# Structured agents for physical construction

*(Wednesday posters, Pacific Ballroom #36)*


**Victor Bapst***

**Alvaro Sanchez-Gonzalez***

Carl Doersch

Kim Stachenfeld

Pushmeet Kohli

Peter Battaglia

Jessica Hamrick

DeepMind

**equal contribution*

# Humans are a “Construction Species”

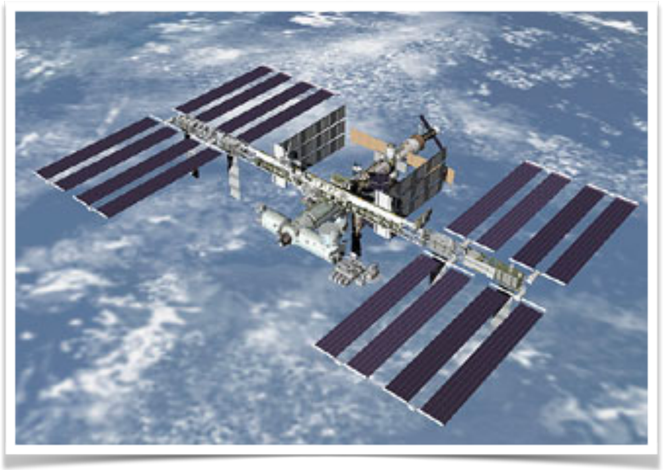

# Humans are a “Construction Species”

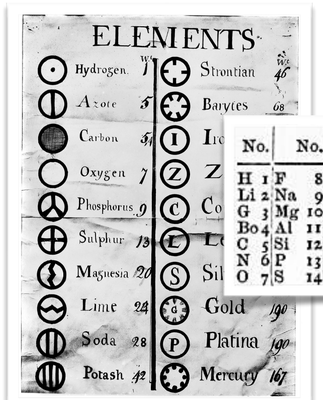


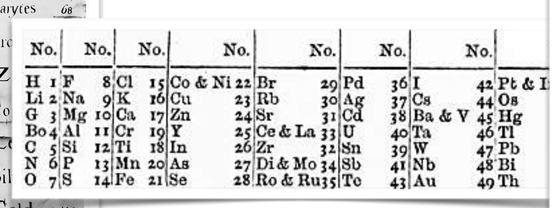

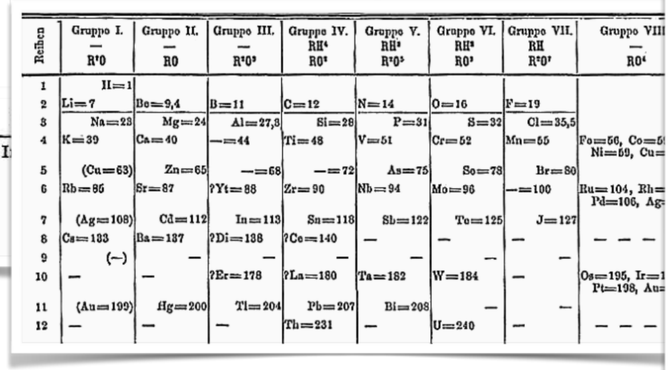

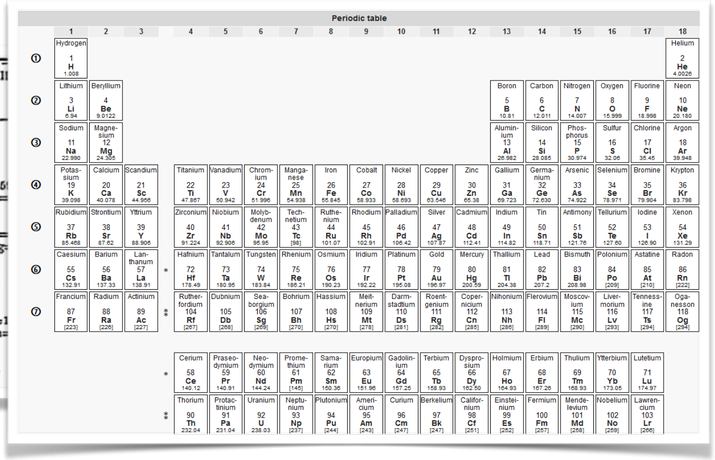

# Humans are a “Construction Species”

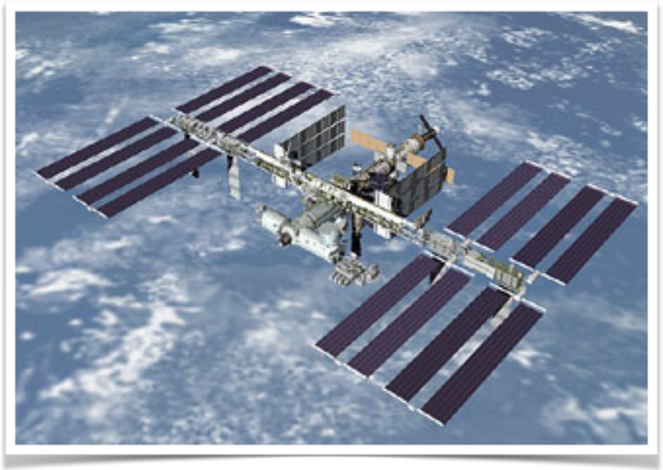

Physical construction: the ability to compose **objects**, subject to physical **dynamics**, in order to serve a **function**.

# Contributions

# Contributions

1. A suite of challenging physical ***construction tasks***

# Contributions

1. A suite of challenging physical ***construction tasks***

2. A new type of ***structured agent*** that uses:

# Contributions

1. A suite of challenging physical ***construction tasks***

2. A new type of ***structured agent*** that uses:

structured representations
object-centric relative actions
combination of model-free and model-based

# Construction Tasks

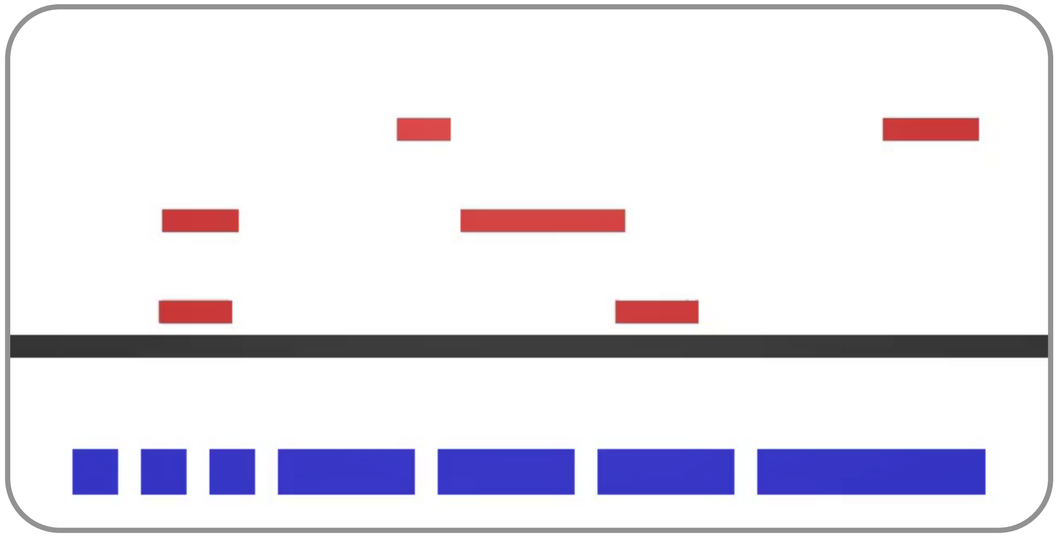

# Construction Tasks

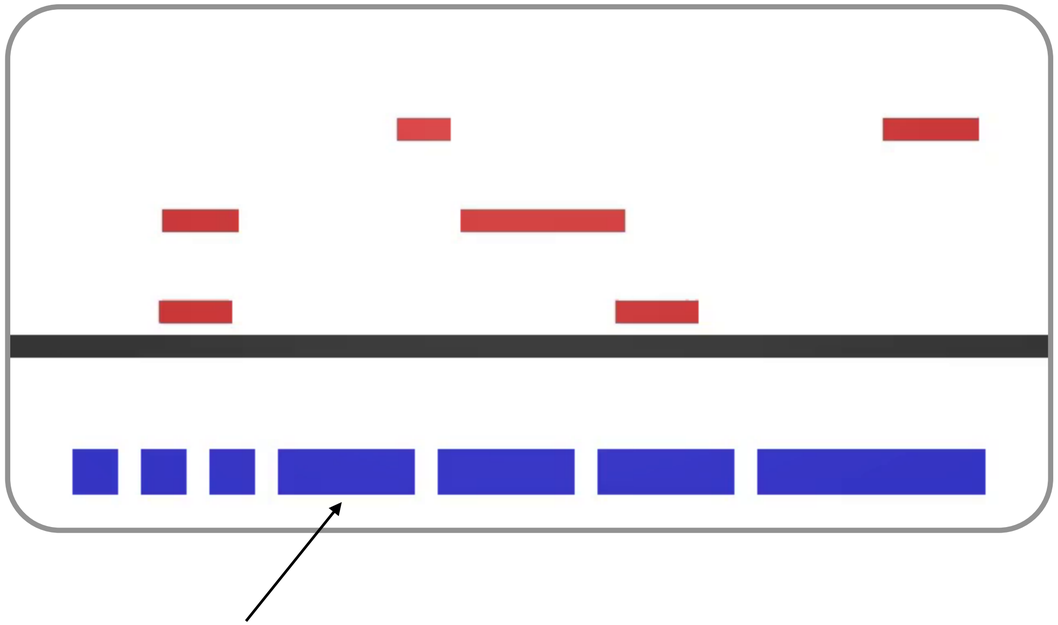

Pick up **blocks** and place them in the scene
(and optionally make them **sticky**)

# Construction Tasks

Avoid touching **obstacles**

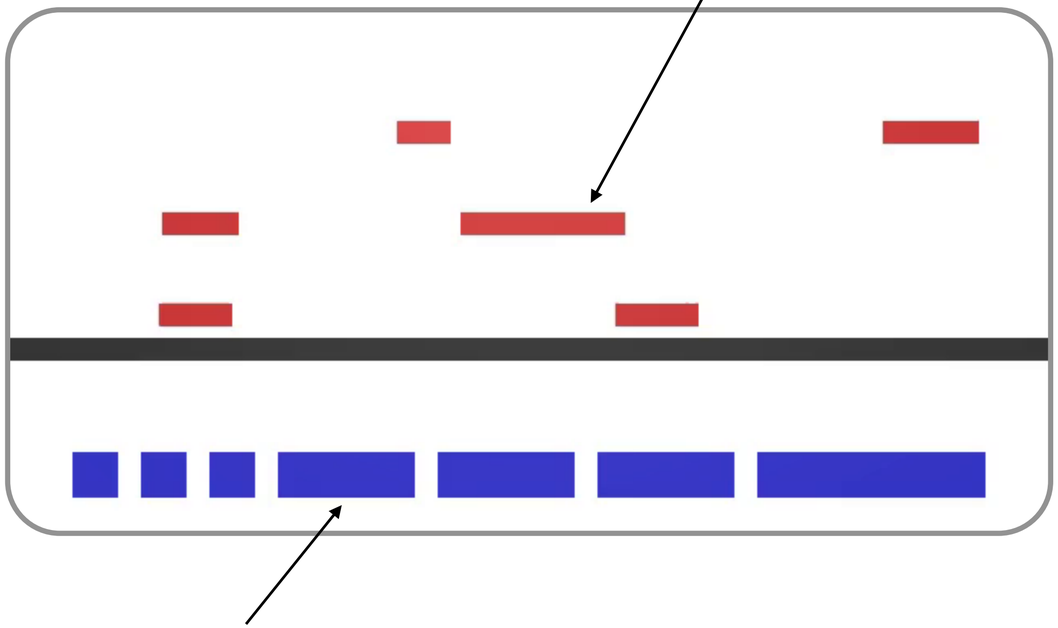

Pick up **blocks** and place them in the scene
(and optionally make them **sticky**)

# Construction Tasks

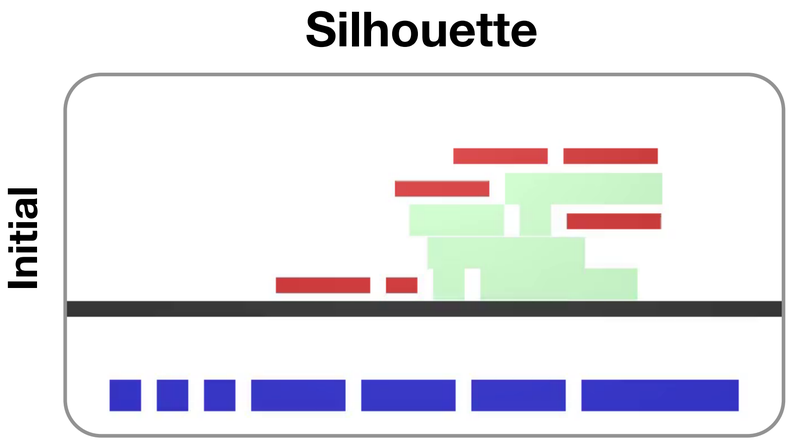

# Construction Tasks

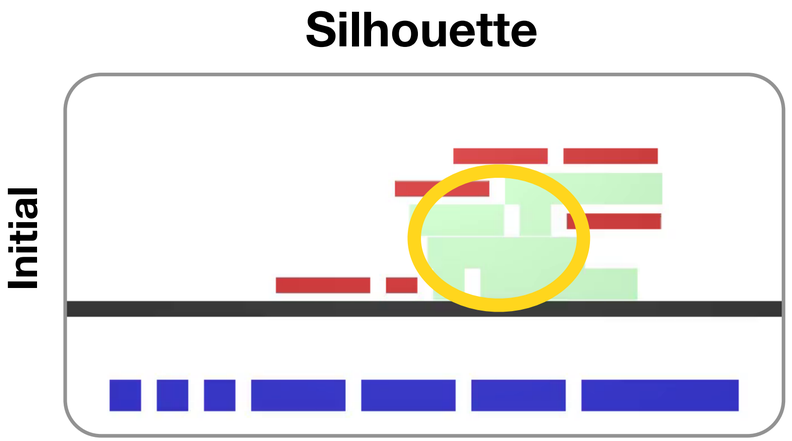

# Construction Tasks

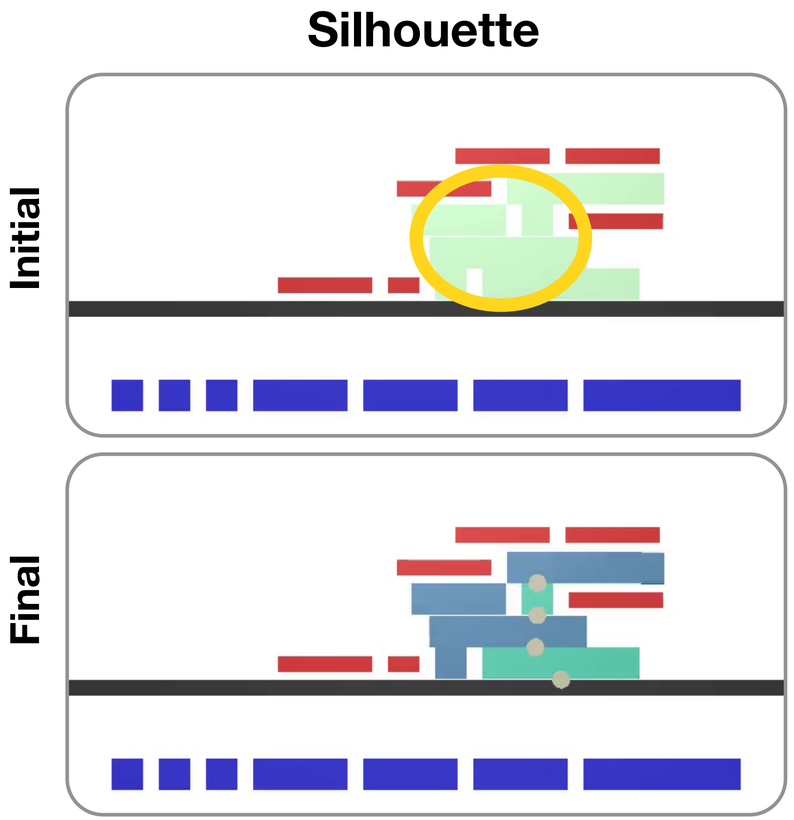

# Construction Tasks

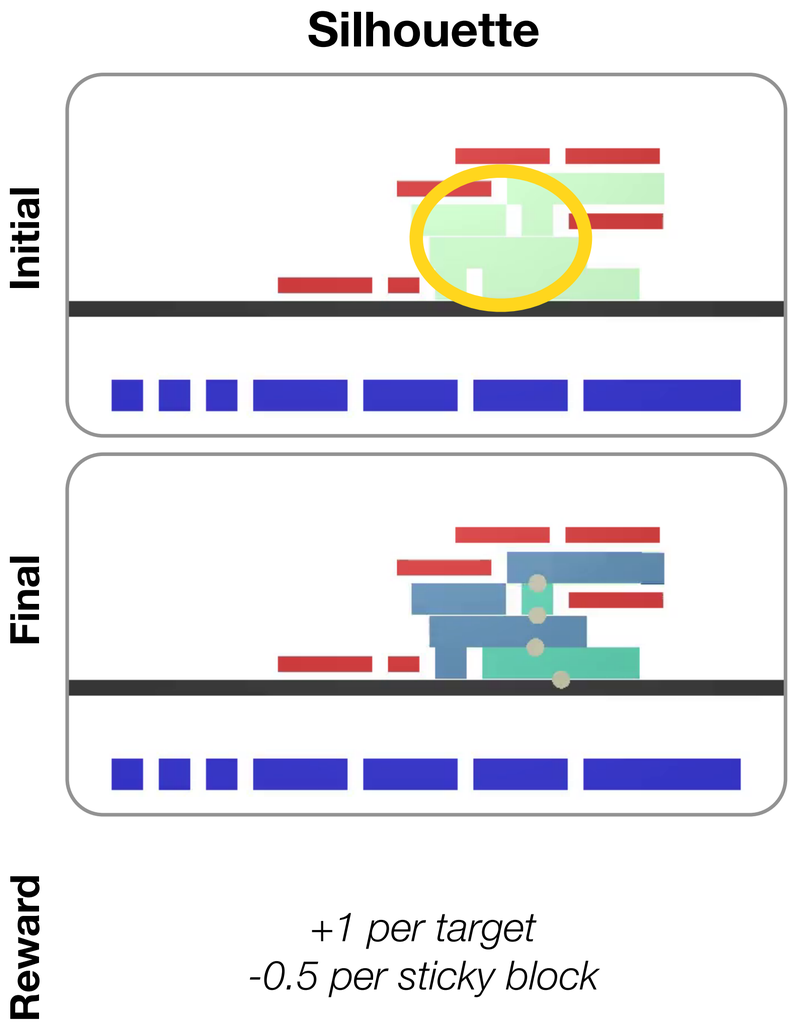

# Construction Tasks

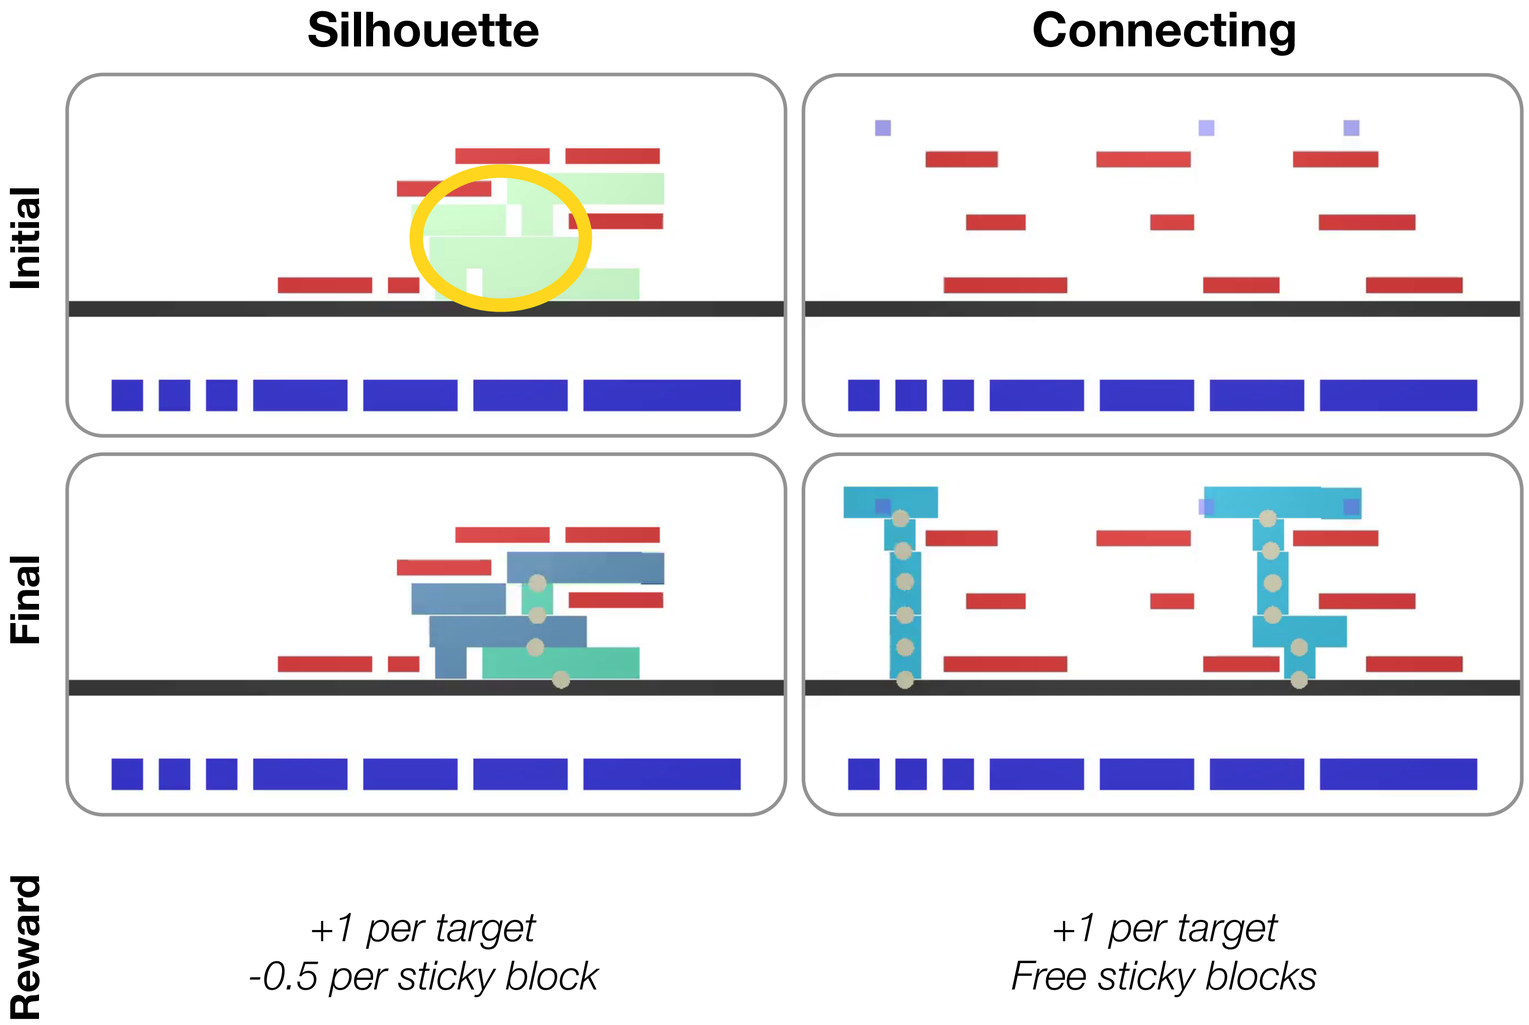

# Construction Tasks

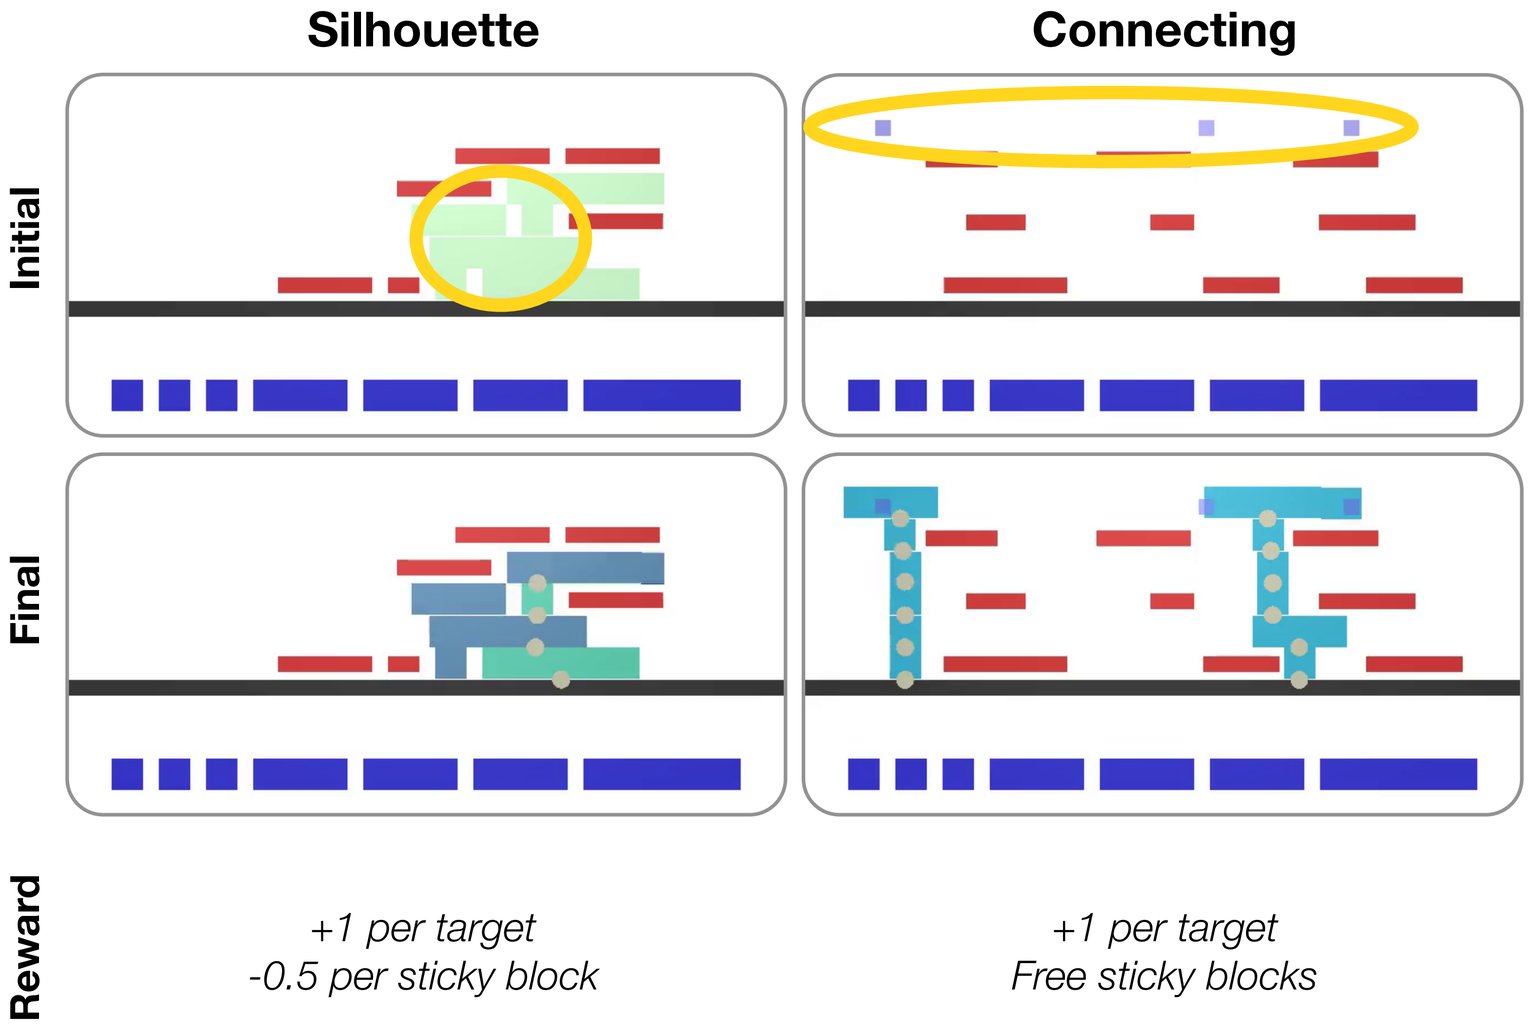

# Construction Tasks

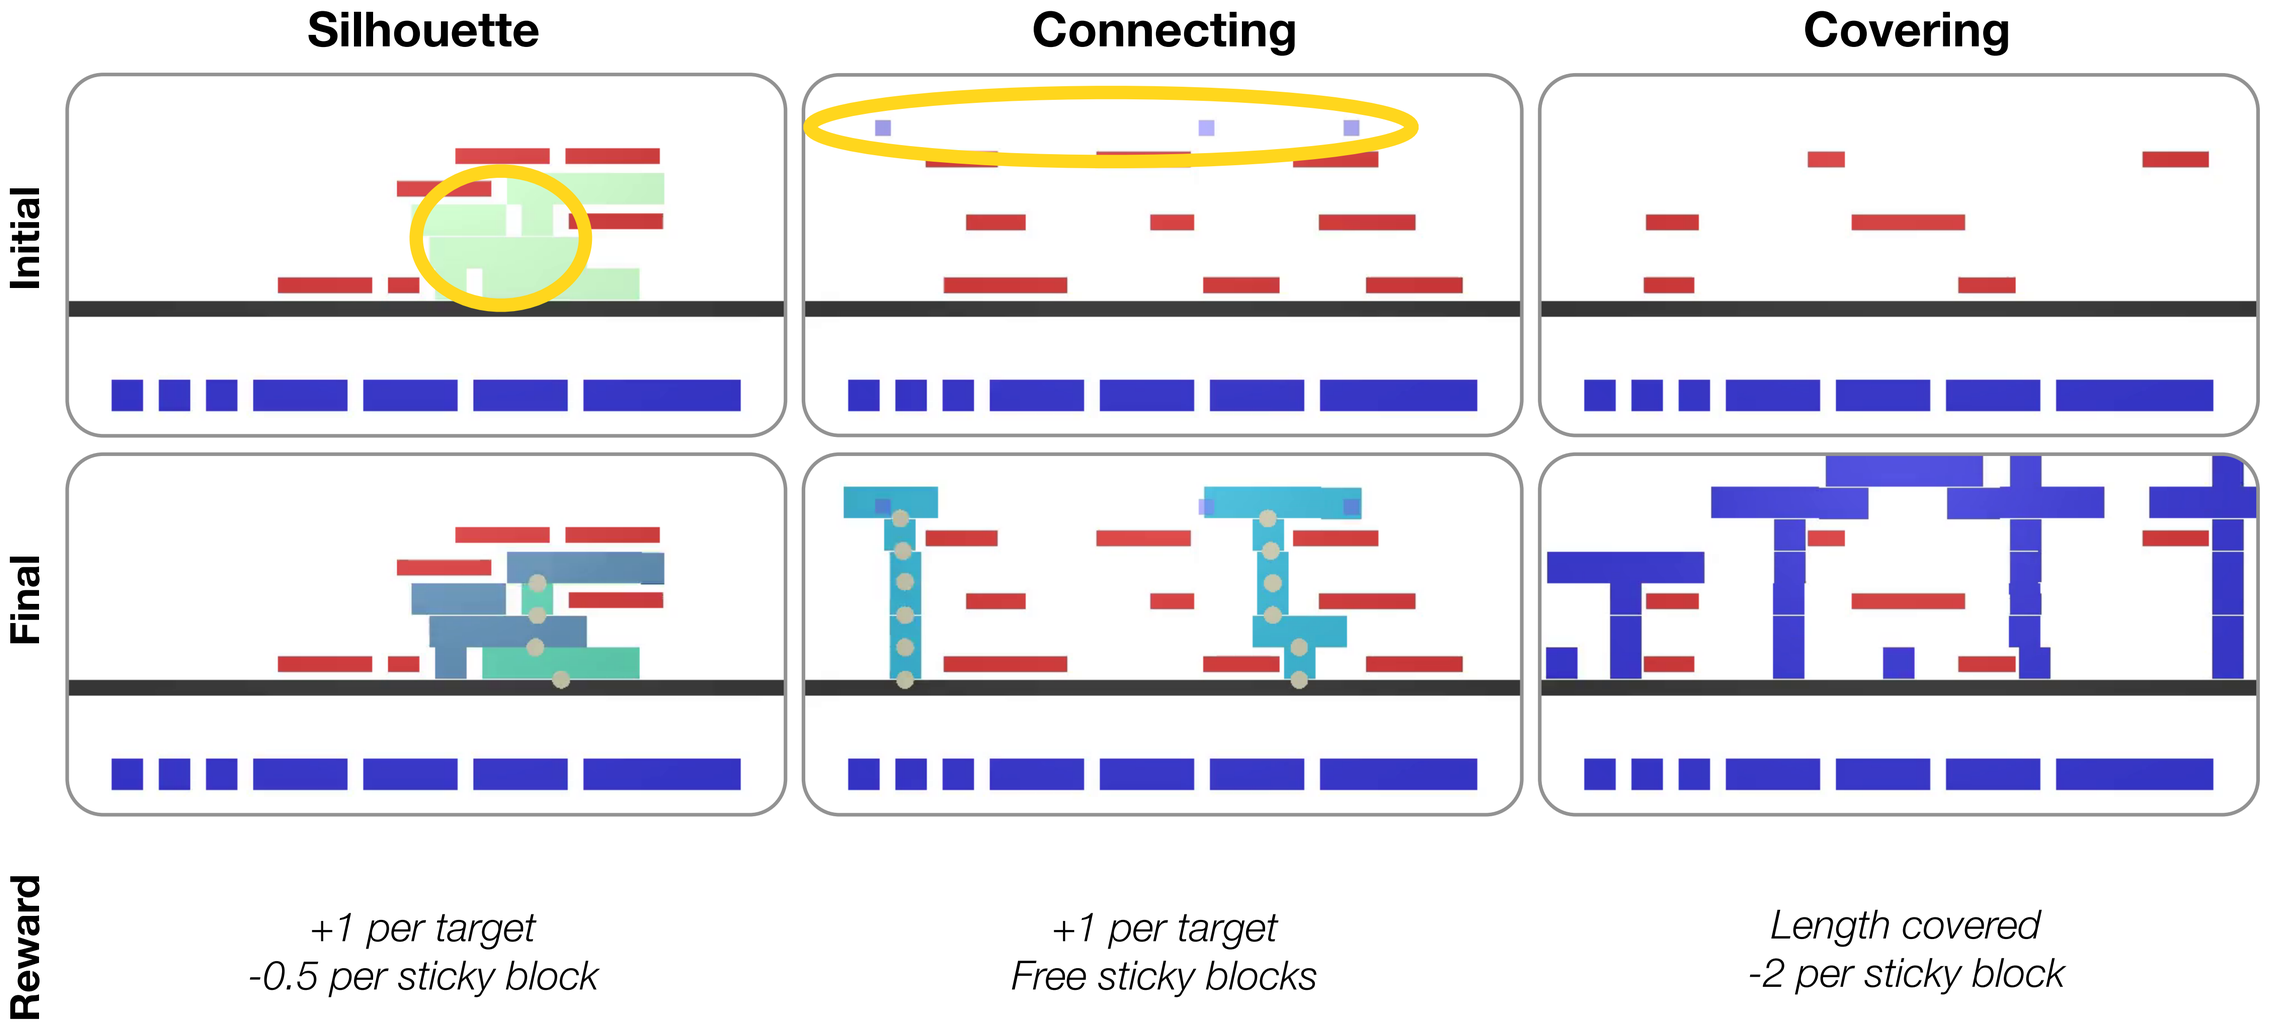

# Construction Tasks

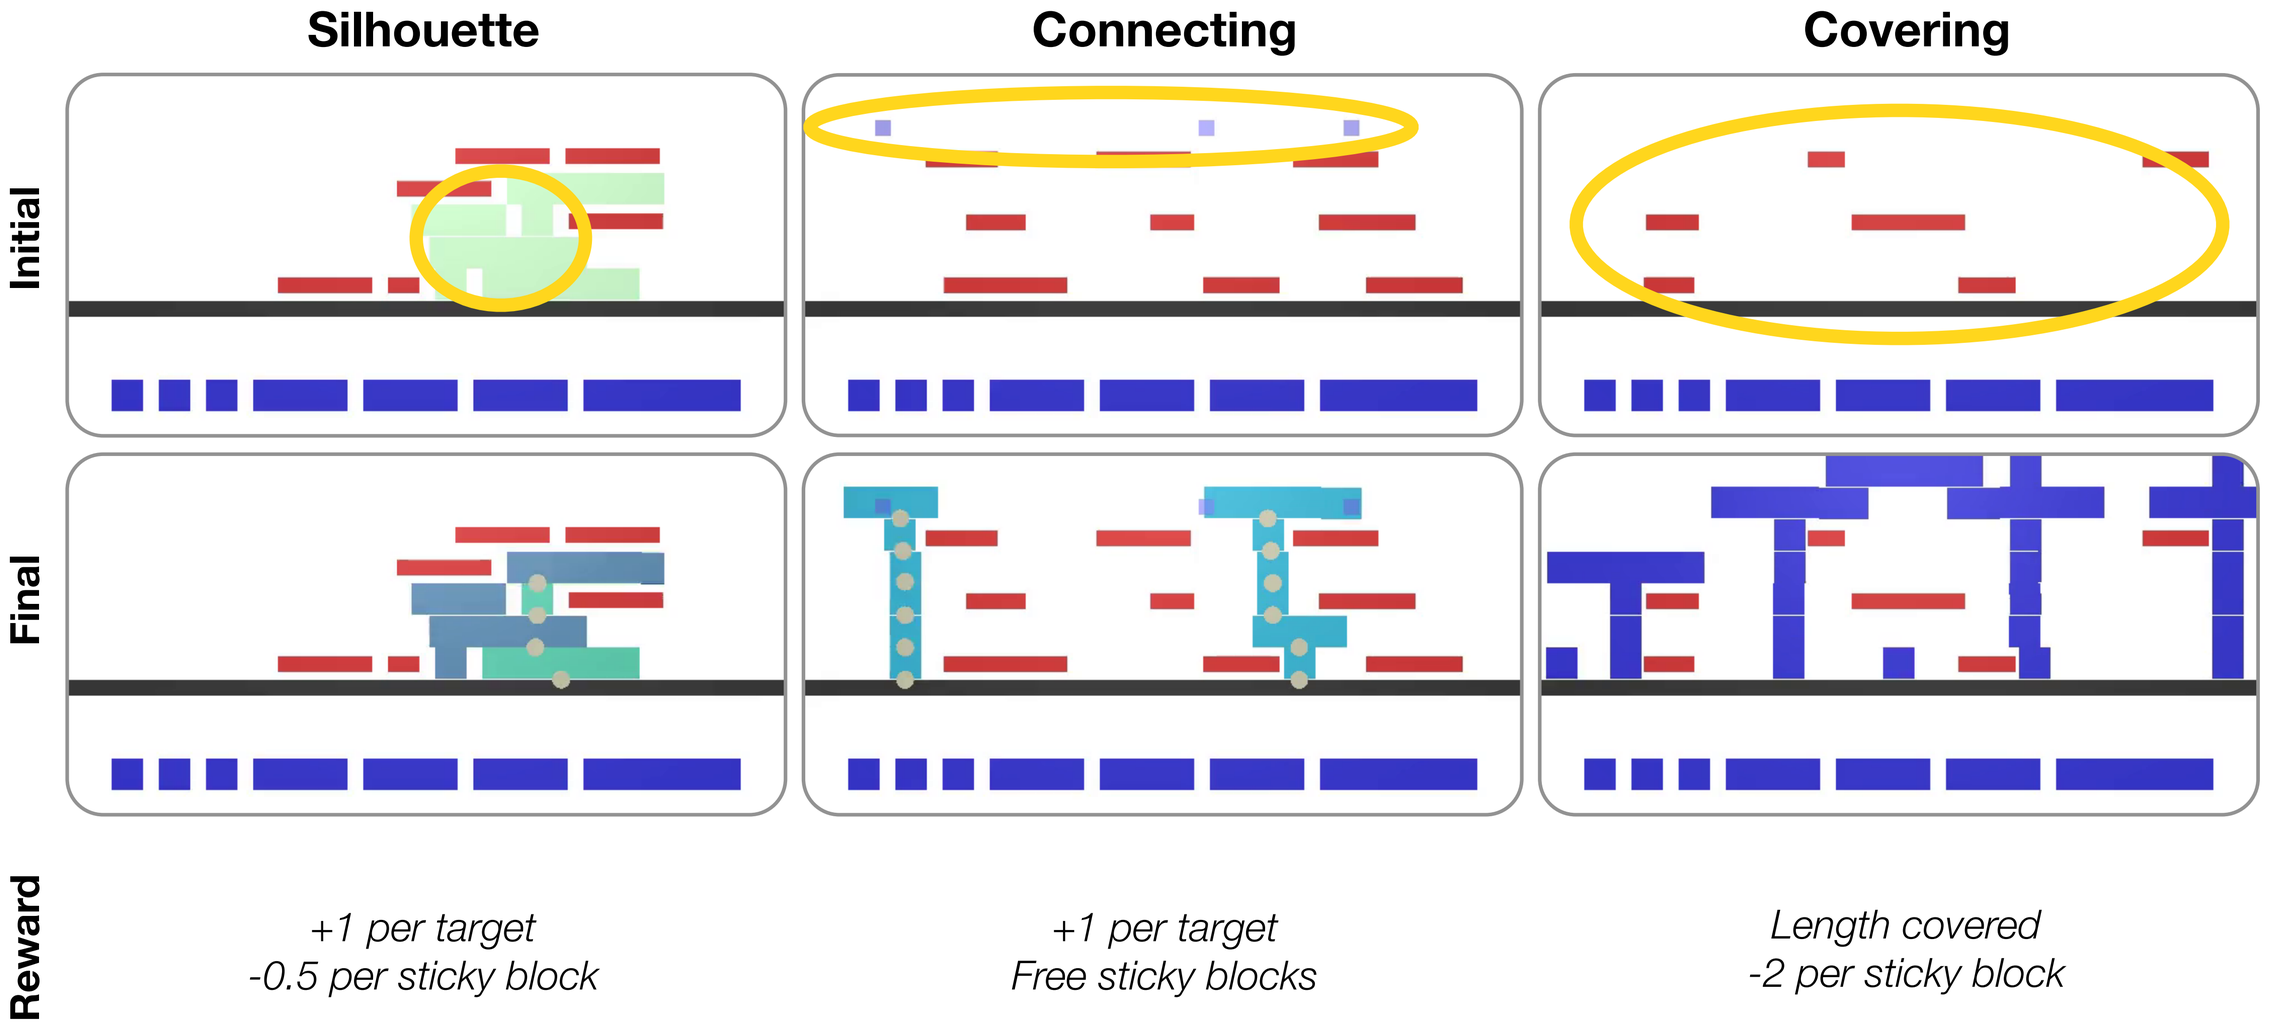

# Action Format: Absolute vs. Relative

# Action Format: Absolute vs. Relative

*“Place block D at position* (7.2, 8.5)”

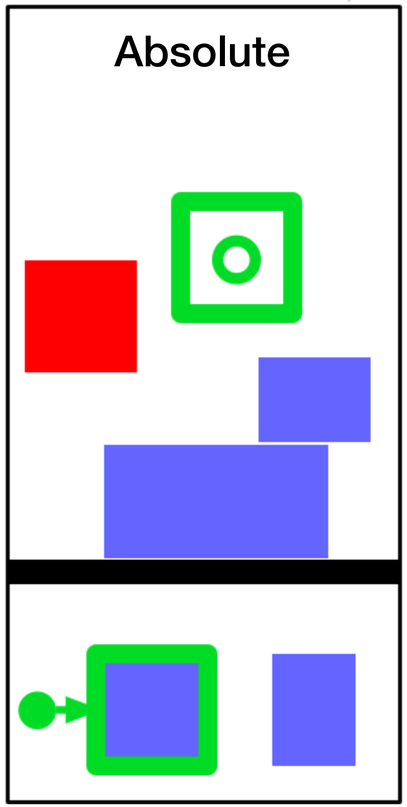

# Action Format: Absolute vs. Relative

*"Place block D at position* (7.2, 8.5)*"*

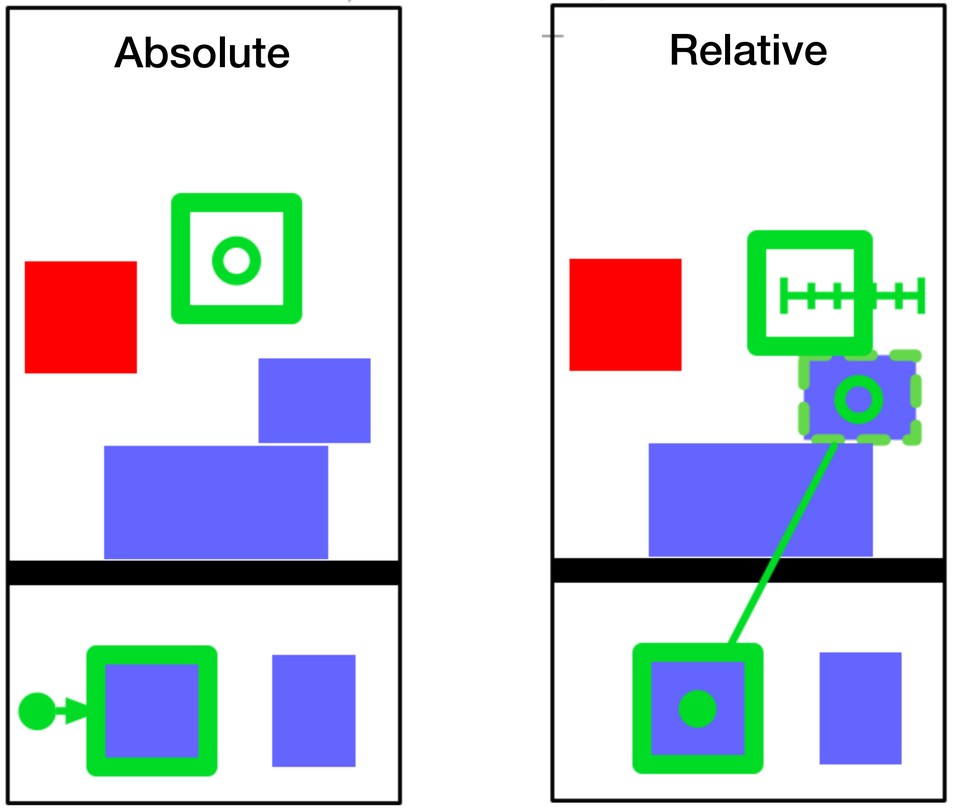


*"Place block D on the top left of block B"*

# Action Format: Absolute vs. Relative

*"Place block D at position* (7.2, 8.5)*"*

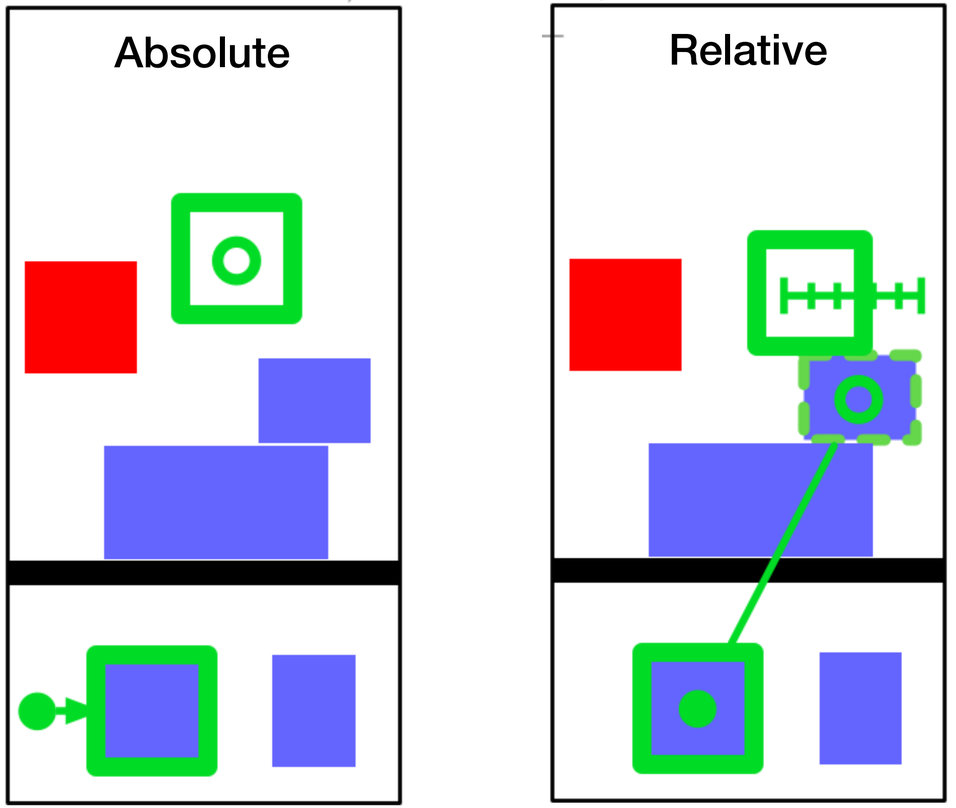


*"Place block D on the top left of block B"*

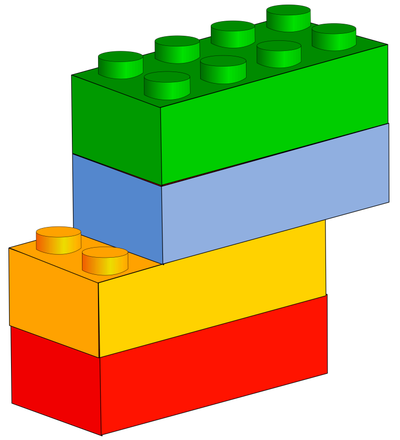

# Action Format: Absolute vs. Relative

*"Place block D at position (7.2, 8.5)"*

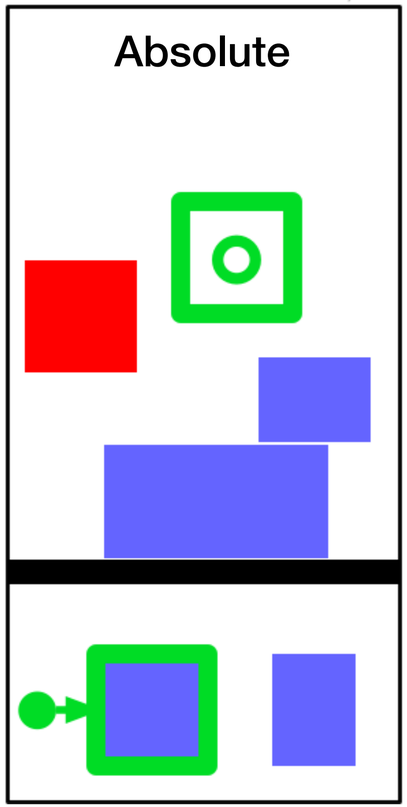


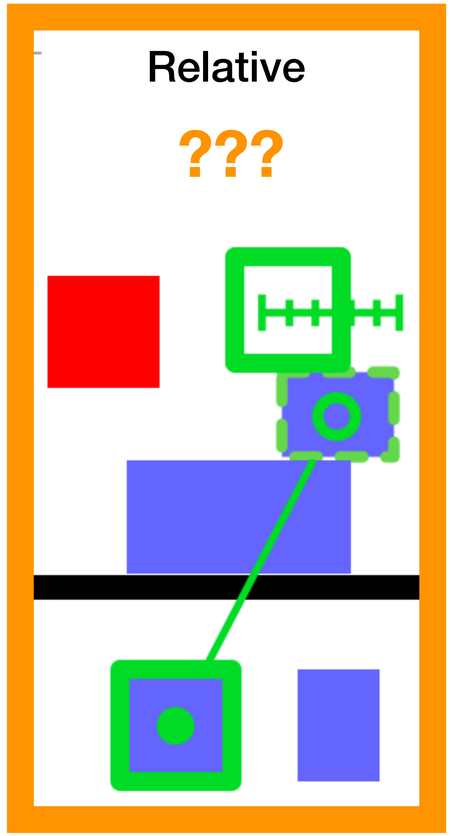


*"Place block D on the top left of block B"*

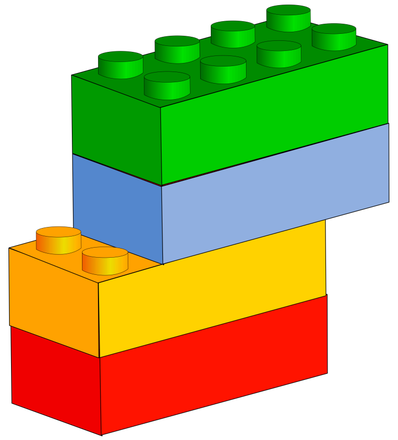

# Graph Network Agent (GN-DQN)

# Graph Network Agent (GN-DQN)

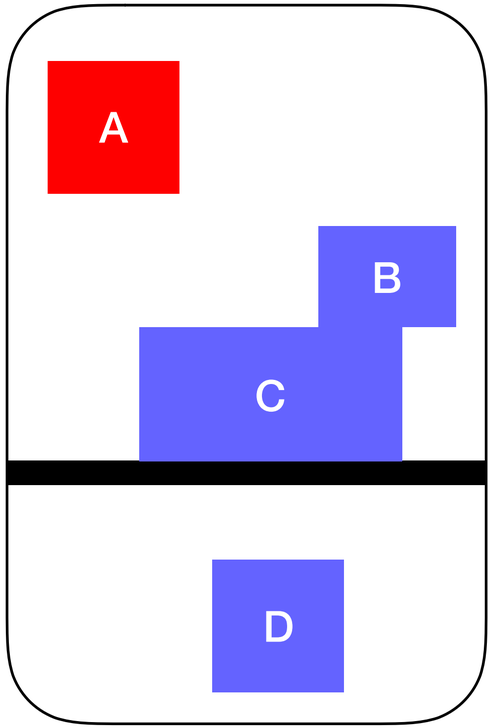


Given
observation

# Graph Network Agent (GN-DQN)

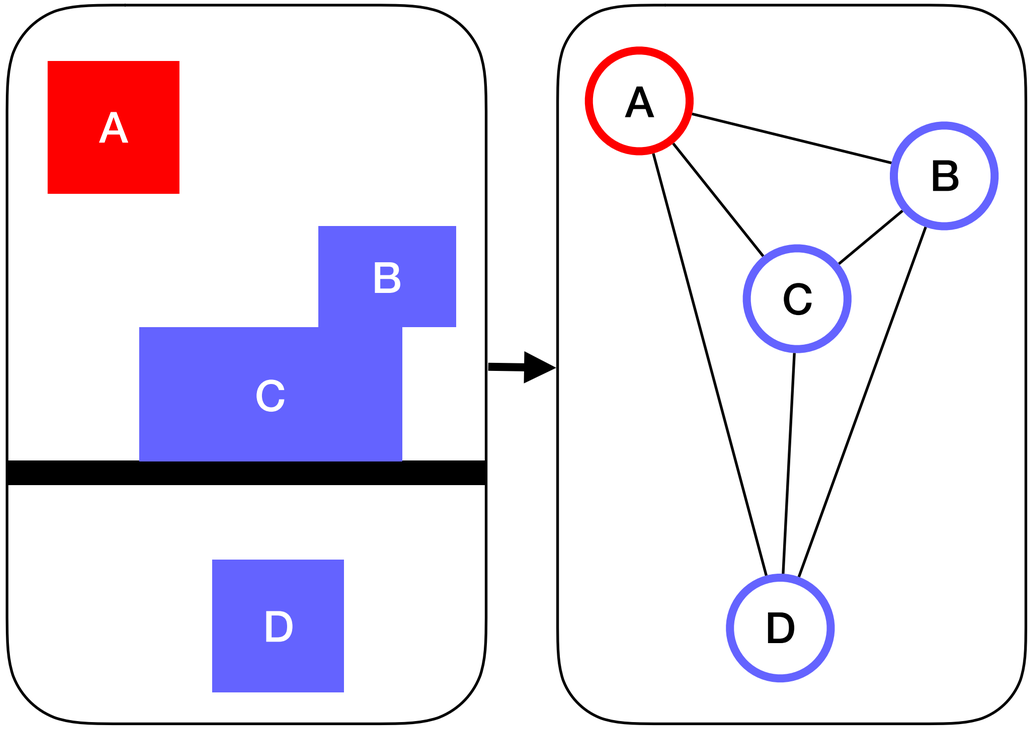


Given observation

Structured representation

# Background: Graph Networks

Graph neural networks: Gori et al. (2005), Scarselli et al. (2005, 2009), Li et al. (2015), Kipf & Welling (2016), Gilmer et al. (2017), many more!

# Background: Graph Networks

Graph neural networks: Gori et al. (2005), Scarselli et al. (2005, 2009), Li et al. (2015), Kipf & Welling (2016), Gilmer et al. (2017), many more!

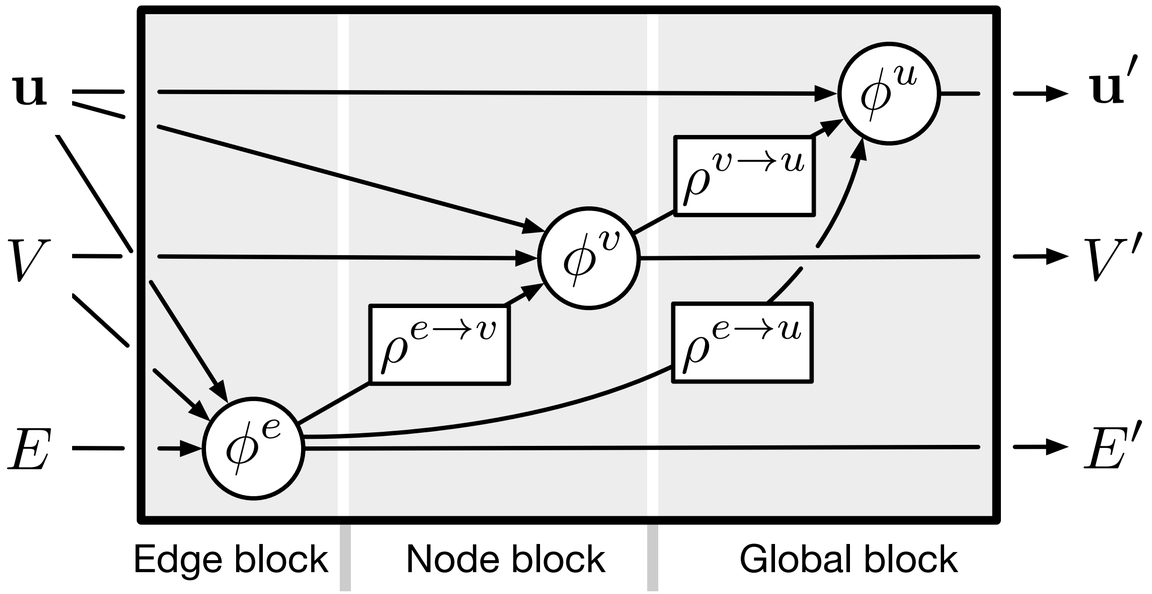


**Battaglia, Hamrick, Bapst, Sanchez-Gonzalez, Zambaldi, et al. (*arXiv* 2018)**

# Background: Graph Networks

Graph neural networks: Gori et al. (2005), Scarselli et al. (2005, 2009), Li et al. (2015), Kipf & Welling (2016), Gilmer et al. (2017), many more!

1. Takes ***graphs*** as input, return graphs as output

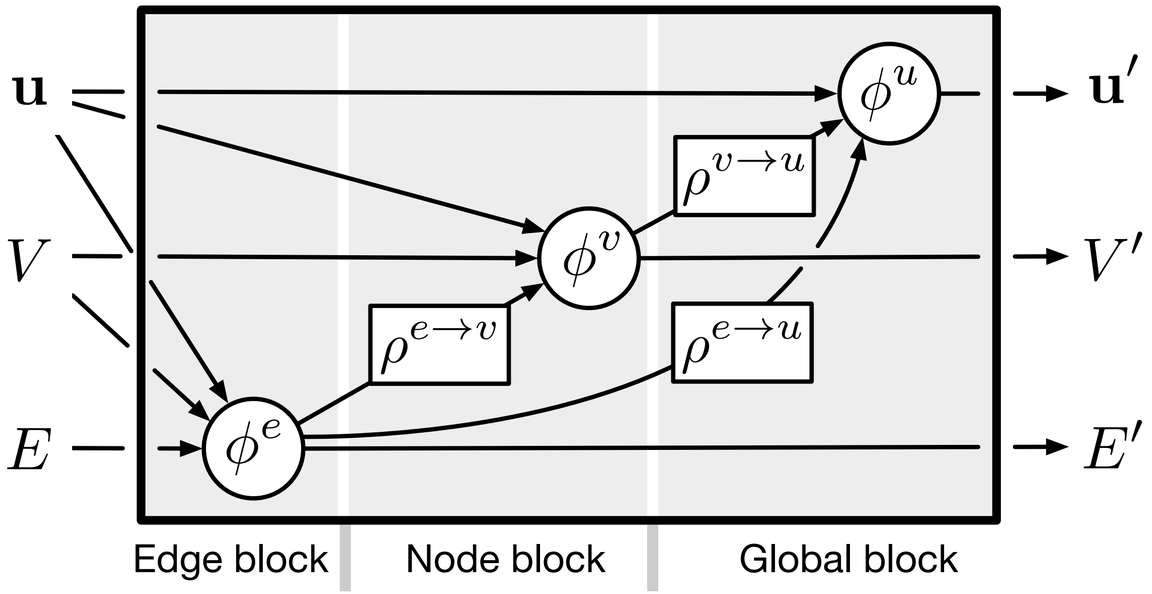


**Battaglia, Hamrick, Bapst, Sanchez-Gonzalez, Zambaldi, et al. (*arXiv* 2018)**

# Background: Graph Networks

Graph neural networks: Gori et al. (2005), Scarselli et al. (2005, 2009), Li et al. (2015), Kipf & Welling (2016), Gilmer et al. (2017), many more!

1. Takes ***graphs*** as input, return graphs as output

2. Invariant to the ***permutation*** of the nodes and edges

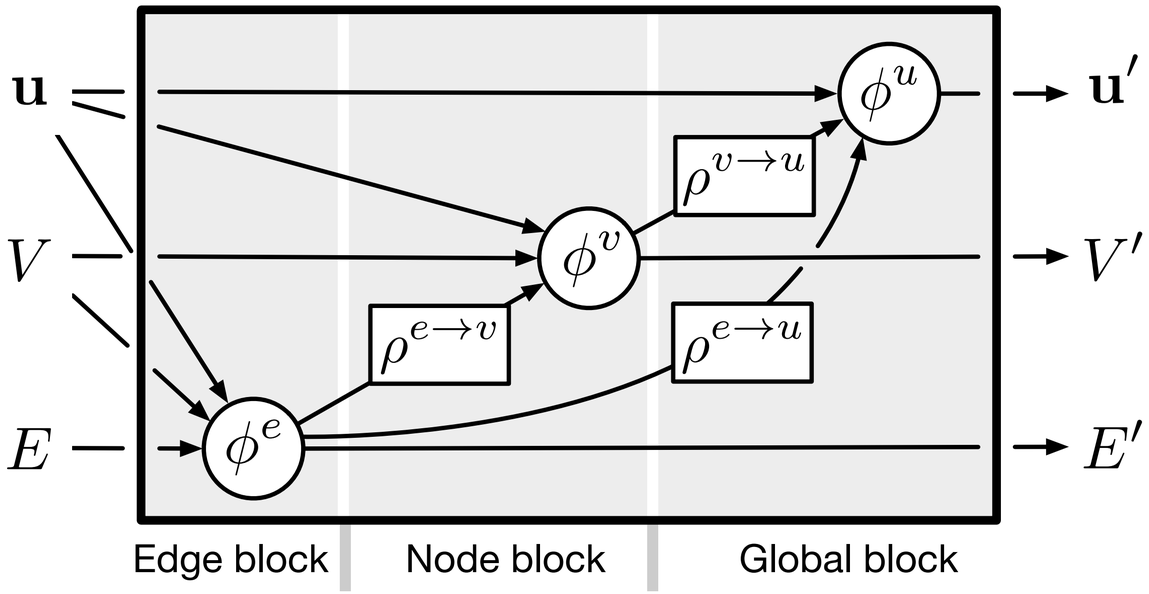


**Battaglia, Hamrick, Bapst, Sanchez-Gonzalez, Zambaldi, et al. (*arXiv* 2018)**

# Background: Graph Networks

Graph neural networks: Gori et al. (2005), Scarselli et al. (2005, 2009), Li et al. (2015), Kipf & Welling (2016), Gilmer et al. (2017), many more!

1. Takes ***graphs*** as input, return graphs as output

2. Invariant to the ***permutation*** of the nodes and edges

3. Scales to different ***numbers*** of nodes and edges

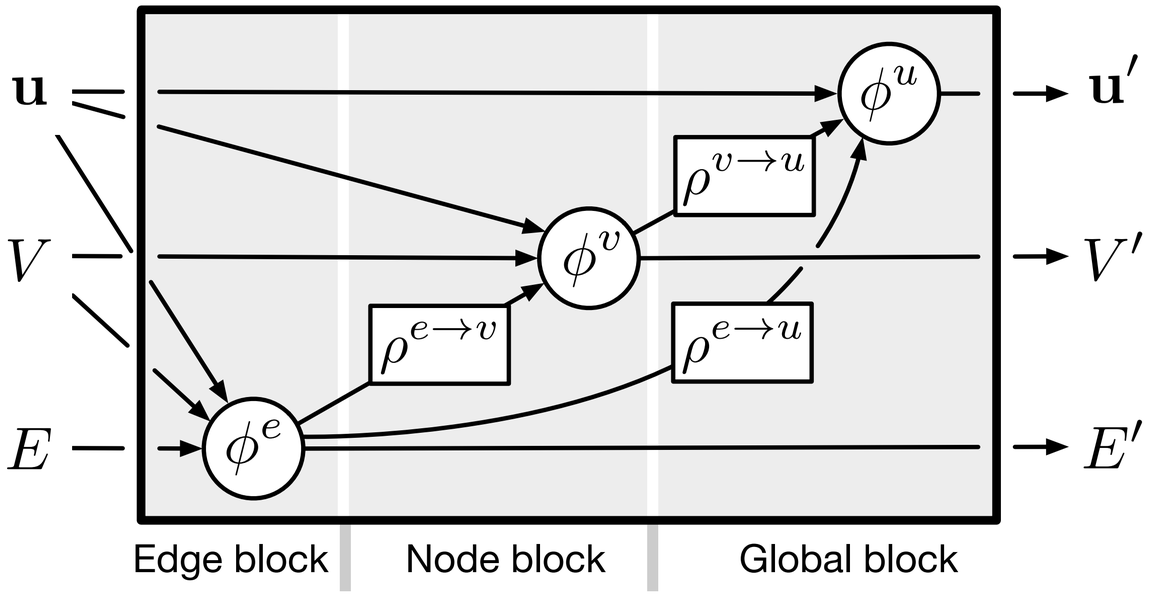


**Battaglia, Hamrick, Bapst, Sanchez-Gonzalez, Zambaldi, et al. (*arXiv* 2018)**

# Graph Network Agent (GN-DQN)

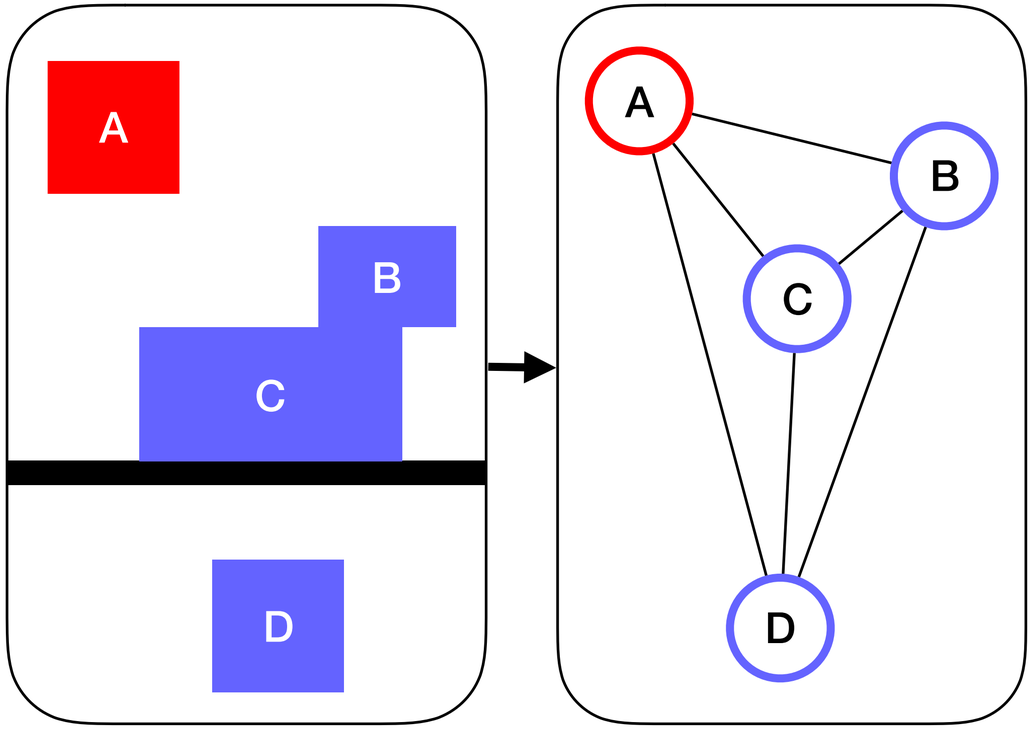


Given observation

Structured representation

# Graph Network Agent (GN-DQN)

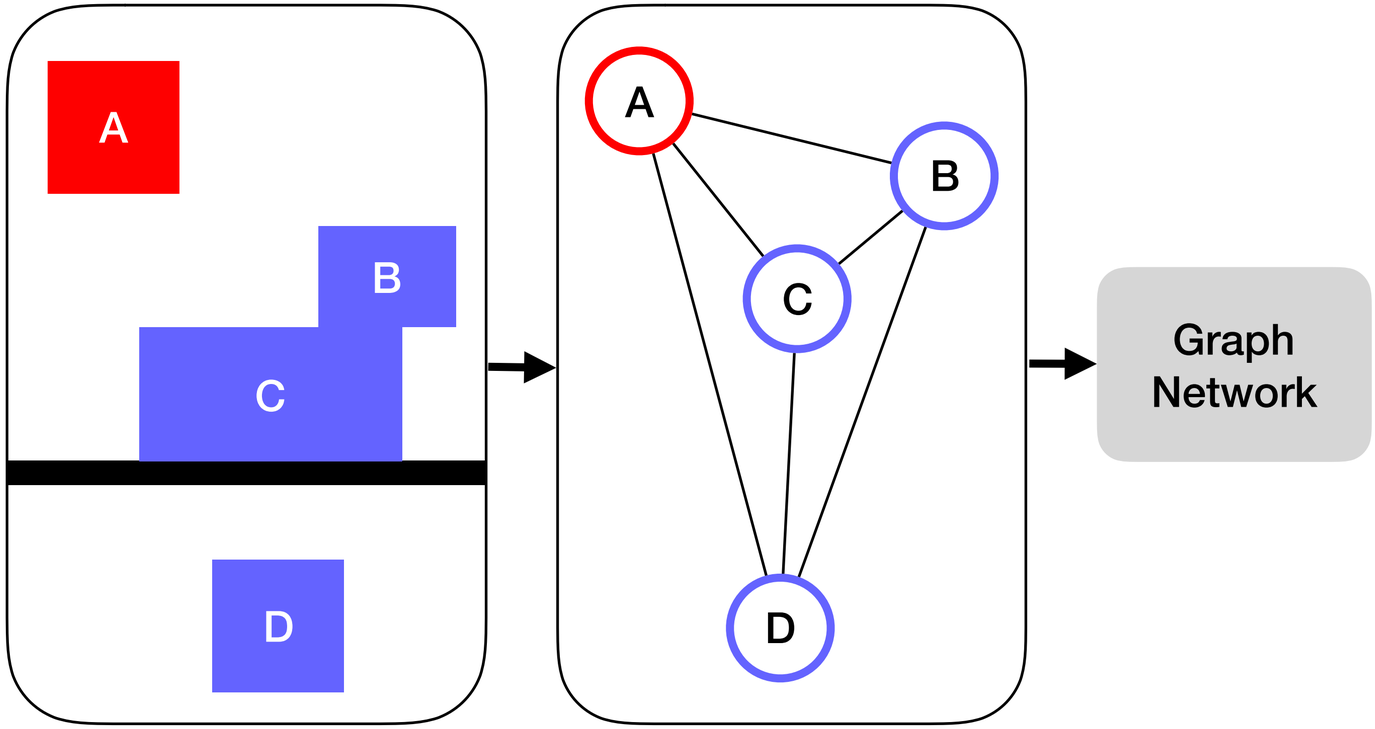

# Graph Network Agent (GN-DQN)

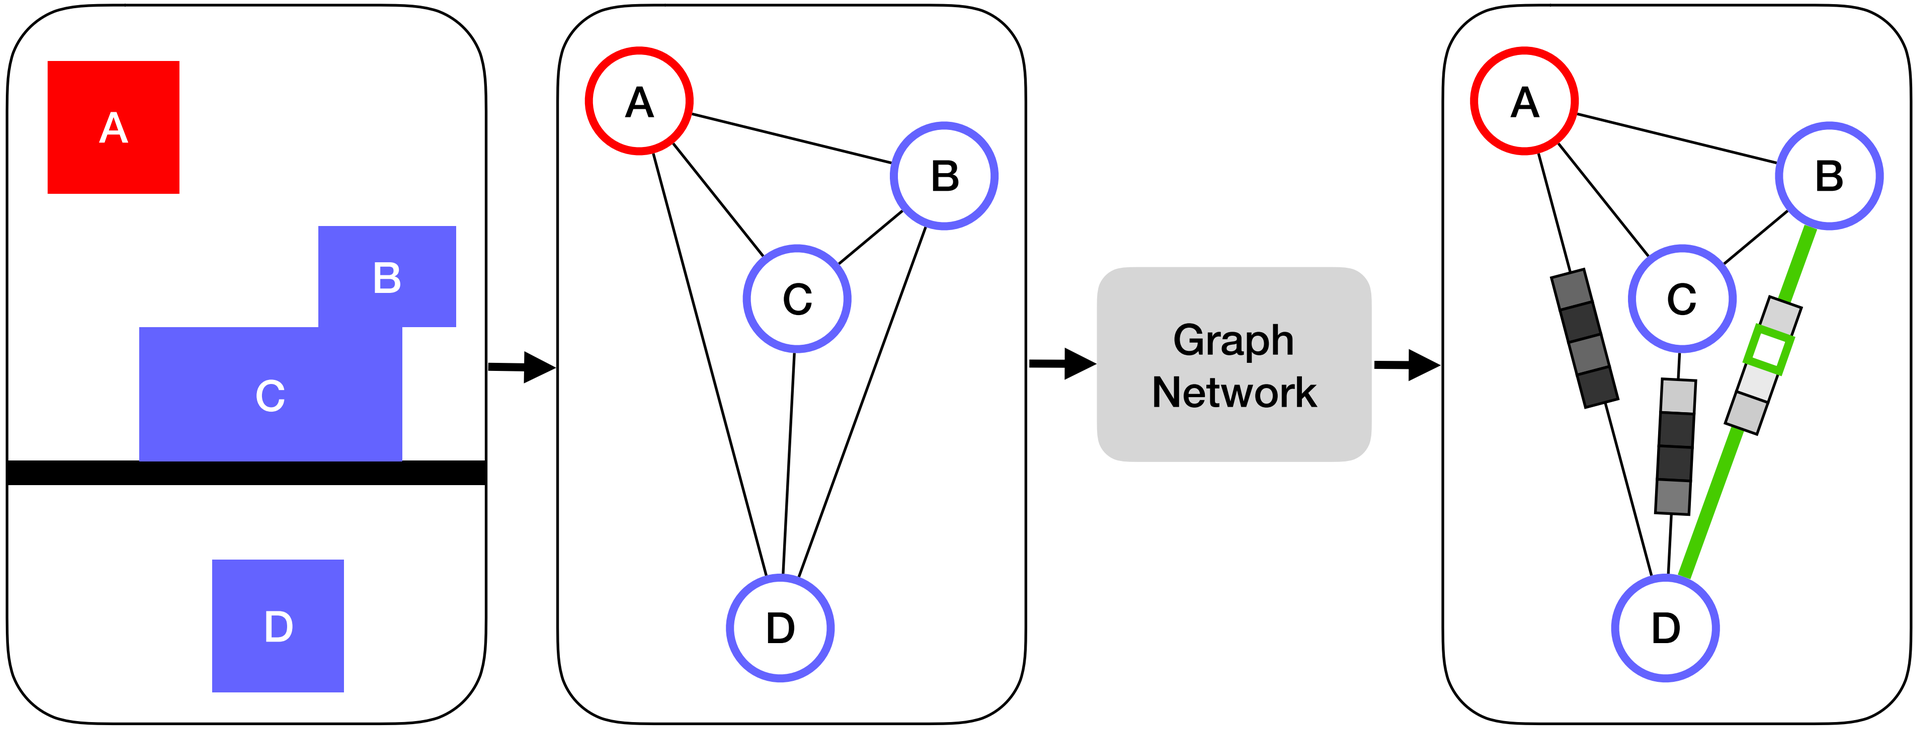


Given observation

Structured representation

Relative action

# Graph Network Agent (GN-DQN)

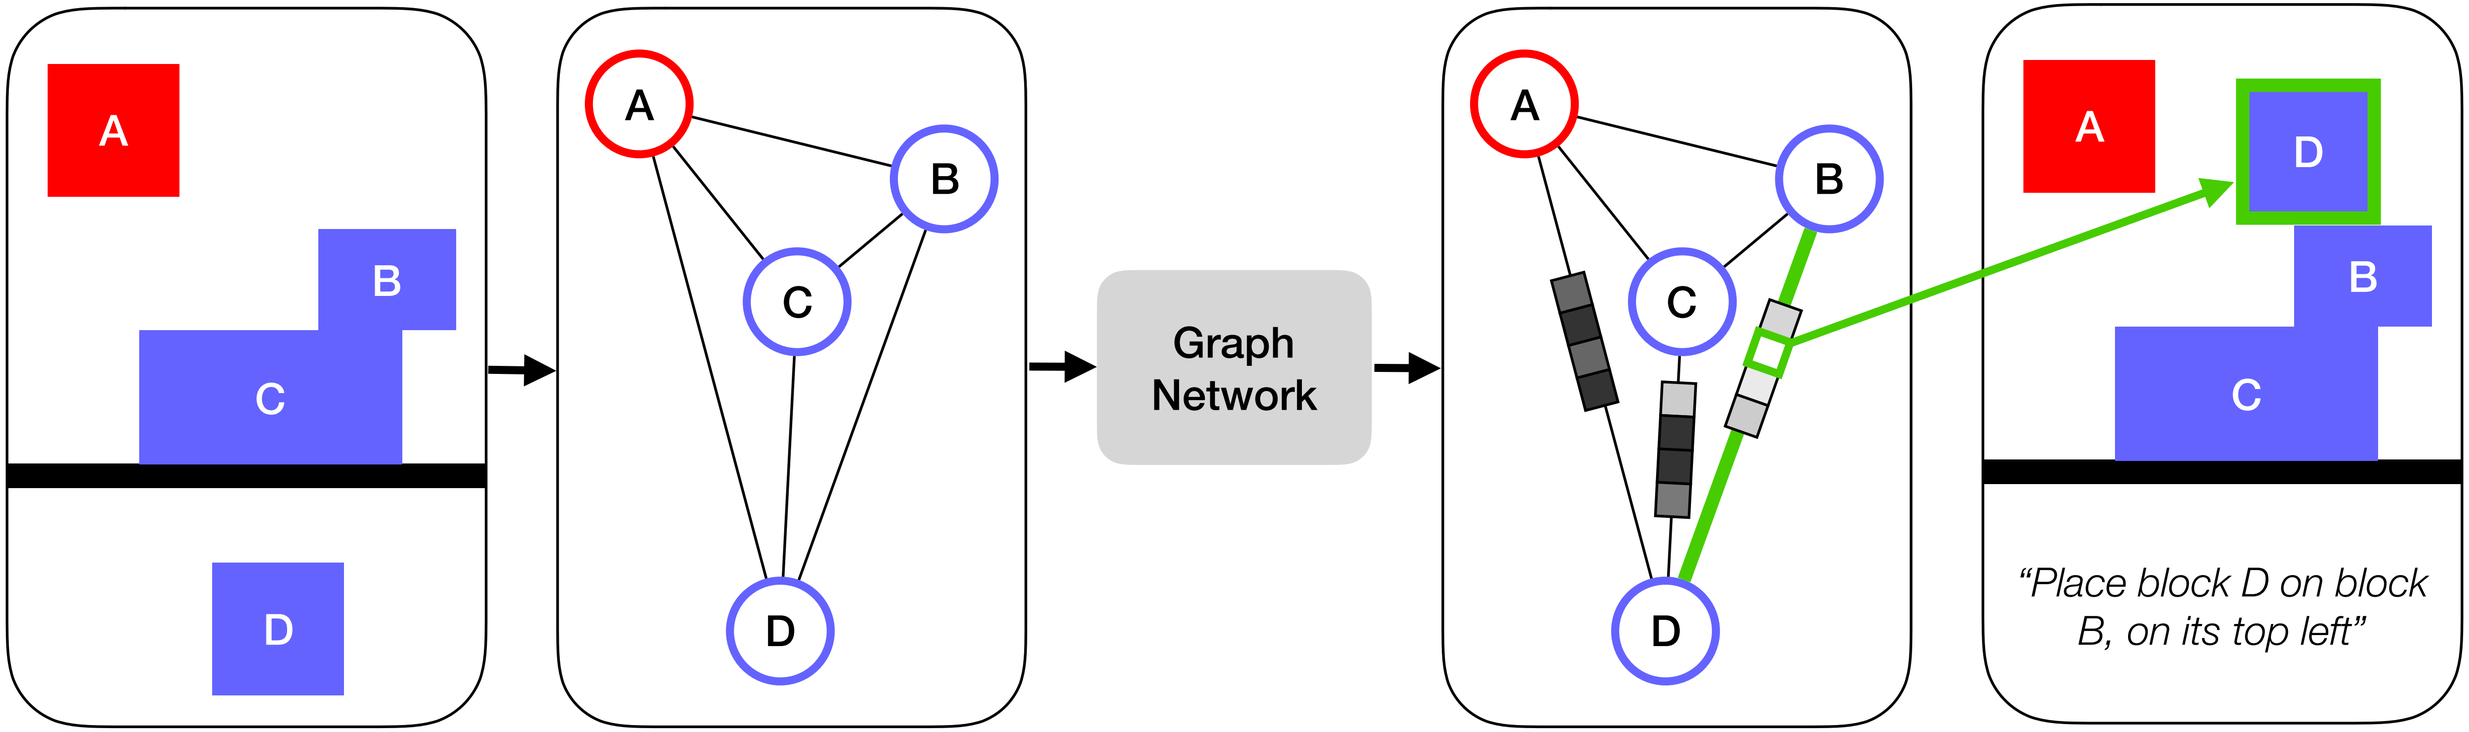


Given observation

Structured representation

Relative action

# Graph Network Agent (GN-DQN)

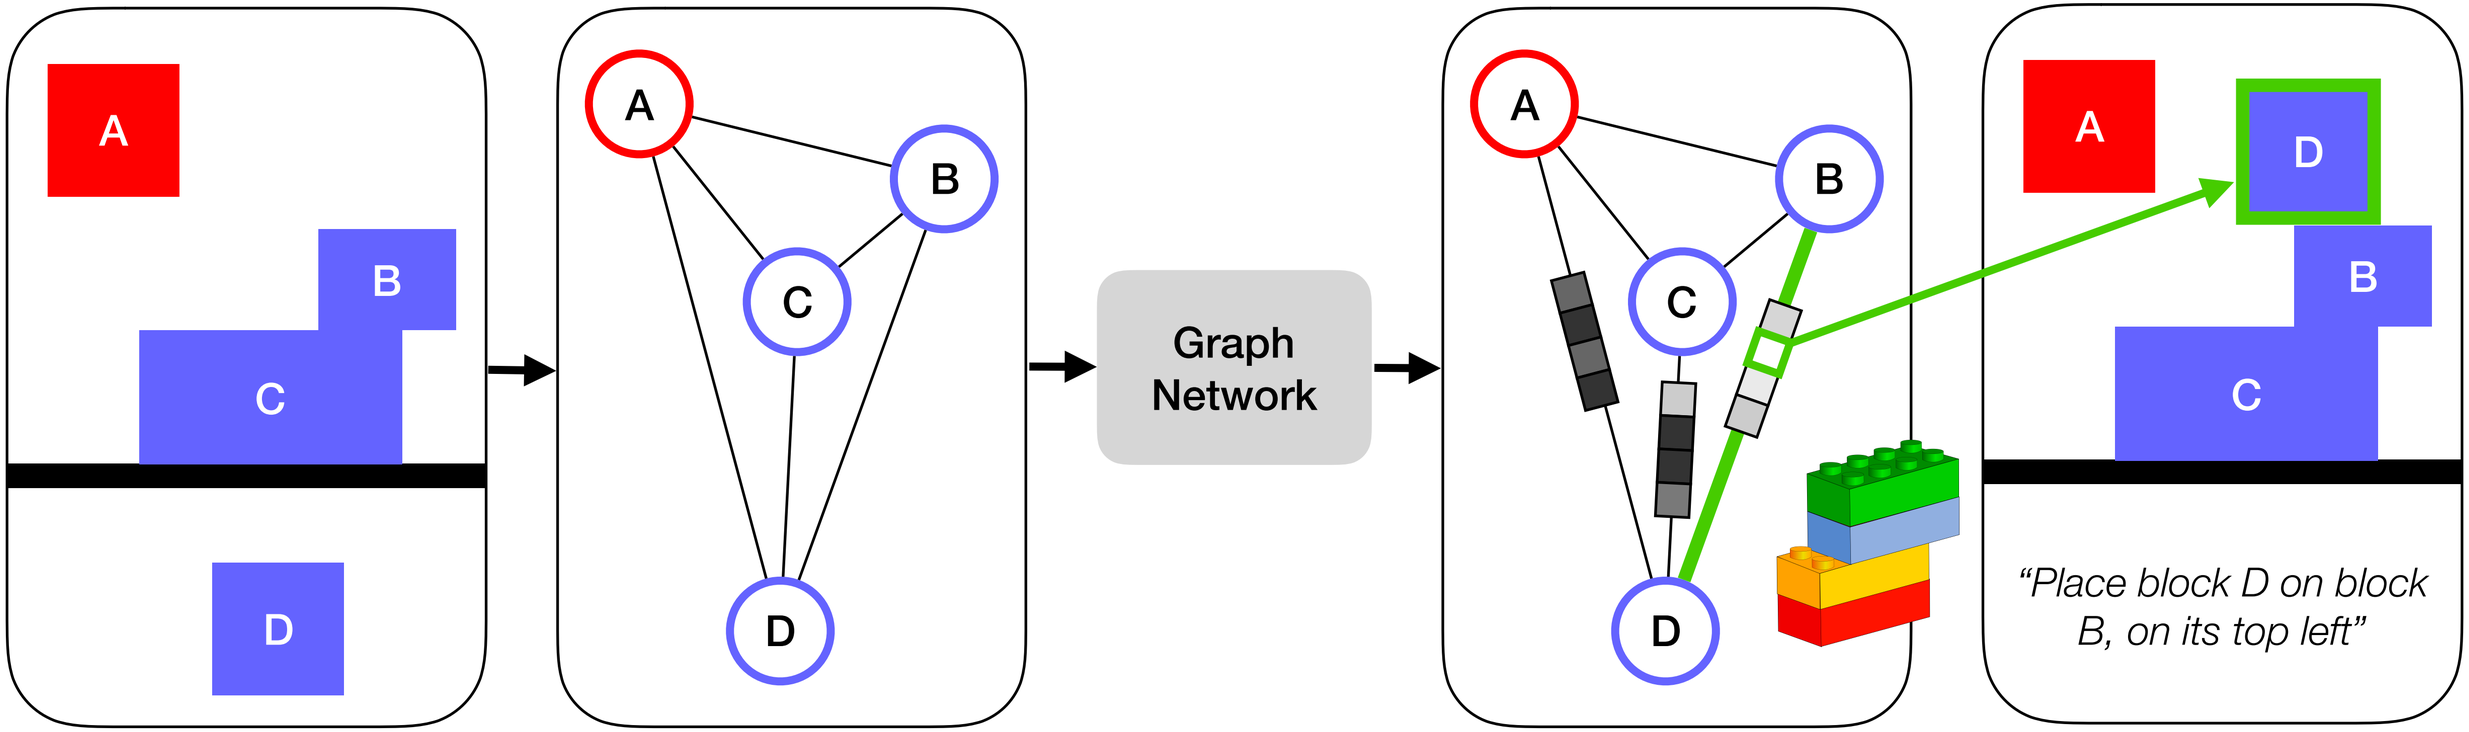

# Graph Network Agent (GN-DQN)

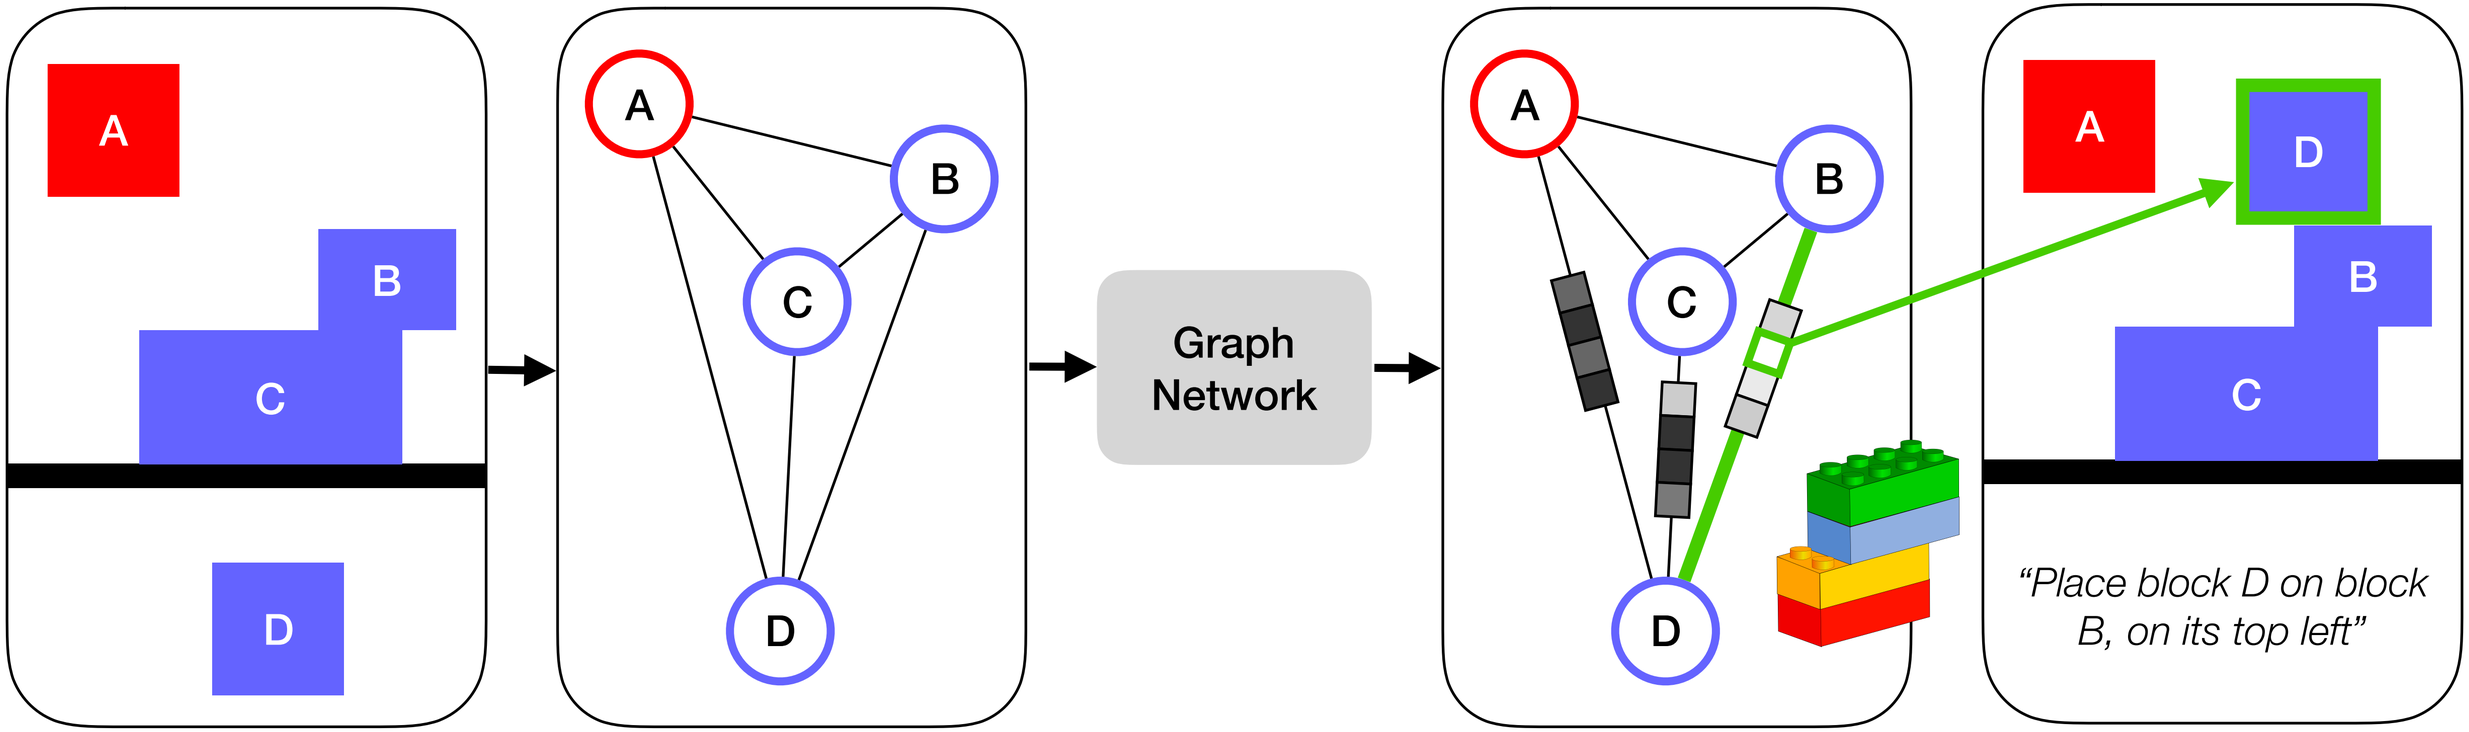


Given observation

Structured representation

Relative action

**10 objects = 1350 actions!**

# Baseline Architectures

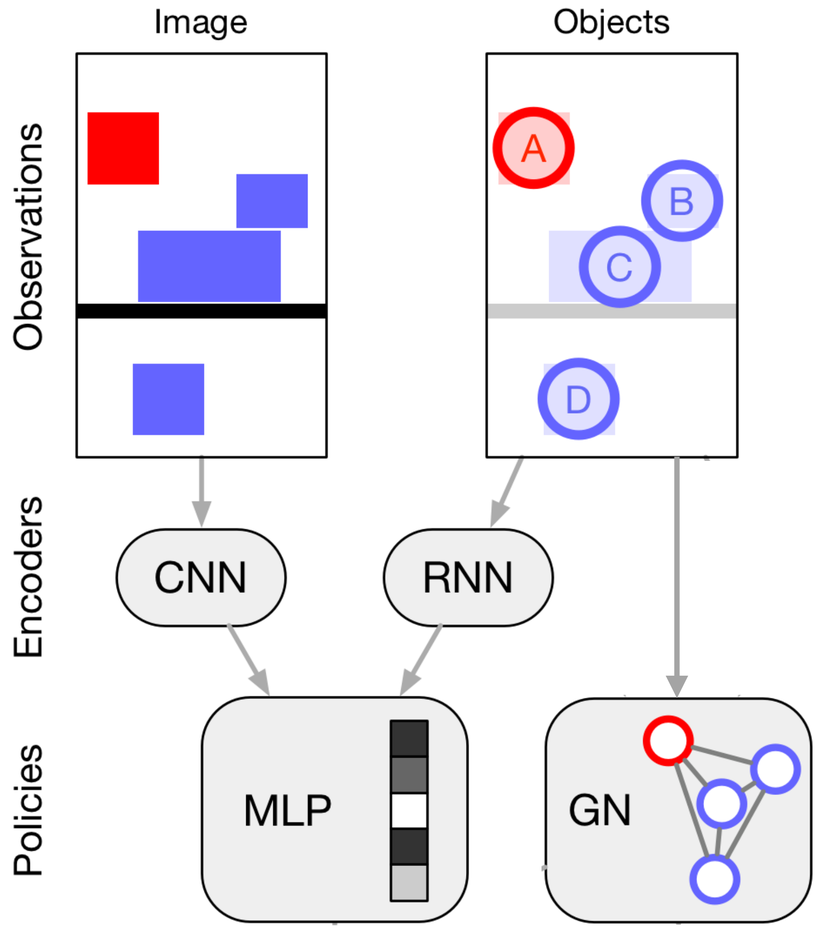


| Agent | Obs. | Encoder | Policy | Learning Algorithm | Actions |
|---|---|---|---|---|---|
| **GN-DQN** | Objects | - | GN | DQN | Discrete |
| **GN-RS0** | Objects | - | GN | RS0 | Continuous |
| **RNN-RS0** | Objects | RNN | MLP | RS0 | Continuous |
| **CNN-RS0** | Image | CNN | MLP | RS0 | Continuous |

# Key Questions

# Key Questions

1. What is the contribution of ***relative*** vs. absolute actions?

# Key Questions

1. What is the contribution of ***relative*** vs. absolute actions?

2. What is the contribution of ***structured*** representations?

# Key Questions

1. What is the contribution of ***relative*** vs. absolute actions?

2. What is the contribution of ***structured*** representations?

3. What is the contribution of ***planning***?

# What is the contribution of ***relative*** vs. absolute actions?

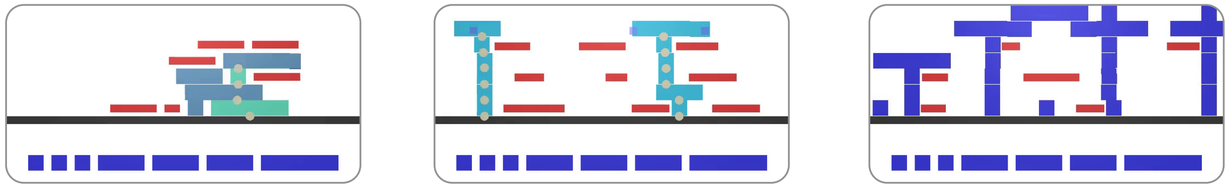

Best absolute agent

Relative agent (GN-DQN)

*(Median across 10 seeds, with error bars for min/max seed)*

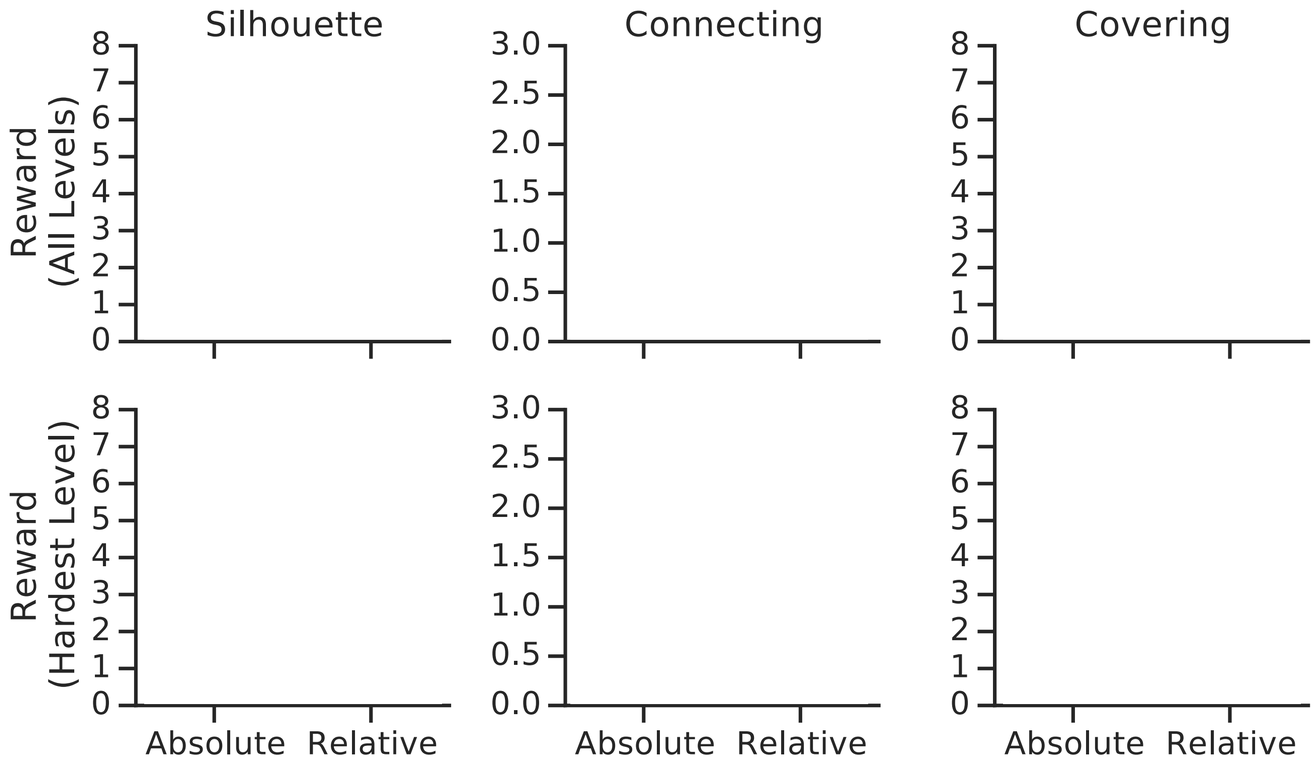

# What is the contribution of ***relative*** vs. absolute actions?

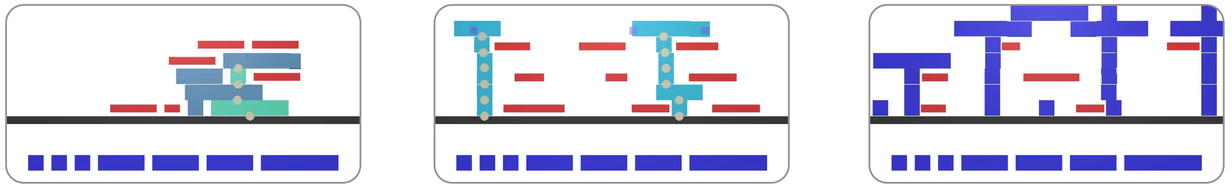

Best absolute agent

Relative agent (GN-DQN)

*(Median across 10 seeds, with error bars for min/max seed)*

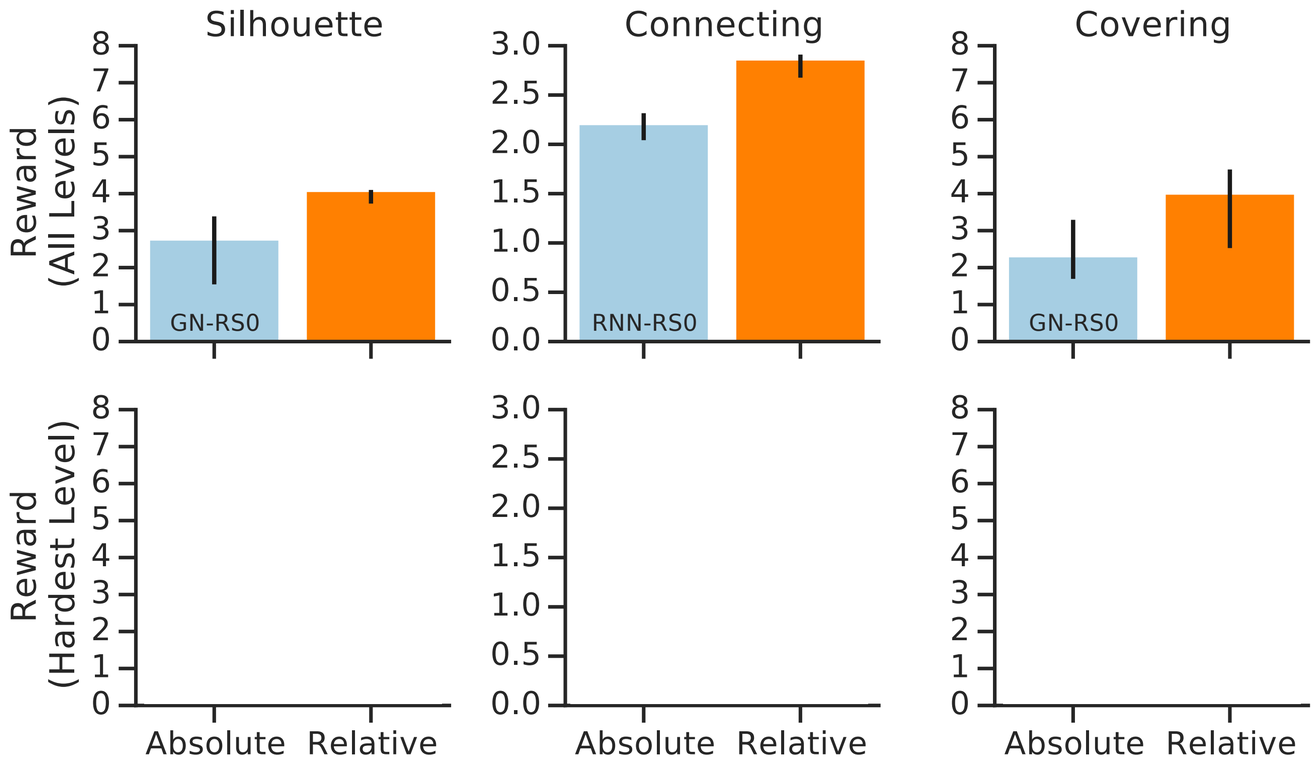

# What is the contribution of ***relative*** vs. absolute actions?

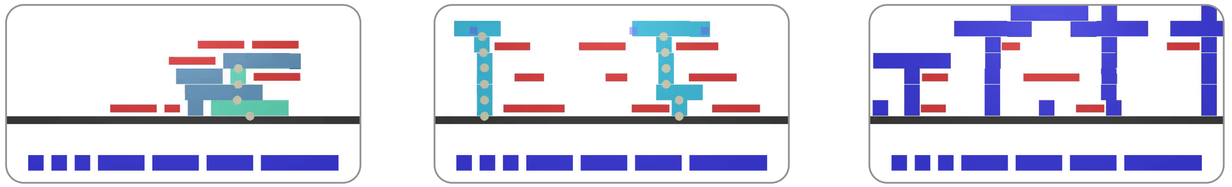

Best absolute agent

Relative agent (GN-DQN)

*(Median across 10 seeds, with error bars for min/max seed)*

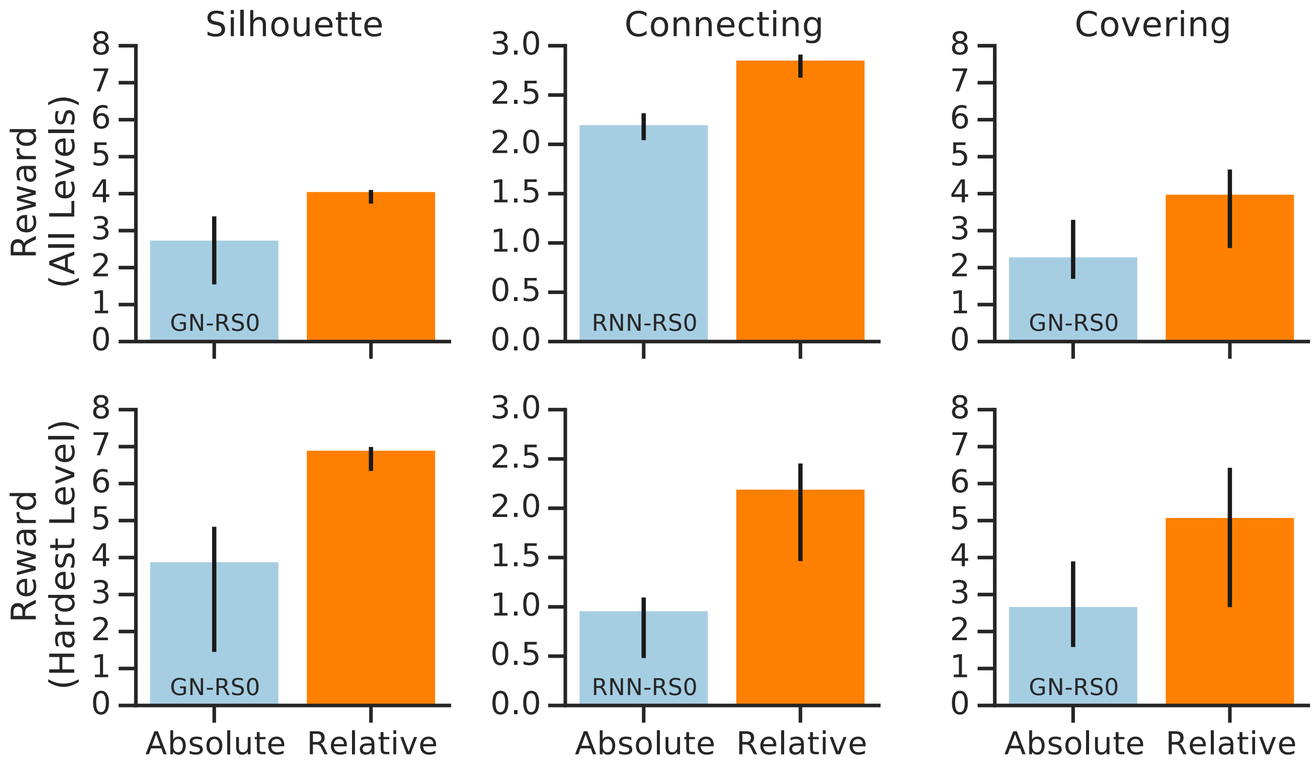

# What is the contribution of ***structured*** representations?

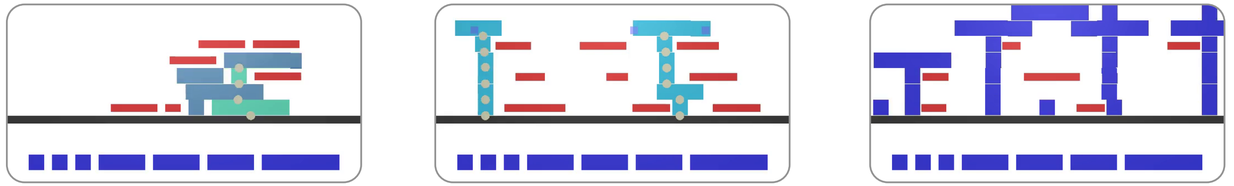

Best unstructured agent

Structured agent (GN-DQN)

(Both with relative actions)

*(Median across 10 seeds, with error bars for min/max seed)*

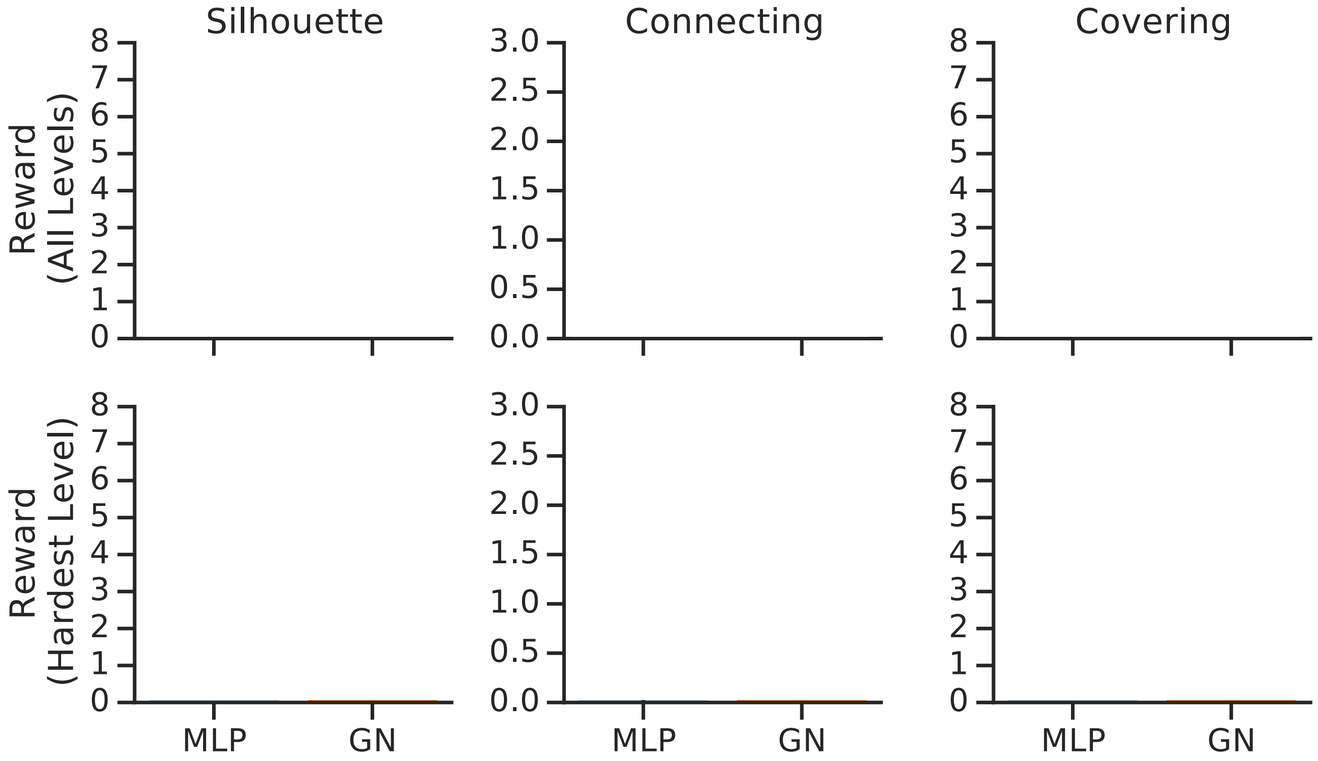

# What is the contribution of ***structured*** representations?

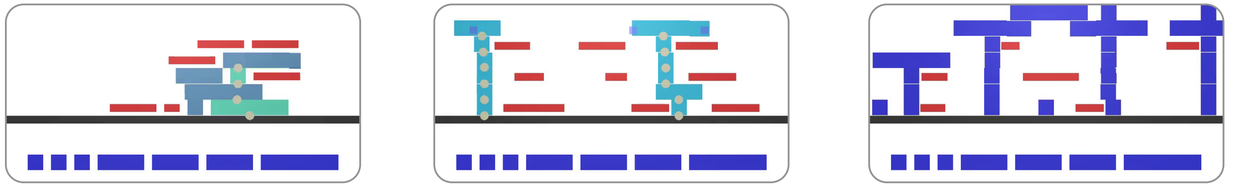

Best unstructured agent

Structured agent (GN-DQN)

(Both with relative actions)

*(Median across 10 seeds, with error bars for min/max seed)*

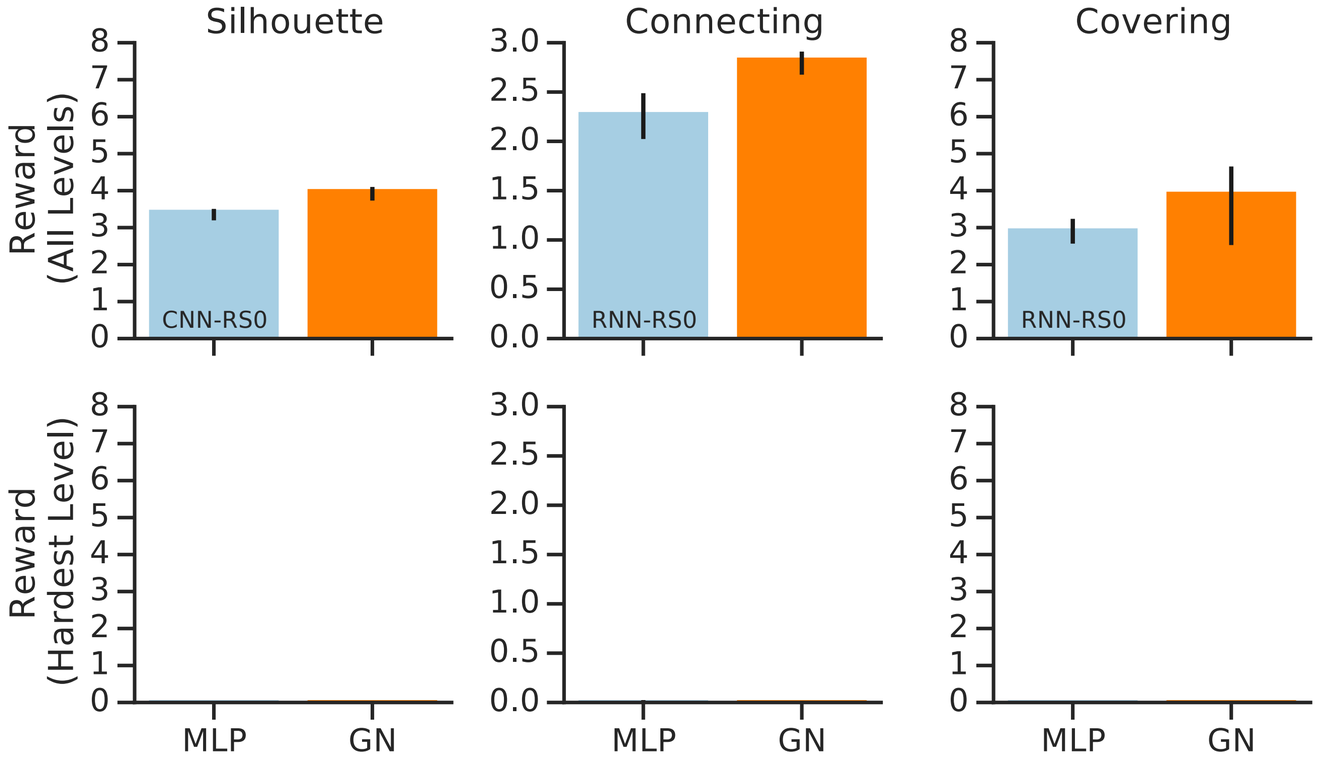

# What is the contribution of ***structured*** representations?

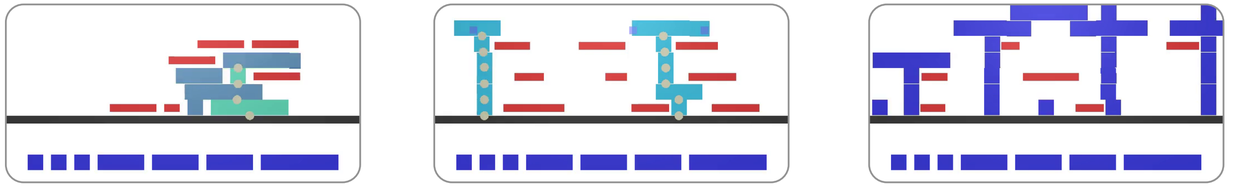

Best unstructured agent

Structured agent (GN-DQN)

(Both with relative actions)

*(Median across 10 seeds, with error bars for min/max seed)*

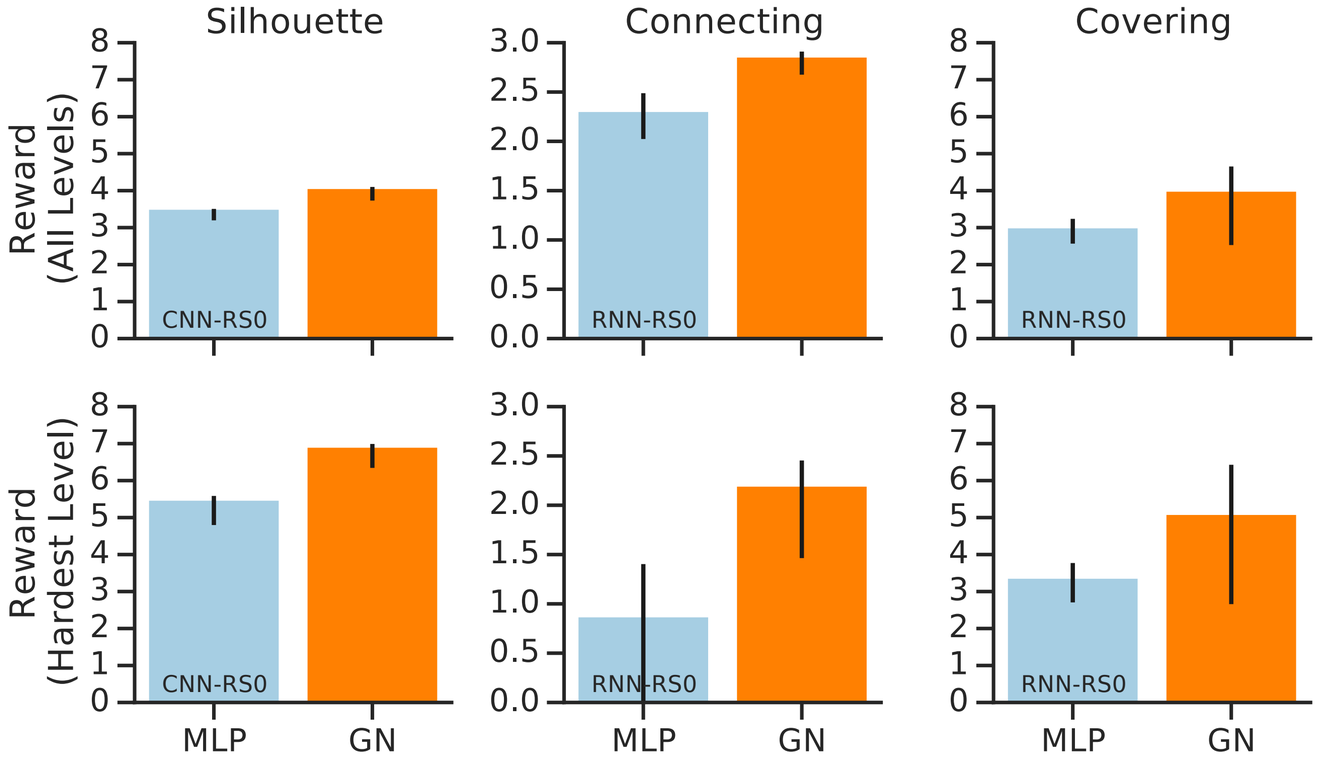

# Silhouette

**Absolute Actions (GN-RS0)**
**(Average reward: 4.83)**

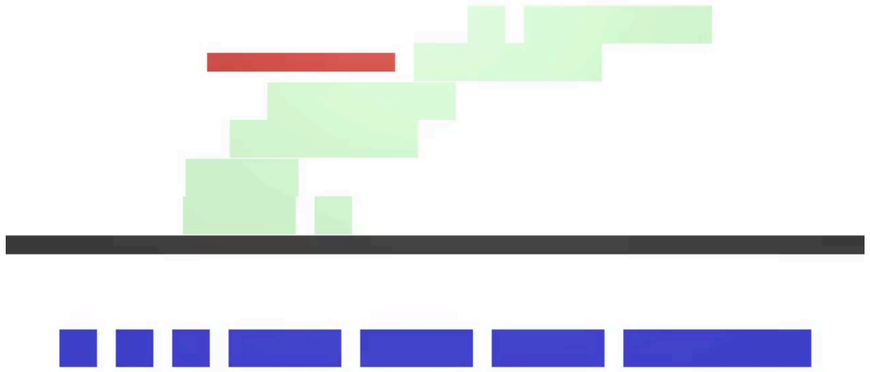

**Relative Actions (GN-DQN)**
**(Average reward: 6.99)**

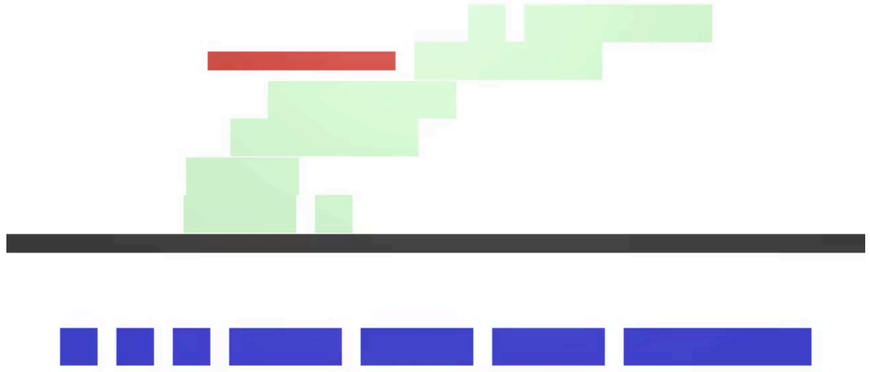

*Reward: +1 per target, -0.5 per sticky block*

# Silhouette

**Absolute Actions (GN-RS0)**
**(Average reward: 4.83)**

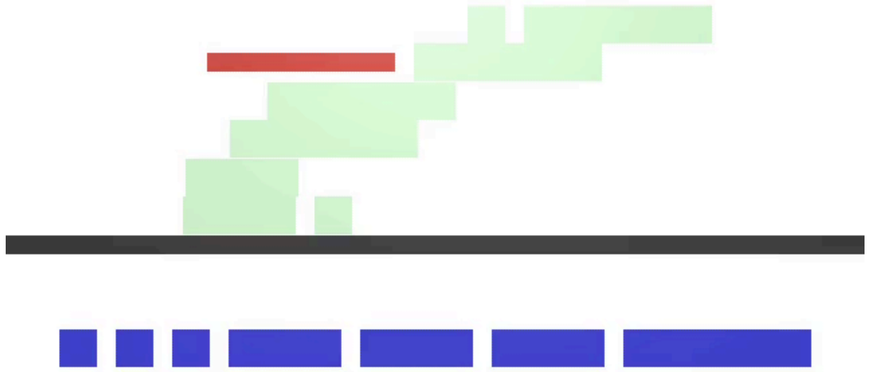

**Relative Actions (GN-DQN)**
**(Average reward: 6.99)**

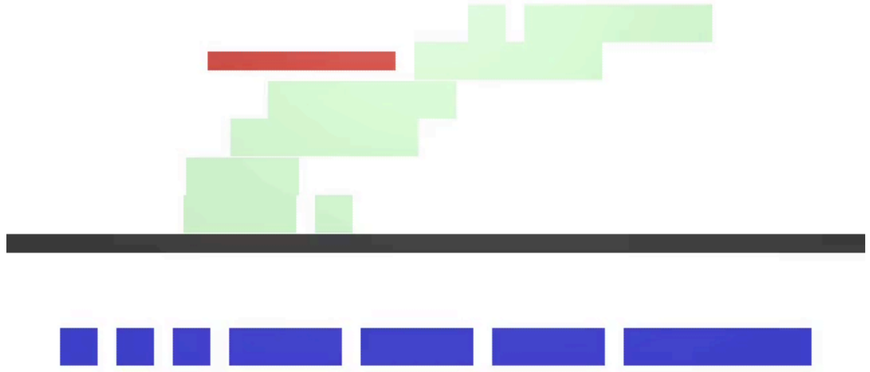

*Reward: +1 per target, -0.5 per sticky block*

# Connecting

**Absolute Actions (RNN-RS0)**
**(Average reward: 1.09)**

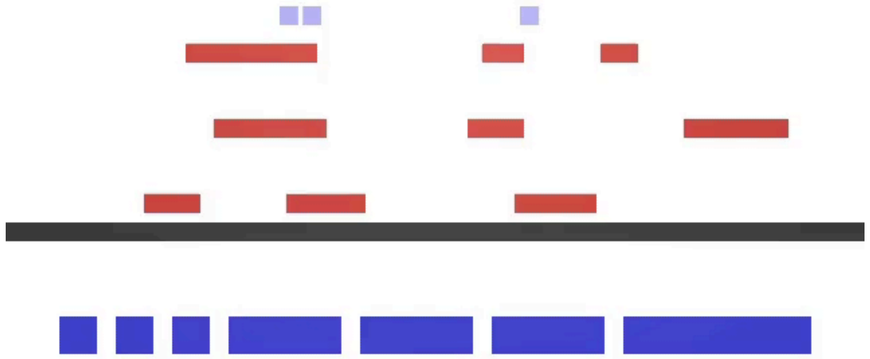

**Relative Actions (GN-DQN)**
**(Average reward: 2.45)**

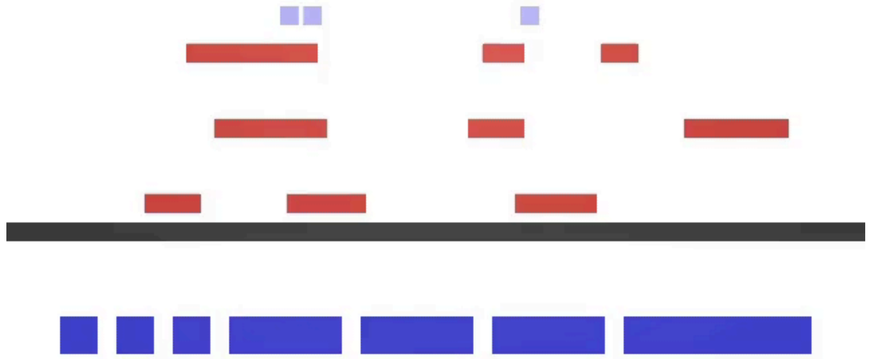

*Reward: +1 per target, free sticky blocks*

# Connecting

**Absolute Actions (RNN-RS0)**
**(Average reward: 1.09)**

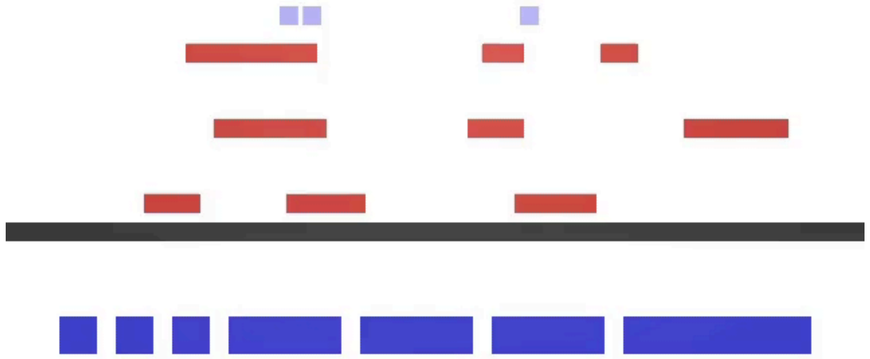

**Relative Actions (GN-DQN)**
**(Average reward: 2.45)**

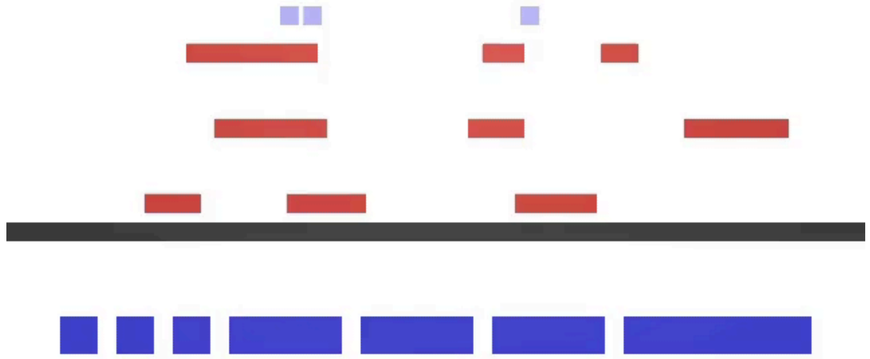

*Reward: +1 per target, free sticky blocks*

# Covering

**Absolute Actions (GN-RS0)**
**(Average reward: 3.90)**

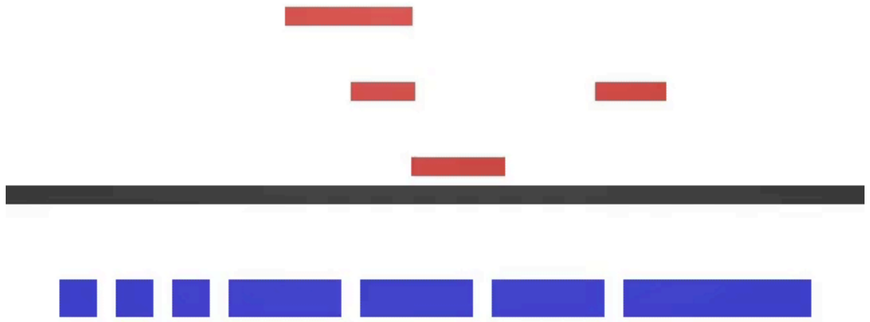

**Relative Actions (GN-DQN)**
**(Average reward: 6.43)**

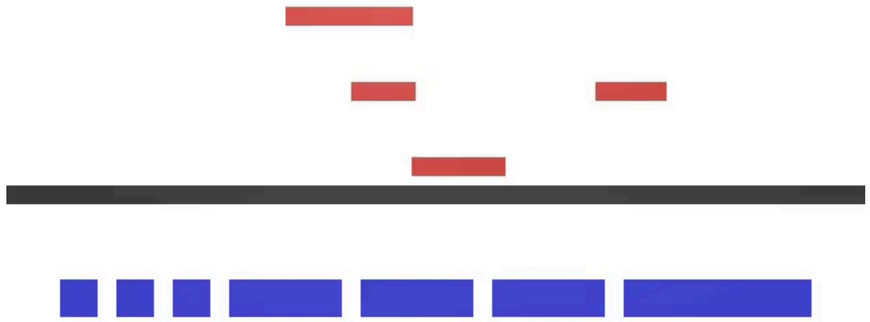

*Reward: proportional to length covered, -2 per sticky block*

# Covering

**Absolute Actions (GN-RS0)**
**(Average reward: 3.90)**

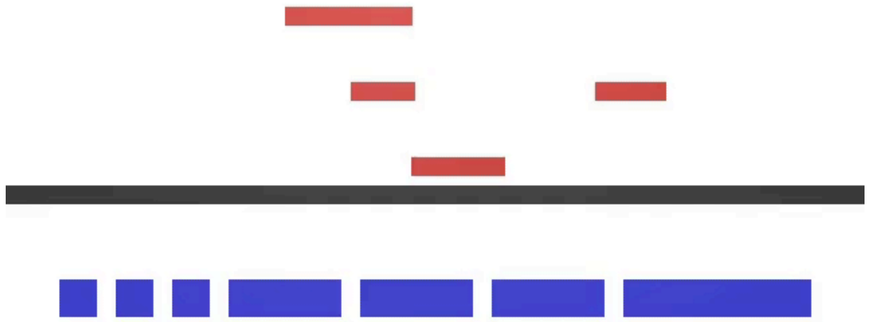

**Relative Actions (GN-DQN)**
**(Average reward: 6.43)**

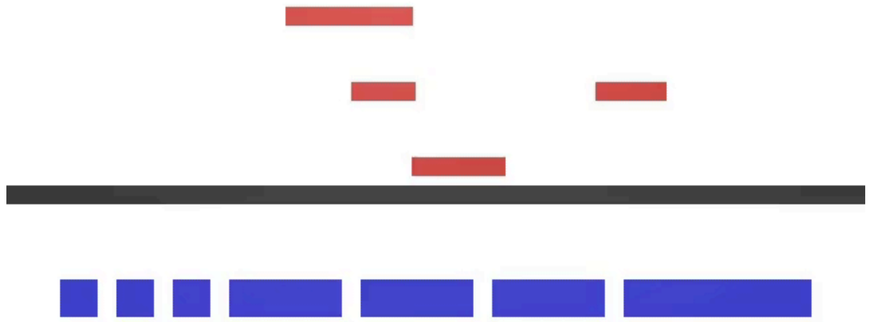

*Reward: proportional to length covered, -2 per sticky block*

# What is the contribution of ***planning***?

# What is the contribution of ***planning***?

Covering Hard

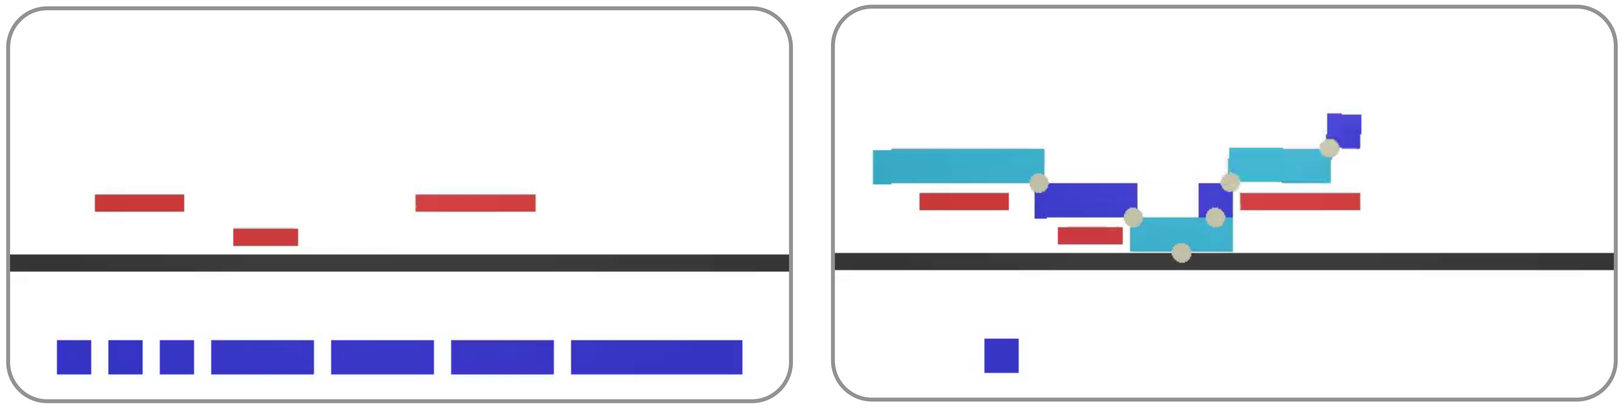

*Length covered*
*-0.5 per sticky block*

# What is the contribution of ***planning***?

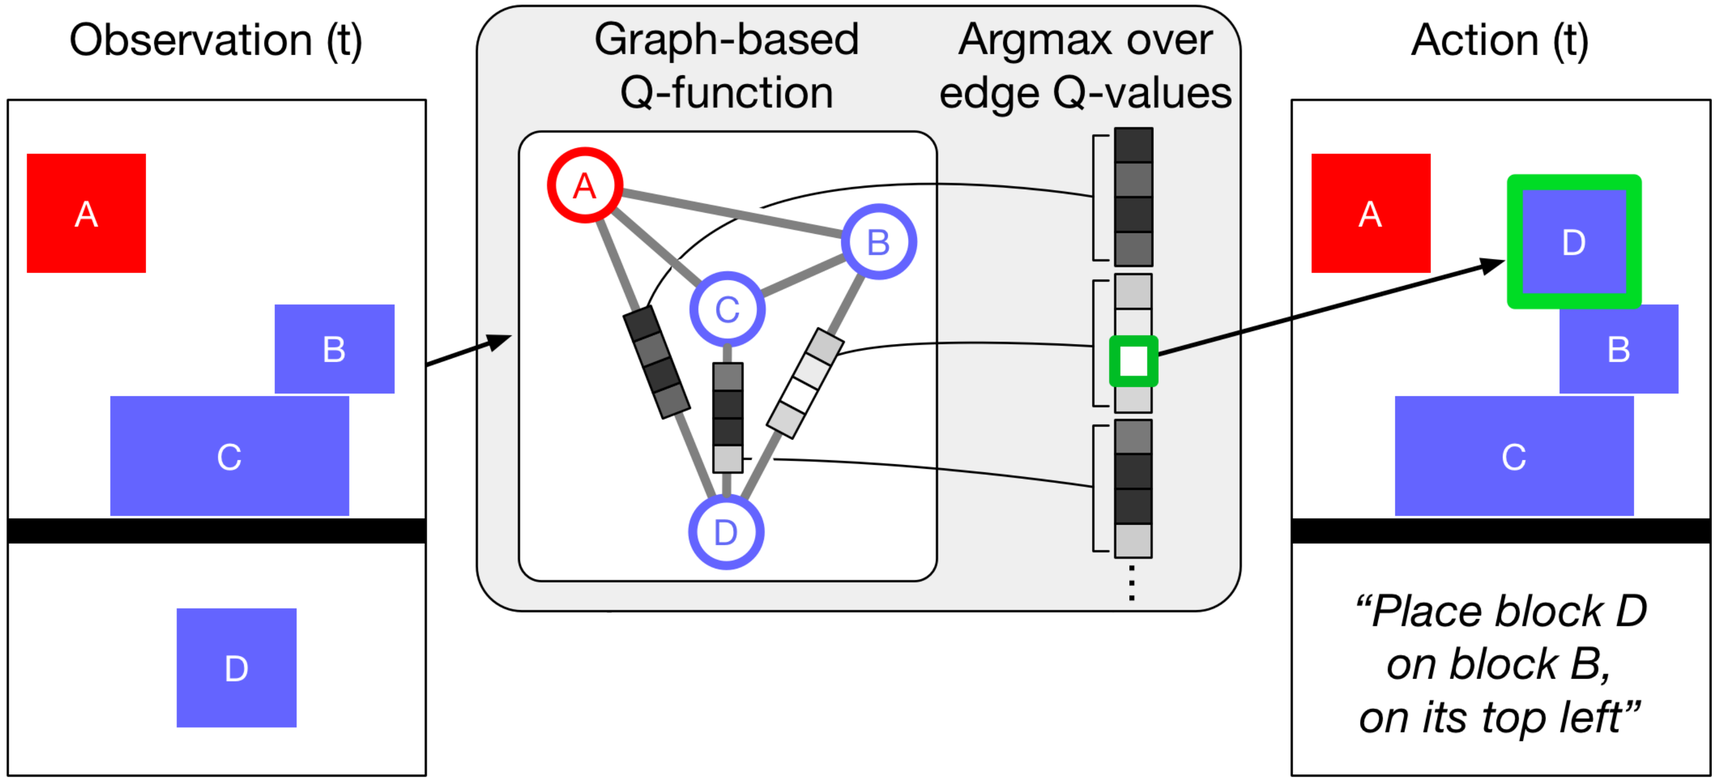

# What is the contribution of ***planning***?

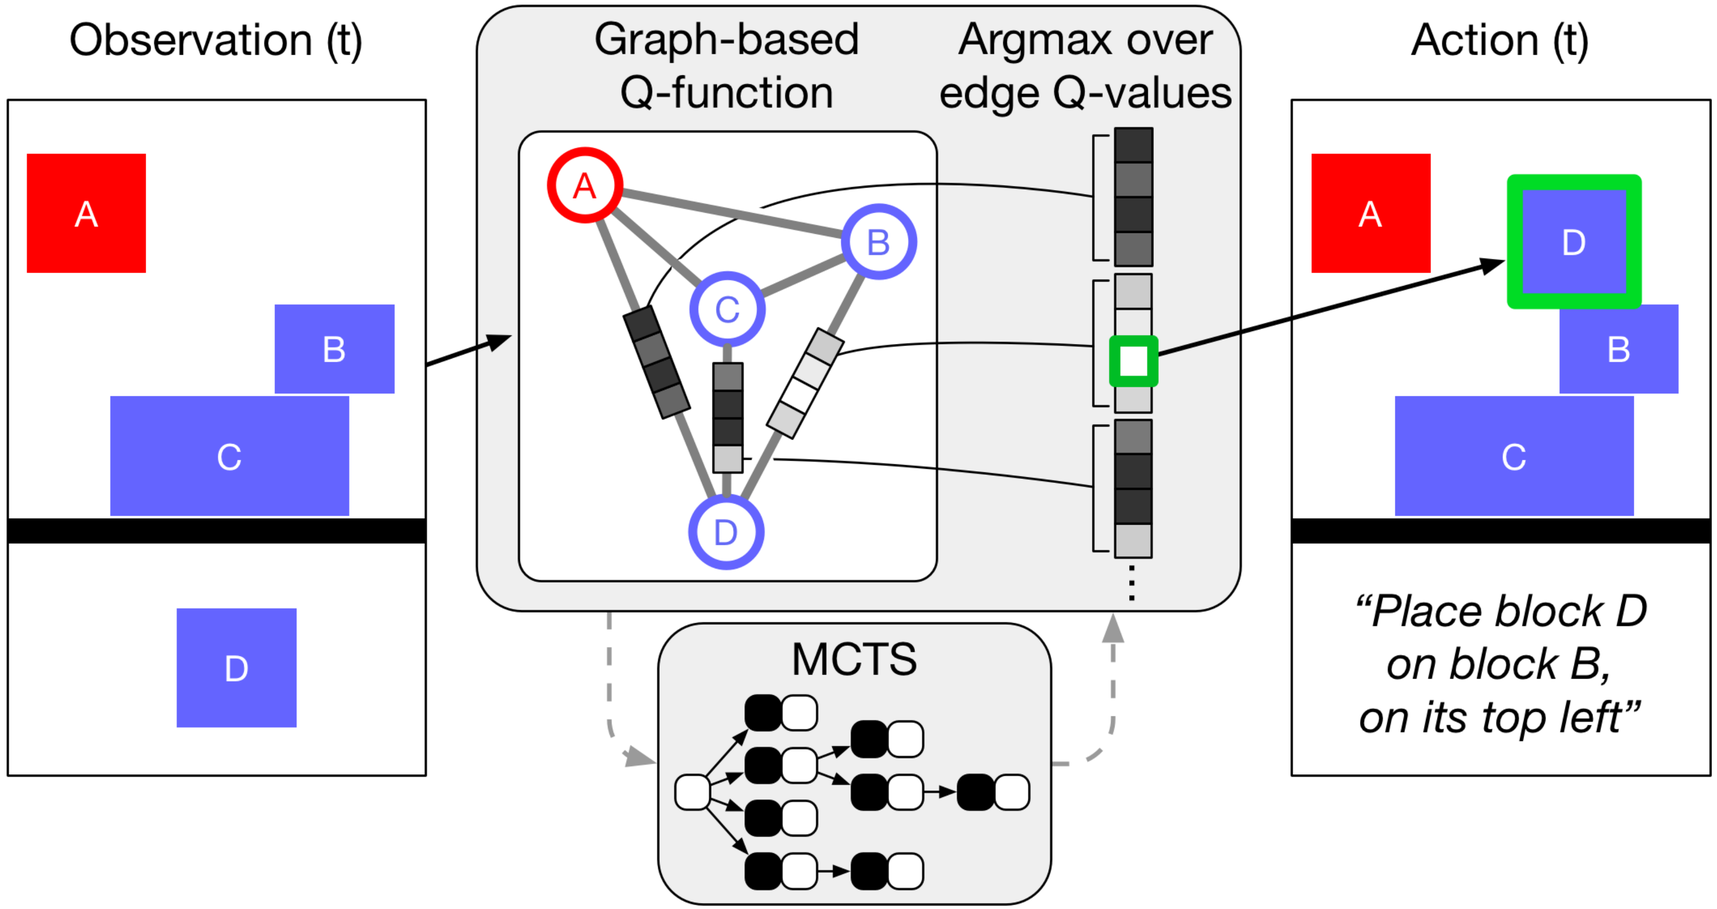

# What is the contribution of ***planning***?

Model-free

MCTS @ test time only

MCTS @ train (10) and test

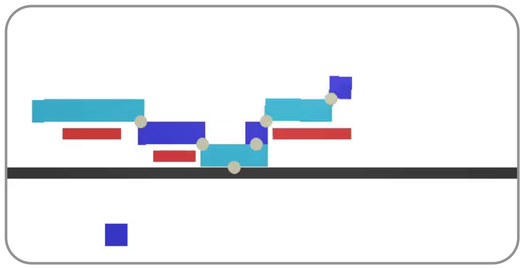

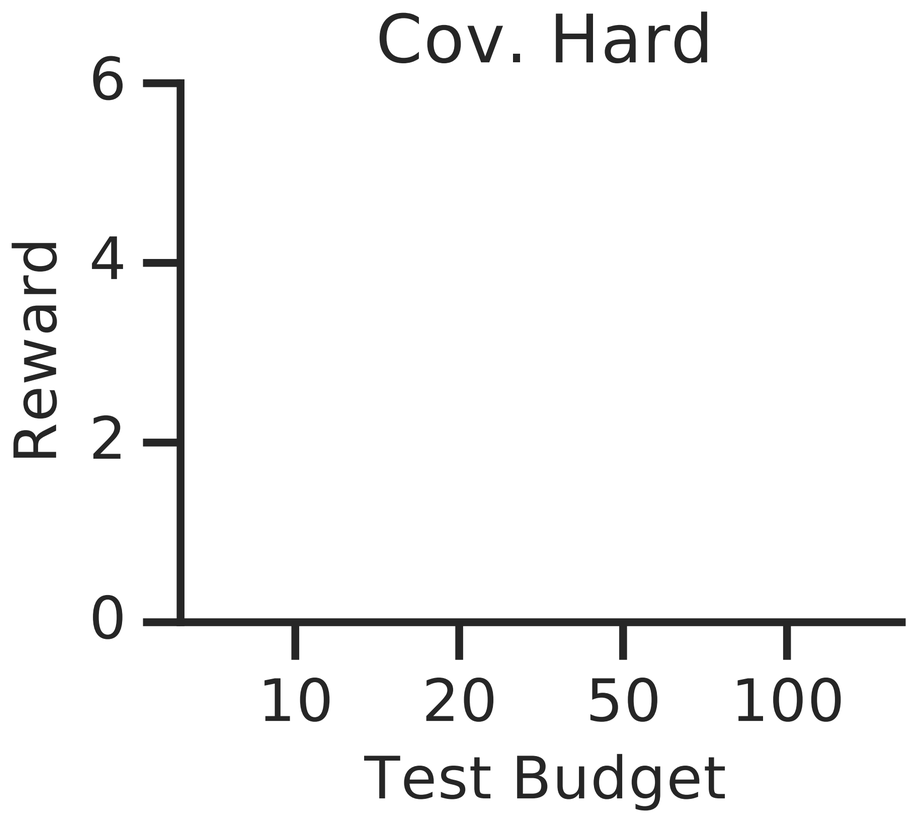


*(Median across 10 seeds, with error bars for min/max seed)*

# What is the contribution of ***planning***?

Model-free

MCTS @ test time only

MCTS @ train (10) and test

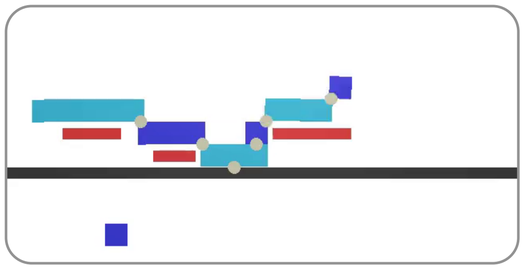

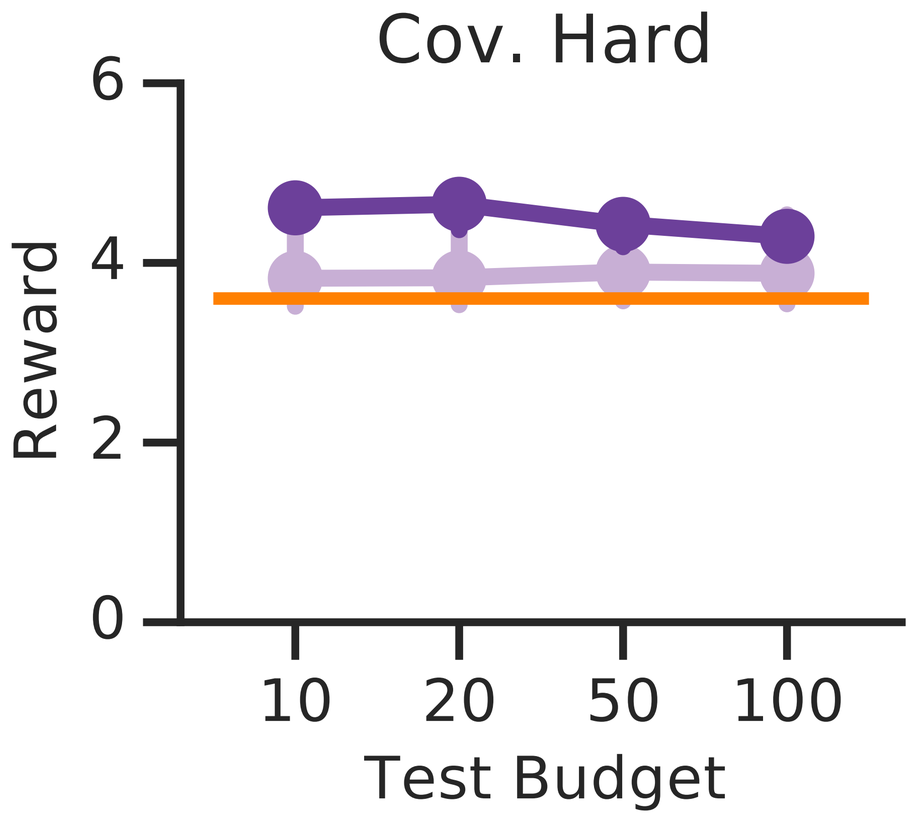


*(Median across 10 seeds, with error bars for min/max seed)*

# What is the contribution of ***planning***?

Model-free

MCTS @ test time only

MCTS @ train (10) and test

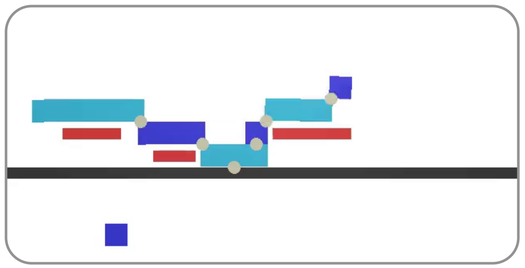

*(Median across 10 seeds, with error bars for min/max seed)*

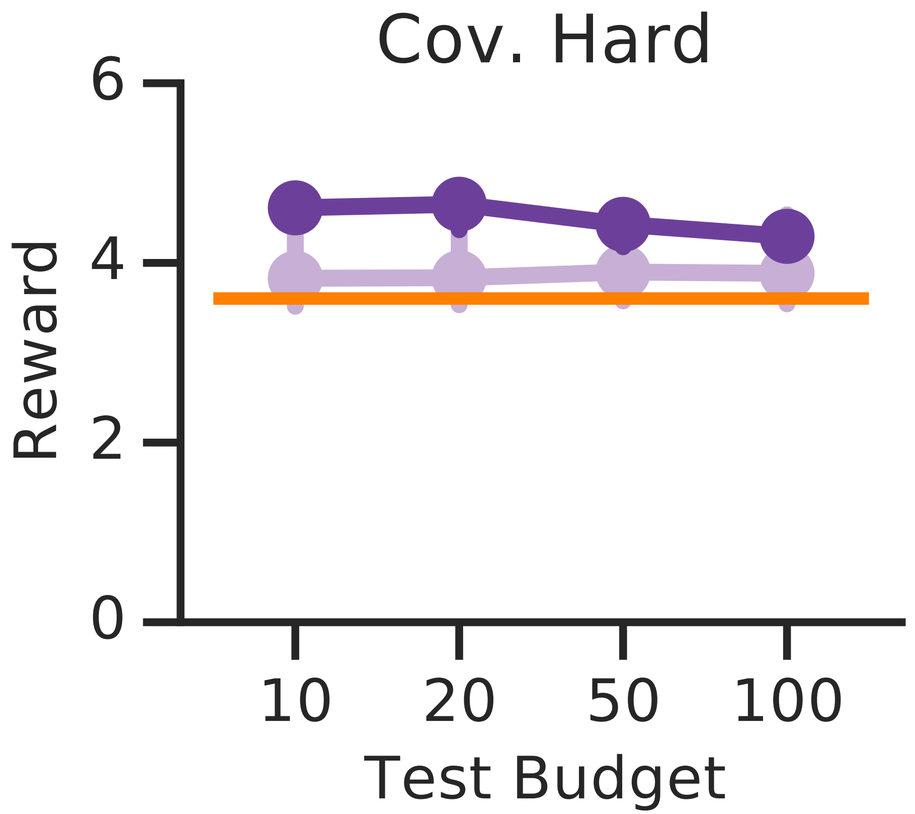


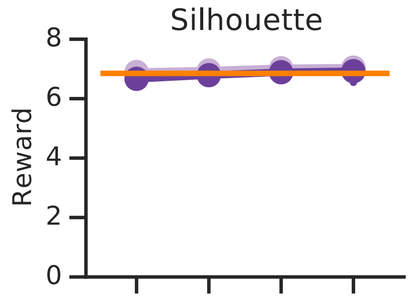


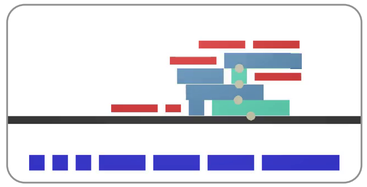

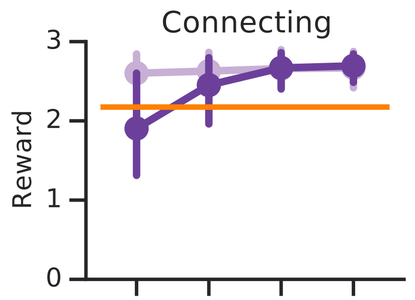


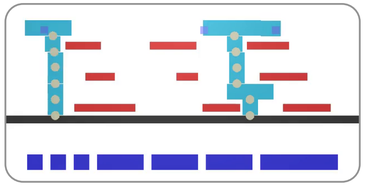

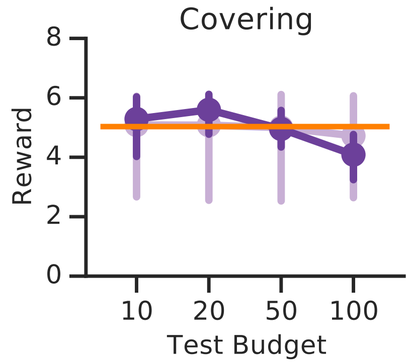


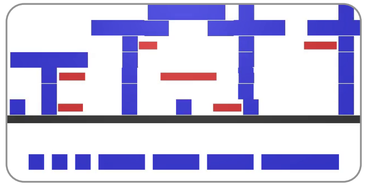

# Covering Hard

**Model-Free (GN-DQN)**
**Reward: 3.1**

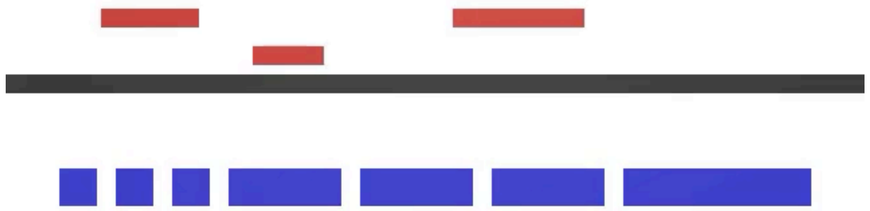

**Model-Based (GN-DQN-MCTS)**
**Reward: 4.1**

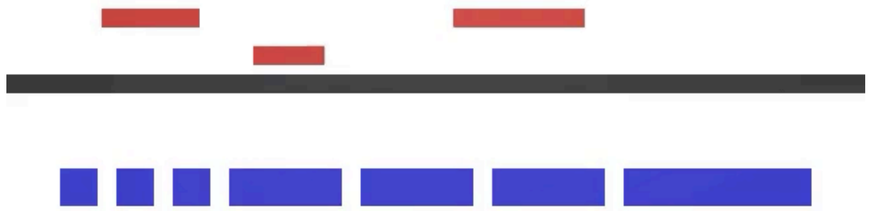

*Reward: proportional to length covered, -0.5 per sticky block*

# Covering Hard

**Model-Free (GN-DQN)**
**Reward: 3.1**

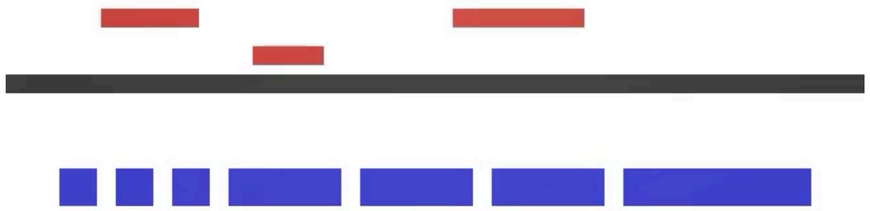

**Model-Based (GN-DQN-MCTS)**
**Reward: 4.1**

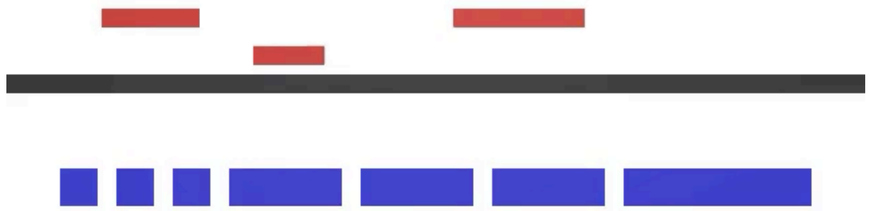

*Reward: proportional to length covered, -0.5 per sticky block*

# Additional Results: Generalization

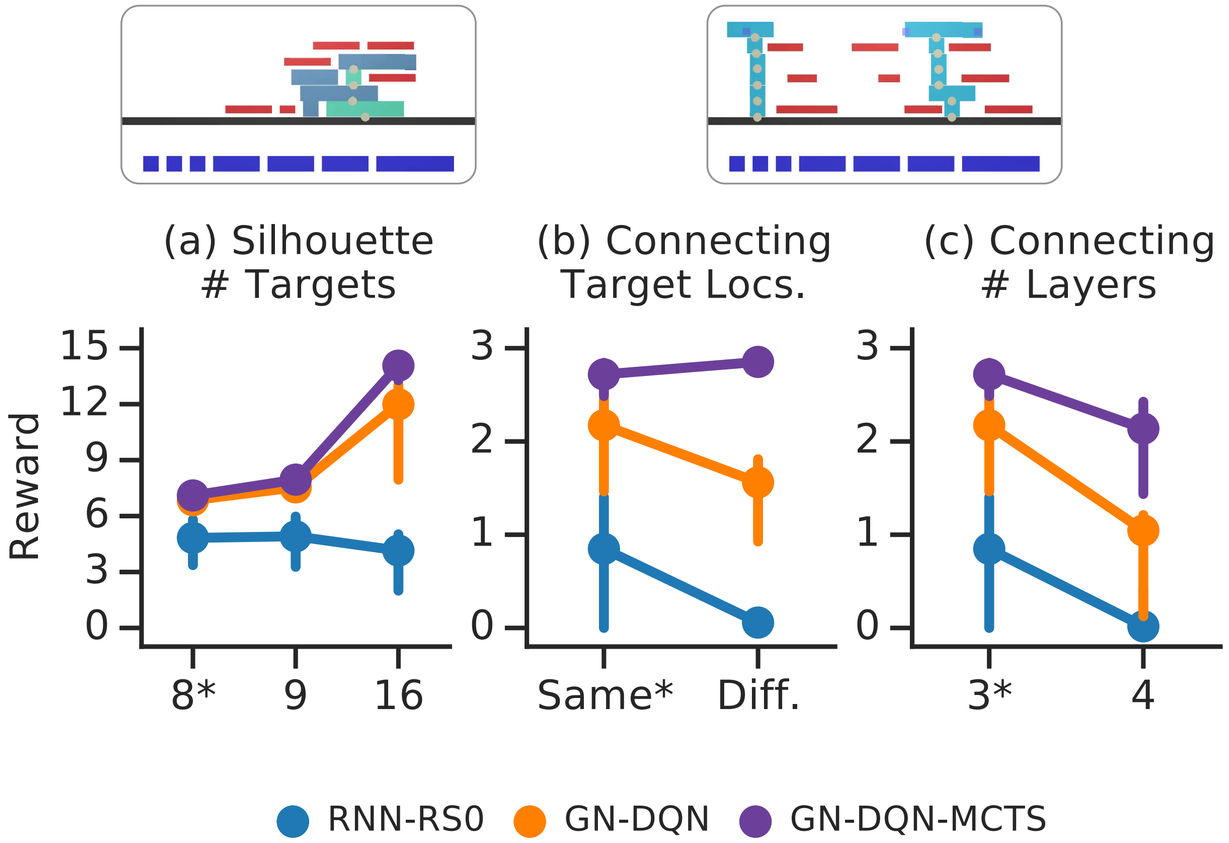

# Additional Results: Generalization

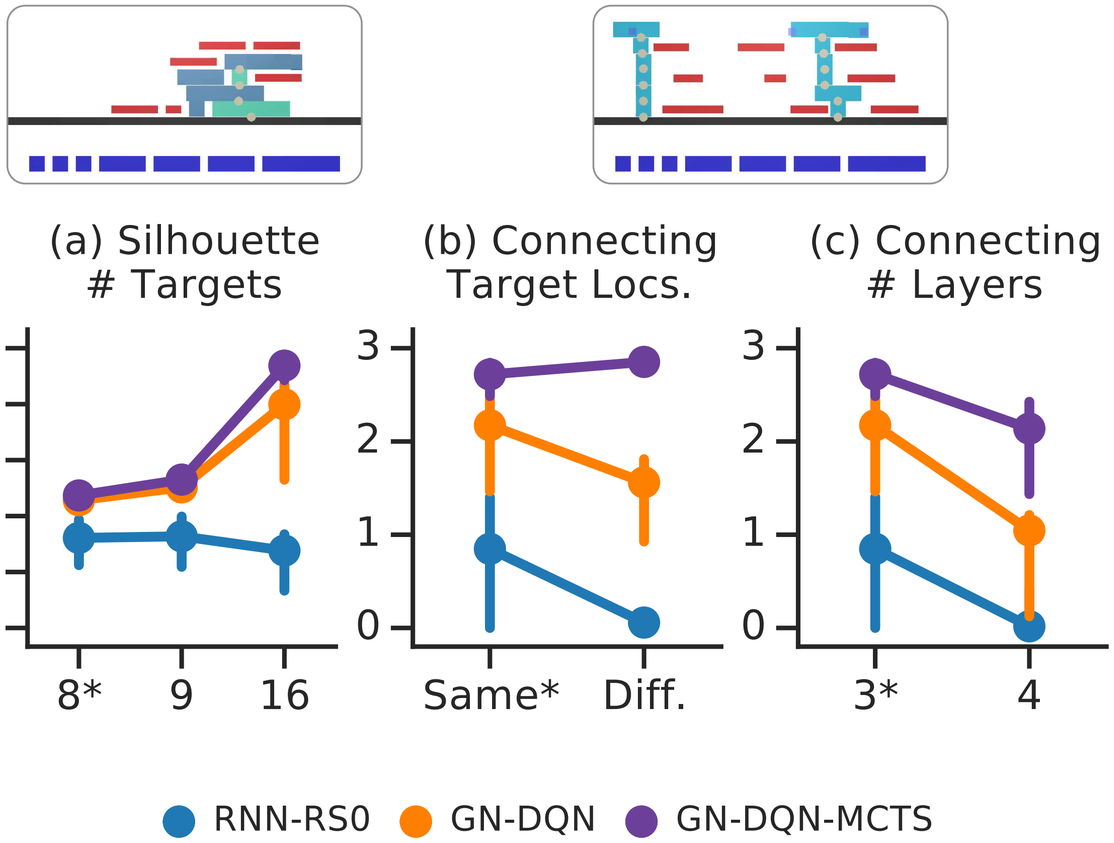


**GN-DQN-MCTS**
**(Average reward: 14.25)**

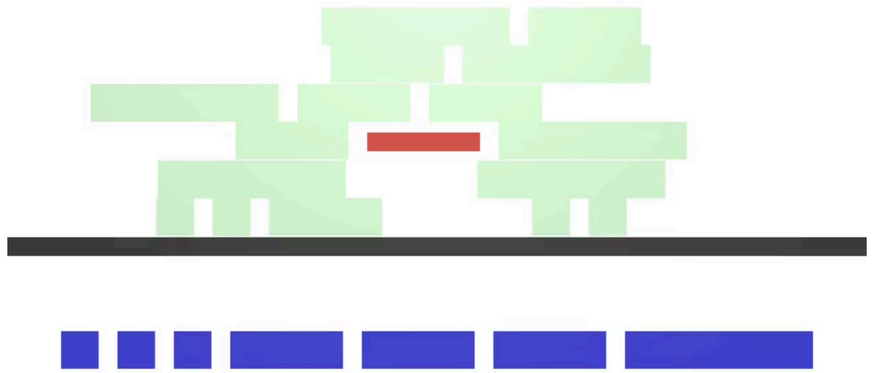

# Additional Results: Generalization

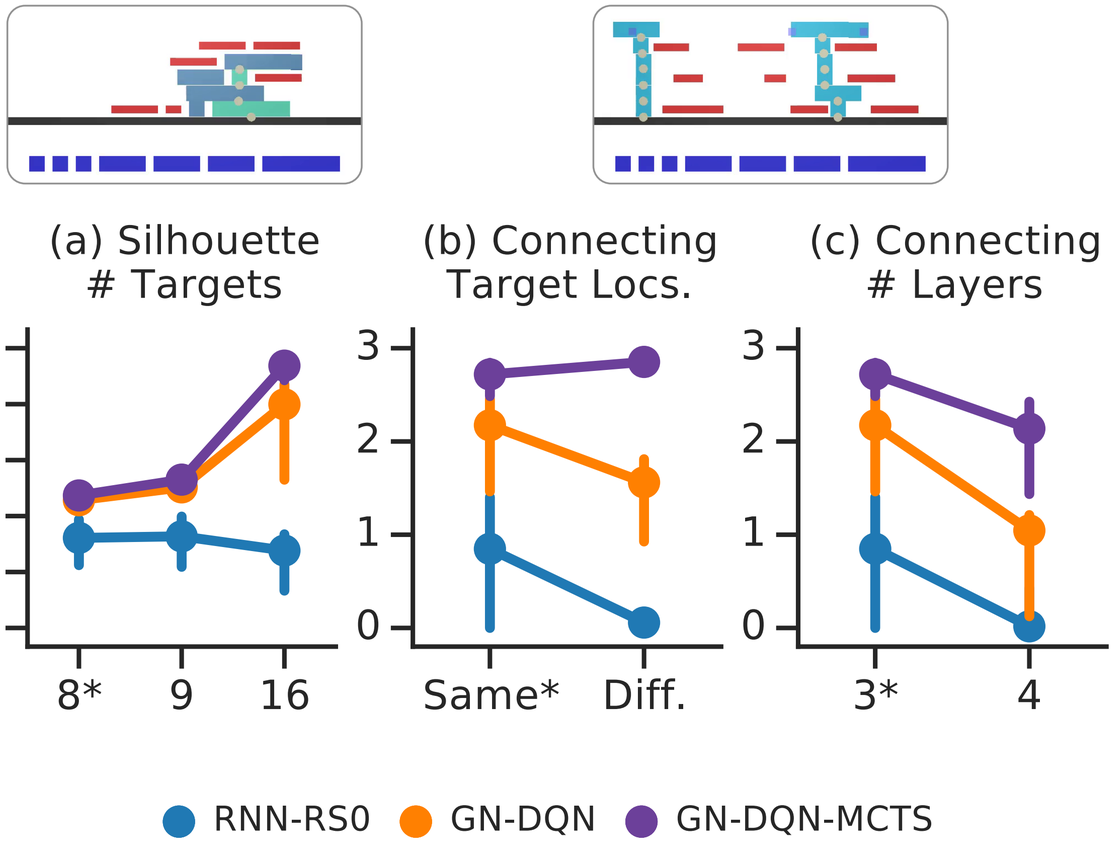


**GN-DQN-MCTS**
**(Average reward: 14.25)**

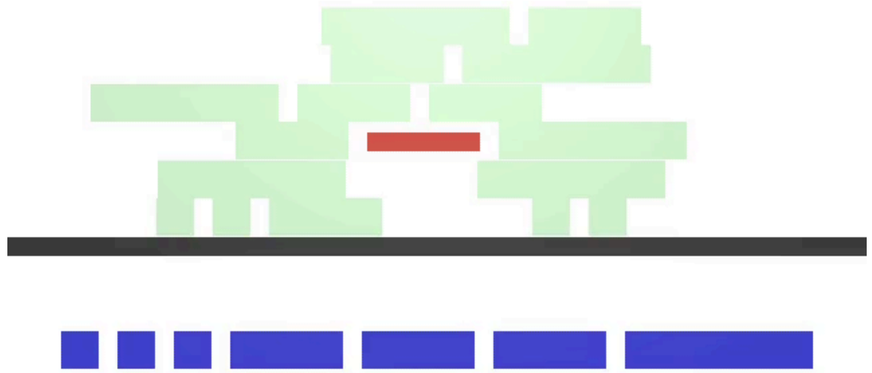

# Key Questions

1. What is the contribution of ***relative*** vs. absolute actions?

2. What is the contribution of ***structured*** representations?

3. What is the contribution of ***planning***?

# Key Questions

1. What is the contribution of ***relative*** vs. absolute actions?

**80-130%** *improvement*

2. What is the contribution of ***structured*** representations?

**25-155%** *improvement*

3. What is the contribution of ***planning***?

**28%** *improvement (Covering Hard)*

1. A suite of challenging ***construction tasks***

2. A new type of ***structured agent*** that uses:

    structured representations
    object-centric relative actions
    combination of model-free and model-based

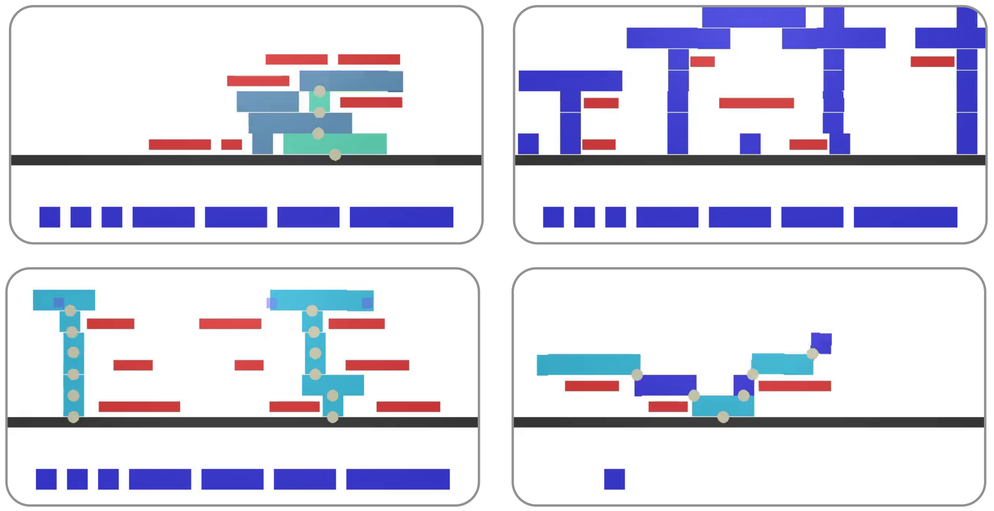

*Only the beginning! Next steps: bigger scenes, harder tasks, more interesting physics, connecting with perception and control, learning environment models, etc.*

**Come see the poster tonight: Pacific Ballroom #36**

# Extra slides

# Graph Networks

Battaglia, Hamrick, Bapst, Sanchez-Gonzalez, Zambaldi, et al. (*arXiv* 2018)

| Edges | Nodes | Globals |
|---|---|---|
| $E$ | $V$ | $\mathbf{u}$ |

**Edge update**

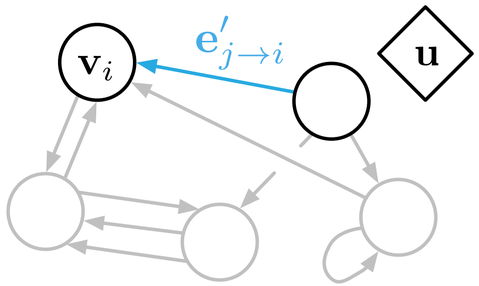


$$\mathbf{e}'_{i\to j} = \phi_e(\mathbf{v}_i, \mathbf{v}_j, \mathbf{e}_{i\to j}, \mathbf{u})$$

**Node update**

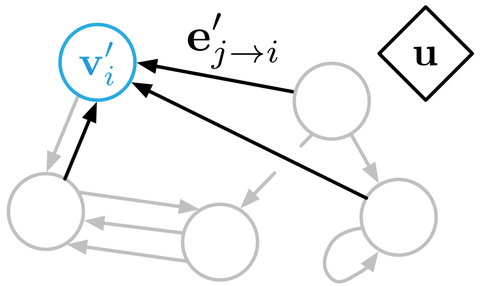


$$\mathbf{v}'_i = \phi_v(\mathbf{v}_i, \sum_j \mathbf{e}'_{j\to i}, \mathbf{u})$$

**Globals update**

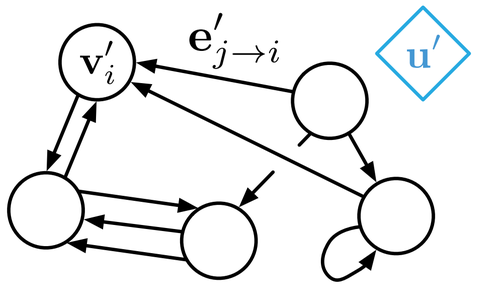


$$\mathbf{u}' = \phi_u(\sum_i \mathbf{v}'_i, \sum_{i,j} \mathbf{e}'_{i\to j}, \mathbf{u})$$

# CNN comparison

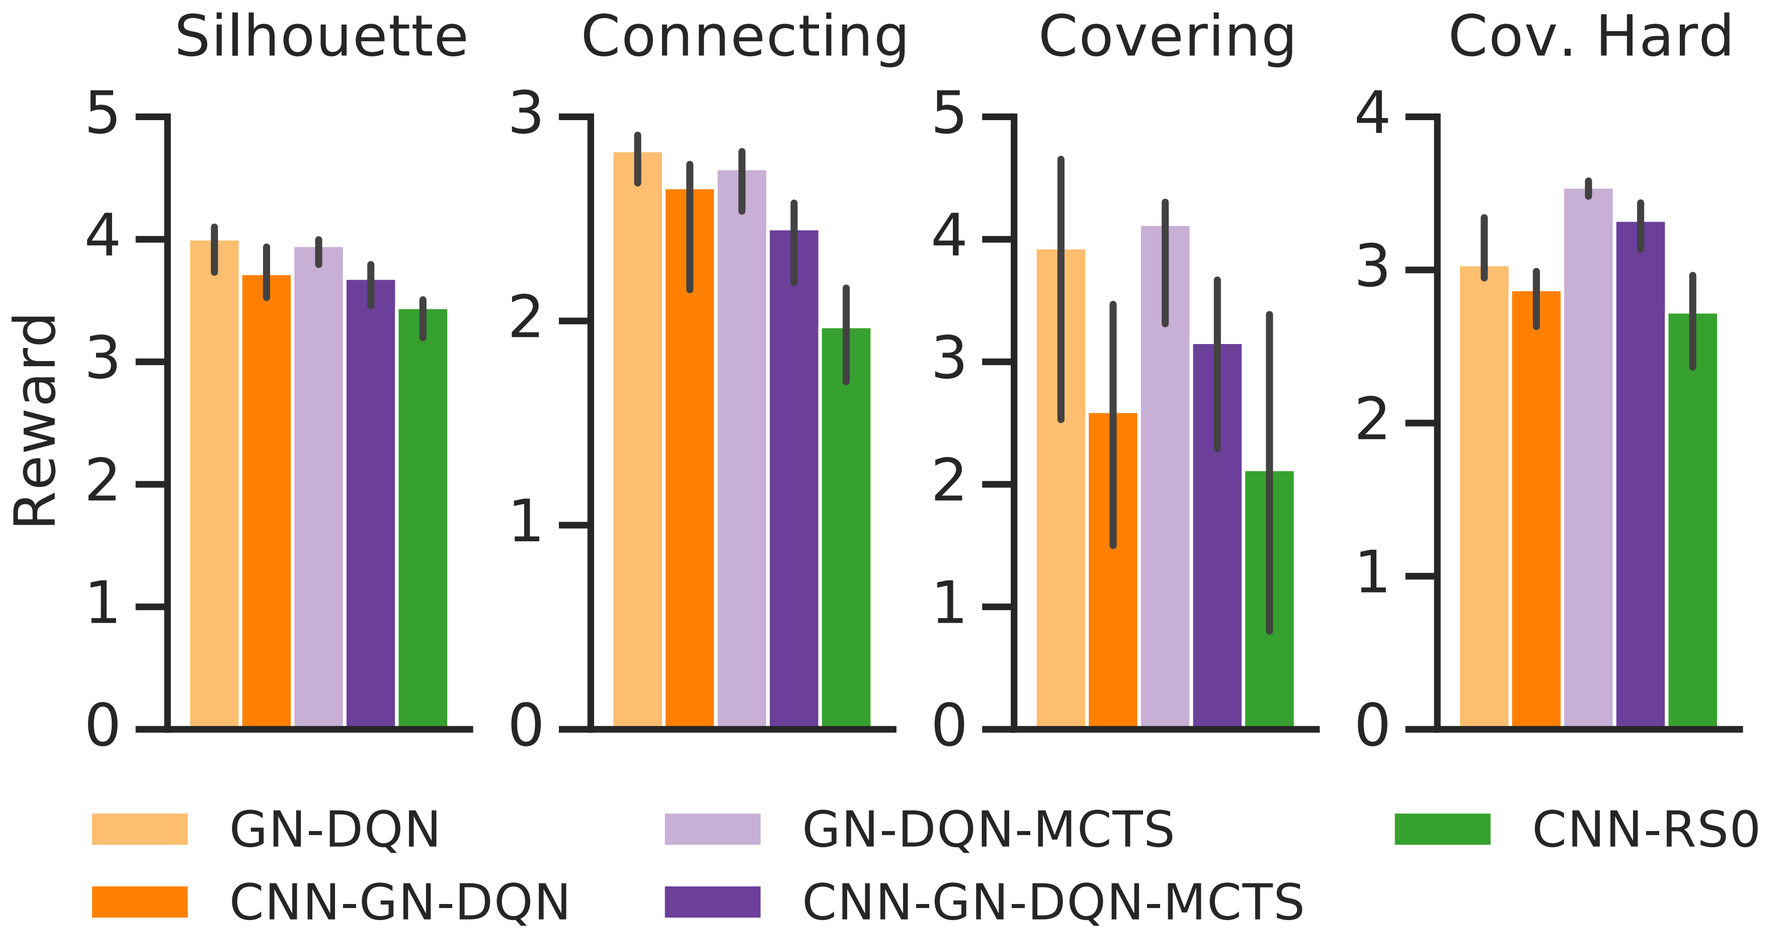

# Full Comparison

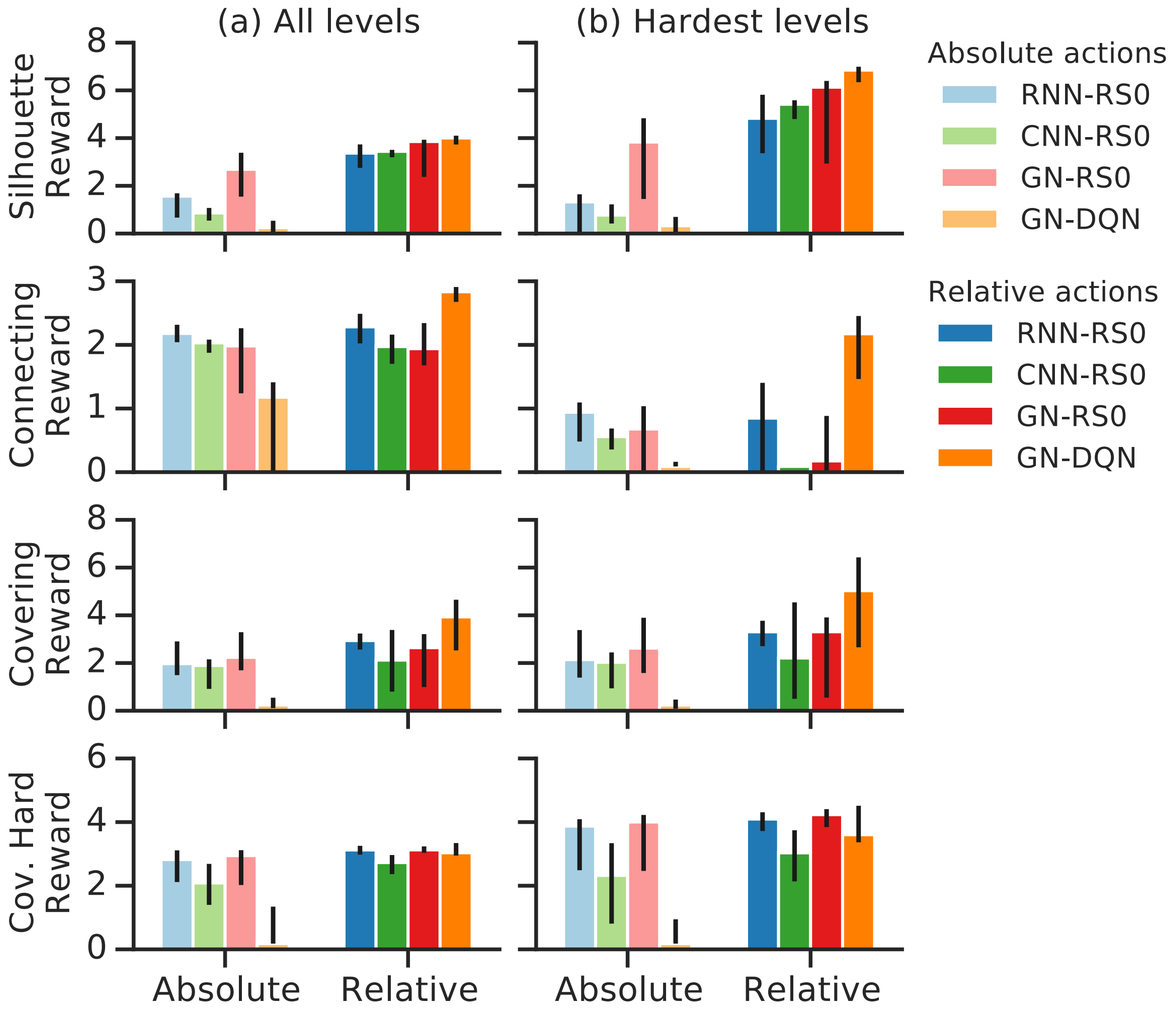

# More Comparisons