

Evolving Graphical Planner: Contextual Global Planning for Vision-and-Language Navigation



Zhiwei Deng



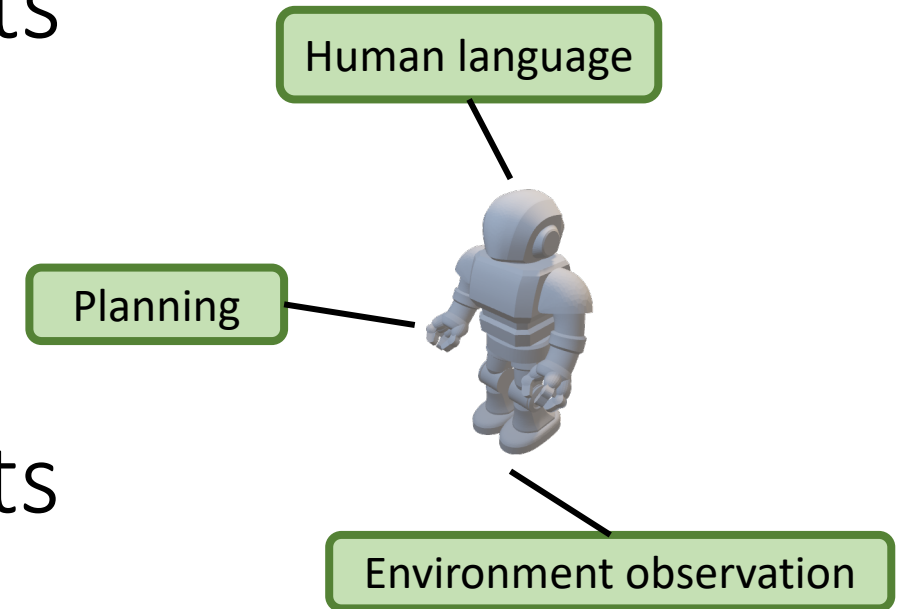
Karthik Narasimhan



Olga Russakovsky

Human communicates with robots
- through language

Robots interact with environments
- perceive visual information
- perform planning, take actions



Vision-and-Language Navigation Task

Unseen environment



Vision-and-Language Navigation Task

Unseen environment




Vision-and-Language Navigation Task

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination

Photorealistic images

Human annotated instructions



Vision-and-Language Navigation Task


Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination

Photorealistic images

Human annotated instructions

Navigation in a room



Start


Target



Vision-and-Language Navigation Task

Unseen environment


Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Vision-and-Language Navigation Task

Unseen environment


Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Vision-and-Language Navigation Task

Unseen environment


Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Challenge 1: Reason over observation and languages

Unseen environment


Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Challenge 1: Reason over observation and languages

Unseen environment


Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Challenge 2: Perform error correction and recovery

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Challenge 2: Perform error correction and recovery

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Deviate from correct path

Challenge 2: Perform error correction and recovery

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Incorrect action



Correct action

Deviate from correct path

Vision-and-Language Navigation Task

Unseen environment



Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Vision-and-Language Navigation Task

Unseen environment


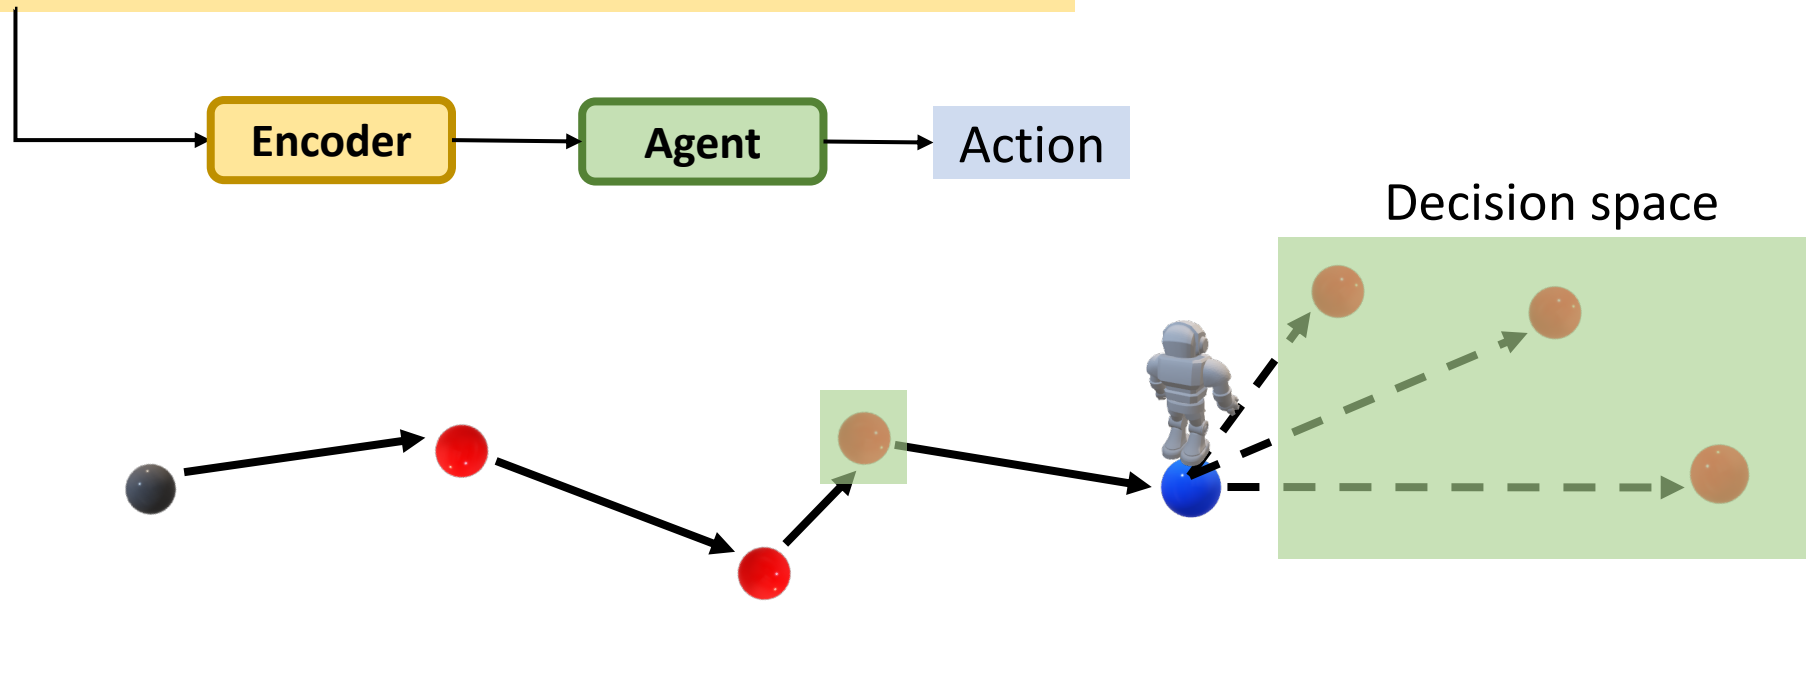
Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Current navigation architectures

Unseen environment

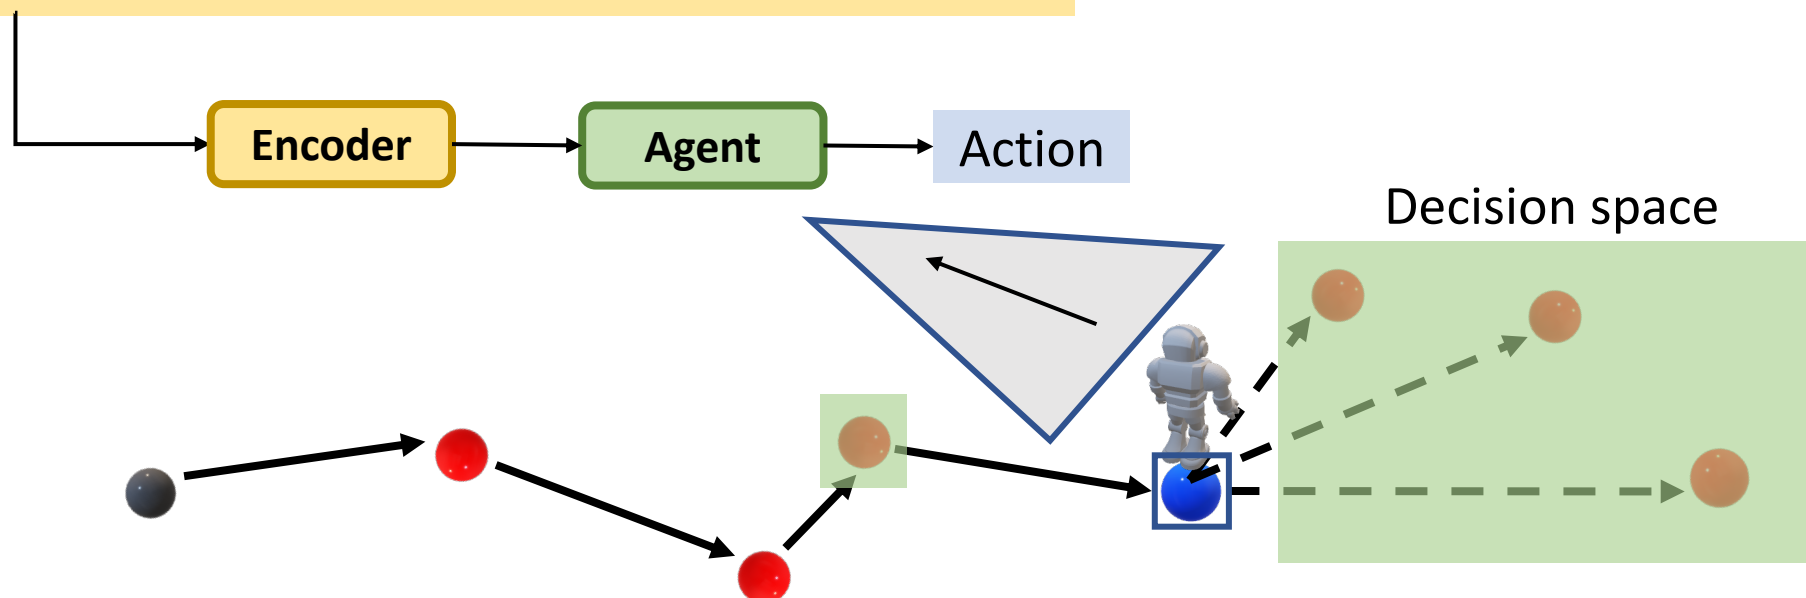

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Current navigation architectures

Unseen environment

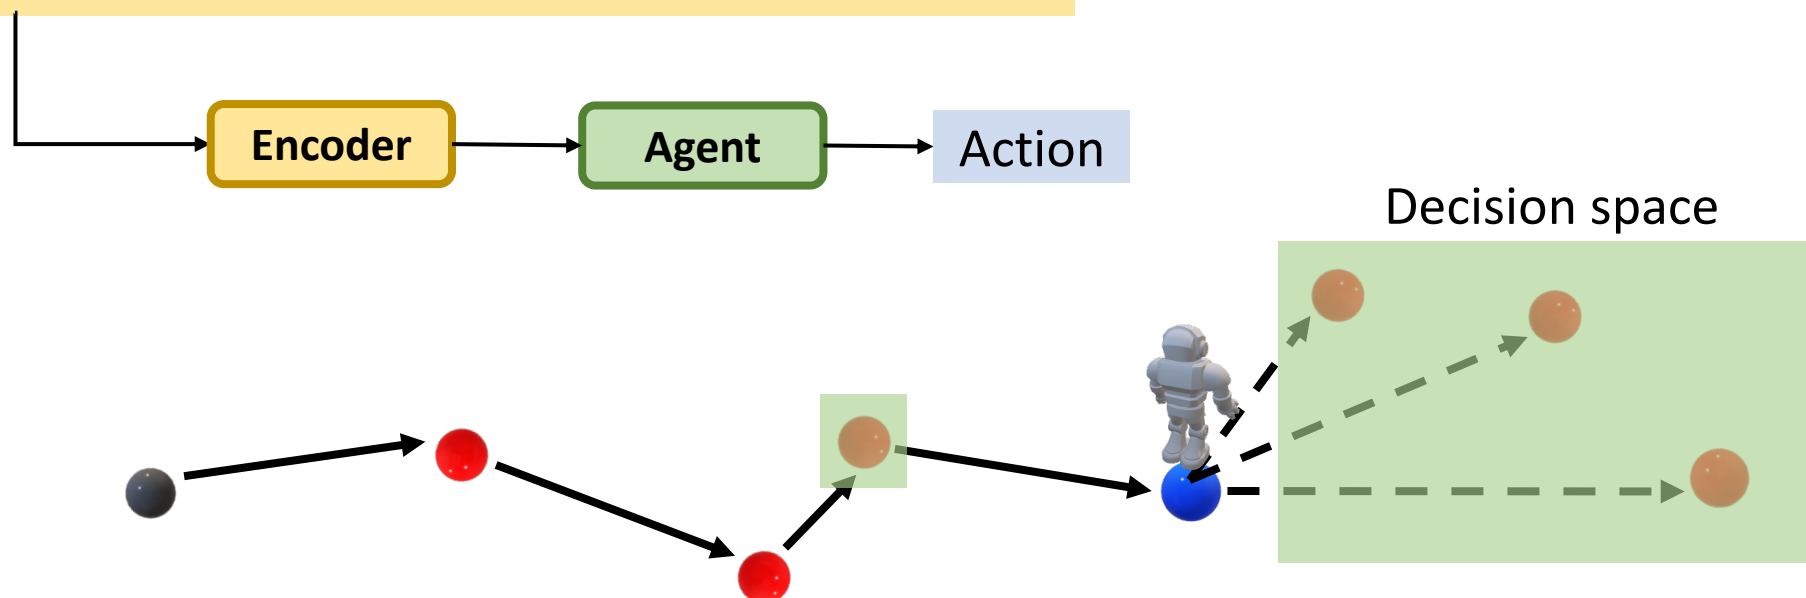
Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Current navigation architectures

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination





**Existing navigation architectures for VLN:
constrained local decision space**

Current navigation architectures

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



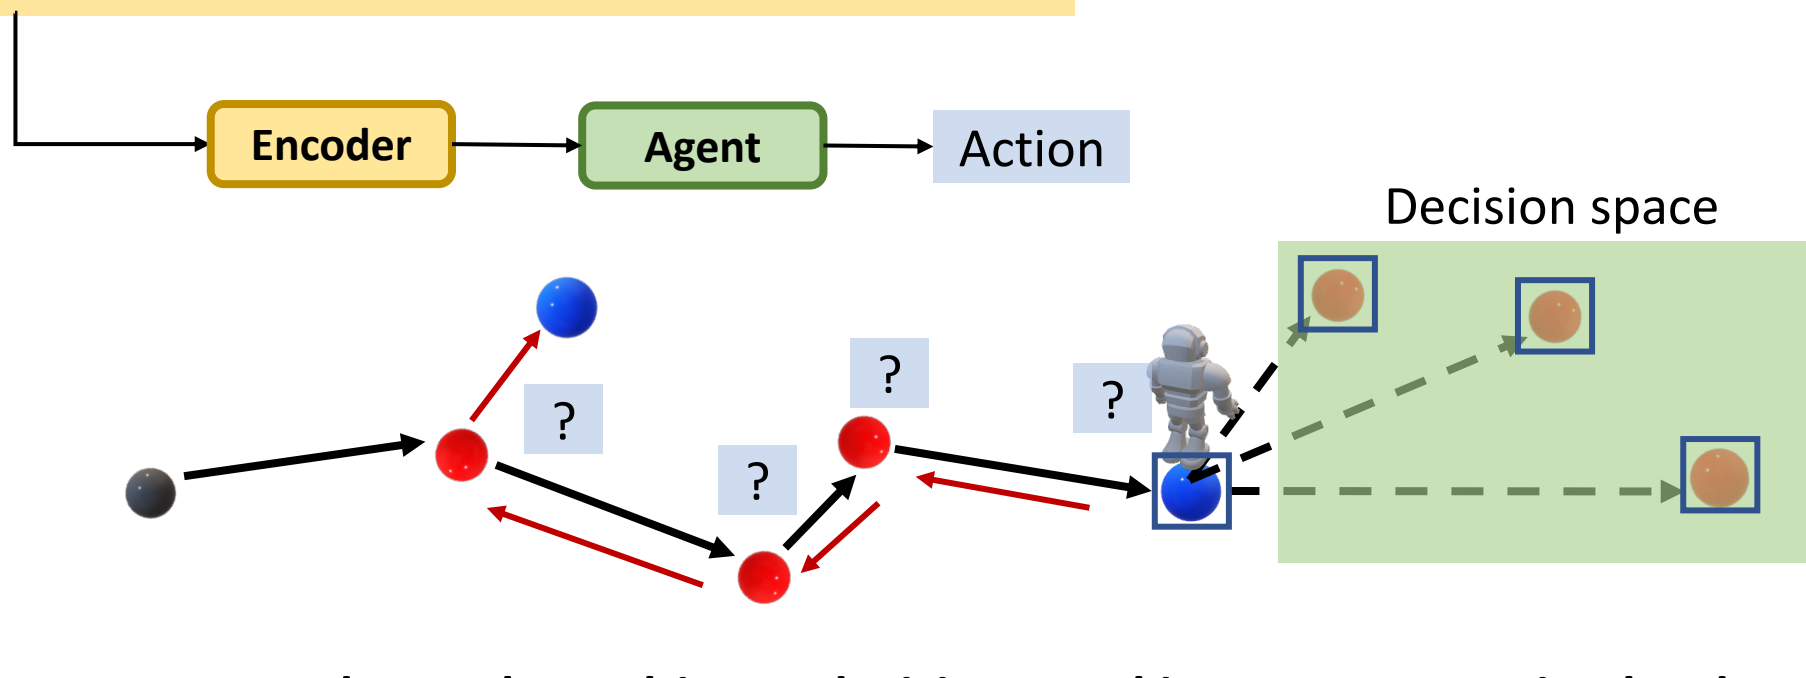
Alignment confusion

Observation + decision space

Current navigation architectures

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination

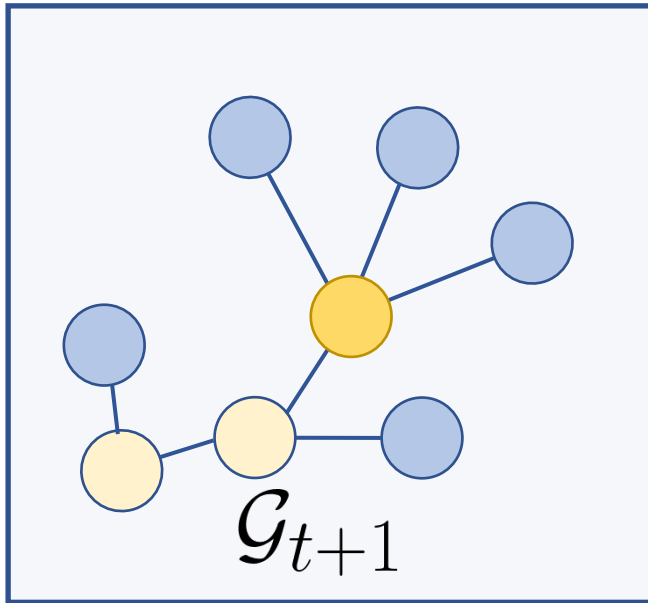


Need to make multi-step decisions, making error correction harder

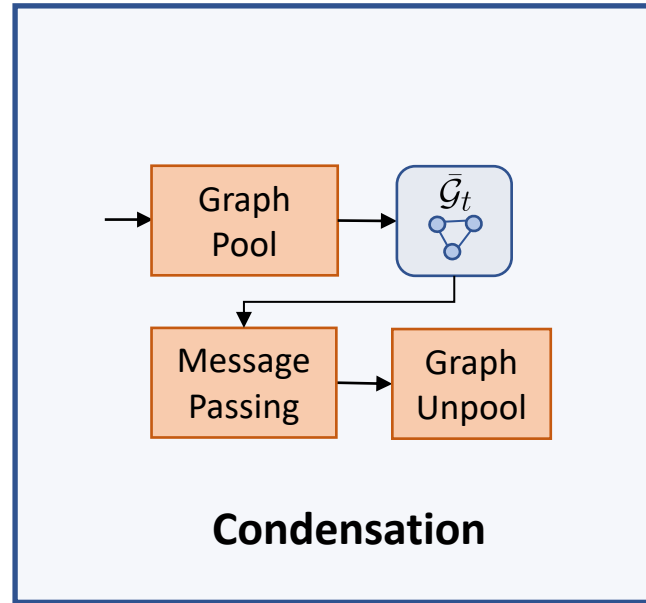
Our work: Evolving Graphical Planner

A differentiable graphical planner


Evolving Graphical Structure



Proxy graphs for planning




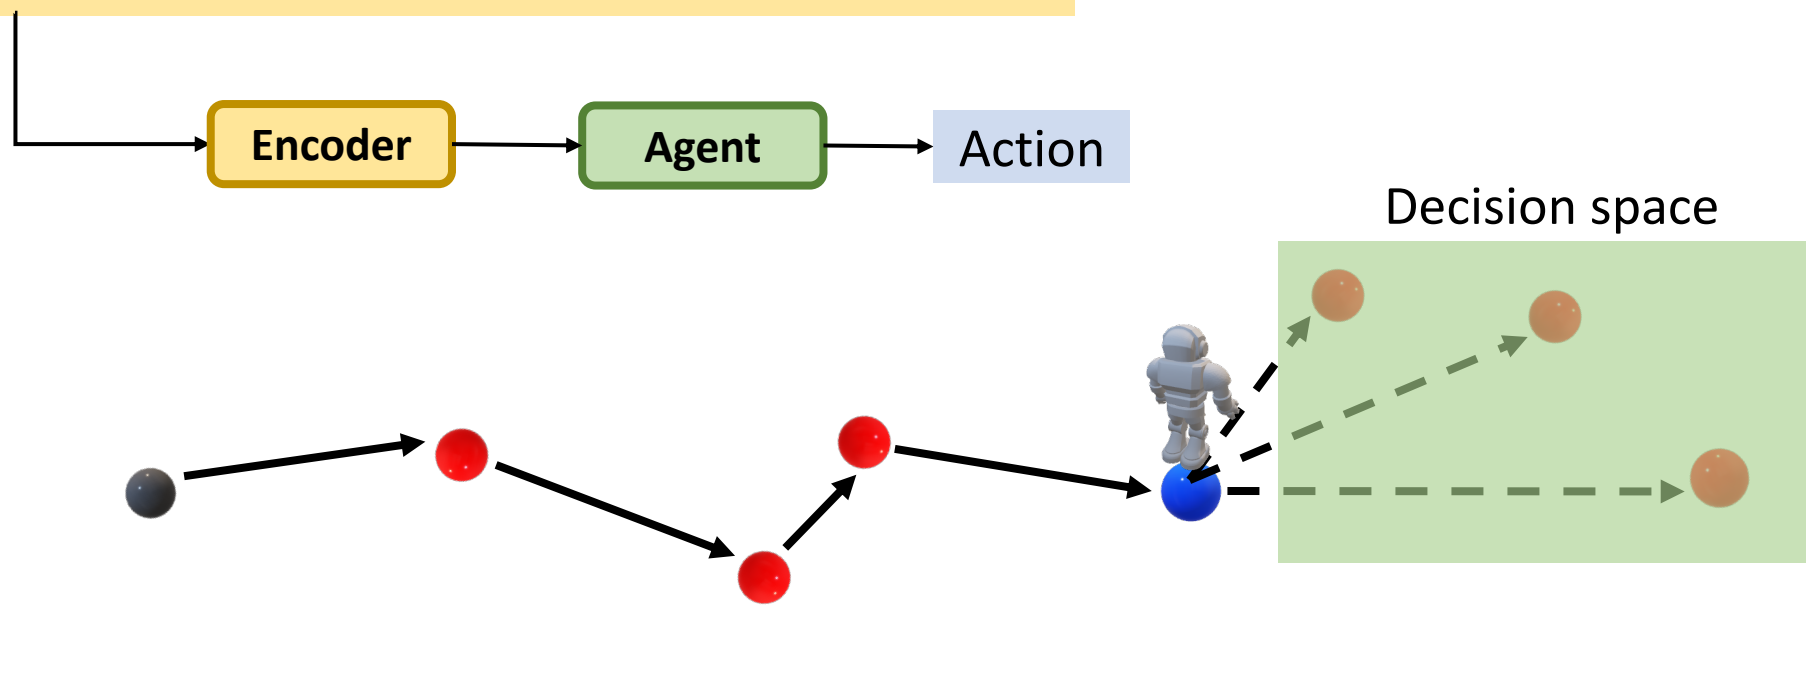
Graph-augmented supervision



Standard VLN navi-agent

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination





Our work: Evolving Graphical Planner

A differentiable graphical planner: global decision space helps


Our work: Evolving Graphical Planner

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination





Topological map




Our work: Evolving Graphical Planner

Unseen environment

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting. Enter the bedroom, you will reach your destination



Topological map





Our work: Evolving Graphical Planner

A differentiable graphical planner: Graphical memory – topological connection + raw feat.

Instructions

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...



Graphical memory

$$G_t = (V_t, E_t)$$

$$v_t^i = (\text{visual}_t^i, \text{angle}_t^i)$$

Observations
(visual + angle)






Our work: Evolving Graphical Planner

A differentiable graphical planner: Graphical memory

Instructions

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting..

Observations
(visual + angle)



Graphical memory

○ Grounding: global alignment


Our work: Evolving Graphical Planner

A differentiable graphical planner: Graphical memory

Instructions

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Observations
(visual + angle)



Graphical memory

- Follow the memorized path
- Decision made in single step
- Easier error correction



Our work: Evolving Graphical Planner

A differentiable graphical planner: Proxy graphs

Instructions

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Observations
(visual + angle)



Ever expanding graph...



Our work: Evolving Graphical Planner

A differentiable graphical planner: Proxy graphs

Instructions

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Observations
(visual + angle)




Ever expanding graph...

**Operate on the full graph:
high planning cost**

Our work: Evolving Graphical Planner


A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$



Ever expanding graph...

$$\tilde{G}_t = (V_t, E_t)$$



Pool
➔

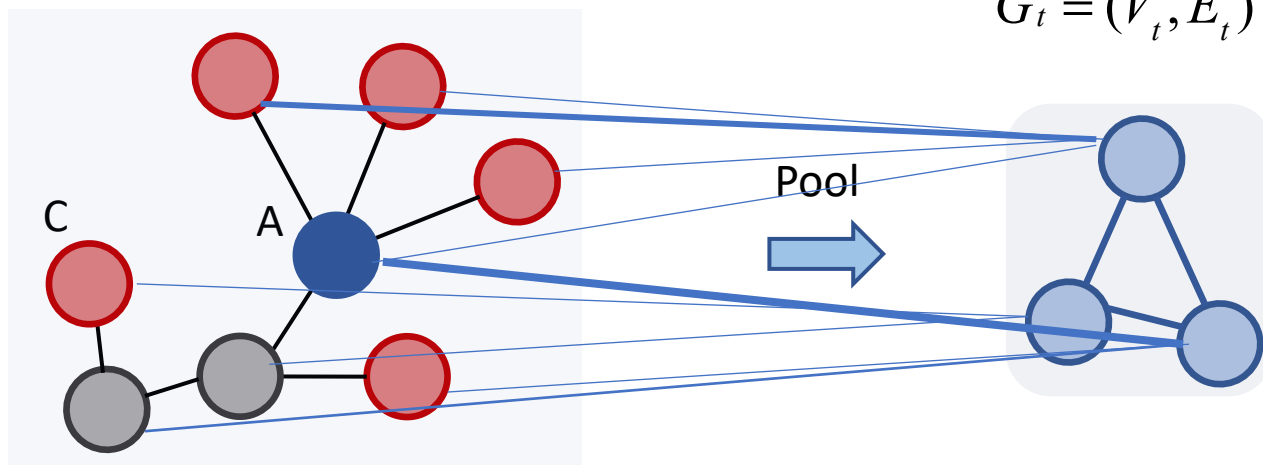
$$\tilde{V}_t = A_t^T V_t$$

$$\tilde{E}_t = A_t^T E_t A_t$$

Our work: Evolving Graphical Planner

A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$




Pooling matrix A_t : soft “attention” or aggregation from the original graph

Our work: Evolving Graphical Planner


A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$



$$\tilde{G}_t = (V_t, E_t)$$

Pool
➔




Pooling matrix A_t : obtained from $f(G_t, \text{language}, \text{agent} - \text{state})$

Our work: Evolving Graphical Planner


A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$



$$\tilde{G}_t = (V_t, E_t)$$

Pool
➔




Neural message passing: $GraphNeuralNetworks(G_t, k = steps)$


Our work: Evolving Graphical Planner

A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$



$$\tilde{G}_t = (V_t, E_t)$$




Un-pool
←

Pooling matrix A_t : transpose as the un-pool matrix

Our work: Evolving Graphical Planner


A differentiable graphical planner: Proxy graphs

$$G_t = (V_t, E_t)$$



Propose next action

$$\tilde{G}_t = (V_t, E_t)$$




Un-pool
←

Our work: Evolving Graphical Planner


A differentiable graphical planner: Proxy graphs – multi-channel

$$G_t = (V_t, E_t)$$

$$\{\tilde{G}_t^k (V_t^k, E_t^k)\}, k = 1, \dots, K$$




Un-pool
←



Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...




Expert trajectories are provided

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...




How to use expert trajectory supervision?

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 1: “teacher forcing”




Expert trajectory dataset: $D = \{(a_1, a_2, \dots, a_{T_i})_i\}$

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 1: “teacher forcing”




$$P(a_1, a_2, \dots, a_T | s) = P(a_1 | s) \prod_{t=2}^T P(a_t | a_1, a_2, \dots, a_{t-1}, s)$$

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 1: “teacher forcing” – drifting issue in unseen data




Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 2: “student forcing”




Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 2: “student forcing” – generate new supervision (shortest path)



$$D^* = \{(a_1^*, a_2^*, \dots, a_{T_i}^*)_i\}$$

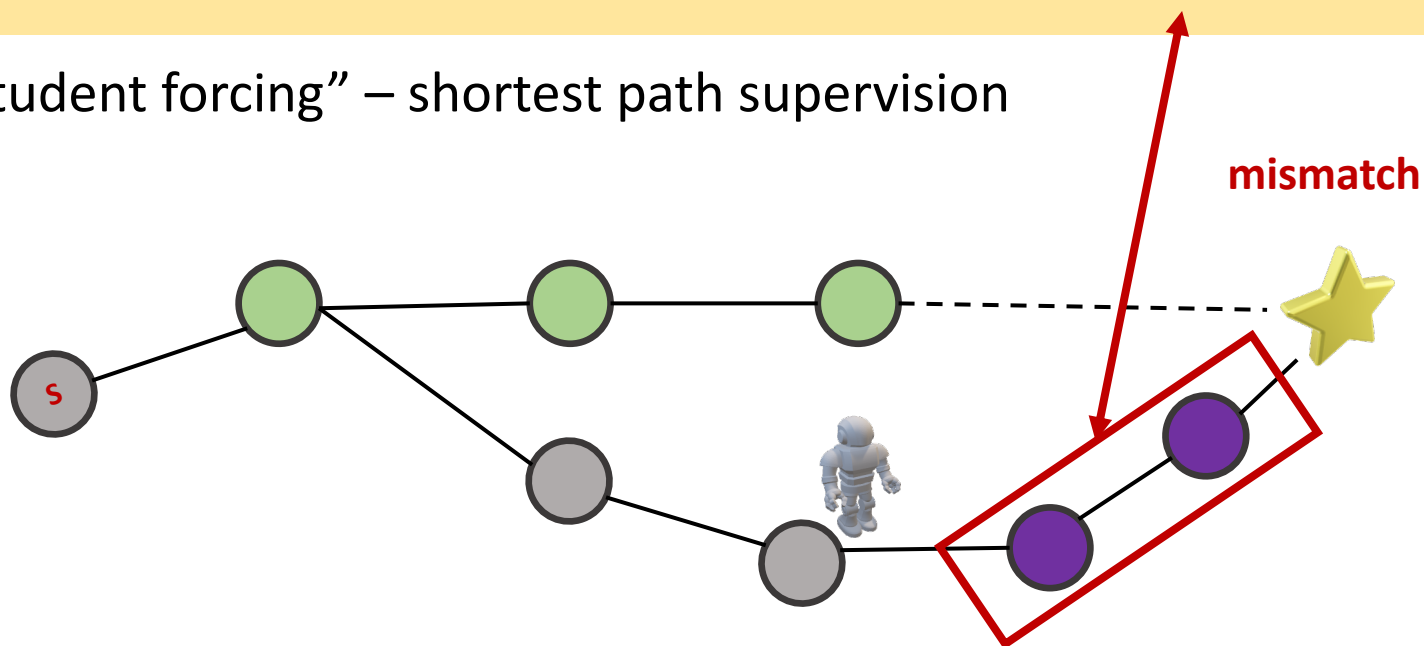
$$D \cup D^*$$

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 2: “student forcing” – shortest path supervision




Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 2: “student forcing” – graph augmented supervision

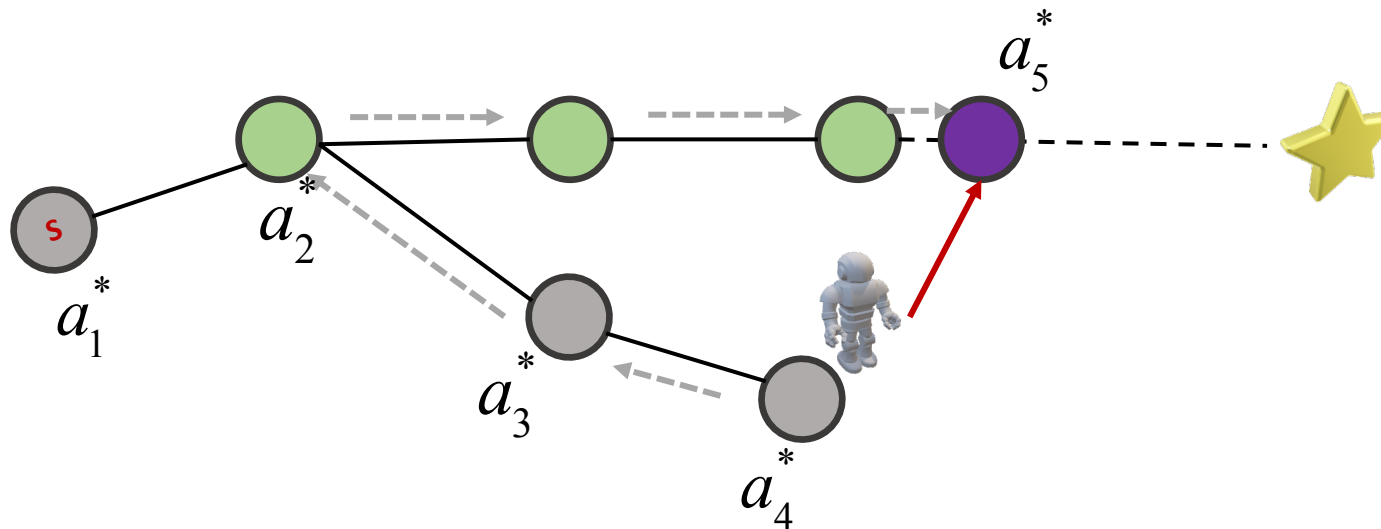


Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

Option 2: “student forcing” – graph augmented supervision



$$D^* = \{(a_1^*, a_2^*, \dots, a_{T_i}^*)_i\}$$


$$D \cup D^*$$

Our work: Evolving Graphical Planner

A differentiable graphical planner: how to supervise the imitation learner?

Facing the end of the bed, take an immediate right and exit the bedroom through the open doorway. Walk straight until you see a large red painting. At the painting make a turn towards and go through the doorway on the right of the painting...

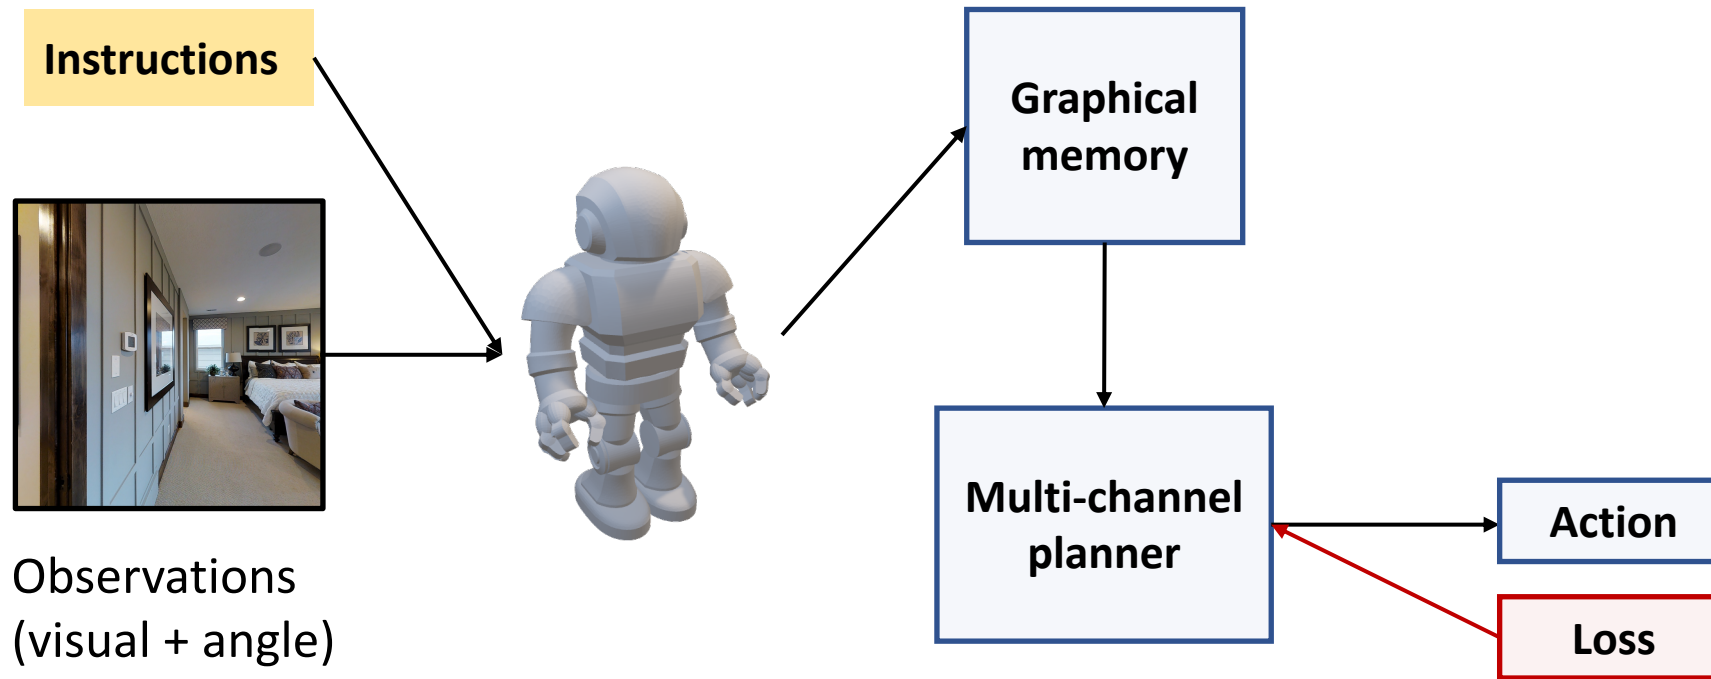
Option 2: “student forcing” – graph augmented supervision



- Ground truth always exists
- No mismatch problem
- No need to access the ENV

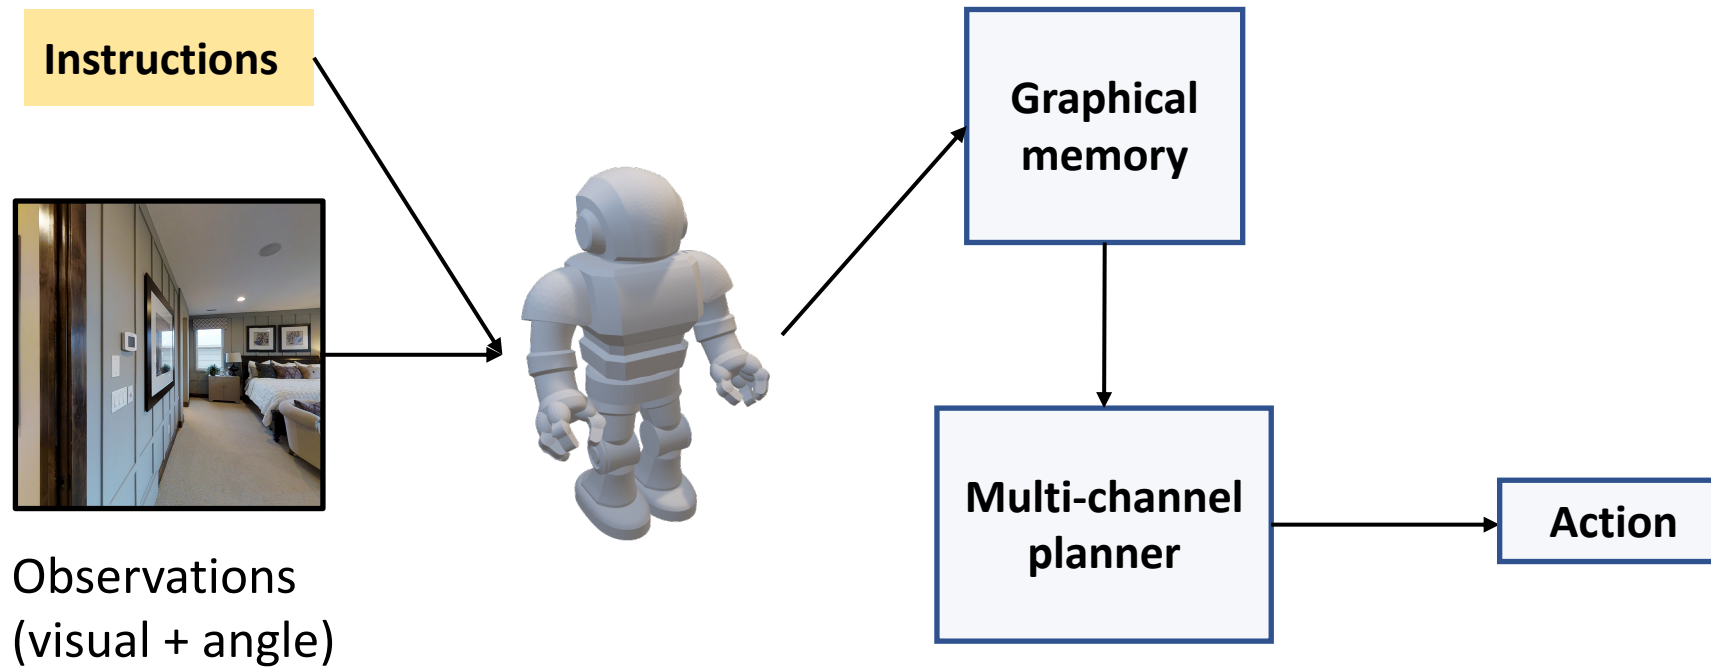
Our work: Evolving Graphical Planner

A differentiable graphical planner: full training process




Our work: Evolving Graphical Planner

A differentiable graphical planner: test inference matches the training




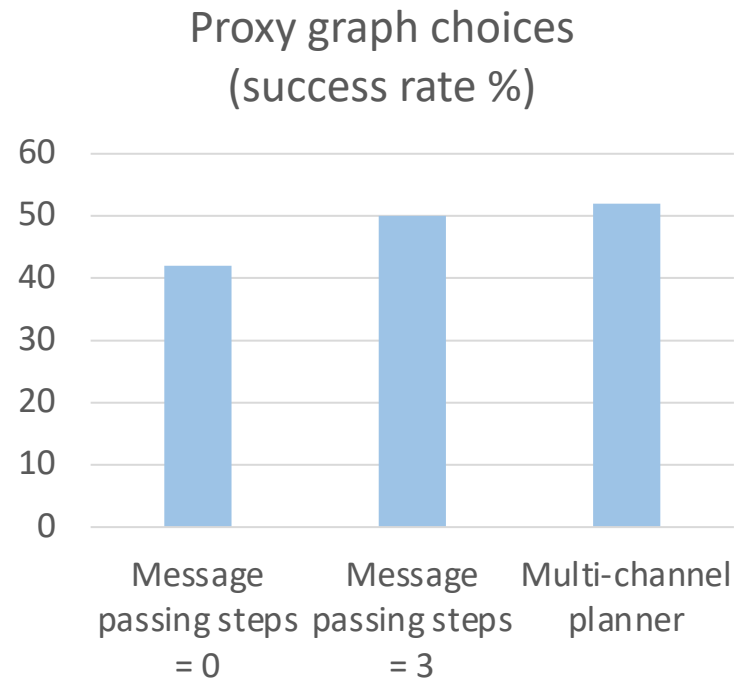

Experiments

- Room-to-Room (R2R): all trajectories are generated through shortest-path, emphasize on **goal reaching**



Contribution of each component

- Room-to-Room (R2R): all trajectories are generated through shortest-path, emphasize on goal reaching



The global decision space, the planner and the new supervision strategy help on navigation success rate

Compare to existing backbones


- Room-to-Room (R2R): all trajectories are generated through shortest-path, emphasize on goal reaching

Models	Type	Val unseen				Test			
		NE ↓	SR% ↑	SPL% ↑	OSR% ↑	NE ↓	SR% ↑	SPL% ↑	OSR% ↑
Seq2Seq [1]	IL	6.01	39	-	53	7.81	22	-	28
Ghost [25]	IL	7.20	35	31	44	7.83	33	30	42
SF* [2]	IL	6.62	36	-	45	6.62	35	28	44
RCM* [47]	IL+RL	5.88	43	-	52	6.12	43	38	50
Monitor [14]	IL	5.98	44	30	58	-	-	-	-
Monitor* [14]	IL	5.52	45	32	56	5.67	48	35	59
Regretful [19]	IL	5.36	48	37	61	-	-	-	-
Regretful* [19]	IL	5.32	50	41	59	5.69	48	40	56
Baseline agent	IL	6.20	43	36	52	-	-	-	-
EGP (ours)	IL	5.34	52	41	65	-	-	-	-
EGP* (ours)	IL	4.83	56	44	64	5.34	53	42	61

We outperform previous backbone architecture

Room-for-room with pure imitation learning

- Room-for-Room (R4R): measured by Coverage weighted by Length Score (CLS), normalized dynamic time warping (DTW), Success rate weighted normalized Dynamic Time Warping (SDTW), emphasize on **path following**



Room-for-room with pure imitation learning

- Room-for-Room (R4R): measured by Coverage weighted by Length Score (CLS), normalized dynamic time warping (DTW), Success rate weighted normalized Dynamic Time Warping (SDTW)

Models	Type	PL	NE ↓	SR [%] ↑	CLS ↑	nDTW ↑	SDTW ↑
Random	-	23.6	10.4	13.8	22.3	18.5	4.1
Speaker-Follower[18]	IL+RL	19.9	8.47	23.8	29.6	-	-
RCM + goal-oriented[18]	IL+RL	32.5	8.45	28.6	20.4	26.9*	11.4*
RCM + fidelity-oriented[18]	IL+RL	28.5	8.08	26.1	34.6	30.4*	12.6*
PTA low-level[53]	IL+RL	10.2	8.19	27.0	35.0	20.0	8.0
PTA high-level[53]	IL+RL	17.7	8.25	24.0	37.0	32.0	10.0
EGP (ours)	IL	18.3	8.0	30.2	44.4	37.4	17.5

We achieve the state-of-the-art using **pure imitation learning**

Contributions

- A differentiable graphical planner that extends the decision space globally
- A new supervision strategy for training imitation agent in navigation
- Introduce proxy graphs for improving the efficiency of planning

Email: zhiweid@cs.princeton.edu