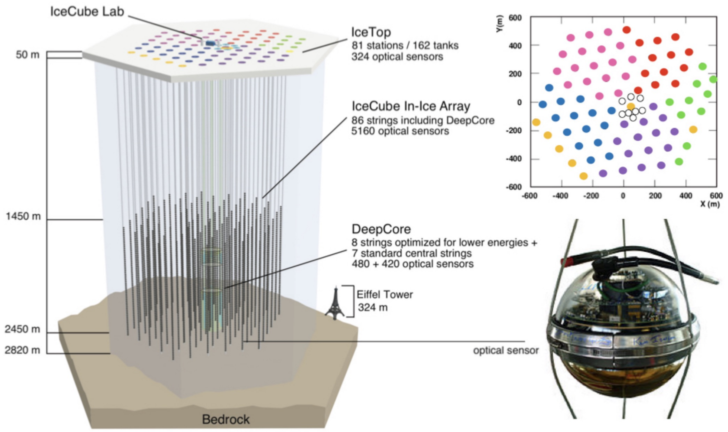

5,160 sensors

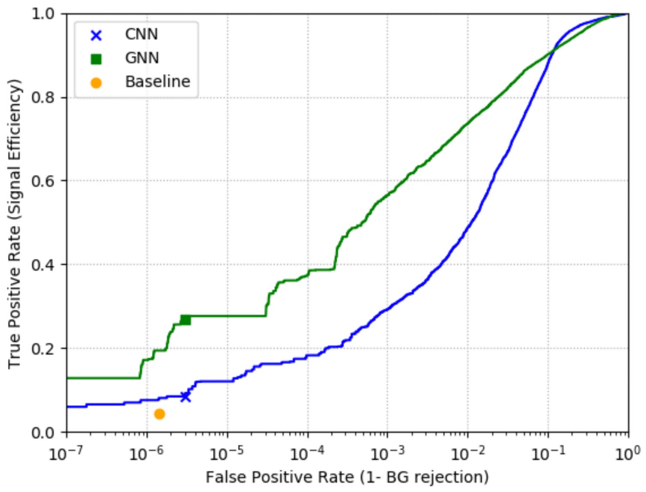

GNNs for physics simulation

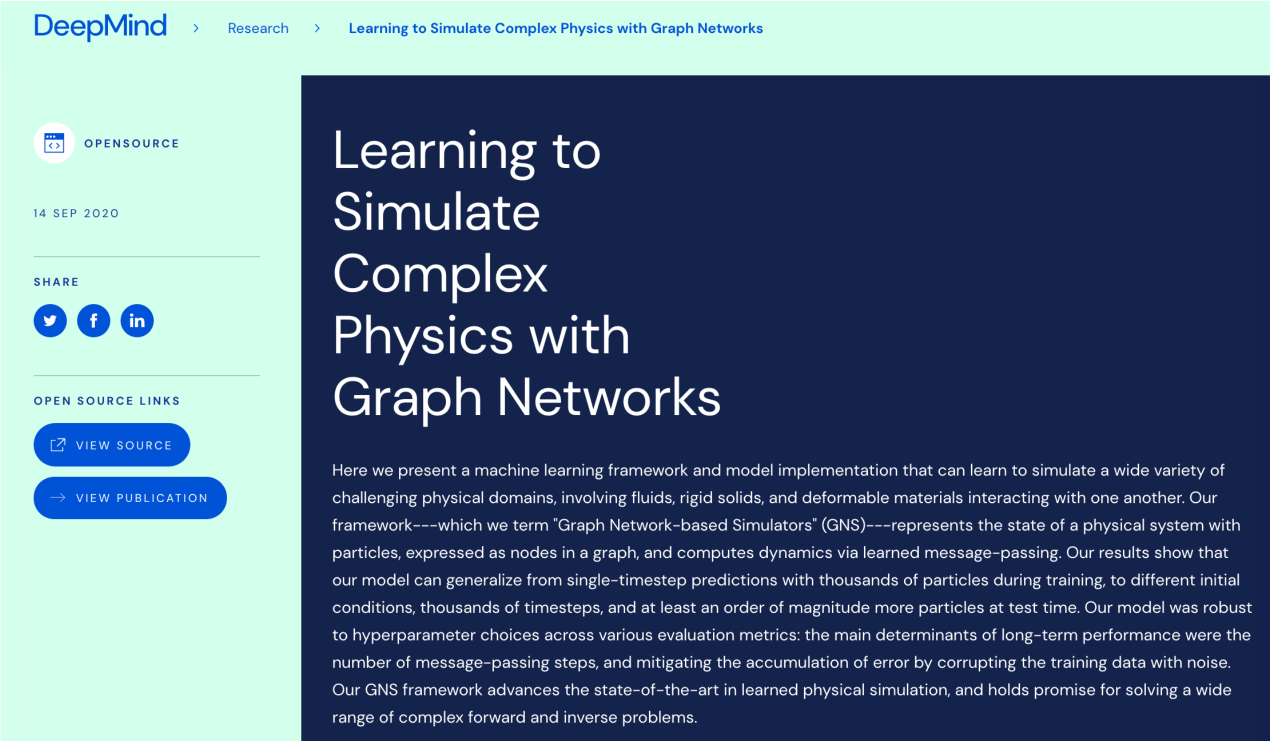
DeepMind › Research › Learning to Simulate Complex Physics with Graph Networks

OPENSOURCE

14 SEP 2020

SHARE

OPEN SOURCE LINKS

VIEW SOURCE

VIEW PUBLICATION

Learning to Simulate Complex Physics with Graph Networks

Here we present a machine learning framework and model implementation that can learn to simulate a wide variety of challenging physical domains, involving fluids, rigid solids, and deformable materials interacting with one another. Our framework---which we term "Graph Network-based Simulators" (GNS)---represents the state of a physical system with particles, expressed as nodes in a graph, and computes dynamics via learned message-passing. Our results show that our model can generalize from single-timestep predictions with thousands of particles during training, to different initial conditions, thousands of timesteps, and at least an order of magnitude more particles at test time. Our model was robust to hyperparameter choices across various evaluation metrics: the main determinants of long-term performance were the number of message-passing steps, and mitigating the accumulation of error by corrupting the training data with noise. Our GNS framework advances the state-of-the-art in learned physical simulation, and holds promise for solving a wide range of complex forward and inverse problems.

https://deepmind.com/research/open-source/Learning-to-Simulate-Complex-Physics-with-Graph-Networks

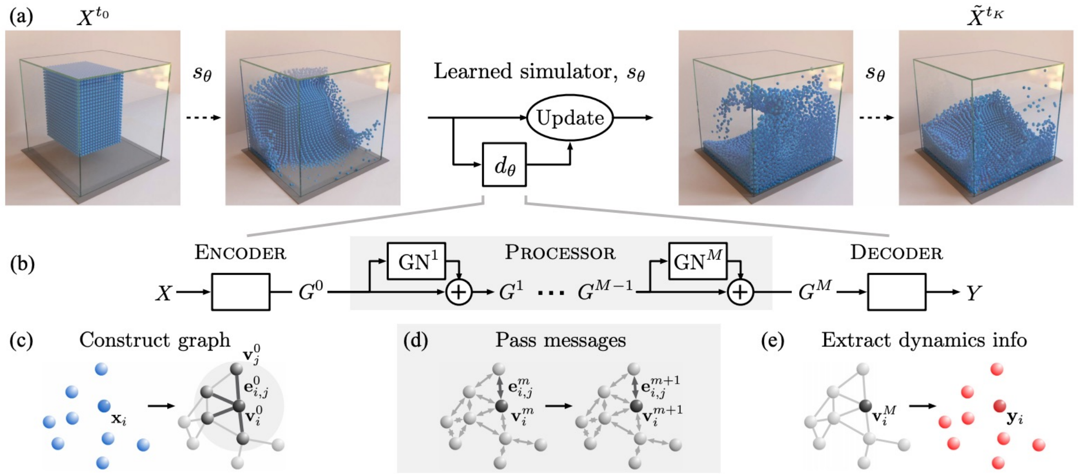

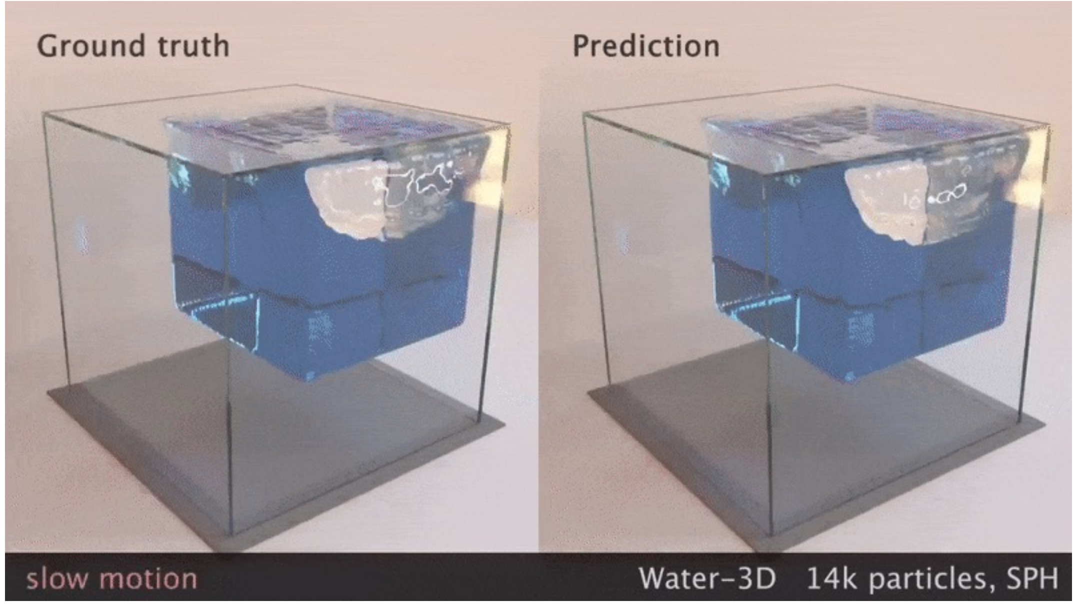

GNNs for code bug detection

- Experiments conducted on a large dataset with 4.9M methods in 92 different project versions show that GNNs have a relative improvement up to 160% on F-score when comparing with the state-of-the-art bug detection approaches.

```
    private void print(InputStream body, String jobName) throws PrintException {
        if (printerOperations.getPrintService().isDocFlavorSupported(printerOperations.
            getFlavor())) {
            PrintDocument printDoc = new PrintDocument(body, printerOperations.getFlavor());
-           printerOperations.print(printDoc, config.getCopies(), config.isSendToPrinter(),
                config.getMimeType(), jobName);
+           printerOperations.print(printDoc, config.isSendToPrinter(), config.getMimeType(),
                jobName);
        }
    }
```

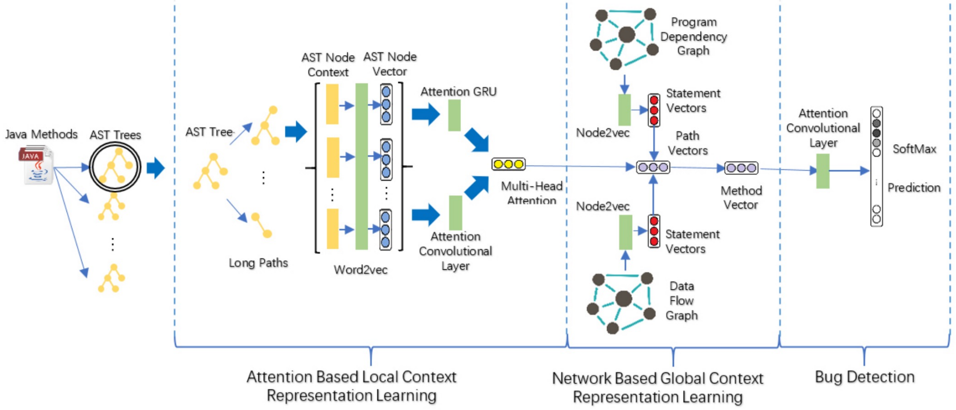

Fig. 2. Overview of our Approach.

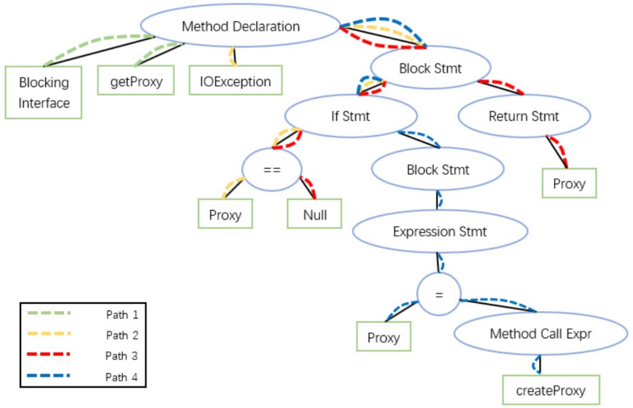

(b) The AST of the code in Figure 3a.

https://dl.acm.org/doi/pdf/10.1145/3360588

GNNs for genomics

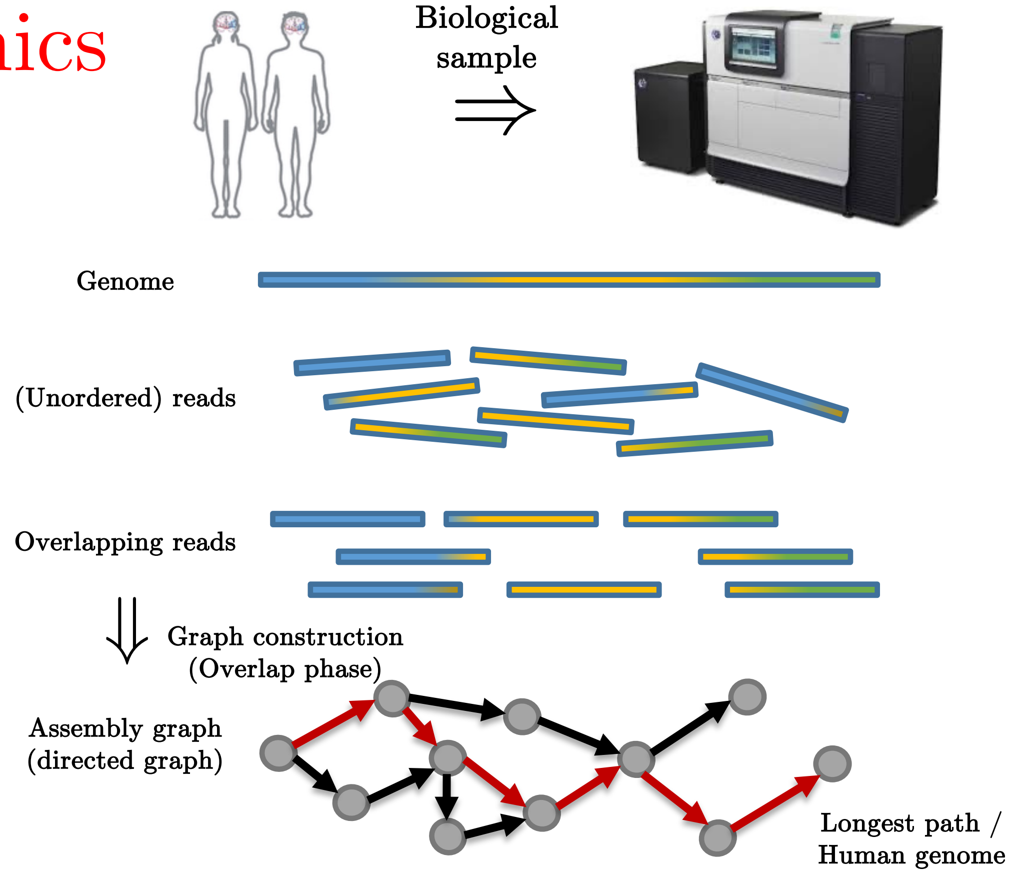

- We proposed to use GNN to replace human heuristics ⇒ AI-based genome assembler
- Advantage : Solve genome assembly independently of any type of sequencing machine and no hand-crafting of genome assemblers.
- Given a state-of-the-art genome assembler (Raven), we demonstrated that learned heuristics with GNN outperforms human engineered rules.

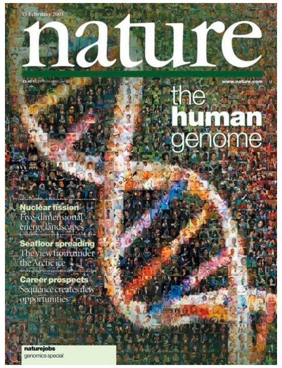

2001

2022

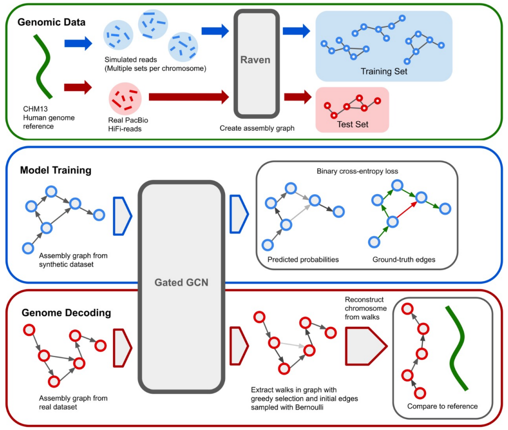

https://arxiv.org/pdf/2206.00668.pdf

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

GNNs for industry

Market potential has a large landscape :

GNN is a general technology that can be applied to several tasks

- Finance: fraud/anomaly detection, forecast prediction for e.g. sales, investments, credit risk
- Recommendation: better targeted ads, improved user experience/loyalty
- Knowledge graphs: enhanced CRM, query relationships/interactions with company and customers
- Healthcare: drug design, new diagnostic tools for doctors s.a. brain analysis
- Robotics: better 3D point representation, planning and reasoning
- NLP: improved Q&A chatbot with contextual graphs
- Resource management: supply chain and warehouse/inventory optimization
- Transportation: more accurate and dynamic delivery time
- Etc.

GNNs for industry

- Recent technology
- Kumo : https://kumo.ai
  - Start-up raised \$18.5 Million in July 2022
  - A co-funder is Jure Leskovec (Stanford)
- Amazon Neptune: https://aws.amazon.com/neptune/machine-learning

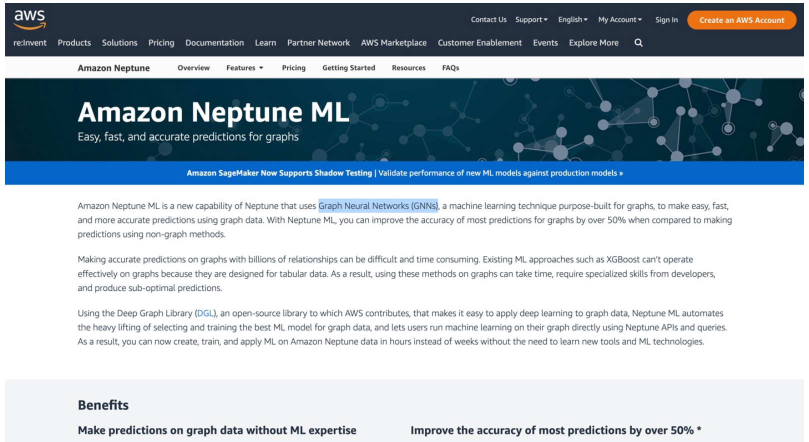
aws
Contact Us Support English My Account Sign In Create an AWS Account
re:Invent Products Solutions Pricing Documentation Learn Partner Network AWS Marketplace Customer Enablement Events Explore More

Amazon Neptune Overview Features Pricing Getting Started Resources FAQs

Amazon Neptune ML

Easy, fast, and accurate predictions for graphs

Amazon SageMaker Now Supports Shadow Testing | Validate performance of new ML models against production models »

Amazon Neptune ML is a new capability of Neptune that uses Graph Neural Networks (GNNs), a machine learning technique purpose-built for graphs, to make easy, fast, and more accurate predictions using graph data. With Neptune ML, you can improve the accuracy of most predictions for graphs by over 50% when compared to making predictions using non-graph methods.

Making accurate predictions on graphs with billions of relationships can be difficult and time consuming. Existing ML approaches such as XGBoost can't operate effectively on graphs because they are designed for tabular data. As a result, using these methods on graphs can take time, require specialized skills from developers, and produce sub-optimal predictions.

Using the Deep Graph Library (DGL), an open-source library to which AWS contributes, that makes it easy to apply deep learning to graph data, Neptune ML automates the heavy lifting of selecting and training the best ML model for graph data, and lets users run machine learning on their graph directly using Neptune APIs and queries. As a result, you can now create, train, and apply ML on Amazon Neptune data in hours instead of weeks without the need to learn new tools and ML technologies.

Benefits

Make predictions on graph data without ML expertise

Improve the accuracy of most predictions by over 50% *

Amazon Neptune

kumo

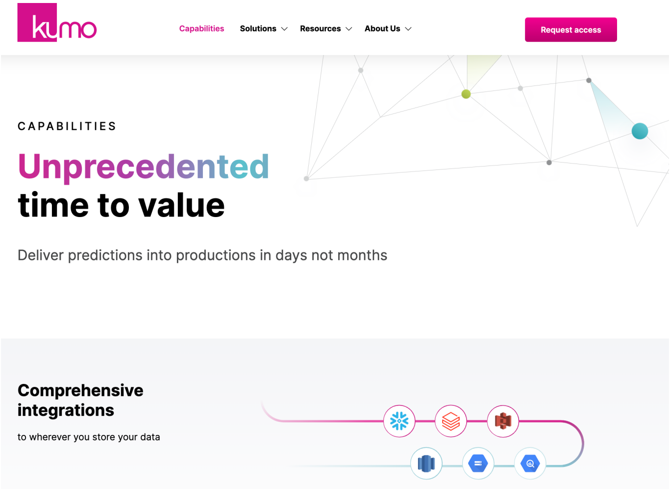

GNN pipeline

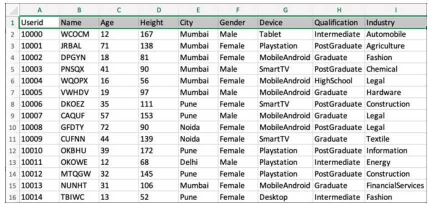

	A	B	C	D	E	F	G	H	I
1	Userid	Name	Age	Height	City	Gender	Device	Qualification	Industry
2	10000	WCOCM	12	167	Mumbai	Male	Tablet	Intermediate	Automobile
3	10001	JRBAL	71	138	Mumbai	Female	Playstation	PostGraduate	Agriculture
4	10002	DPGYN	18	81	Mumbai	Female	MobileAndroid	Graduate	Fashion
5	10003	PNSQX	41	90	Mumbai	Male	SmartTV	PostGraduate	Chemical
6	10004	WQOPX	16	56	Mumbai	Female	MobileAndroid	HighSchool	Legal
7	10005	VWHDV	19	97	Mumbai	Male	MobileAndroid	Graduate	Hardware
8	10006	DKOEZ	35	111	Pune	Female	SmartTV	PostGraduate	Construction
9	10007	CAQUF	57	153	Pune	Male	MobileAndroid	Graduate	Legal
10	10008	GFDTY	72	90	Noida	Female	MobileAndroid	PostGraduate	Legal
11	10009	CUFNN	44	139	Noida	Female	SmartTV	Graduate	Textile
12	10010	OKBHU	39	172	Pune	Female	Playstation	PostGraduate	Information
13	10011	OKOWE	12	68	Delhi	Male	Playstation	Intermediate	Energy
14	10012	MTQGW	32	145	Pune	Female	Playstation	PostGraduate	Construction
15	10013	NUNHT	31	106	Mumbai	Female	MobileAndroid	Graduate	FinancialServices
16	10014	TBIWC	13	52	Pune	Female	Desktop	Intermediate	Fashion

- Collect data (user/customer features, product features, etc.)

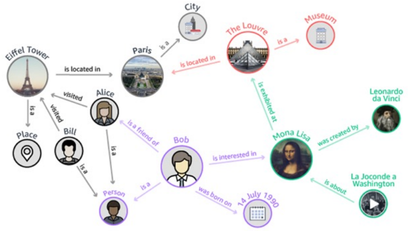

- Generate graphs from data relationships and features

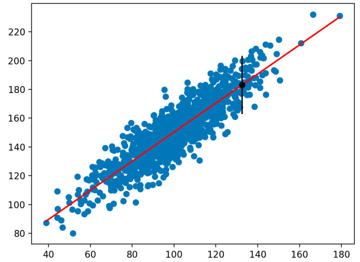

- Train/validate GNN for given predictive tasks

- Cloud storage, security, computing and deployment

- Customers query an API to get predictions
  Easy to use for non-expert DL/GNN customers

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN books and libraries
- Conclusion

Books

- Graph Representation Learning Book, Springer, 2020
  - William Hamilton, McGill University
  - https://link.springer.com/book/10.1007/978-3-031-01588-5
  - https://github.com/RHxW/CV-DL-Docs/blob/master/GRL_Book.pdf

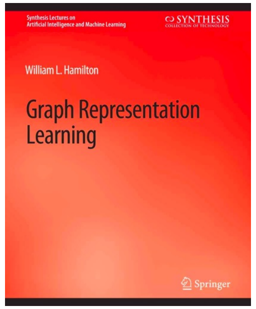

- Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges, 2021
  - Michael M. Bronstein (Oxford), Joan Bruna (NYU), Taco Cohen (Qualcomm), Petar Veličković (DeepMind)
  - https://arxiv.org/pdf/2104.13478.pdf

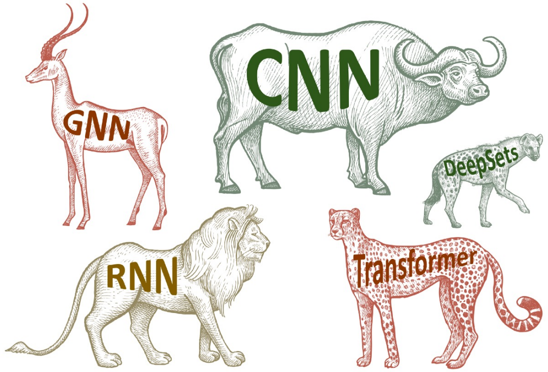

- Graph Neural Networks: Foundations, Frontiers, Applications, Springer, 2022
  - Lingfei Wu (Pinterest), Peng Cui (Tsinghua), Jian Pei, (Duke), Liang Zhao (Emory University)
  - https://graph-neural-networks.github.io (English and Chinese versions)

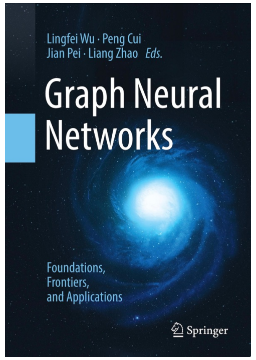

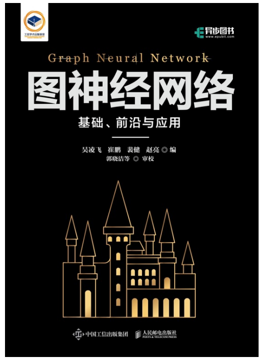

Libraries

- Amazon DGL (Deep Graph Library)
  - First released in Dec 2018 (11k+ stars)
  - PyTorch / TensorFlow / MxNet
  - https://www.dgl.ai

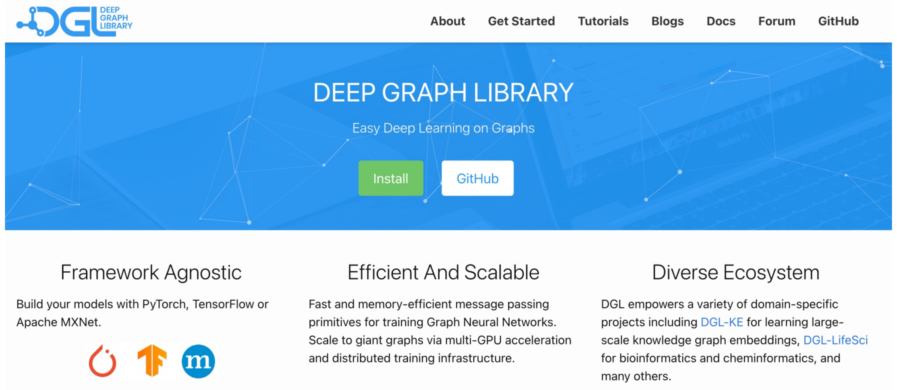

- Kumo PyG (PyTorch Geometric)
  - First released in Mar 2019 (16k+ stars)
  - PyTorch
  - https://www.pyg.org

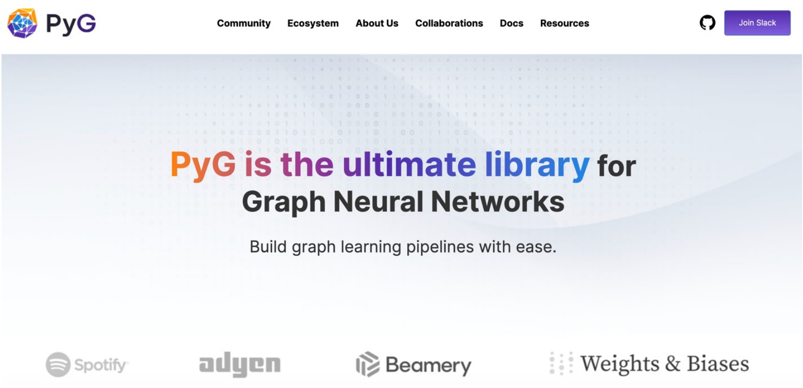

- TensorFlow GNNs
  - First released in Nov 2021 (1k+ stars)
  - https://github.com/tensorflow/gnn

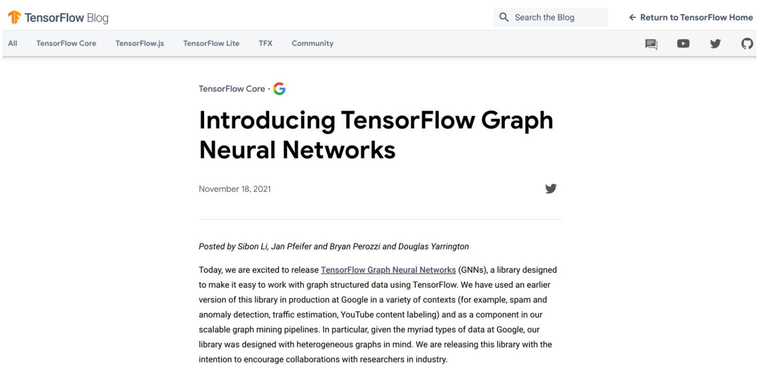

Outline

- The Deep Learning (DL) revolution
- Limitations of DL
- Graph-Structured data
- Graph Neural Networks (GNNs)
- GNN case studies
- GNN for industry
- GNN libraries
- Conclusion

Conclusion

- Graph Neural Networks are a breakthrough in Machine Learning.
  - GNNs have become the standard toolkit for analyzing graph-structured data.
  - They generalize CNNs/RNNs/Transformers from grids/sequences to complex relational data structures.
  - Graphs are everywhere because everything is connected.
  - GNNs are highly flexible and have been/will be applied to a large variety of applications.
  - Large-scale training (linear complexity) with distributed computing.
  - Supervised, reinforcement and self-supervised training.
  - GNNs will boost business analytics.

Tentative Lectures

- → Introduction to Graph Machine Learning
- Part 1: GML without feature learning (before 2014)
  - Introduction to Graph Science
  - Graph Analysis Techniques without Feature Learning
    - Graph clustering
    - Classification
    - Recommendation
    - Dimensionality reduction
- Part 2 : GML with shallow feature learning (2014-2016)
  - Shallow graph feature learning
- Part 3 : GML with deep feature learning, a.k.a. GNNs (after 2016)
  - Graph Convolutional Networks (spectral and spatial)
  - Weisfeiler-Lehman GNNs
  - Graph Transformer & Graph ViT/MLP-Mixer
  - Benchmarking GNNs
  - Molecular science and generative GNNs
  - GNNs for combinatorial optimization
  - GNNs for recommendation
  - GNNs for knowledge graphs
  - Integrating GNNs and LLMs

Questions?