

Deep Learning with Knowledge Graphs

Jure Leskovec



Includes joint work with M. Fey, J. You, S. Jegelka, K. Xu, R. Ying, W. Hu,
M. Zitnik, P. Battaglia, K. Subbian, N. Rao, P. Li, A. Wang

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

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


New frontiers beyond classic neural networks that learn on images and sequences



Grids



Groups



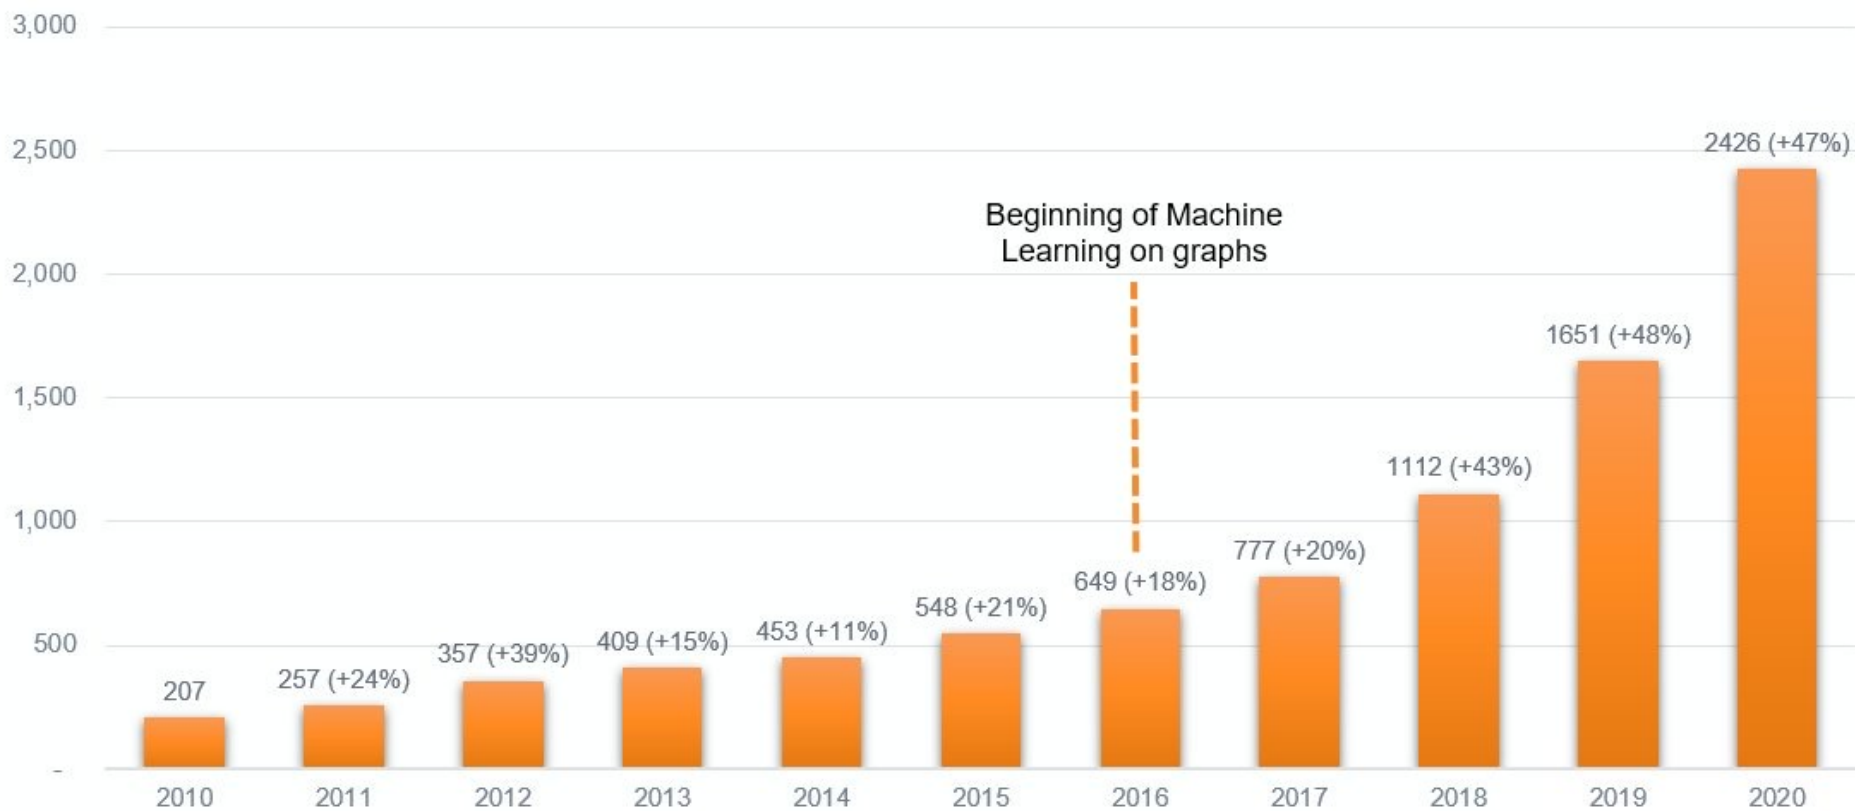
Graphs



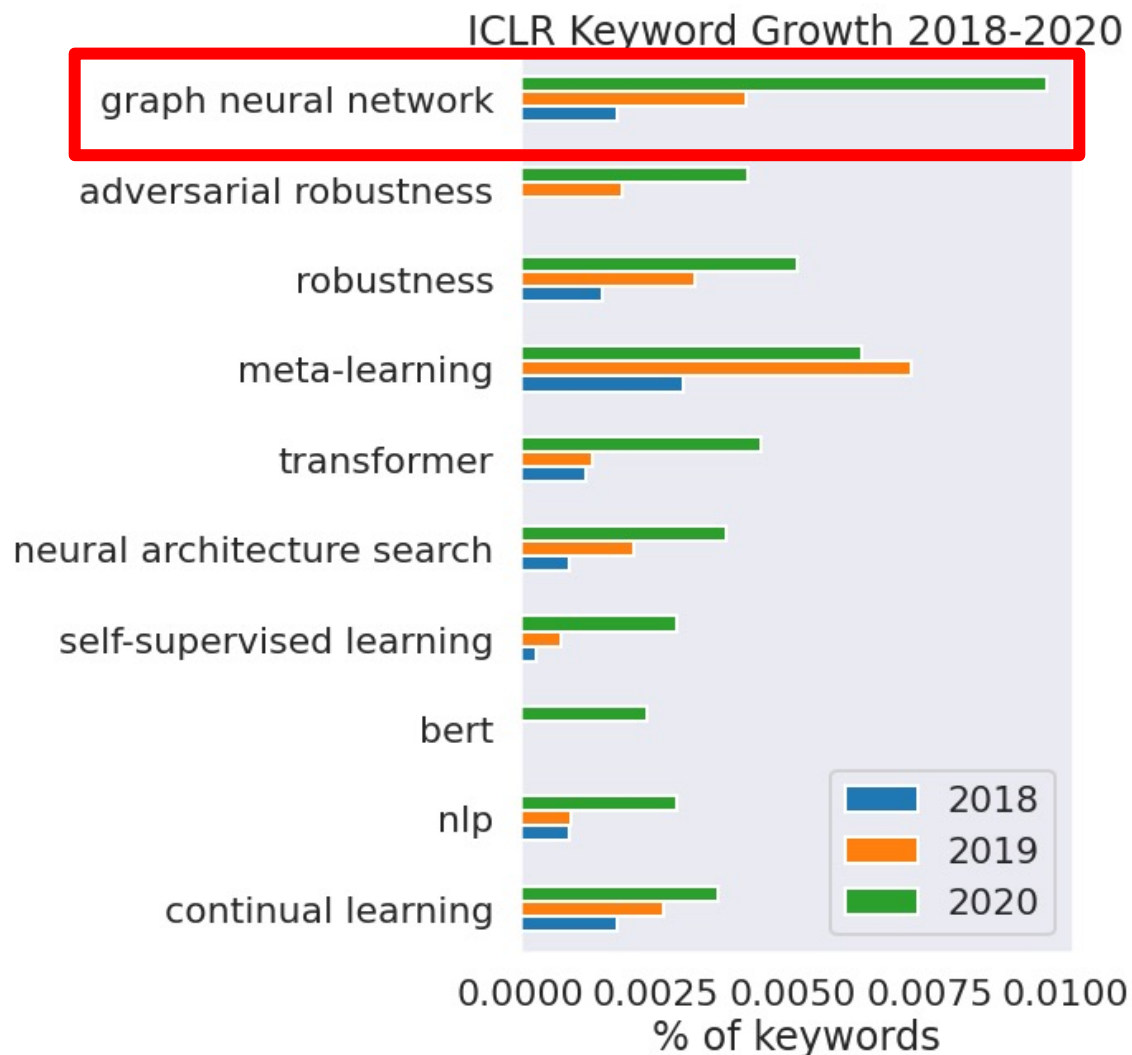
**Geodesics &
Gauges**

Graph ML is on Fire

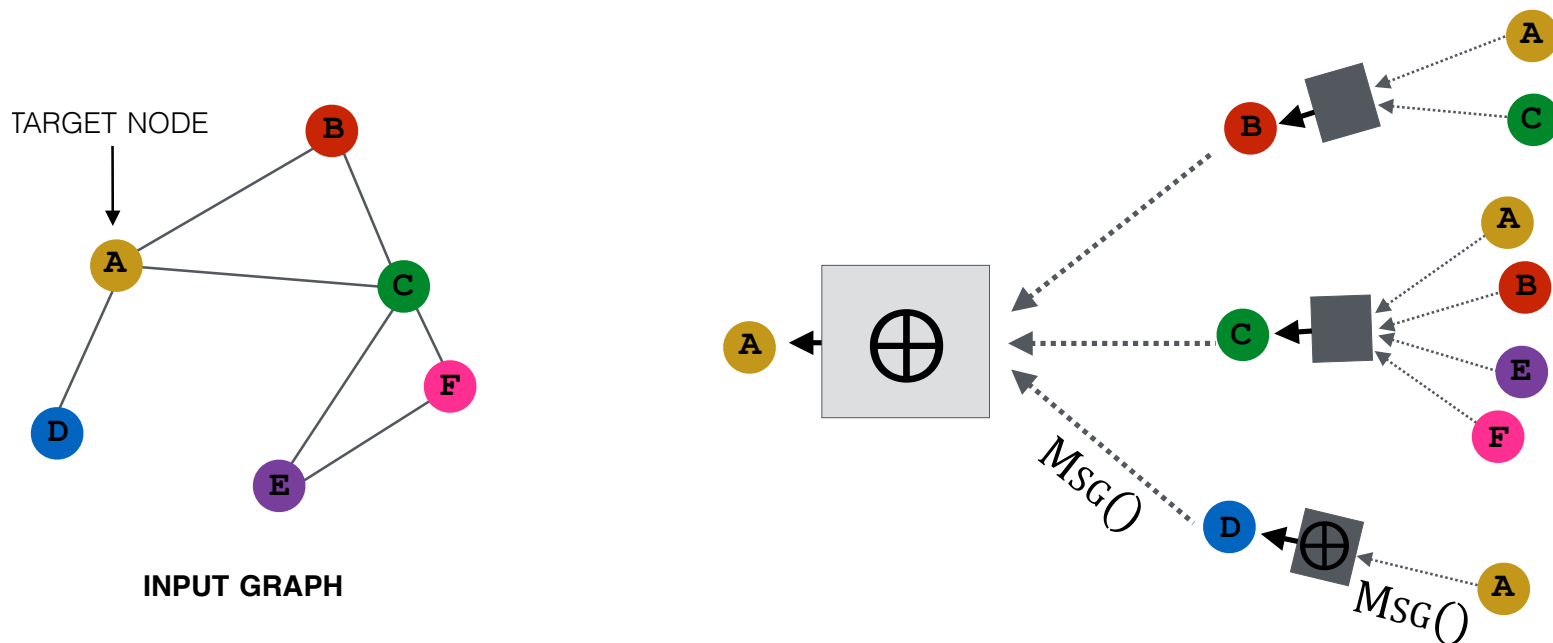
Number of papers with 'graph' in title (ArXiv).



Graph ML is on Fire




Graph Neural Network



Each node defines a computation graph

- Each edge in this graph is a MSG function
- \oplus is a message **aggregation** function

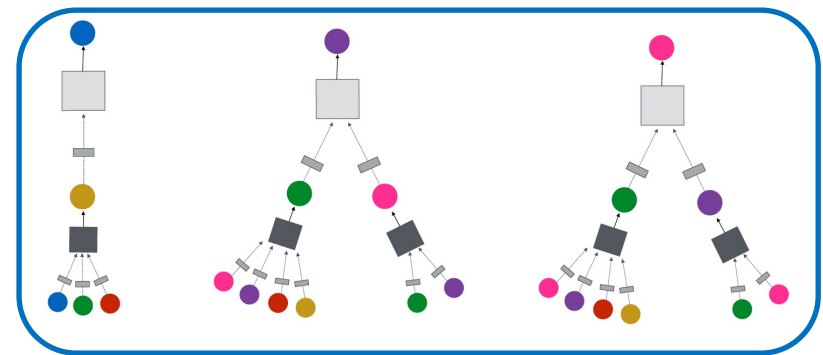

Inductive Capability



INPUT GRAPH

Train the model on a subset of nodes

Make predictions for nodes we never trained on!




Nodes have different computation graphs

Two Views of GNNs

The GNN does two things:

- **1) Powerful feature transformer/smoother**
 - Node features get passed and transformed around node's L -hop neighborhood
- **2) Each node can have a different computation graph and the network is also able to capture/learn its structure**




Key Benefits of GNNs

- GNNs adapt to the **shape** of data
 - Other Deep Learning architectures assume fixed input (matrix, sequence)
- GNNs can integrate multimodal data of various cardinalities and shapes

Key Benefits of GNNs

- GNNs subsume CNNs and Transformers as special cases

But scalability remains a challenge!



Image



Graph



$$\text{GNN formulation: } h_v^{(l+1)} = \sigma(\mathbf{W}_l \sum_{u \in \mathcal{N}(v)} \frac{h_u^{(l)}}{|\mathcal{N}(v)|} + \mathbf{B}_l h_v^{(l)}), \forall l \in \{0, \dots, L-1\}$$

$$\text{CNN formulation: } h_v^{(l+1)} = \sigma(\sum_{u \in \mathcal{N}(v)} \mathbf{W}_l^u h_u^{(l)} + \mathbf{B}_l h_v^{(l)}), \forall l \in \{0, \dots, L-1\}$$

Scalable & Deep GNNs

Deep GNNs on large-scale graphs result in neighbor explosions

- **Exponentially increasing dependency graph of nodes over layers**




degree

d

d^2

d^3

Our Work: GNNAutoScale


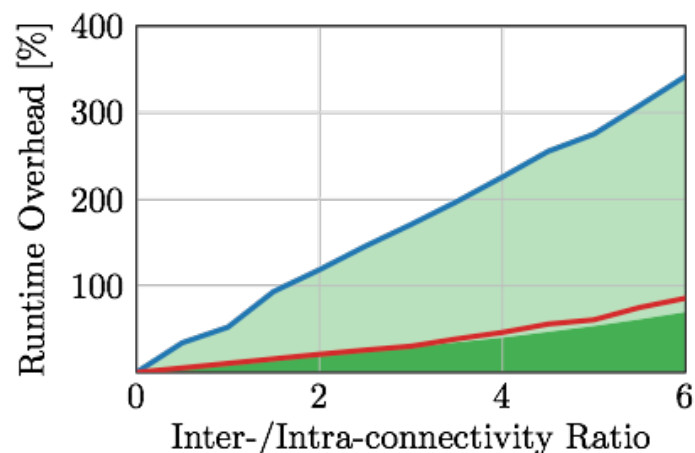


- Pull histories for out-of-minibatch nodes
- Push estimated embeddings to histories
- **Theory:**
 - (1) Bounded approximation error
 - (2) Provably as expressive as the WL-test


GNNAutoScale

Fast history access via **asynchronous device transfers** to and from the GPU

— Serial Access — Concurrent Access
■ Computational Overhead ■ I/O Overhead



(a) Serial execution



(b) Concurrent execution

Experimental Results

Dataset	Runtime (s)		Memory (MB)	
	GTTF	GAS	GTTF	GAS
CORA	0.077	0.006	18.01	2.13
PUBMED	0.071	0.006	28.79	2.19
PPI	0.976	0.007	134.86	12.37
FLICKR	1.178	0.007	325.97	16.32

Method	#nodes	717K	169K	2.4M
	#edges	7.9M	1.2M	61.9M
	Method	YELP	ogbn-arxiv	ogbn-products
2-layer	Full-batch	6.64GB/100%	1.44GB/100%	21.96GB/100%
	GRAPHSAGE	0.76GB/ 9%	0.40GB/ 27%	0.92GB/ 2%
	CLUSTER-GCN	0.17GB/ 13%	0.15GB/ 40%	0.16GB/ 16%
	GAS	0.51GB/100%	0.22GB/100%	0.36GB/100%

- 10-100x speedup over GTTF [ICLR '21]
 - GTTF 10x better than naïve implementation
- 10-20x less memory over GTTF [ICLR '21]
 - GTTF 4x better than naïve implementation
- GNNAutoScale can be applied with any GNN architecture (GCN, SAGE, GAT)



PyG: GNN Library



PyG bundles the state-of-the-art in Graph Representation Learning

- 50+ GNN architectures
- 200+ benchmark datasets
- Extendable via a message passing interface
- Dedicated sparsity-aware CUDA kernels
- Support for various scalability techniques
- Heterogeneous Graph Support
- GNN Design Space Exploration



PyG: GNN Library



PyG is  PyTorch-on-the-rocks

- Keeps **design principles** close to vanilla PyTorch
- Fits nicely into the **PyTorch ecosystem**

```
from torch.nn import Conv2d

class CNN(torch.nn.Module):
    def __init__(self):
        self.conv1 = Conv2d(3, 64)
        self.conv2 = Conv2d(64, 64)

    def forward(self, input):
        h = self.conv1(input)
        h = h.relu()
        h = self.conv2(h)
        return h
```

```
from torch_geometric.nn import GCNConv

class GNN(torch.nn.Module):
    def __init__(self):
        self.conv1 = GCNConv(3, 64)
        self.conv2 = GCNConv(64, 64)

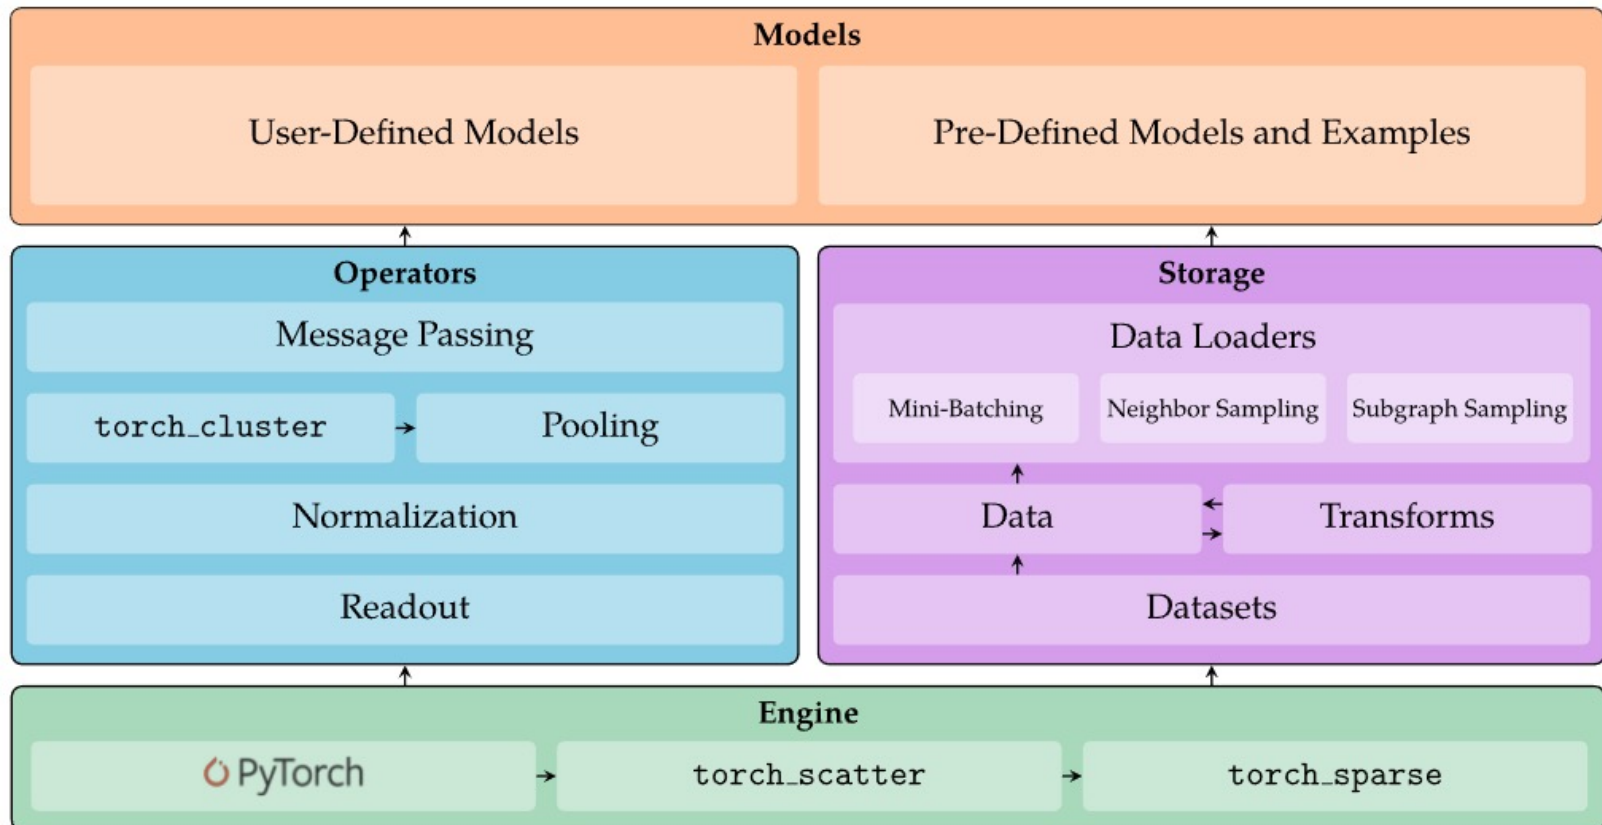
    def forward(self, input, edge_index):
        h = self.conv1(input, edge_index)
        h = h.relu()
        h = self.conv2(h, edge_index)
        return h
```




PyG: GNN Library



PyG provides the **state-of-the-art** in **Graph Representation Learning**





PyG: GNN Library

[PYG.ORG](https://pyg.org):

- ~800 research papers written using  PyG
- ~40K monthly downloads
- ~250 external contributors
- ~1800 members on Slack



PyG



MIT license



Cite this repository ▾



14.5k stars



241 watching



2.6k forks



Stanford
University



Pinterest

amazon



NVIDIA®

AIRBUS

pwc

intel®


Uber

PayPal




What are some
applications of PyG?


Graphs are Everywhere



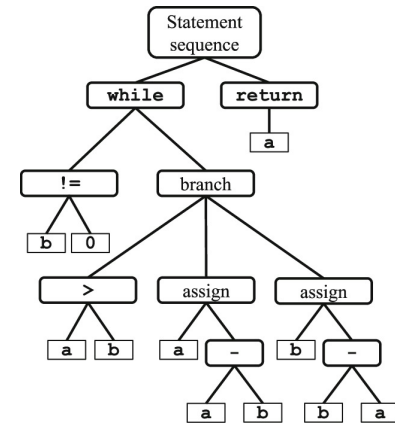
Molecular graphs




Social networks




Knowledge graph




AST



Biological networks



Citation networks



3D meshes

Real-world data is full of relations.

Graphs model how “stuff” is connected.

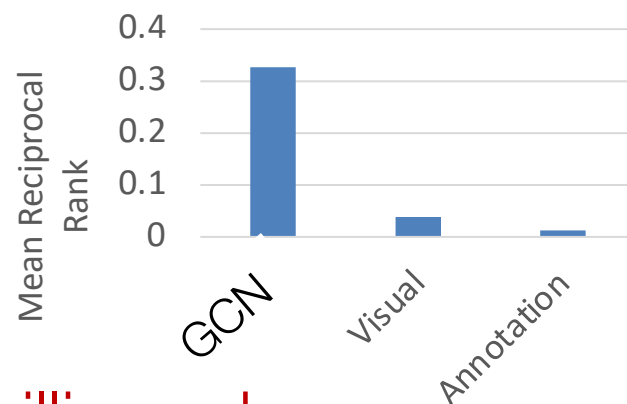
(1): Recommender Engine



Task: Recommend related pins



Task: Learn node embeddings z_i s.t.
 $d(z_{cake1}, z_{cake2}) < d(z_{cake1}, z_{sweater})$



Challenges:




- Massive size: 3 billion nodes, 20 billion edges
- Heterogeneous data: Rich image and text features

(2) Fraud & Intrusion Detection

Fraud and intrusion detection in dynamic transaction graphs

Financial networks

Communication networks



(3) Text + Graph Reasoning

Question

If it is not used for **hair**, a **round brush** is an example of what?
A. **hair brush** B. **bathroom** C. **art supplies*** D. **shower**

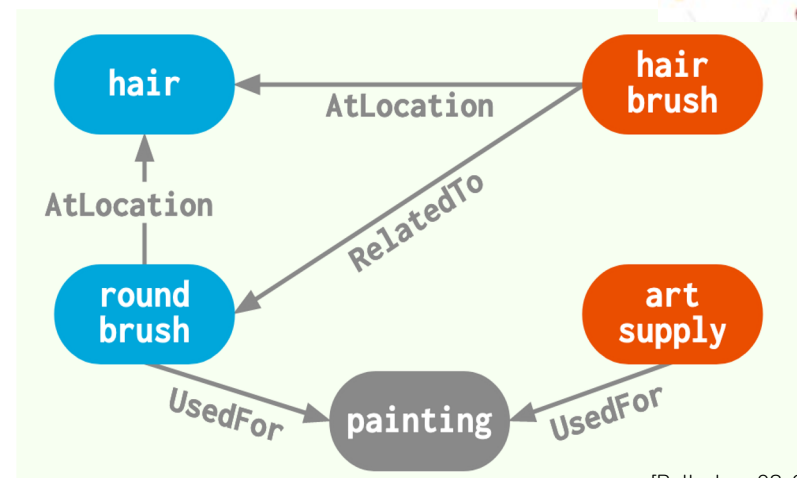
Knowledge sources

Pre-trained language model (LM)



[Devlin+19; Liu+19; Brown+20; ...]

Knowledge Graph (KG)



[Bollacker+08; Speer+16]


ConceptNet
An open, multilingual knowledge graph

Freebase

Interpretability

(a) Attention visualization direction: BFS from **Q**


Where would you find a **basement** that can be accessed with an **elevator**? A. **closet** B. **church** C. **office building***




Structured Reasoning

Original Question

If it is **not** used for **hair**, a **round brush** is an example of what?
A. **hair brush** B. **art supply***



GNN 1st Layer



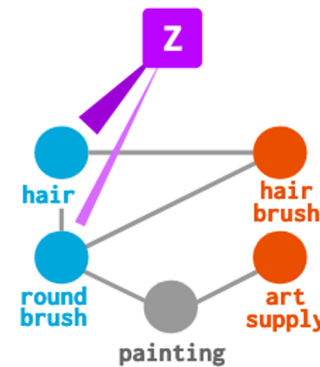
GNN Final Layer

A. hair brush (0.38)
B. art supply (0.64)

Model Prediction

(a) Negation Flipped

If it is used for **hair**, a **round brush** is an example of what?
A. **hair brush** B. **art supply**

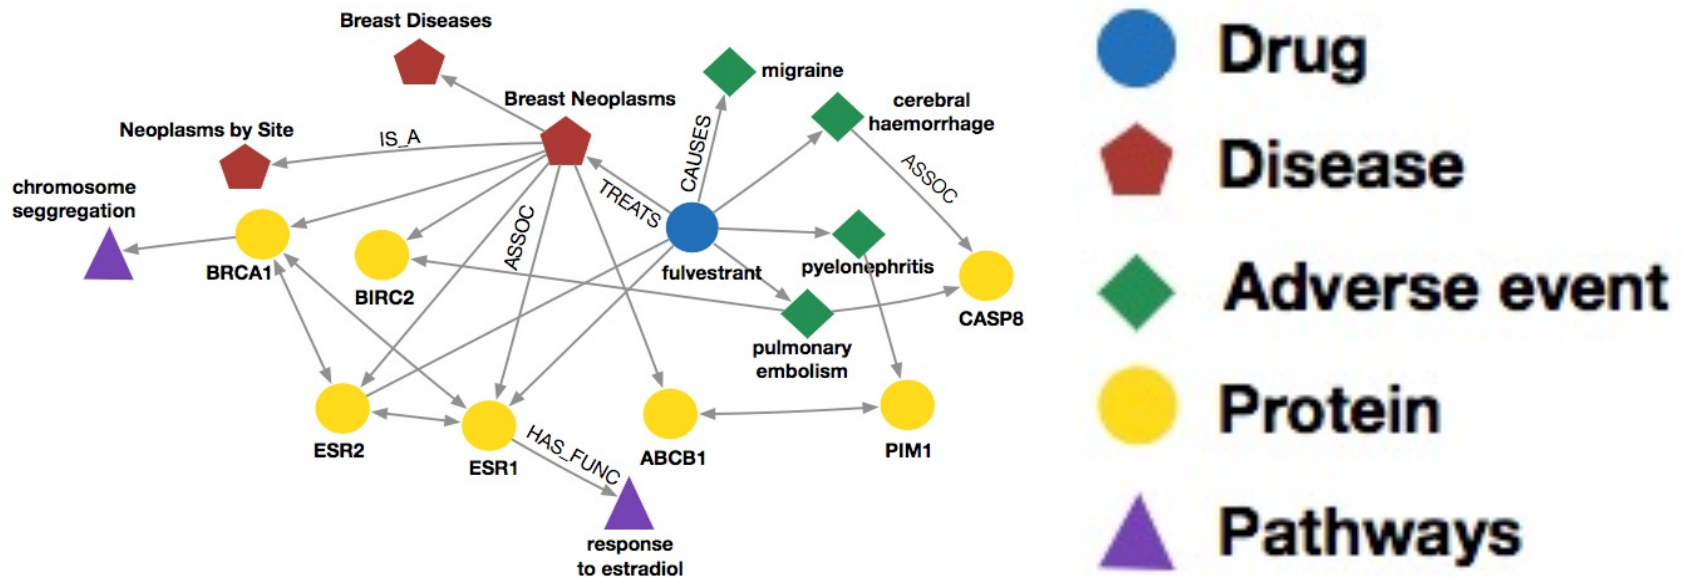


GNN Final Layer

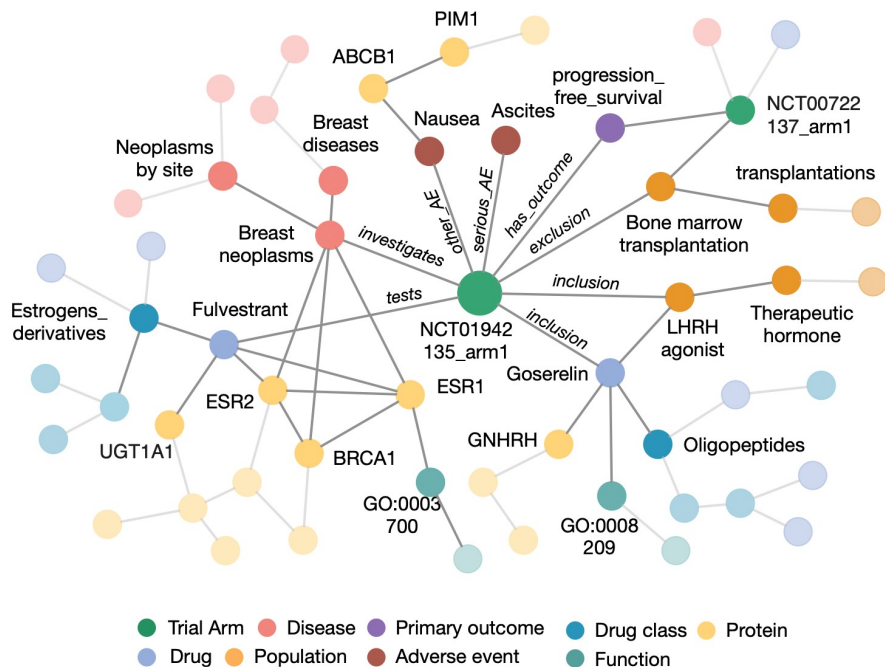
A. **hair brush (0.81)**
B. art supply (0.19)

Model Prediction

(4) Data in Biomedicine



TrialNet Knowledge Graph

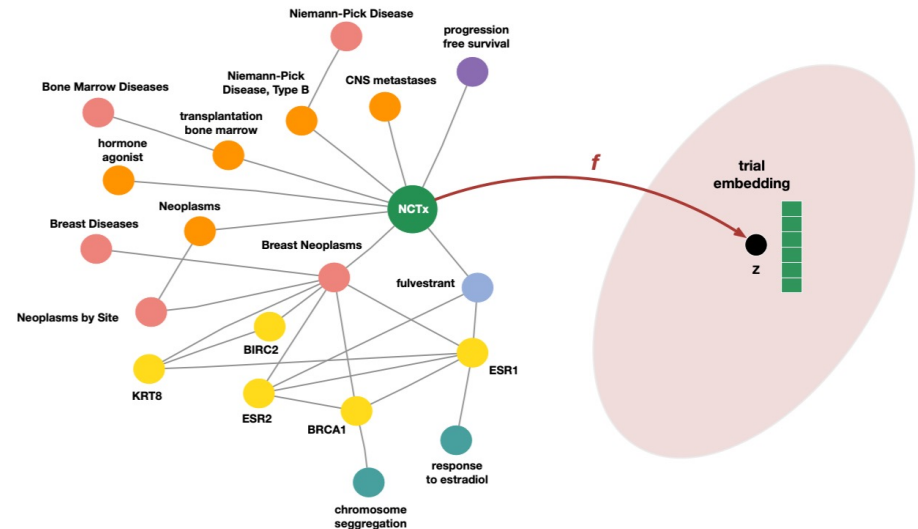


- Covers **70k interventional clinical trials**
- **300k nodes** (clinical trials protocol entities, biological and chemical entities)
- **10.8 millions** edges across 18 relation types

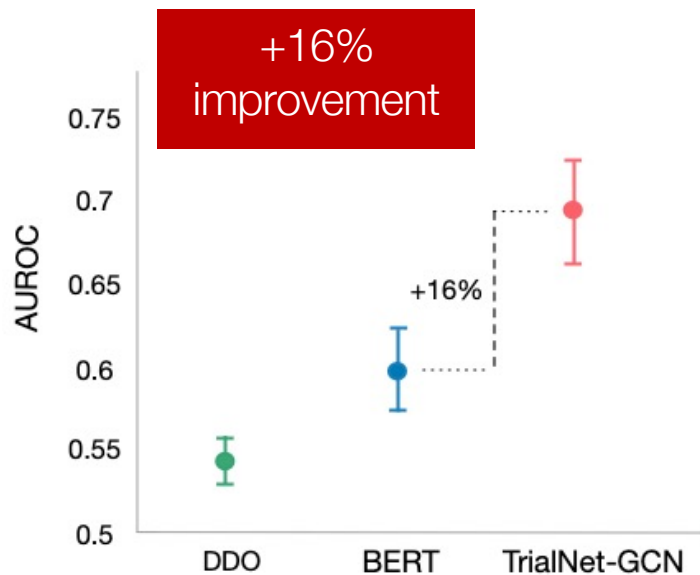
Learning Embeddings

How can we leverage TrialNet and build predictive models?

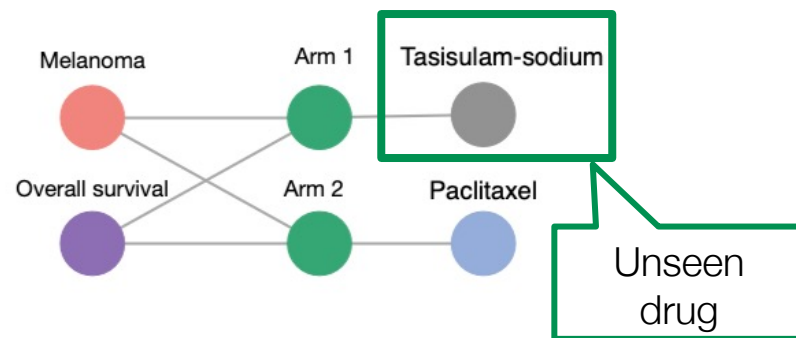
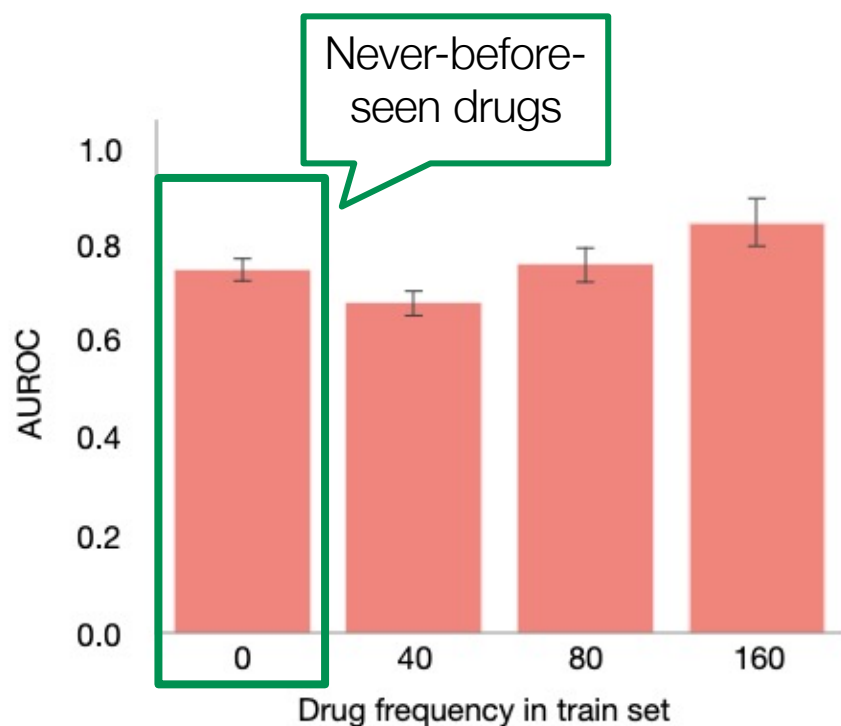
Key idea: Learn to embed trials, drugs, diseases, population, proteins...



Predicting Drug Efficacy



Generalization to New Drugs



Which arm has better outcome??

Model prediction

Tasisulam: 0.25

Paclitaxel: 0.75

Ground truth

Tasisulam: 6.8 months

Paclitaxel: 9.4 months

Safety terminated study

(9) ML on Business data

Production ML

An on-call engineer's biggest nightmare

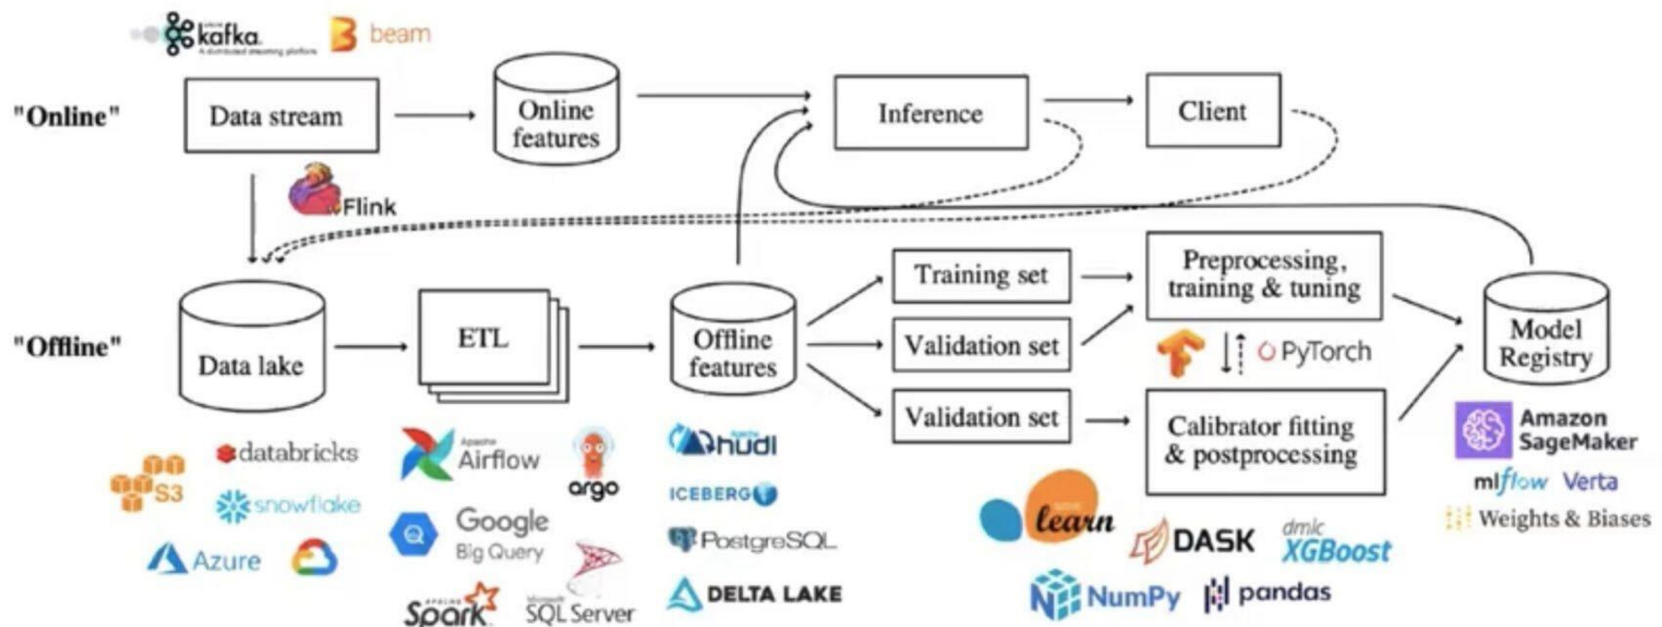
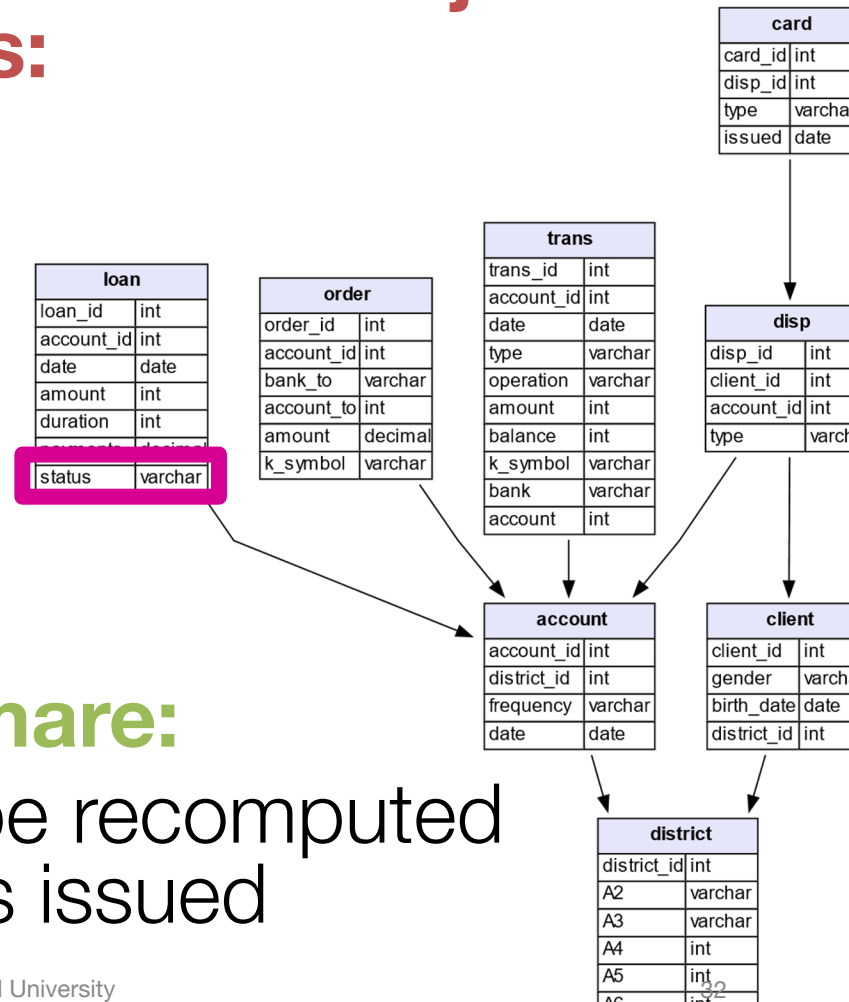


Figure 1: High-level architecture of a generic end-to-end machine learning pipeline. Logos represent a sample of tools used to construct components of the pipeline, illustrating heterogeneity in the tool stack. Shankar et al. 2021

Example: Predict Credit Risk

Data scientists writes workflows to join tables, generate features:


- account age
- avg. balance
- total # transactions
- # cards
- # bad transactions
- ...




Time consistency nightmare:

Every feature needs to be recomputed for the time the loan was issued

Representation Learning

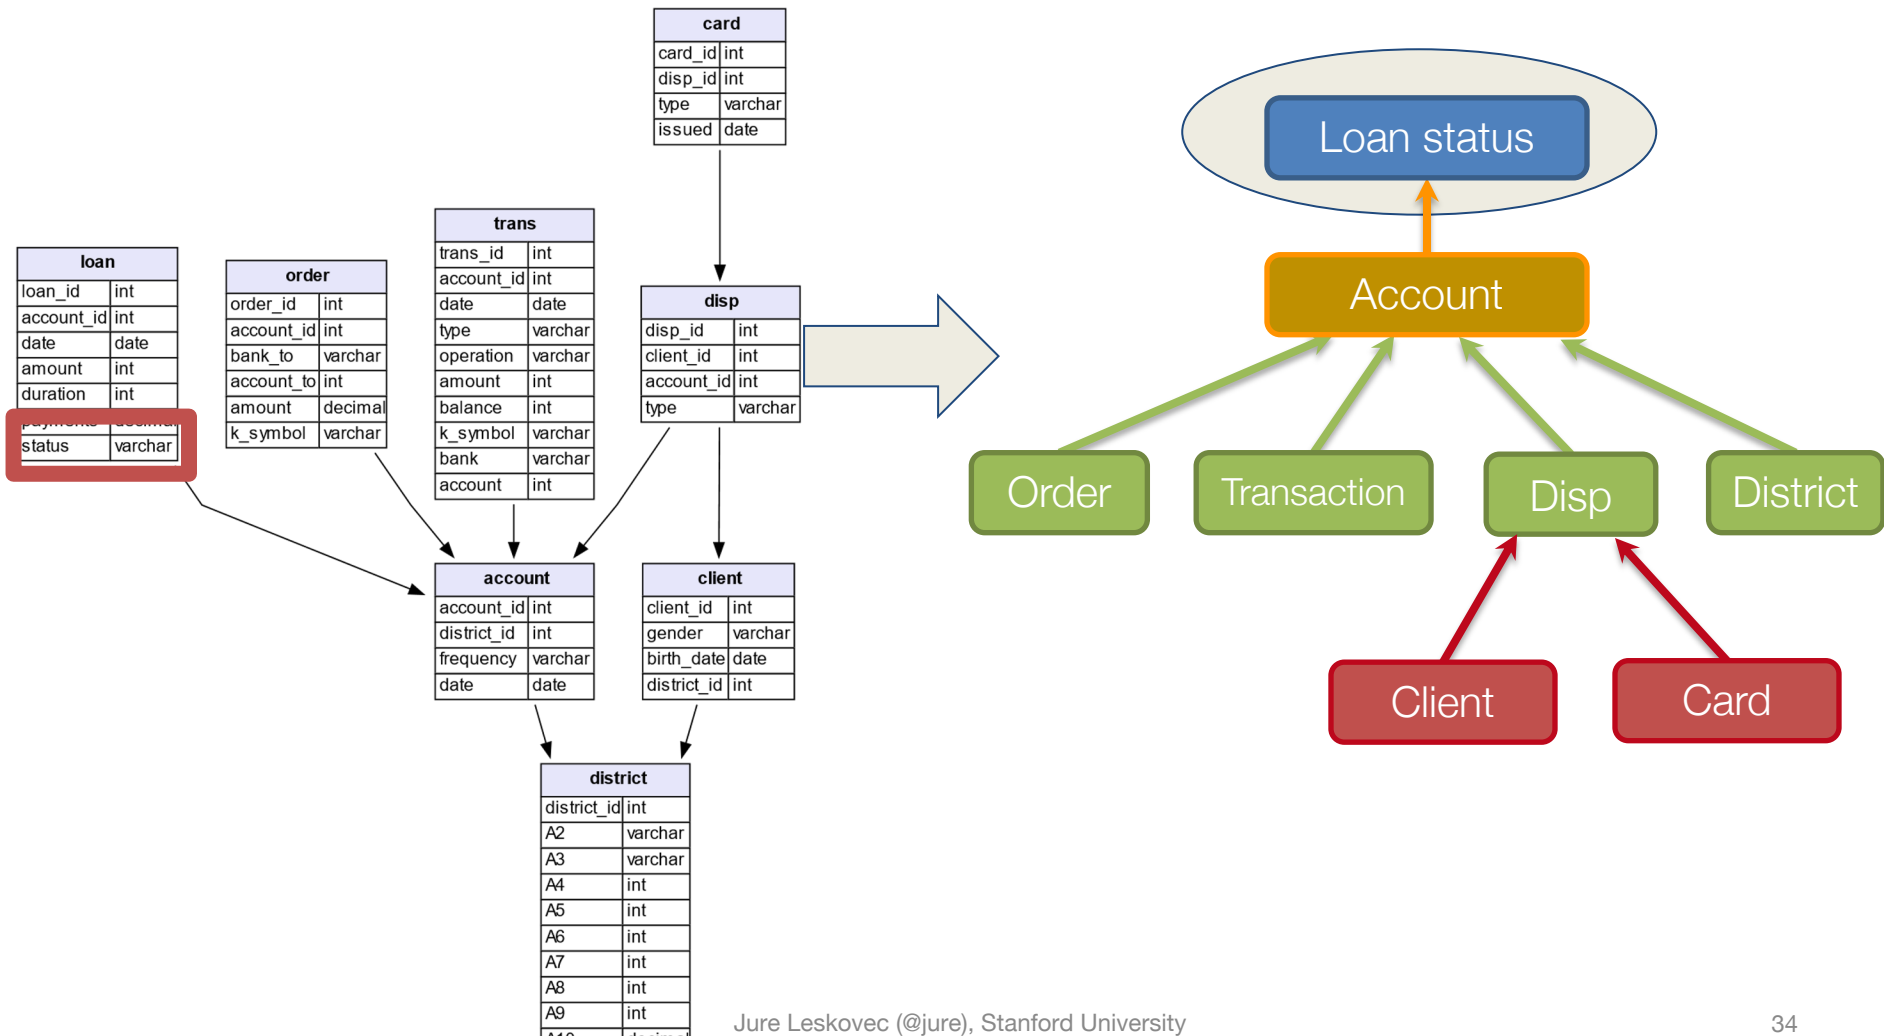


Classical computer vision: hand-crafted features (e.g. SIFT)
+ simple classifier (e.g. SVM)

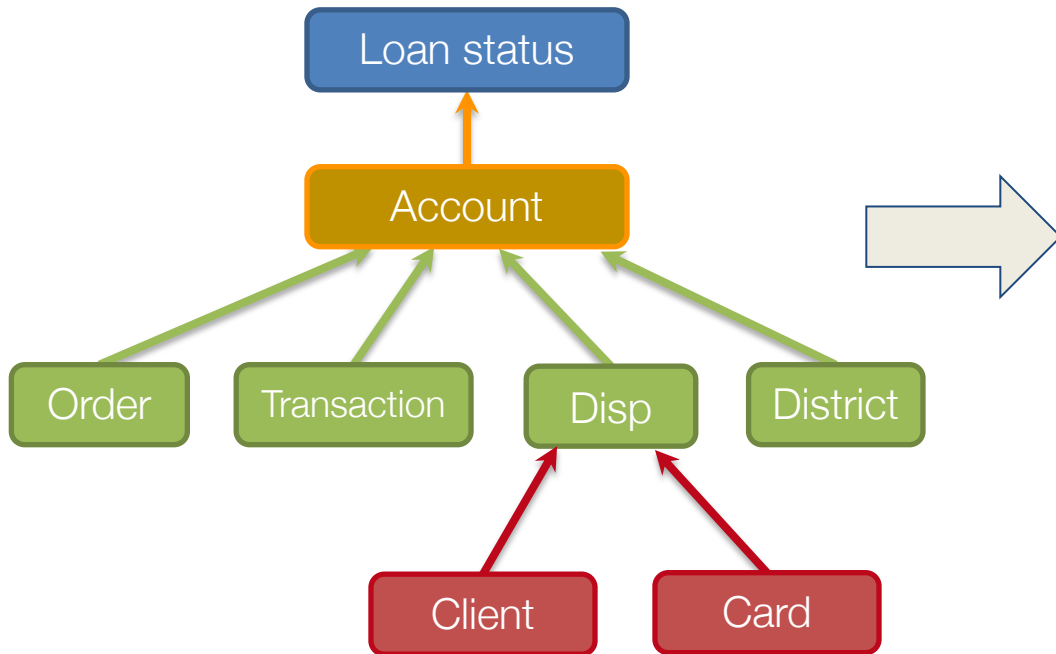


Modern computer vision: data-driven end-to-end systems

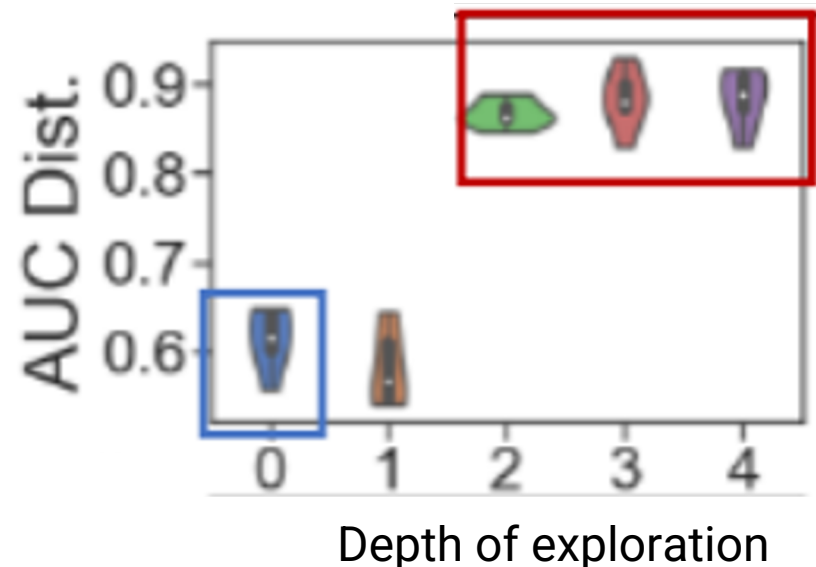
kumo: Example



kumo: State of the Art Results



Method	AUC
Naïve Neural Net baseline	0.65
Stanford PhD (XGBoost)	0.88
Best research paper	0.92
KUMO	0.97



kumo : Benefits

Learned representation (features) is optimal for a given task

State of the art model performance

Drastic simplification of the ML stack

No feature store, no backfills, no aggregations, no feature crosses

Easy productionisation

Faster time to value

Amplifies data scientists to be more productive

Conclusion


GNNs allow for representation on complex relational data

- Fuse node features & relations
 - State-of-the-art accuracy graph machine learning tasks
- Model size independent of graph size; can scale to billions of nodes
 - Largest embedding to date (5B nodes, 20B edges)
- Leads to significant performance gains

Conclusion

- GNNs are a very general type of NNs
- GNNs subsumes CNNs and Transformers as special cases
- GNNs have a huge range of applications

PhD Students




Alexandra Porter Camilo Ruiz Hongyu Ren Hamed Nilforoshan Michi Yasunaga




Jared Davis Serina Chang Weihua Hu Yusuf Roohani Kaidi Cao

Post-Doctoral Fellows




Maria Brbic Tailin Wu

Industrial Visitors



Yanan Wang Jing Huang Yoshitaka Inoue Shigeru Maya

Research Staff



Adrijan Bradaschia Rok Susic

Industry Partnerships



Learn more:

<http://pyg.org>

<http://cs.stanford.edu/~jure>

Contact: jure@cs.stanford.edu