## Graph networks for learning about complex systems





- Peter Battaglia
  - DeepMind

Workshop on Graph Neural Networks Siena, Italy July 22, 2019



### Many complex systems are structured

#### Molecule





Rigid Body System

Sentence



The brown dog jumped.

Mass-Spring System

#### n-body System



Image



### Many complex systems are structured



Rigid Body System

Sentence and Parse Tree





The brown dog jumped.



Image and Fully-Connected Scene Graph





### The standard deep learning toolkit...







... is not well-suited to reasoning over structured representations.

State

### The standard deep learning toolkit...



But graph neural networks, which can learn a form of message-passing on graphs, are.

... is not well-suited to reasoning over structured representations.

State

*n*-body System









*n*-body System



#### **Edge function**

 $\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$ 

Compute "message" from lacksquarenode and edge attributes associated with an edge









#### **Edge function**

 $\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$ 

Compute "message" from ulletnode and edge attributes associated with an edge

### *n*-body System







*n*-body System



#### **Edge function**

 $\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$ 

Compute "message" from lacksquarenode and edge attributes associated with an edge







*n*-body System



#### **Edge function**

#### **Node function**

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

Compute "message" from node and edge attributes associated with an edge

- $\mathbf{v}'_k \leftarrow \phi^v(\bar{\mathbf{e}}'_i, \mathbf{v}_i, \mathbf{u})$

Update node info from previous node state and aggregated "messages"







#### **Edge function**

#### **Node function**

$$\mathbf{e}'_k \leftarrow \phi^e(\mathbf{e}_k, \mathbf{v}_{r_k}, \mathbf{v}_{s_k})$$

Compute "message" from node and edge attributes associated with an edge

- $\mathbf{v}'_k \leftarrow \phi^v (\bar{\mathbf{e}}$

### *n*-body System



$$ar{\mathbf{e}}_i', \mathbf{v}_i, \mathbf{u})$$

Update node info from previous node state and aggregated "messages"

Trained to predict node states at  $t_1$  from states at  $t_0$ 







## Physical systems as graphs

### n-body









### Nodes: bodies Edges: gravitational forces

### Balls

- Nodes: balls
- Edges: rigid collisions between balls, and walls

### String





Nodes: masses

Edges: springs and rigid collisions





## 1000-step rollouts of true (top row) vs predicted (bottom row)

n-body







#### Balls

### String









### Zero shot generalisation to larger systems

n-body







True

#### Balls

### String







## Interaction Network: Predicting potential energy



#### **Global function**

 $\mathbf{u}' \leftarrow \phi^u(\bar{\mathbf{v}}')$ 

Rather than making node-wise lacksquarepredictions, edge and node updates can be used to make global predictions.

#### *n*-body System



#### Trained to predict system's potential energy







### Visual interaction network: Simulate from input images

#### Multi-frame encoder (conv net-based)



Interaction network



Watters et al., 2017, NeurIPS



### Visual interaction network: Simulate from input images

#### Mass-springs



True

Model

#### Bouncing balls



True

Model

Can even predict invisible objects, inferred from how they affect visible ones

Watters et al., 2017, NeurIPS





## Relation Network (RN) for scene understanding



- Classify scenes

- Infer novel scene structures • Learn object factorizations from input states or images Support one-shot learning

No node updates, where instead all per-edge outputs are summed and passed to a global function

Our experiments showed that RNs can learn to:

Raposo et al., 2017, ICLR workshop



### CLEVR Visual Question-Answering (VQA) dataset



There is a sphere with the same size as the metal cube; is it made of the same material as the small red sphere?





**RN** architecture for VQA

### Results: RN applied to CLEVR





### Results: RN applied to bAbl Q-A dataset

Sandra moved to the garden. John moved to the office. Sandra journeyed to the bathroom. Mary moved to the hallway. Daniel travelled to the office. John went back to the garden. John moved to the bedroom.

Where is Sandra? **bathroom** 

Model

- Entities are LSTM-encoded sentences
- Input to an RN

Results

Solves 18/20 bAbl tasks



### Results: RN can infer relations in dot motion

#### Trained on mass-spring systems





### Results: RN can infer relations in dot motion

#### Trained on mass-spring systems





## Graph Networks (GNs)

#### Why do we need another Graph Neural Network variant?

We designed GNs to be both expressive, and easy to implement ullet

## Graph Networks (GNs)

#### Why do we need another Graph Neural Network variant?

- We designed GNs to be both expressive, and easy to implement
- A GN block is a "graph-to-graph" function approximator
  - The output graph's structure (number of nodes and edge connectivity) matches the input graph's The output graph-, node-, and edge-level attributes will be functions of the input graph's



## What kind of graphs do GNs operate on?

#### "Graph": directed, attributed multi-graph with a global attribute

- "Directed": one-way edges, from a "sender" node to a "receiver" node
- "Multi-graph": there can be more than one edge between vertices, including self-edges
- "Attribute": properties that can be encoded as a vector, set, or even another graph
- "Attributed": edges and vertices have attributes associated with them
- "Global attribute": a graph-level attribute



### Graph: $G = (\mathbf{u}, V, E)$

- Global attribute: u
- Node attributes:  $V = {\mathbf{v}_i}_{i=1:N^v}$
- Edge attributes, sender node indices, receiver node indices:  $E = \{(\mathbf{e}_k, s_k, r_k)\}_{k=1:N^e}$  $s_k, r_k \in \{1, \dots, N^v\}$

### How does a GN block process a graph?



For each edge,  $\mathbf{e}_k, \mathbf{v}_{s_k}, \mathbf{v}_{r_k}, \mathbf{u}$ , are passed to an "edge-wise function":

 $\mathbf{e}'_{k} \leftarrow \phi^{e} \left( \mathbf{e}_{k}, \mathbf{v}_{r_{k}}, \mathbf{v}_{s_{k}}, \mathbf{u} \right)$ 

### **Node block**

For each node,  $\bar{\mathbf{e}}_i', \mathbf{v}_i, \mathbf{u}$ , are passed to a "node-wise function":  $\mathbf{v}_i' \leftarrow \phi^v \left( \mathbf{\bar{e}}_i', \mathbf{v}_i, \mathbf{u} \right)$ 

#### **Global block**

Across the graph,  $\, \bar{e}', \bar{v}', u \,$  , are passed to a "global function":

 $\mathbf{u}' \leftarrow \phi^u \left( \mathbf{\bar{e}}', \mathbf{\bar{v}}', \mathbf{u} \right)$ 











## Composing GN blocks

The GN's graph-to-graph interface promotes stacking GN blocks, passing one GN's output to another GN as input











## Non-local Neural Nets / Self-Attention in the GN formalism

attention weight, which is used to rescale the pooled edge outputs ( $\rho^{e \rightarrow v}$  below)



$$egin{aligned} & \mathcal{O}^{\circ}\left(\mathbf{e}_{k},\mathbf{v}_{r_{k}},\mathbf{v}_{s_{k}},\mathbf{u}
ight)\coloneqq J^{\circ}\left(\mathbf{v}_{r_{k}},\mathbf{v}_{s_{k}}
ight)\ & \phi^{v}\left(ar{\mathbf{e}}_{i}',\mathbf{v}_{i},\mathbf{u}
ight)\coloneqq f^{v}(ar{\mathbf{e}}_{i}')\ & & 
ho^{e
ightarrow v}\left(E_{i}'
ight)\coloneqq rac{1}{\sum_{\{k:\,r_{k}=i\}}a_{k}'}\sum_{\{k:\,r_{k}=i\}}a_{k}'\mathbf{b}_{k}' \end{aligned}$$

The edge function has two components — a scalar-valued pairwise interaction among nodes ( $\alpha^{e}$  below), and a vector-valued function of the sender node ( $\beta^e$  below). The scalar component is the unnormalized



 $= (\alpha^{e} (\mathbf{v}_{r_{k}}, \mathbf{v}_{s_{k}}), \beta^{e} (\mathbf{v}_{s_{k}})) = (a'_{k}, \mathbf{b}'_{k}) = \mathbf{e}'_{k}$ 



## Systems: "DeepMind Control Suite" (Mujoco) & real JACO



### DeepMind Control Suite (Tassa et al., 2018)

### Kinematic tree of the actuated system as a graph

Controllable physical system as a graph:

- Bodies  $\rightarrow$  Nodes
- Joints  $\rightarrow$  Edges
- Global properties





### Forward model: supervised, 1-step training w/ random control inputs

Input graph (t)



#### Chained 100-step predictions

### **Prediction Fixed Swimmer6**

### Expected

Next graph (t+1)







## GN forward model: Multiple systems & zero-shot generalization

Single model trained:

• Pendulum, Cartpole, Acrobot, Swimmer6 & Cheetah

#### **Zero-shot generalization**: Swimmer

- # training links: {**3**, **4**, **5**, **6**, -, **8**, **9**, -, -, ...}
- # testing links: {-, -, -, -, 7, -, -, **10-14**}







### GN forward model: Real JACO data

**Recurrent GN** 

### **Prediction Fixed Real JACO**









### System identification: GN-based inference, under diagnostic control inputs



## Prediction System ID Cartpole ID phase



### Unobserved system parameters (e.g. mass, length) are implicitly inferred







### Control: Model-based planning

Trajectory optimization: the GN-based forward model is differentiable, so we can backpropagate through it, and find a sequence of actions that maximize reward





### Control: Multiple systems via a single model





### Control: Zero-shot control



### Control: Multiple reward functions









## Relational forward models for multi-agent RL





Tacchetti et al., 2019, ICLR

Stag hunt

#### Step 0 with last actions



### Graph representation

#### Forward prediction performance





### Interpretable learned representations



Tacchetti et al., 2019, ICLR

### Agents learn faster with model-augmented observations

- Train a set of agents to perform a game.
- 2. Train an RFM to predict the agents' future actions.
- 3. Train a new agent, whose observations are augmented with the RFM's message magnitudes.

The new agent (blue curve) trains faster in all environments.





Tacchetti et al., 2019, ICLR

### Humans are a "construction species"









### Humans are a "construction species"









#### Our construction behaviors are:

- Combinatorial
- + Structured
- + Use rich physical knowledge

Graph networks





Pick up **blocks** and place them in the scene (and optionally make them sticky)





+1 per target -0.5 per sticky block





+1 per target -0.5 per sticky block

+1 per target Free sticky blocks





+1 per target -0.5 per sticky block

+1 per target Free sticky blocks

Length covered -2 per sticky block





+1 per target -0.5 per sticky block

+1 per target Free sticky blocks

Length covered -2 per sticky block

Length covered -0.5 per sticky block





### Graph net-based agent: model-free



Can be thought of as "graph building" agent



### Graph net-based agent: model-based

Can be thought of as "graph building" agent





### Absolute vs relative actions



# Actions



### Results: "Silhouette" task

**Absolute Actions** 





Reward: +1 per target, -0.5 per sticky block

**Object-Centric Actions** 



### Results: "Connecting" task

**Absolute Actions** 



**Object-Centric Actions** 



Reward: +1 per target, free sticky blocks



### Results: "Covering" task

**Absolute Actions** 





Reward: proportional to length covered, -2 per sticky block

**Object-Centric Actions** 



### Results: "Covering hard" task

**Absolute Actions** 





Reward: proportional to length covered, -0.5 per sticky block

**Object-Centric Actions** 



### Results: Absolute vs relative actions







### Results: Planning agent (using MCTS)









## **Build Graph Nets in Tensorflow**

### github.com/deepmind/graph nets

#### **IPython Notebook demos** (All use same architecture)

Shortest path:



True

Sort: item-to-item connections





Time 0

Sorting:

Predicting physics:

Shortest path: predictions at each message-passing step





#### Predicted

Physics: rollout of mass-spring system pinned at ends

## Build Graph Nets in Tensorflow

Example code for getting started:

```
# Provide your own functions to generate graph-structured data.
input_graphs = get_graphs()
# Create the graph network.
graph_net_module = gn.modules.GraphNetwork(
    edge_model_fn=lambda: snt.nets.MLP([32, 32]),
    node_model_fn=lambda: snt.nets.MLP([32, 32]),
    global_model_fn=lambda: snt.nets.MLP([32, 32]))
```

# Pass the input graphs to the graph network, and return the output graphs. output\_graphs = graph\_net\_module(input\_graphs)

For GNN libraries in PyTorch, check out:

- $\bullet$
- Deep Graph Library: <u>github.com/dmlc/dgl</u>  $\bullet$

github.com/deepmind/graph nets

pytorch\_geometric: github.com/rusty1s/pytorch\_geometric (for a GN analog, see MetaLayer)



Graph neural networks: a first-class member of the deep learning toolkit.

- simulation, but also more sophisticated patterns of reasoning.

Graph neural networks: a first-class member of the deep learning toolkit.

Learned message-passing in graphs can support classification/regression,

- simulation, but also more sophisticated patterns of reasoning.
- are effective in domains for which less structured methods struggle.

• Graph neural networks: a first-class member of the deep learning toolkit.

Learned message-passing in graphs can support classification/regression,

• RL agents with graph-structured inputs, outputs, and/or latent representations

- simulation, but also more sophisticated patterns of reasoning.
- are effective in domains for which less structured methods struggle.

• Graph neural networks: a first-class member of the deep learning toolkit.

Learned message-passing in graphs can support classification/regression,

RL agents with graph-structured inputs, outputs, and/or latent representations

• Build Graph Nets in Tensorflow: github.com/deepmind/graph nets

### Key collaborators

Razvan Pascanu Jess Hamrick Alvaro Sanchez-Gonzalez Victor Bapst Kim Stachenfeld Carl Doersch Nick Watters Andrea Tacchetti Theophane Weber Daniel Zoran Mateusz Malinowski David Raposo Adam Santoro Nicholas Heess Koray Kavukcuoglu

### References

Battaglia et al., 2016, NIPS Watters et al., 2017, NIPS Raposo et al., 2017, ICLR workshop Santoro et al., 2017, NIPS Battaglia et al. 2018 arXiv Sanchez-Gonzalez et al., 2018, ICML Tacchetti et al., 2019, ICLR Bapst et al., 2019, ICML