

Vision-and-Dialog Navigation

Jesse Thomason
University of Washington



Daniel
Gordon



Yonatan
Bisk



Michael
Murray



Maya
Cakmak



Luke
Zettlemoyer

Bringing Robots from Industrial to Human Spaces

Industrial



Bringing Robots from Industrial to Human Spaces

Industrial



Human



Bringing Robots from Industrial to Human Spaces

Industrial



Natural Language



Navigation



Additional Safety



Robust Perception



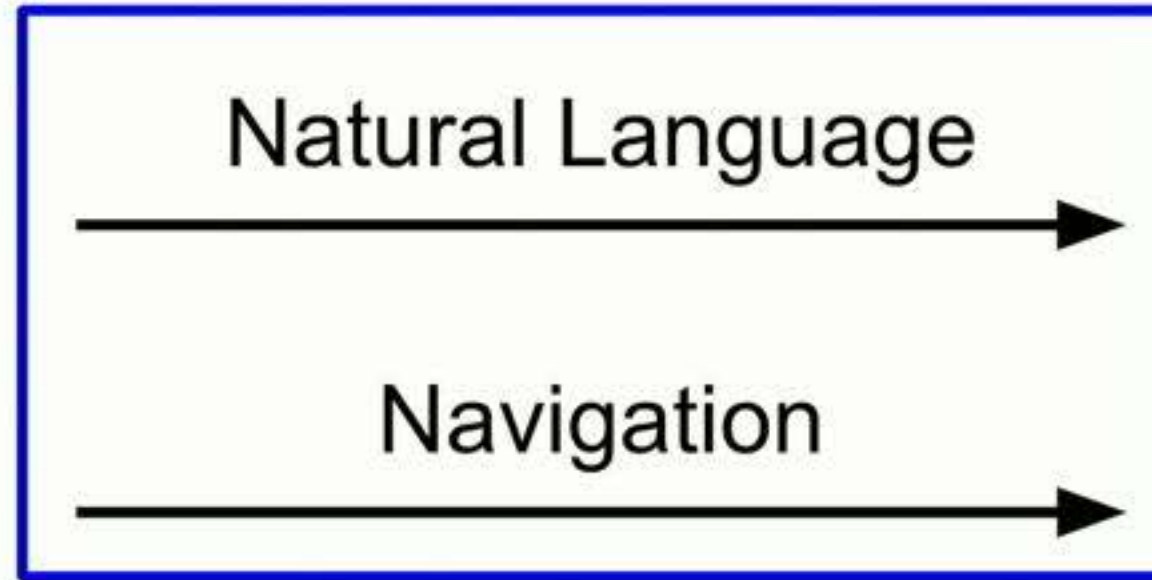
⋮

Human




Bringing Robots from Industrial to Human Spaces

Industrial




Human




⋮



1) In Human Spaces, We Use Natural Language.



1) In Human Spaces, We Use Natural Language.



2) Human Spaces Are *Dynamic* and can be *Unseen*



2) Human Spaces Are *Dynamic* and can be *Unseen*



2) Human Spaces Are *Dynamic* and can be *Unseen*



Outline

Outline

- Language grounding in visual environments

Outline

- Language grounding in visual environments
 - For navigation

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - Navigation from dialog history

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - Navigation from dialog history
- Next steps

Connecting Language and Vision

Connecting Language and Vision

- Common paradigm:

Connecting Language and Vision

- Common paradigm:
 - **Inputs:**
 - Language tokens (e.g., question)

Connecting Language and Vision

- Common paradigm:
 - **Inputs:**
 - Language tokens (e.g., question)
 - Visual context (e.g., photograph)

Connecting Language and Vision

- Common paradigm:
 - **Inputs:**
 - Language tokens (e.g., question)
 - Visual context (e.g., photograph)
- Visual contexts differ in quality across datasets.

Connecting Language and Vision

- Common paradigm:
 - **Inputs:**
 - Language tokens (e.g., question)
 - Visual context (e.g., photograph)
- Visual contexts differ in quality across datasets.
- Output can be a single classification or a sequence.

Visual Fidelity

Rendered



Photorealistic

Visual Context

Static



Dynamic



CLEVR

[Johnson et al., CVPR'17]

Visual Fidelity

Rendered



Photorealistic



Visual Context

Static



Dynamic



Visual Fidelity

Rendered



Photorealistic



CLEVR

[Johnson et al., CVPR'17]

Static



Dynamic



VQA

[Antol et al., CVPR'15]

Visual Context

Visual Fidelity


Rendered

Photorealistic



CLEVR

[Johnson et al., CVPR'17]



Instruction Following

[Chen and Mooney, AAAI'11]

Visual Context

Static

Dynamic



VQA

[Antol et al., CVPR'15]

Visual Fidelity

Rendered



Photorealistic




Visual Context

Static




Dynamic



CLEVR

[Johnson et al., CVPR'17]



Instruction Following

[Chen and Mooney, AAAI'11]



VQA

[Antol et al., CVPR'15]



Room-to-Room

[Anderson et al., CVPR'18] 8

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - Navigation from dialog history
- Next steps


Vision-and-Language Navigation



Exit the bathroom. Turn left and exit the room using the door on the left. Wait there.


“Turn around and exit the library, head down the...”

- Low-level instructions for moving through the environment.

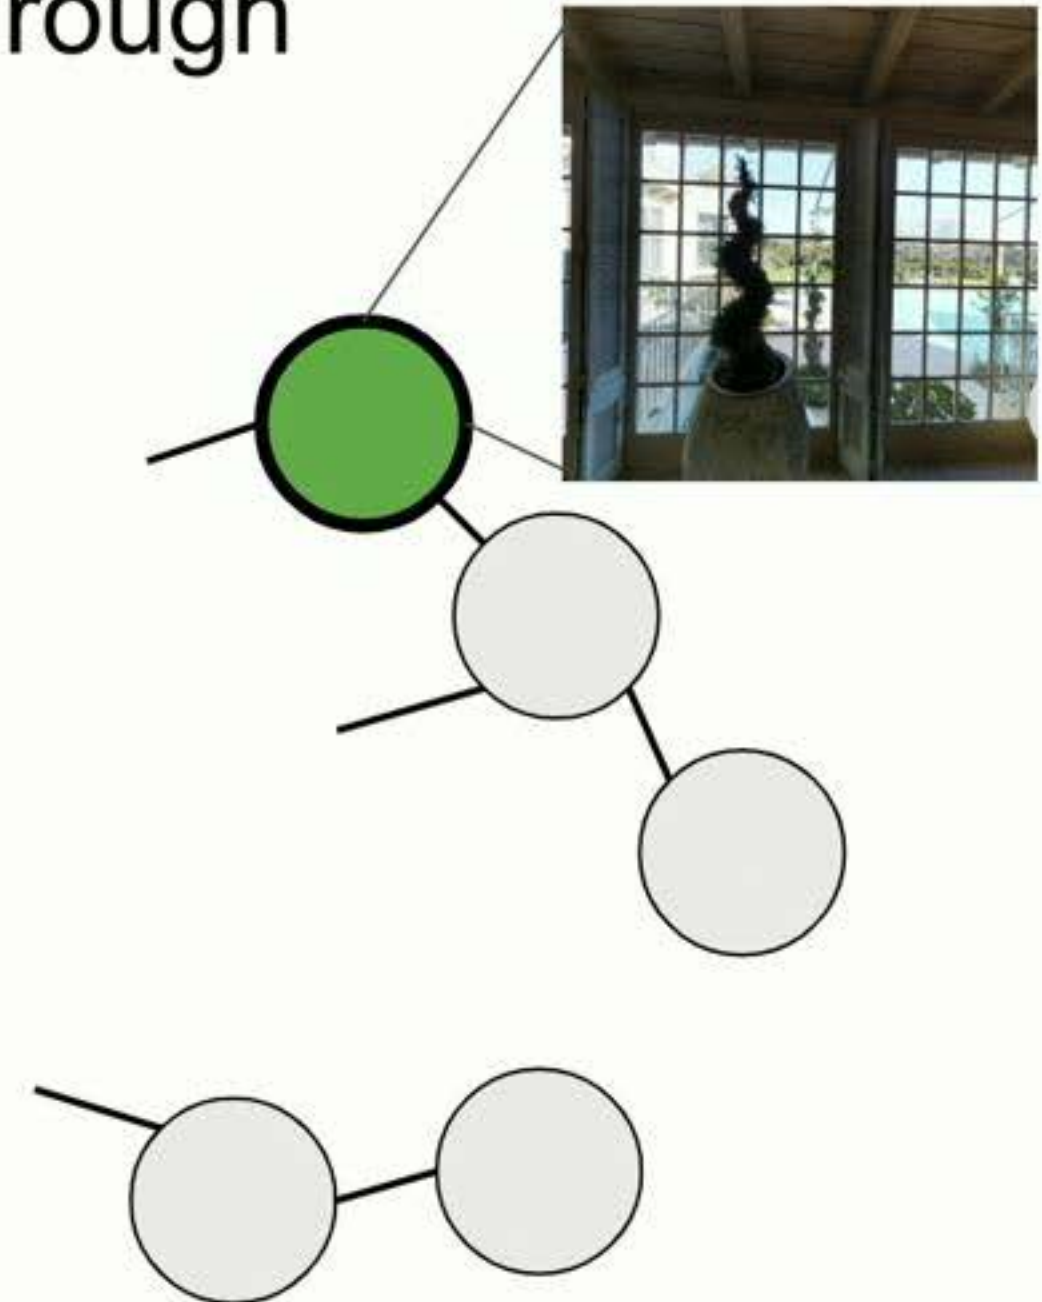


“Turn around and exit the library, head down the...”

- Low-level instructions for moving through the environment.

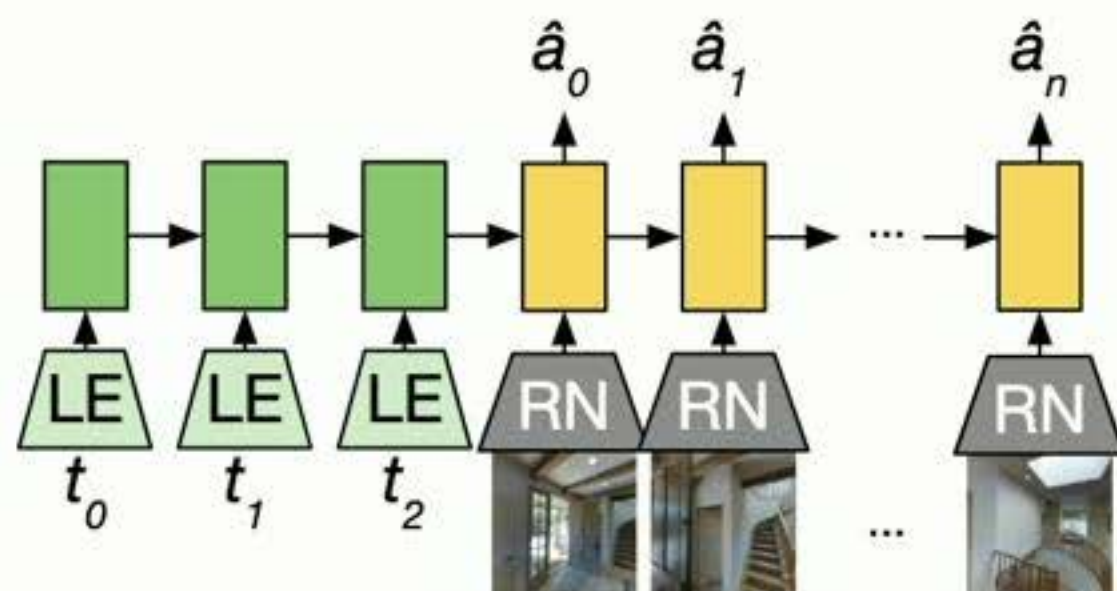


...




Sequence-to-Sequence Model

- Encode the language tokens.
- Decode a sequence of actions to take in the environment.
- At every timestep, receive a new visual observation.




Sequence-to-Sequence Model

- Encode the language tokens.
- Decode a sequence of actions to take in the environment.
- At every timestep, receive a new visual observation.



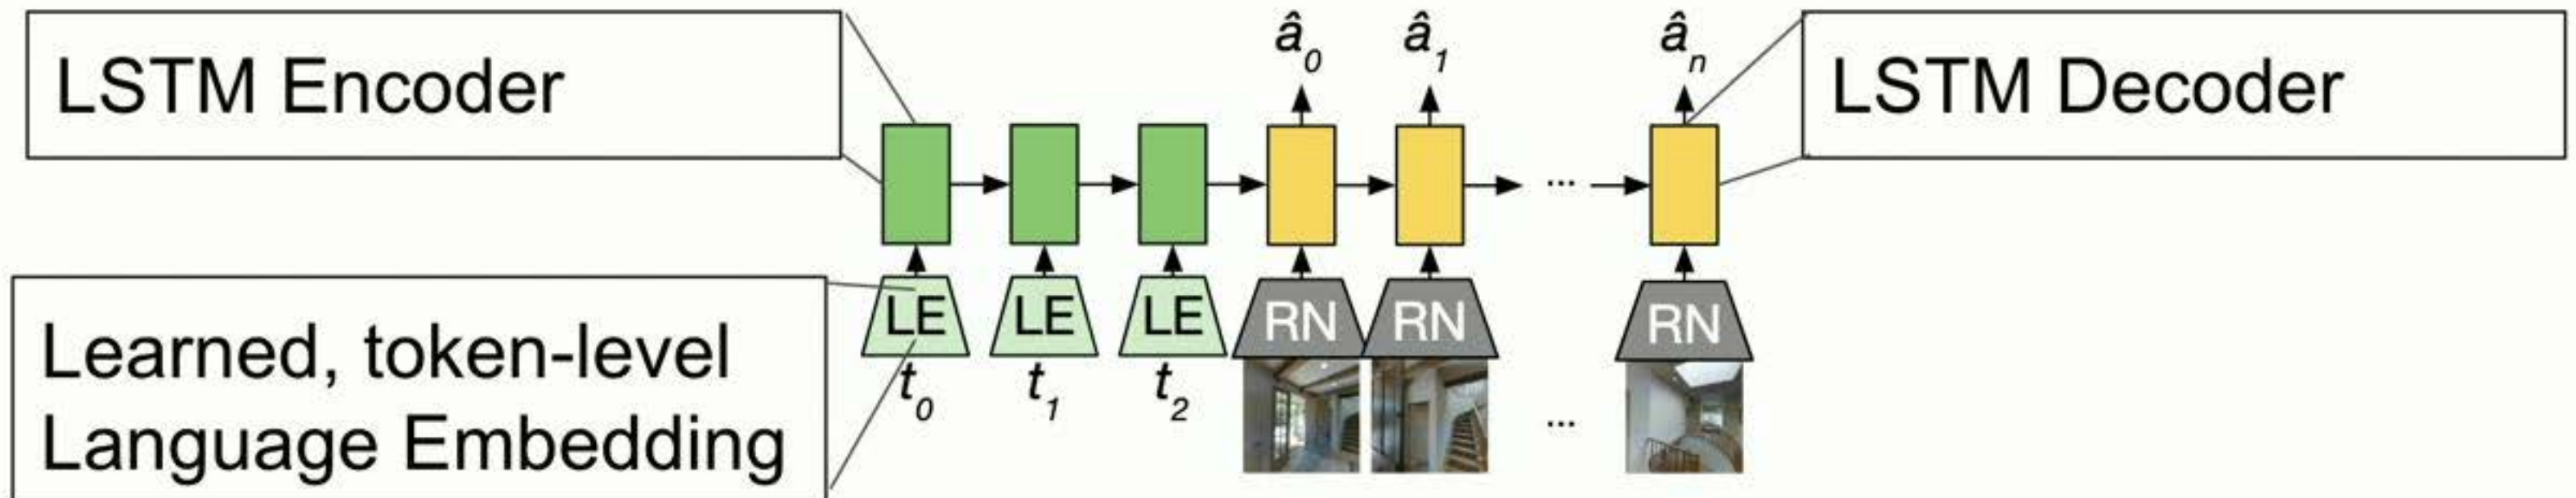
Sequence-to-Sequence Model

- Encode the language tokens.
- Decode a sequence of actions to take in the environment.
- At every timestep, receive a new visual observation.



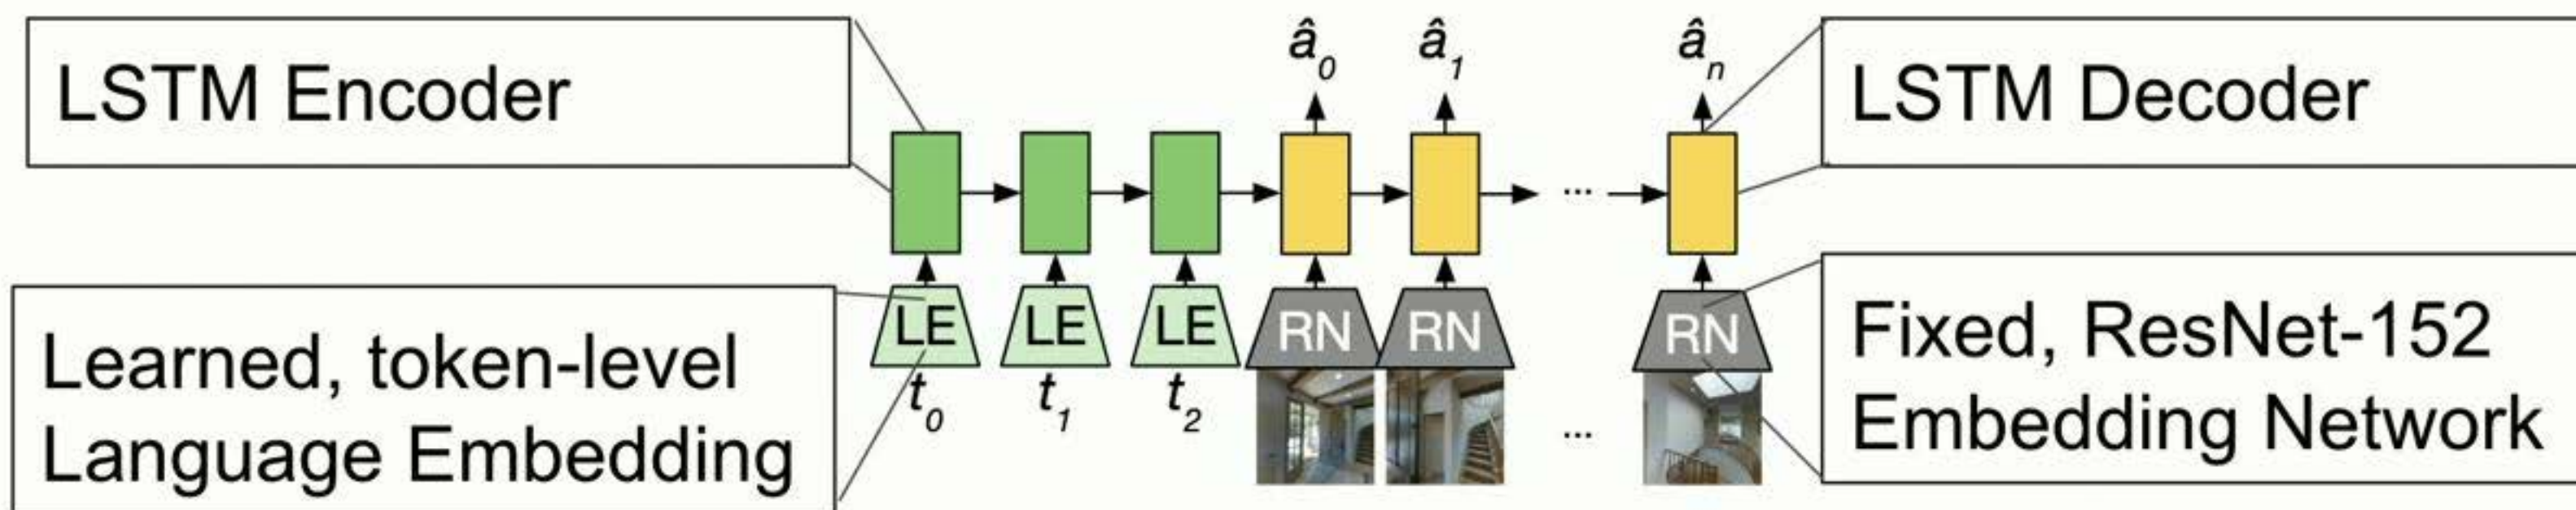
Sequence-to-Sequence Model

- Encode the language tokens.
- Decode a sequence of actions to take in the environment.
- At every timestep, receive a new visual observation.



Sequence-to-Sequence Model

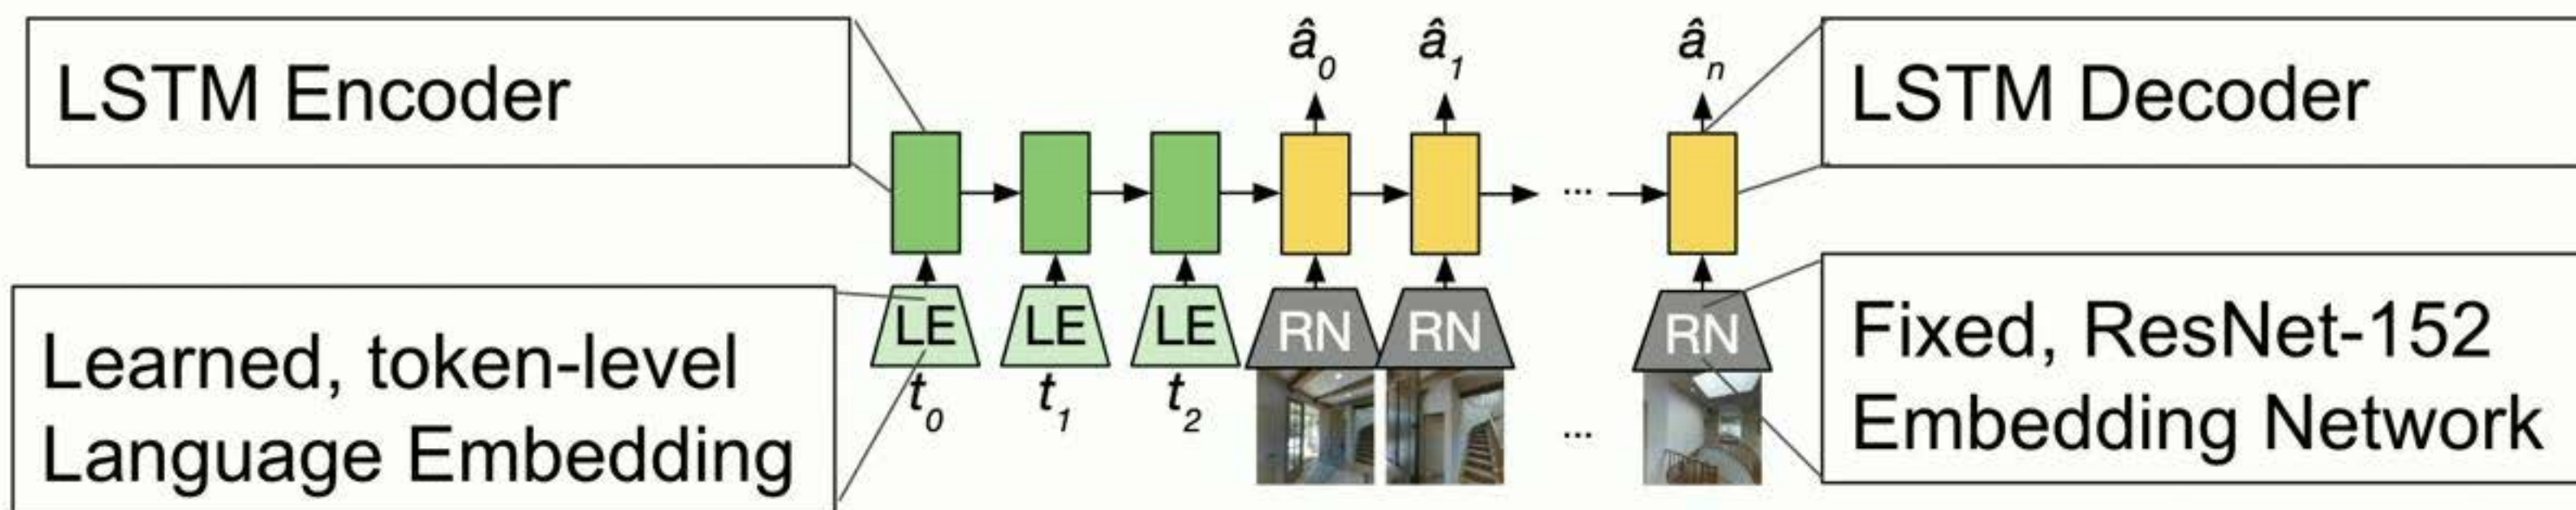
- Encode the language tokens.
- Decode a sequence of actions to take in the environment.
- At every timestep, receive a new visual observation.



Sequence-to-Sequence Model

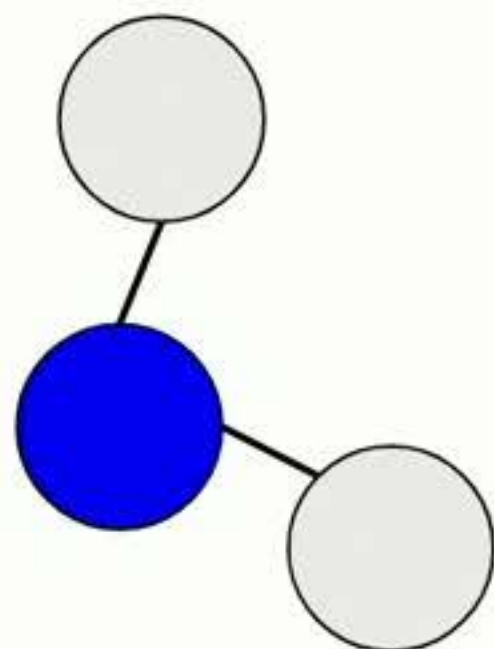

- Encode the language
- Decode a sequence in the environment.
- At every timestep, receive a new visual observation.

We will build on this model.



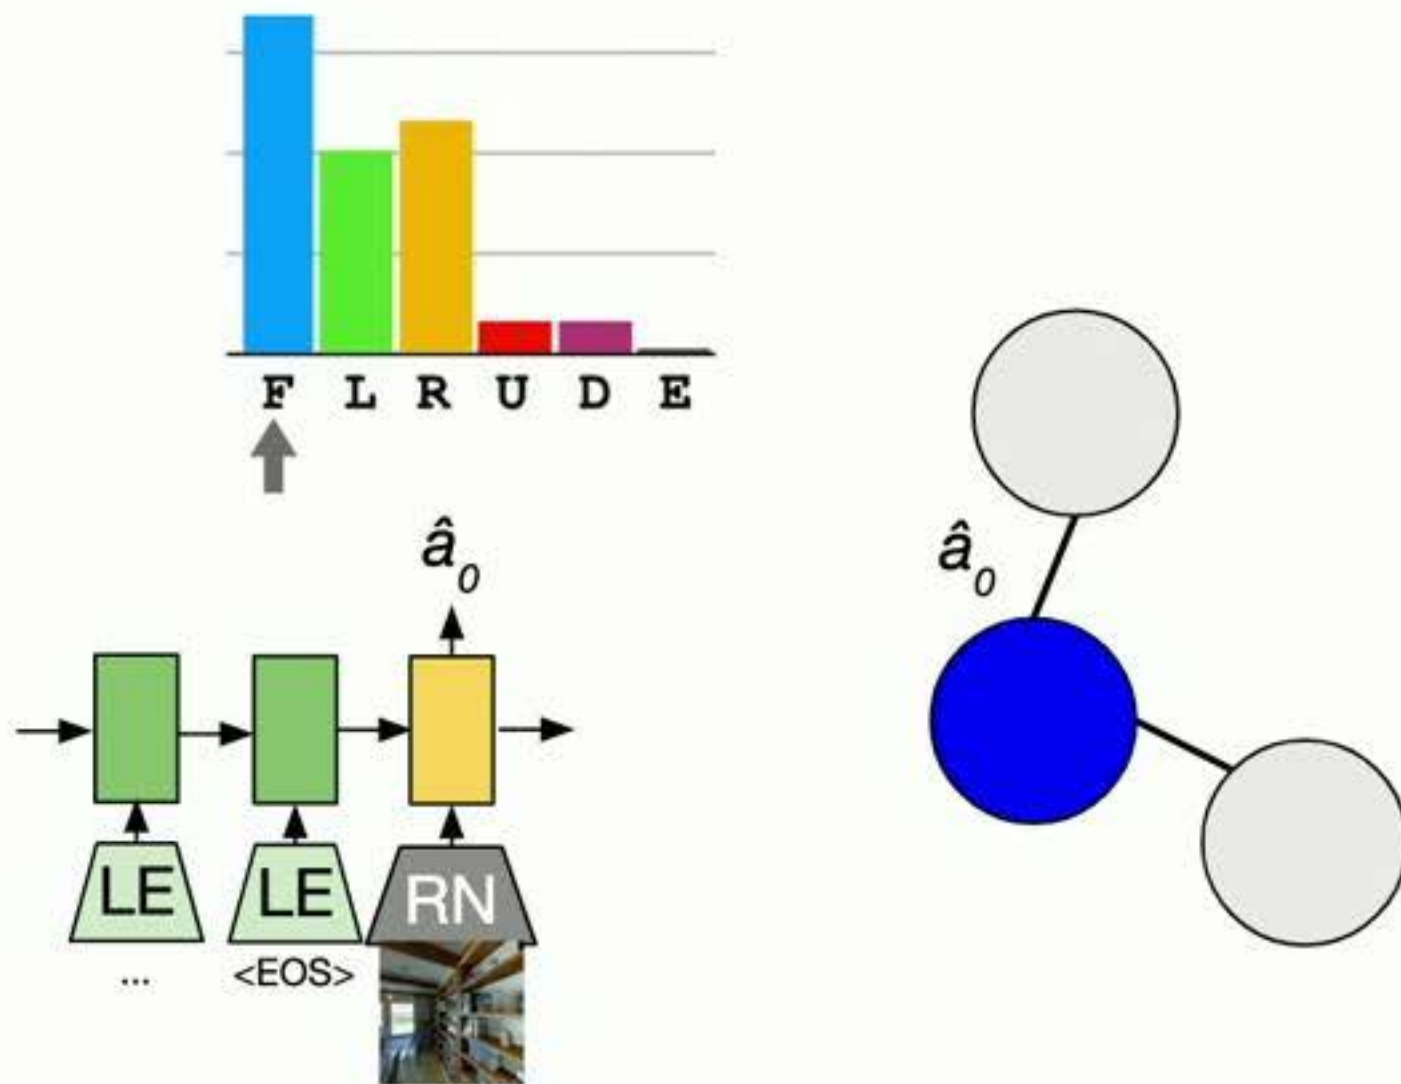
Sequence-to-Sequence Model

- Train to predict the action a shortest-path planner would take from the current state.



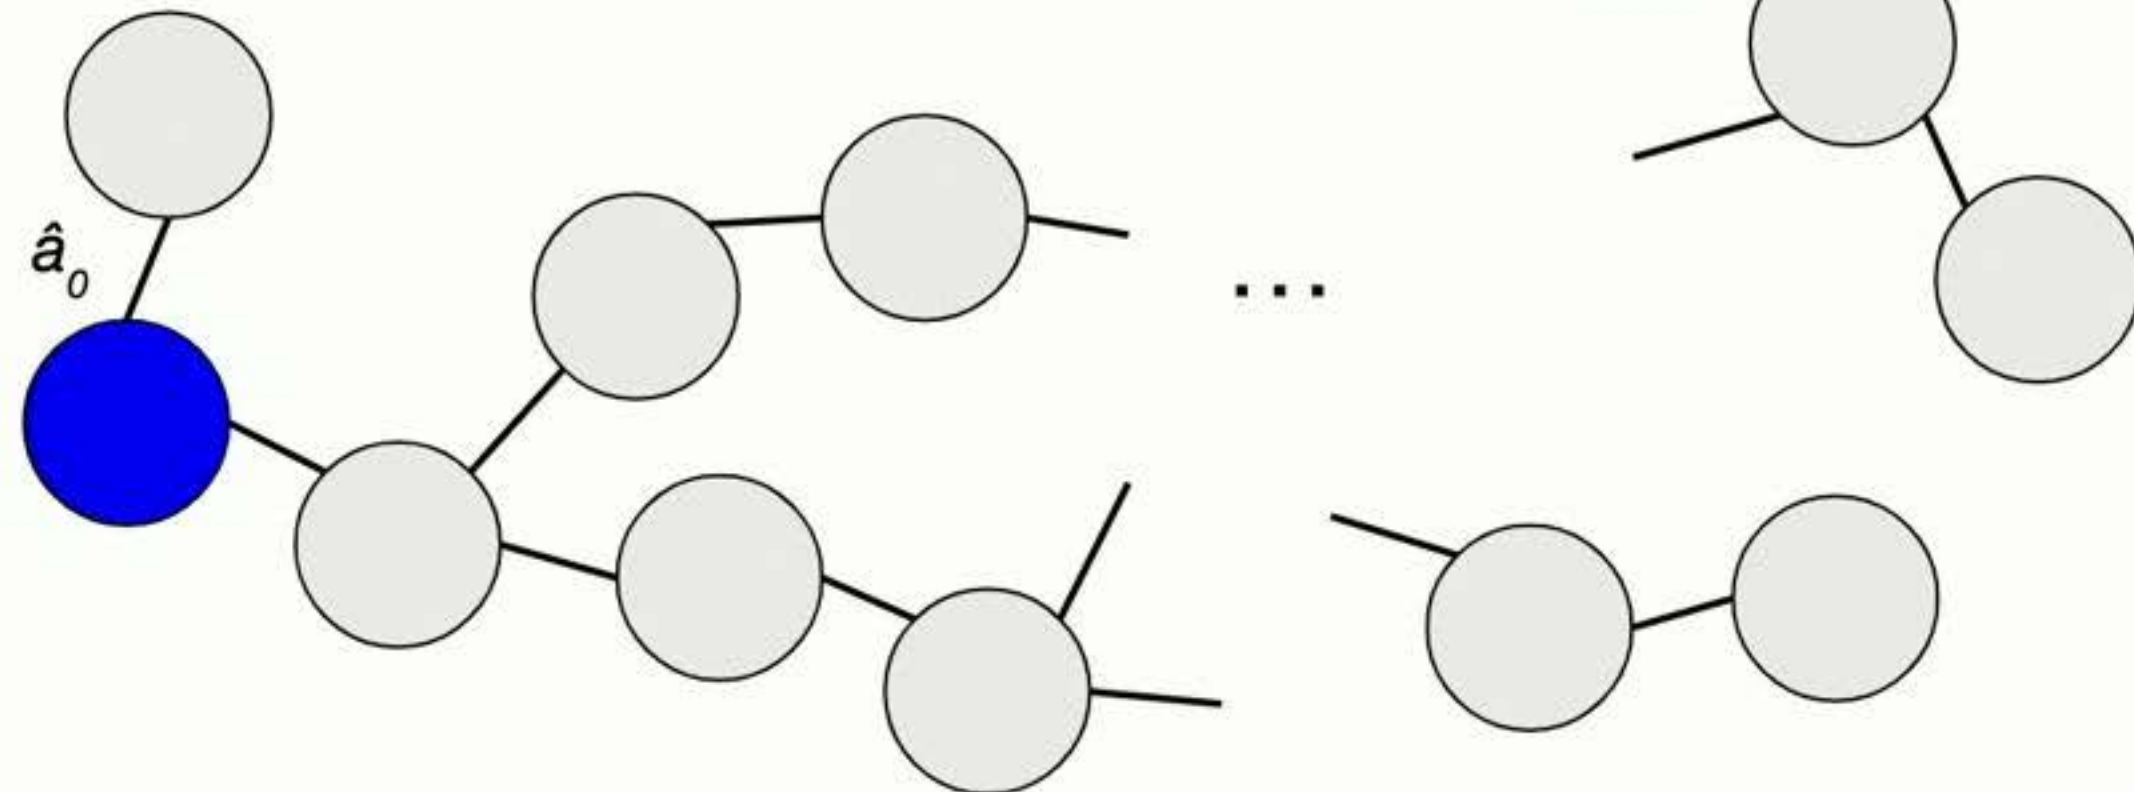
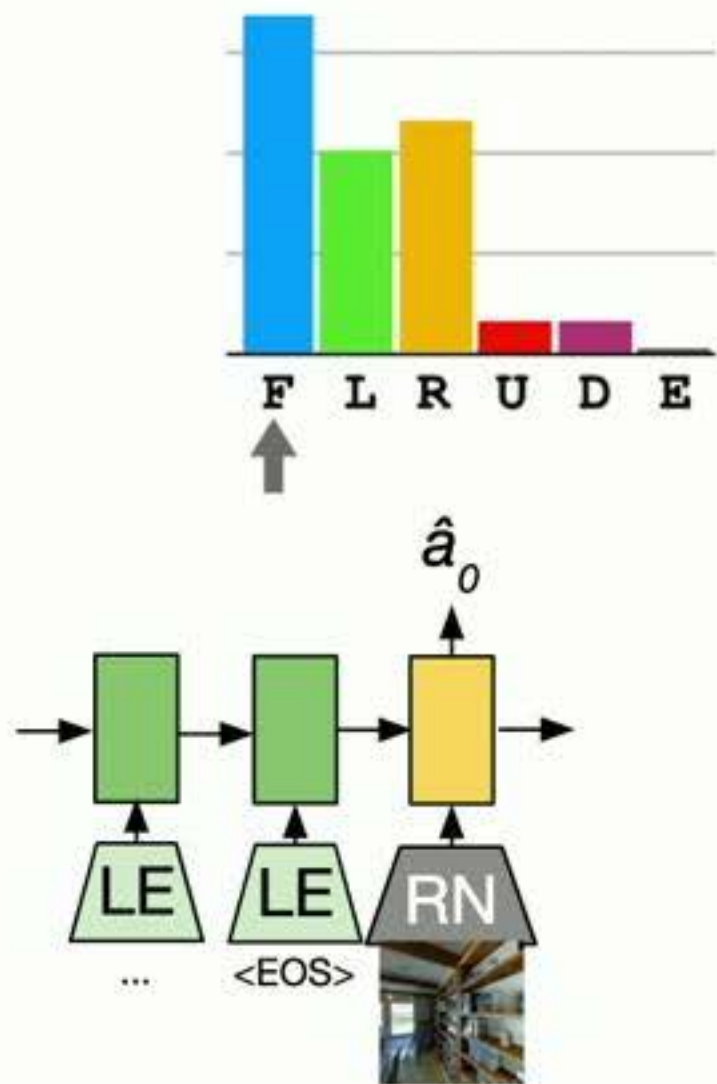
Sequence-to-Sequence Model

- Train to predict the action a shortest-path planner would take from the current state.



Sequence-to-Sequence Model

- Train to predict the action a shortest-path planner would take from the current state.



Language Grounding for QA and Navigation

Input:

Instruction + Frame



Room-2-Room
[Anderson et al., CVPR'18]

Question + Frame



Embodied QA
[Das et al., CVPR'18]

Question + Frame



Interactive QA
[Gordon et al., CVPR'18]

Output:

Navigation Actions

**Navigation Actions +
Answer Action**

**Navigation Actions +
Answer Action**

Language Grounding for QA and Navigation

Input:

Instruction + Frame



Room-2-Room
[Anderson et al., CVPR'18]

Question + Frame



Embodied QA
[Das et al., CVPR'18]

Question + Frame



Interactive QA
[Gordon et al., CVPR'18]

Output:

Navigation Actions

**Navigation Actions +
Answer Action**

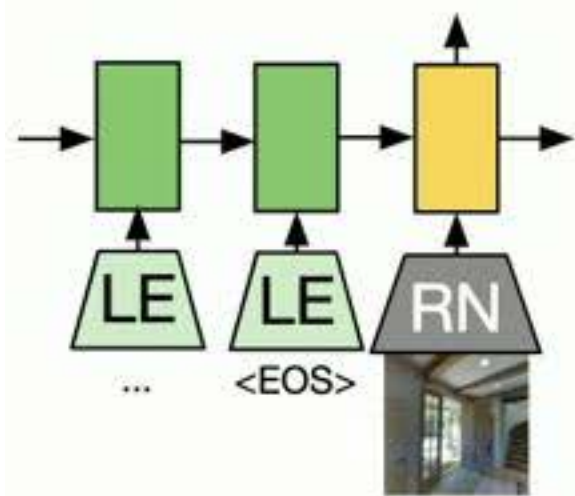
**Navigation Actions +
Answer Action**

Outline

- Language grounding in visual environments
 - For navigation
 - **Unimodal bias [Thomason et al., NAACL'19]**
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - Navigation from dialog history
- Next steps

Inputs and Outputs

Language: *“Walk past the bar and turn right.”*



t_3



t_4




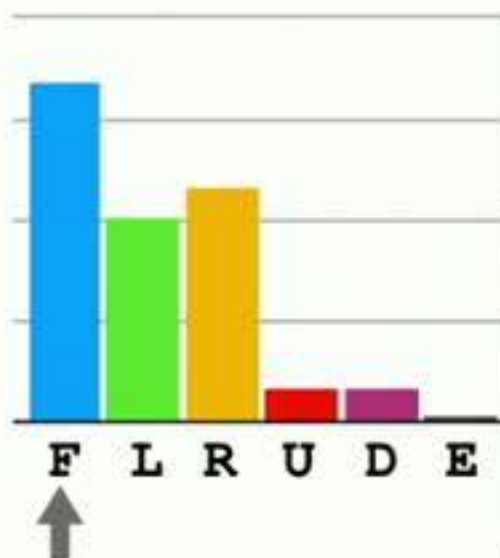
Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Inputs and Outputs

Language: *“Walk past the bar and turn right.”*


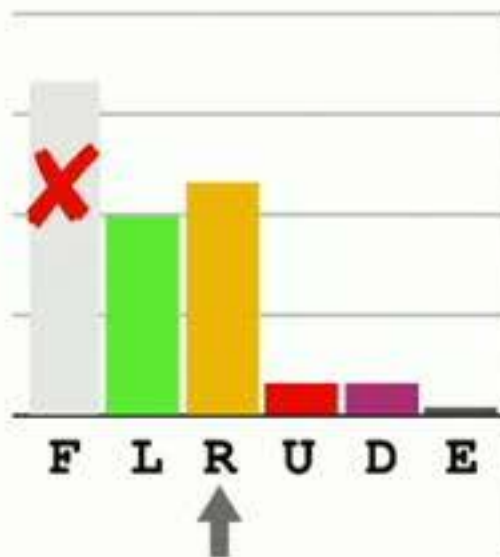
t_1

t_2



t_3


t_4



Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Inputs and Outputs

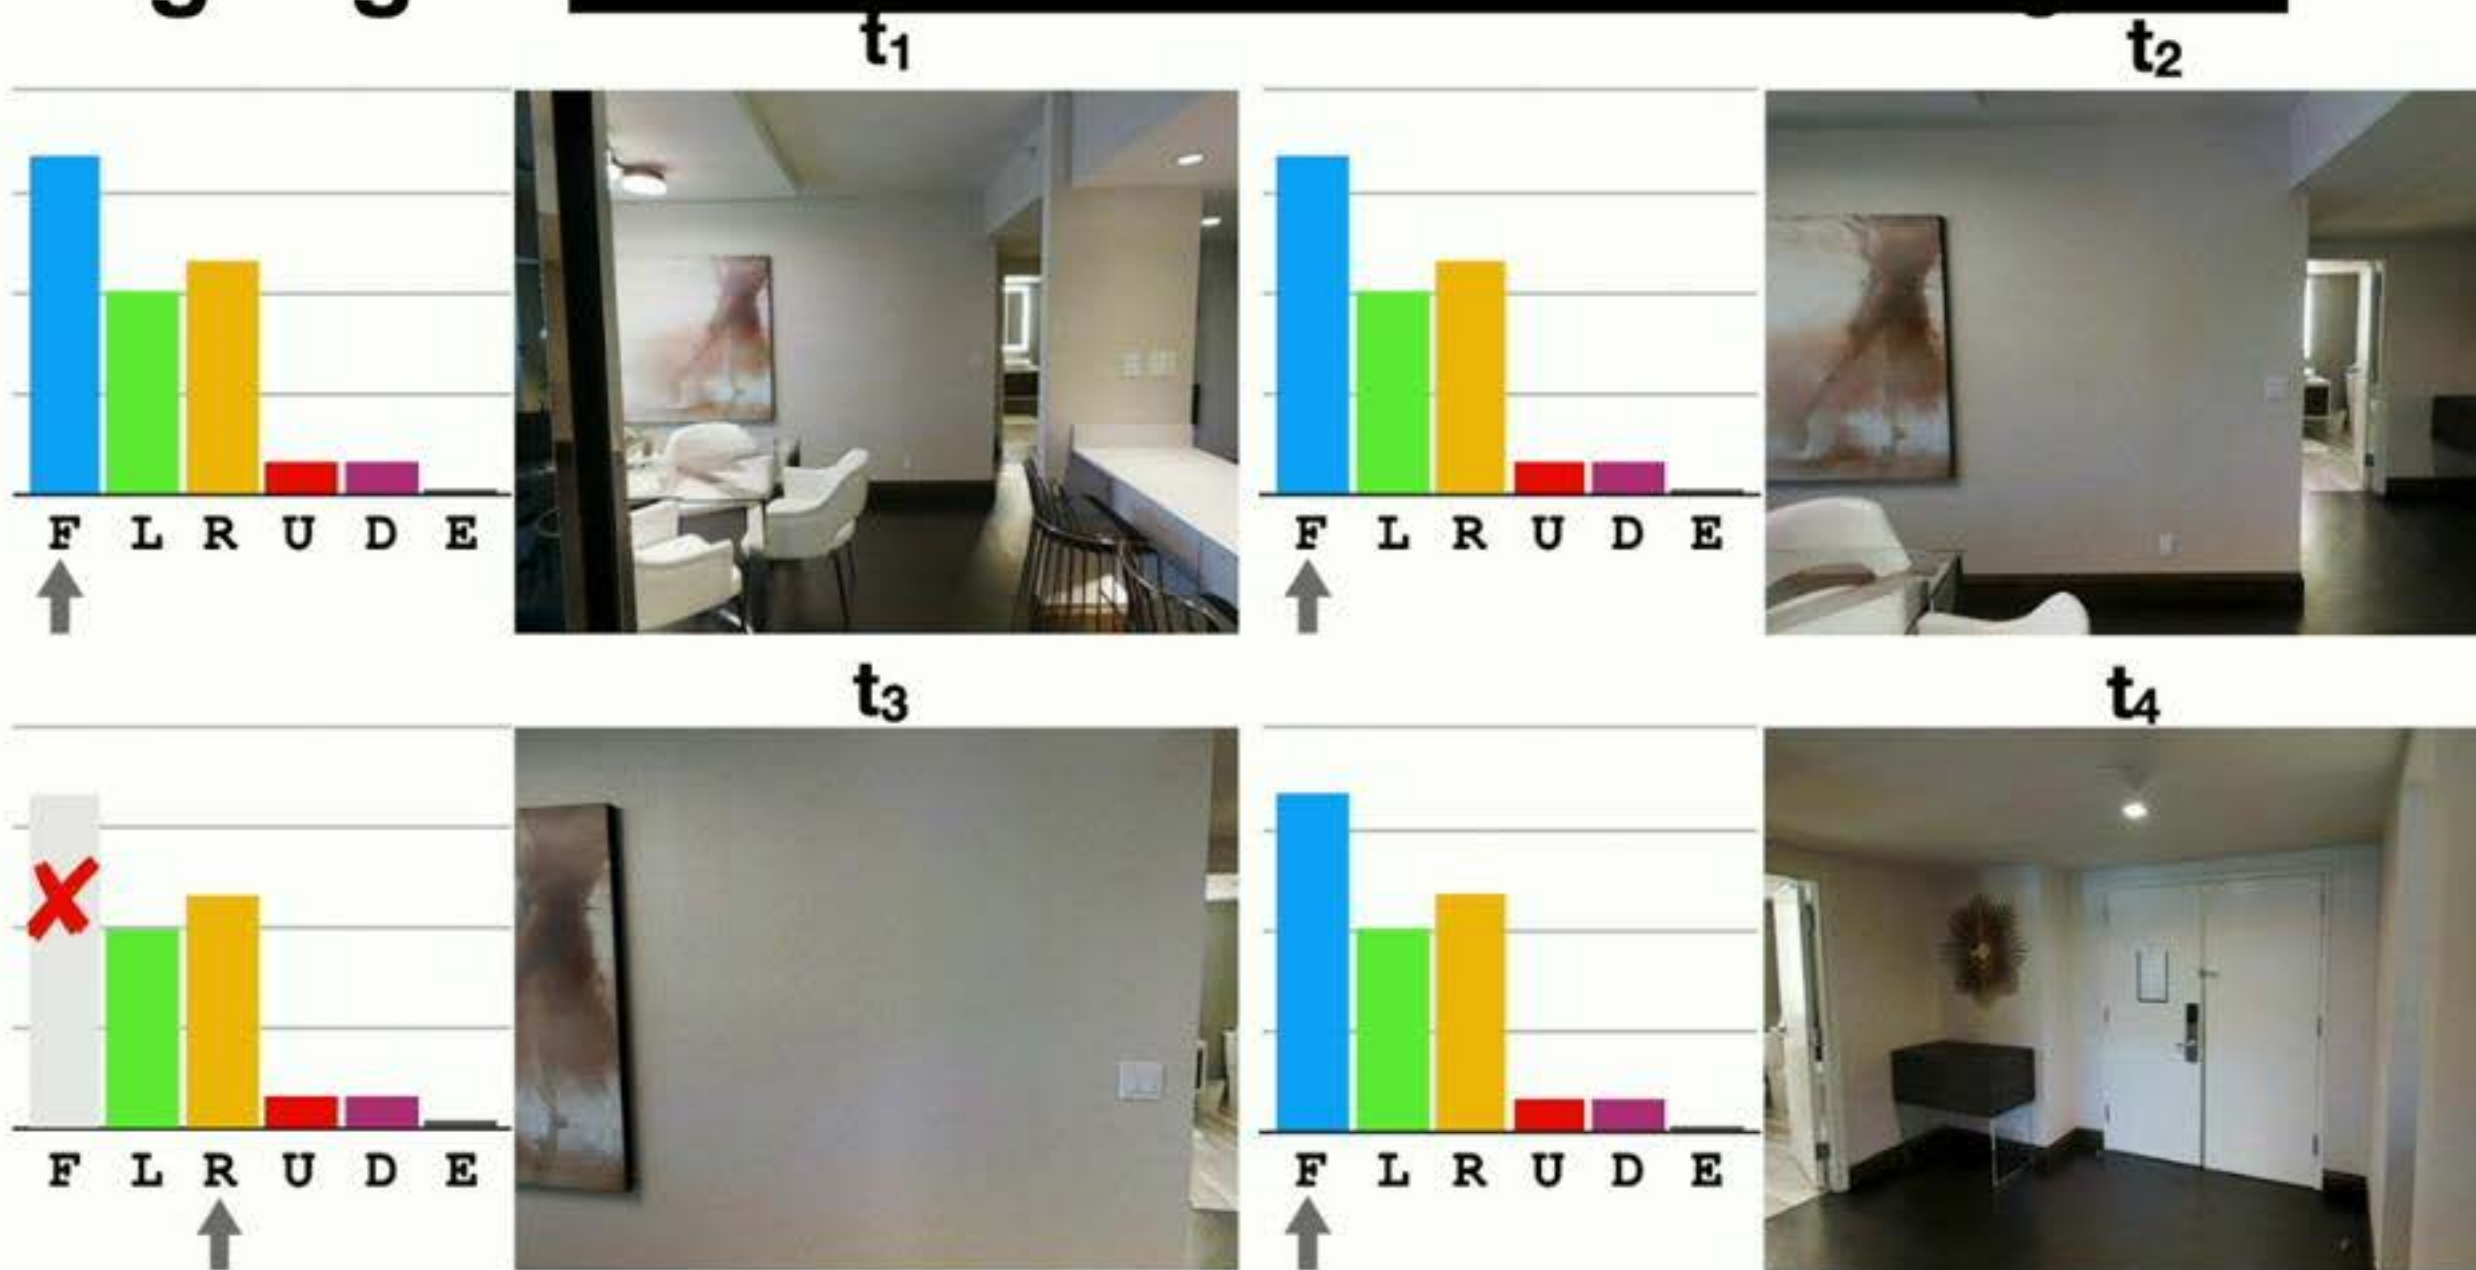
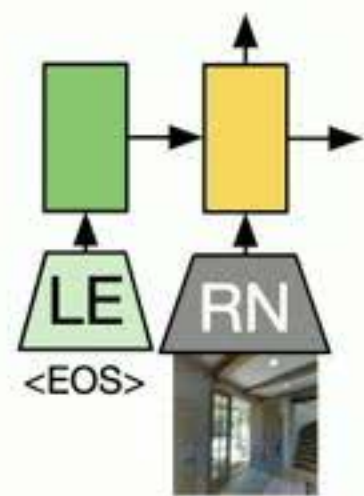
Language: *“Walk past the bar and turn right.”*
 t_1 t_2



Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Unimodal Ablation - *Vision Only*


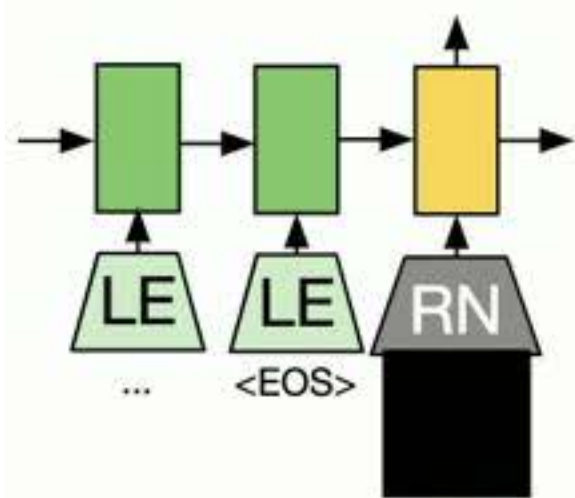
Language: [REDACTED]



Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Unimodal Ablation - *Language Only*



Language: *“Walk past the bar and turn right.”*



Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Unimodal Ablation - *Action Only*

Language: [REDACTED]





Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Unimodal Model Ablations

Unimodal Ablation - *Action Only*

Language: [REDACTED]



Actions: **F**orward, turn **L**eft & **R**ight, tilt **U**p & **D**own, **E**nd

Unimodal Model Ablations

Unimodal Model Ablations



Room-2-Room
[Anderson et al., CVPR'18]

**Beats
baseline**

**Vision-, language-, and
action-only models.**

Unimodal Model Ablations



Room-2-Room
[Anderson et al., CVPR'18]



Embodied QA
[Das et al., CVPR'18]

**Beats
baseline**

**Vision-, language-, and
action-only models.**

**Vision-, language-, and
action-only models.**

Unimodal Model Ablations



Room-2-Room
[Anderson et al., CVPR'18]



Embodied QA
[Das et al., CVPR'18]



Interactive QA
[Gordon et al., CVPR'18]

**Beats
baseline**

**Vision-, language-, and
action-only models.**

**Vision-, language-, and
action-only models.**

Vision-only model.

Unimodal Model Ablations



Room-2-Room
[Anderson et al., CVPR'18]



Embodied QA
[Das et al., CVPR'18]



Interactive QA
[Gordon et al., CVPR'18]

**Beats
baseline**

**Vision-, language-, and
action-only models.**

**Vision-, language-, and
action-only models.**

Vision-only model.

**Beats initial
model!**

Language-only model.

Unimodal Model Ablations



Room-2-Room
[Anderson et al., CVPR'18]



Embodied QA
[Das et al., CVPR'18]



Interactive QA
[Gordon et al., CVPR'18]

**Beats
baseline**

**Vision-, language-, and
action-only models.**

**Vision-, language-, and
action-only models.**


Vision-only model.

**Beats initial
model!**

Language-only model.

Vision-only model.

Unimodal Model Ablations



Low-level Language

[Andersson et al., CVPR'18]

Vision-, language-, and action-only models.

Language-only model.



Embodied QA
[Das et al., CVPR'18]

Vision-, language-, and action-only models.

Vision-only model.



Interactive QA
[Gordon et al., CVPR'18]


Vision-only model.

(none)

Beats baseline

Beats initial model!

Unimodal Model Ablations



[Andersson et al., CVPR'18]

Vision-, language-, and action-only models.

Language-only model.



Embodied QA
[Das et al., CVPR'18]

Vision-, language-, and action-only models.

Vision-only model.



Interactive QA
[Gordon et al., CVPR'18]

Vision-only model.

(none)

Beats baseline

Beats initial model!

Unimodal Model Ablations



Photorealistic
Environments

Low-level
Language

[Andersson et al., CVPR'18]

Vision-, language- and
action-only models.



Simple Visual
Environments

[Das et al., CVPR'18]

Vision-, language- and
action-only models.



Interactive QA

[Gordon et al., CVPR'18]

Vision-only model.

**Beats
baseline**

**Beats initial
model!**

Language-only model.

Vision-only model.

(none)

Unimodal Model Ablations



Photorealistic
Environments

Low-level
Language

[Andersson et al., CVPR'18]

Vision-, language- and
action-only models.



High-level
Language

Simple Visual
Environments

[Das et al., CVPR'18]

Vision-, language- and
action-only models.



Interactive QA
[Gordon et al., CVPR'18]

Vision-only model.

**Beats
baseline**


**Beats initial
model!**

Language-only model.

Vision-only model.

(none)

Unimodal Model Ablations



[Andersson et al., CVPR 2018]

Vision-, language- and action-only models.

Language-only model.



[Das et al., CVPR 2018]

Vision-, language- and action-only models.

Vision-only model.



Interactive [Gordon et al., CVPR 2018]

Vision-only model.

(none)

Beats baseline

Beats initial model!

Lessons

Lessons

- Unimodal baselines expose dataset bias.

Lessons

- Unimodal baselines expose dataset bias.
 - More appropriate than non-learning baselines.

Lessons

- Unimodal baselines expose dataset bias.
 - More appropriate than non-learning baselines.
- Rich visual context prevents visual overfitting.

Lessons

- Unimodal baselines expose dataset bias.
 - More appropriate than non-learning baselines.
- Rich visual context prevents visual overfitting.
- *Underspecified* language context prevents language overfitting.

Lessons

- Unimodal baselines expose dataset bias.
 - More appropriate than non-learning baselines.
- Rich visual context prevents visual overfitting.
- *Underspecified* language context prevents language overfitting.
 - Also, low-level instructions are somewhat unnatural.


Lessons

- Unimodal baselines expose dataset bias.
 - More appropriate than non-learning baselines.
- Rich visual context prevents visual overfitting.
- *Underspecified* language context prevents language overfitting.
 - Also, low-level instructions are somewhat unnatural.
- Why not both?


Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- **Vision-and-Dialog Navigation [Thomason et al., *in sub*]**
 - New dataset - CVDN
 - Navigation from dialog history
- Next steps

“Turn around and exit the library, head down the...”




...




“Turn around and exit the library, head down the...”

- *Room-to-Room* uses low-level language.




...




“Turn around and exit the library, head down the...”

- *Room-to-Room* uses low-level language.
- 24% → 80% of human performance since ‘18.




...




“Go to the room with a plant.”

- This instruction is *underspecified*.




...




“Go to the room with a plant.”

- This instruction is *underspecified*.
- This instruction is *ambiguous*.




...





“Go to the room with a plant.”

- This instruction is *underspecified*.
- This instruction is *ambiguous*.

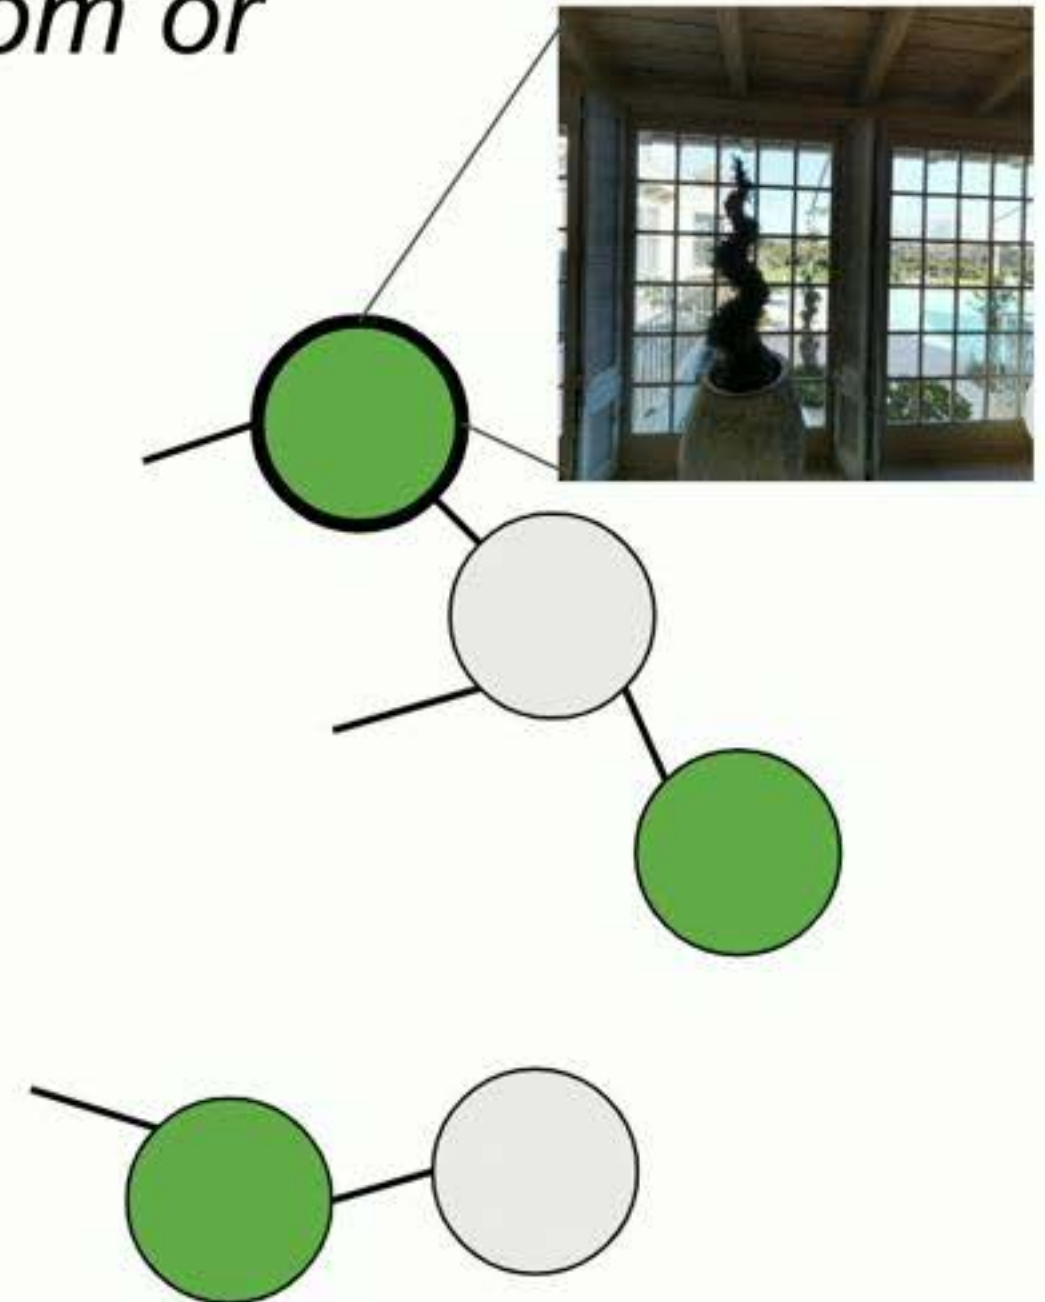
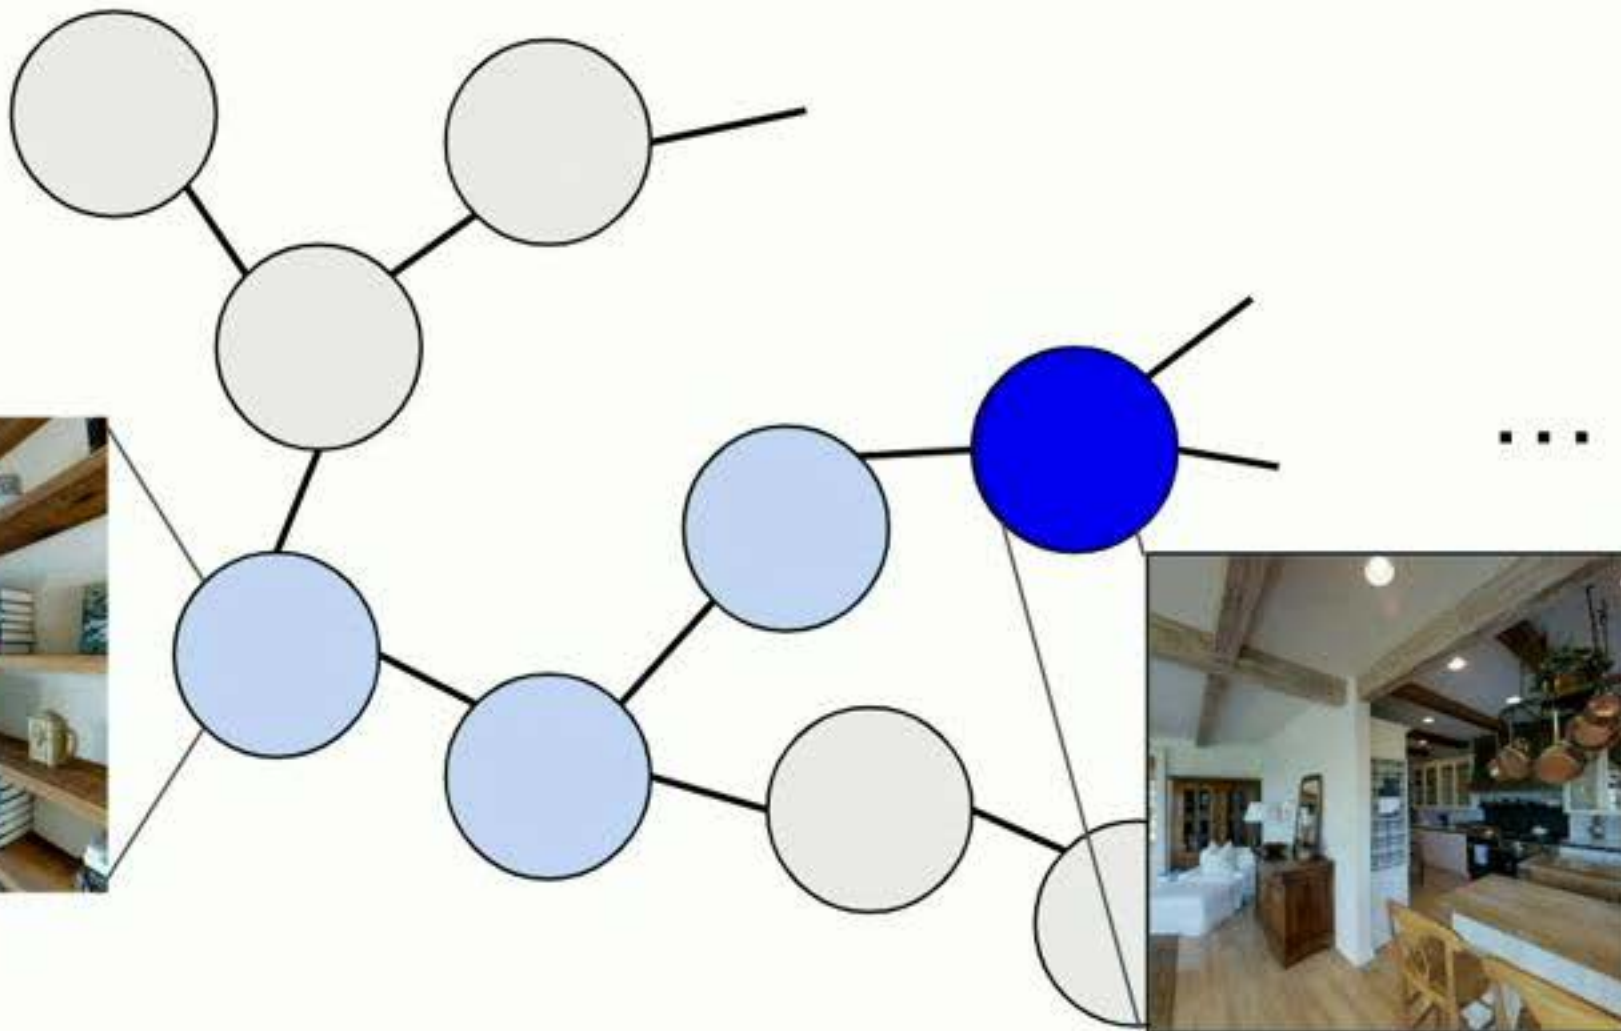


Too Hard!



What if we could just... ask?

- *“Should I continue into the living room or go right towards the kitchen?”*



Visual Fidelity

Rendered



Photorealistic



Visual Context

Static




Dynamic



CLEVR

[Johnson et al., CVPR'17]




Instruction Following

[Chen and Mooney, AAAI'11]



VQA

[Antol et al., CVPR'15]



Room-to-Room

[Anderson et al., CVPR'18]₂₉

Visual Fidelity

Rendered



Photorealistic

CLEVR-Dialog
[Kottur et al.,
NAACL'19]

VisDial
[Das et al.,
CVPR'17]

A **cylinder** is next to a **yellow object**.

Q1 : What shape is **the object**?


A1 : Sphere

Q2 : And material?

A2 : Metal

Q3 : What about **that cylinder**?

A3 : Rubber



Visual Dialog

Q: What is the gender of the one in the white shirt ?

A: She is a woman

Q: What is she doing ?

A: Playing a Wii game

Q: Is that a man to her right

A: No, it's a woman

Visual Context

Static



Dynamic



Language Source

Templates



**Guidance
Abstraction**

Semantic



Visual



Humans



Language Source

Templates

Guidance Abstraction

Semantic

Visual





Talk the Walk
[de Vries et al., arXiv'18]

Humans

Language Source

Templates



VLNA [Nguyen et al., CVPR'19]

Guidance Abstraction

Semantic

Visual



Talk the Walk [de Vries et al., arXiv'18]

Humans

Language Source

Templates



VLNA [Nguyen et al., CVPR'19]

Guidance Abstraction

Semantic

Visual



Talk the Walk
[de Vries et al., arXiv'18]

Vision-and-Dialog Navigation

- Human-human dialogs
- Both participants get an egocentric scene view.

Humans

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - **New dataset**
 - **Cooperative Vision-and-Dialog Navigation (CVDN)**
 - **2k human-human dialogs**
 - Navigation from dialog history
- Next steps

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal [L'19]
- Vision-and-Dialog Navigation [et al., *in sub*]
 - **New dataset**
 - **Cooperative Vision-and-Dialog Navigation (CVDN)**
 - **2k human-human dialogs**
 - Navigation from dialog history
- Next steps

Initial model adapted from R2R does not exhibit unimodal overfitting.


Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.

-- this target object is present in at least two rooms, but only one is correct.

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



-- this target object is present in at least two rooms, but only one is correct.

Visible to both **Navigator** and **Oracle**

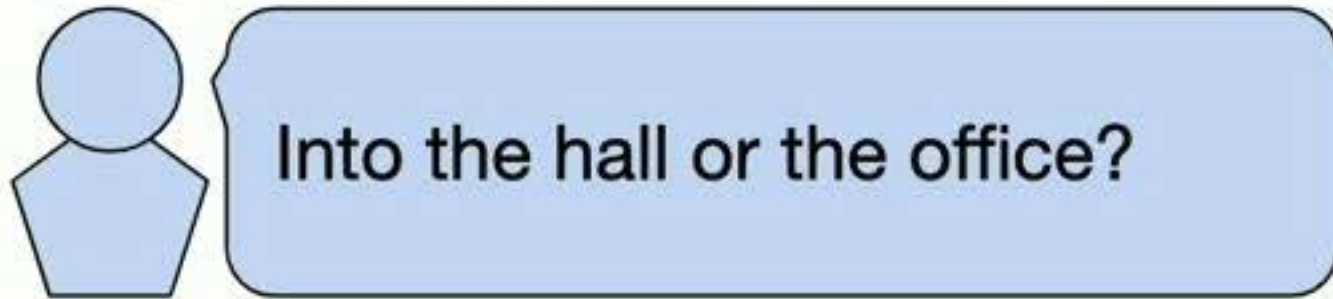
Hint: The goal room contains a *mat*.



-- this target object is present in at least two rooms, but only one is correct.

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



-- this target object is present in at least two rooms, but only one is correct.

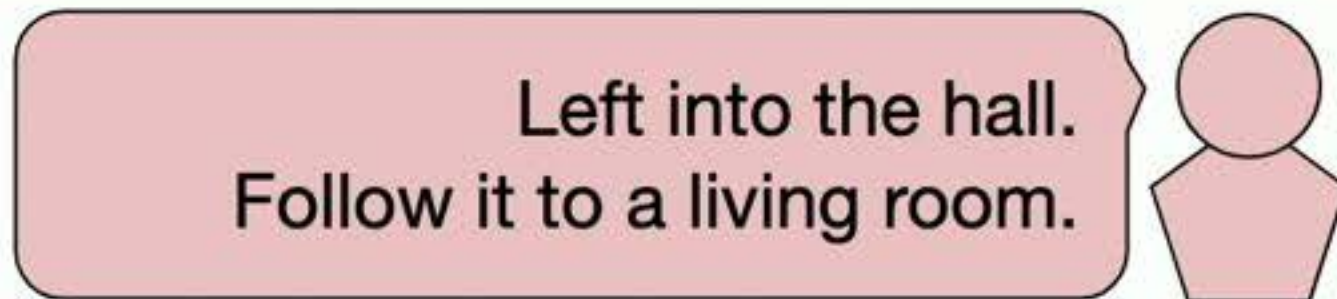
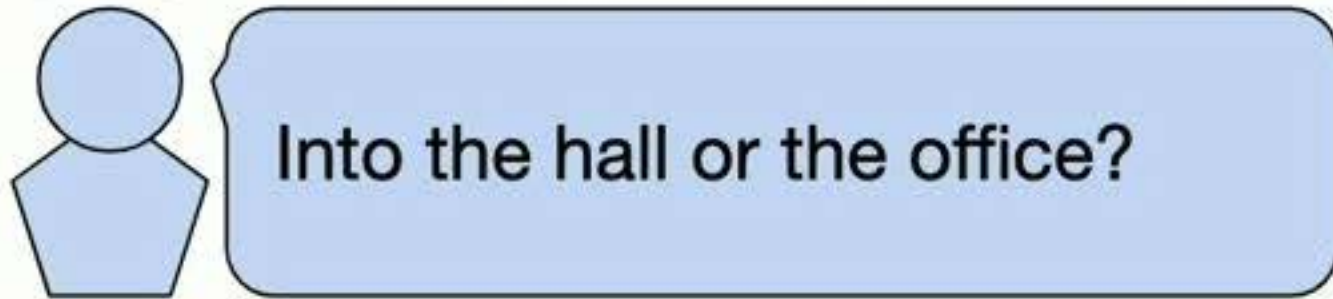
Visible Only to the **Oracle**



-- A shortest-path planner's next steps; up to 5 navigation nodes in the direction of the goal.

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



-- this target object is present in at least two rooms, but only one is correct.

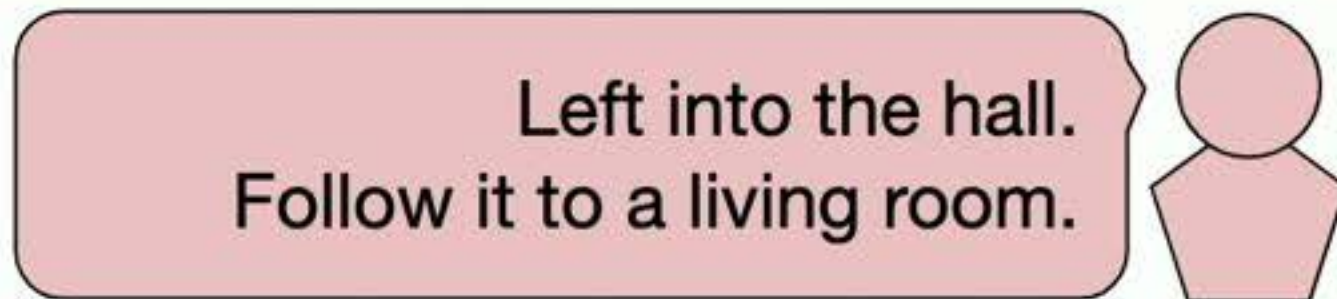
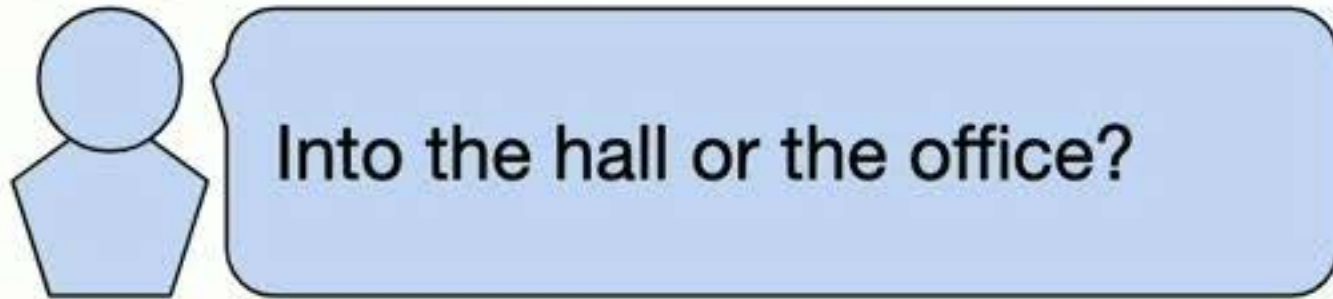
Visible Only to the **Oracle**



-- A shortest-path planner's next steps; up to 5 navigation nodes in the direction of the goal.

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



-- this target object is present in at least two rooms, but only one is correct.


Visible Only to the **Oracle**



-- A shortest-path planner's next steps; up to 5 navigation nodes in the direction of the goal.




Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Into the hall or the office?

Left into the hall.
Follow it to a living room.



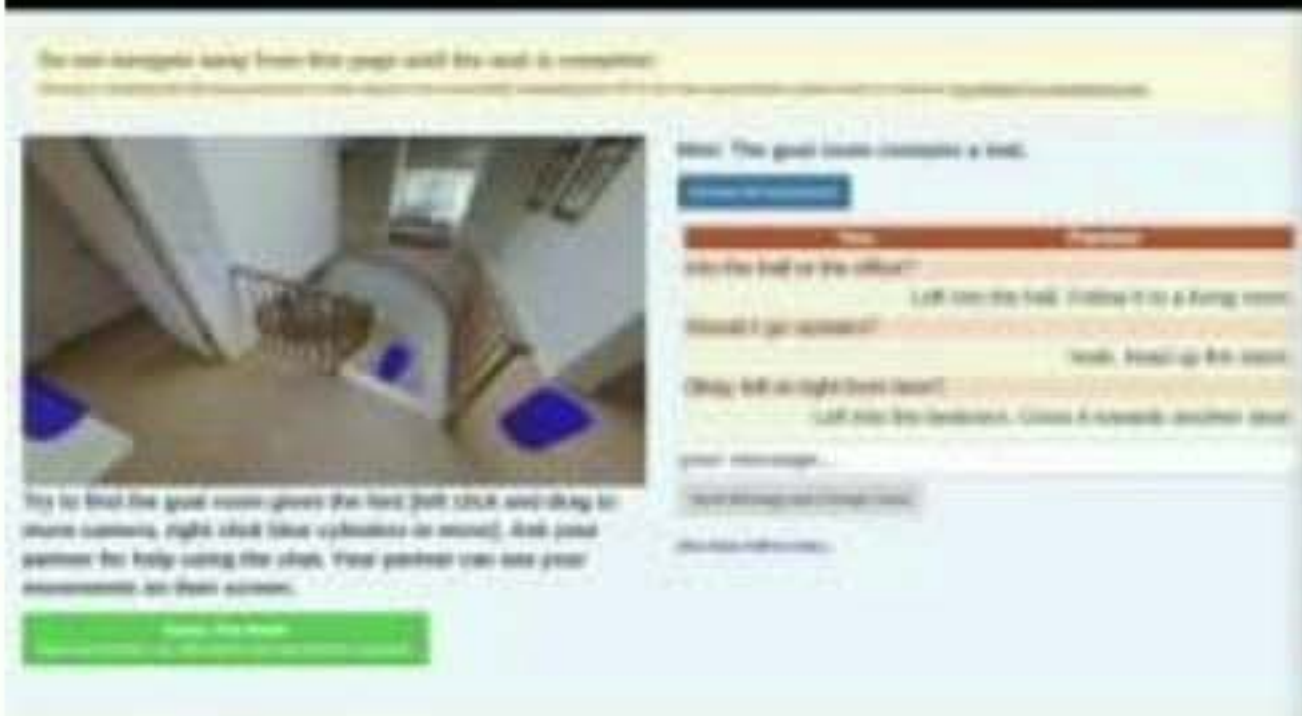
Should I go upstairs?

-- this target object is present in at least two rooms, but only one is correct.

Visible Only to the **Oracle**



-- A shortest-path planner's next steps;
up to 5 navigation nodes in the direction
of the goal.




Navigator View


Oracle View

ON AVERAGE, PATHS IN VDN ARE OVER THREE TIMES LONGER THAN R2R PATHS.

Dialog Leads to Long Paths and Rich Language





Dialog Leads to Long Paths and Rich Language





- Path Length Average:
 - Human (25.0); Planner (17.4)
 - Room-to-Room (6.0)

Dialog Leads to Long Paths and Rich Language



- Path Length Average:
 - Human (25.0); Planner (17.4)
 - Room-to-Room (6.0)

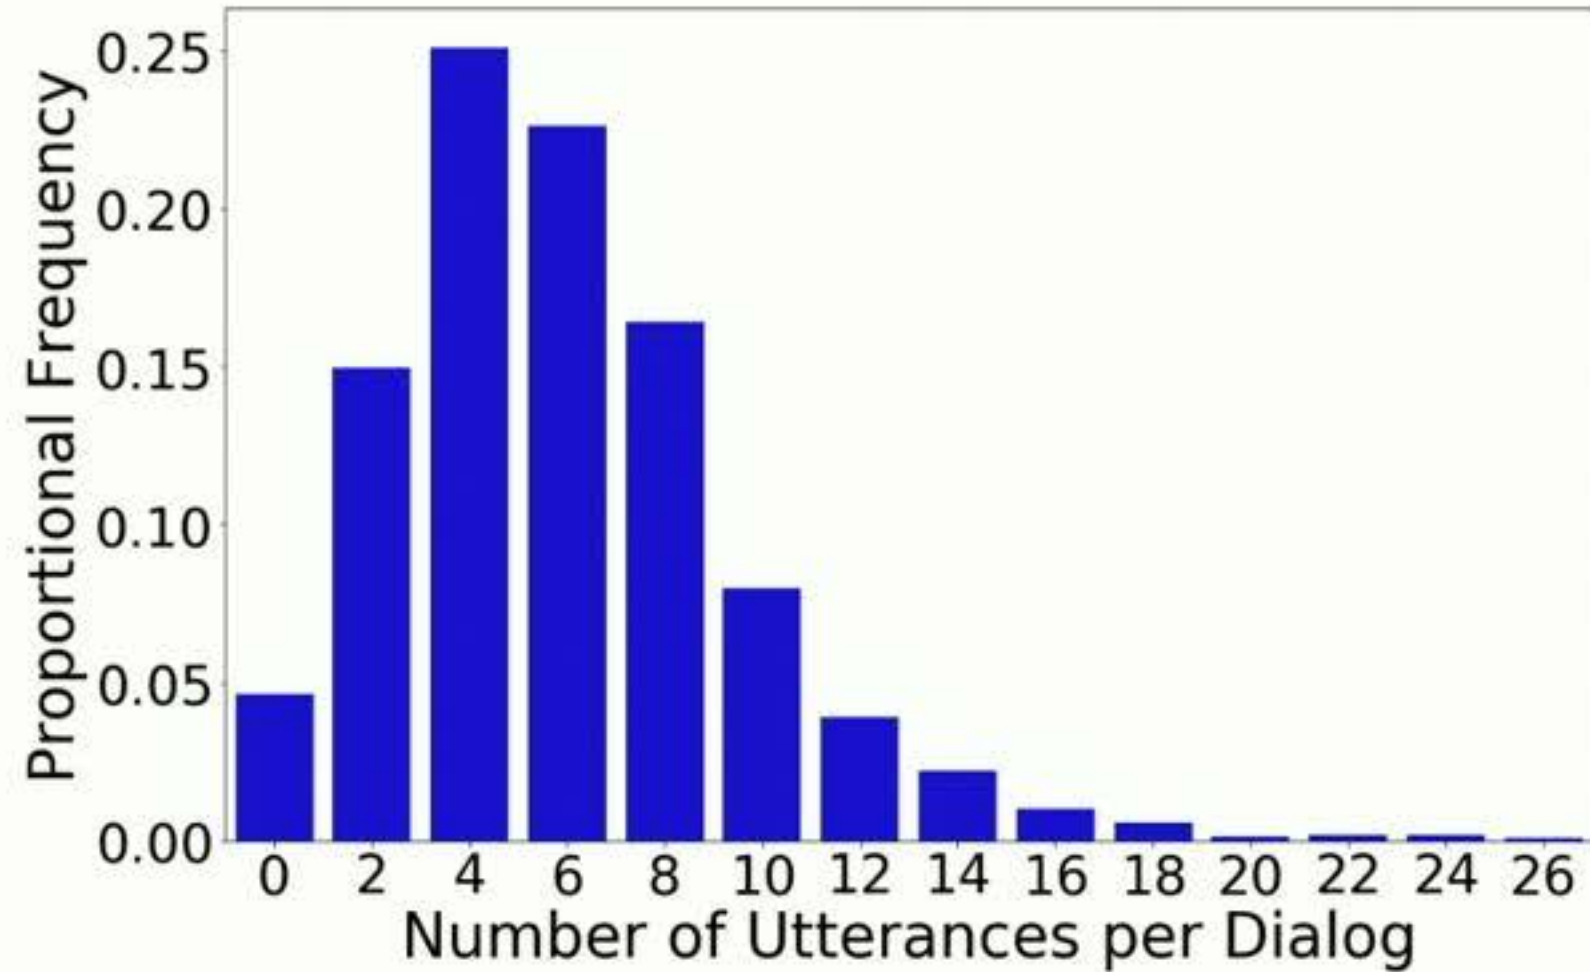
Dialog Leads to Long Paths and Rich Language




- Path Length Average:
 - Human (25.0); Planner (17.4)
 - Room-to-Room (6.0)

- Average total words:
 - CVDN (82)
 - Room-to-Room (29)

Dialog Leads to Long Paths and Rich Language



Dialog Leads to Long Paths and Rich Language



The goal room contains a *rug*.

➔ 3 steps

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

➔ 4 steps

Navigator: Do I need to go in the room with the run or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

➔ 7 steps

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

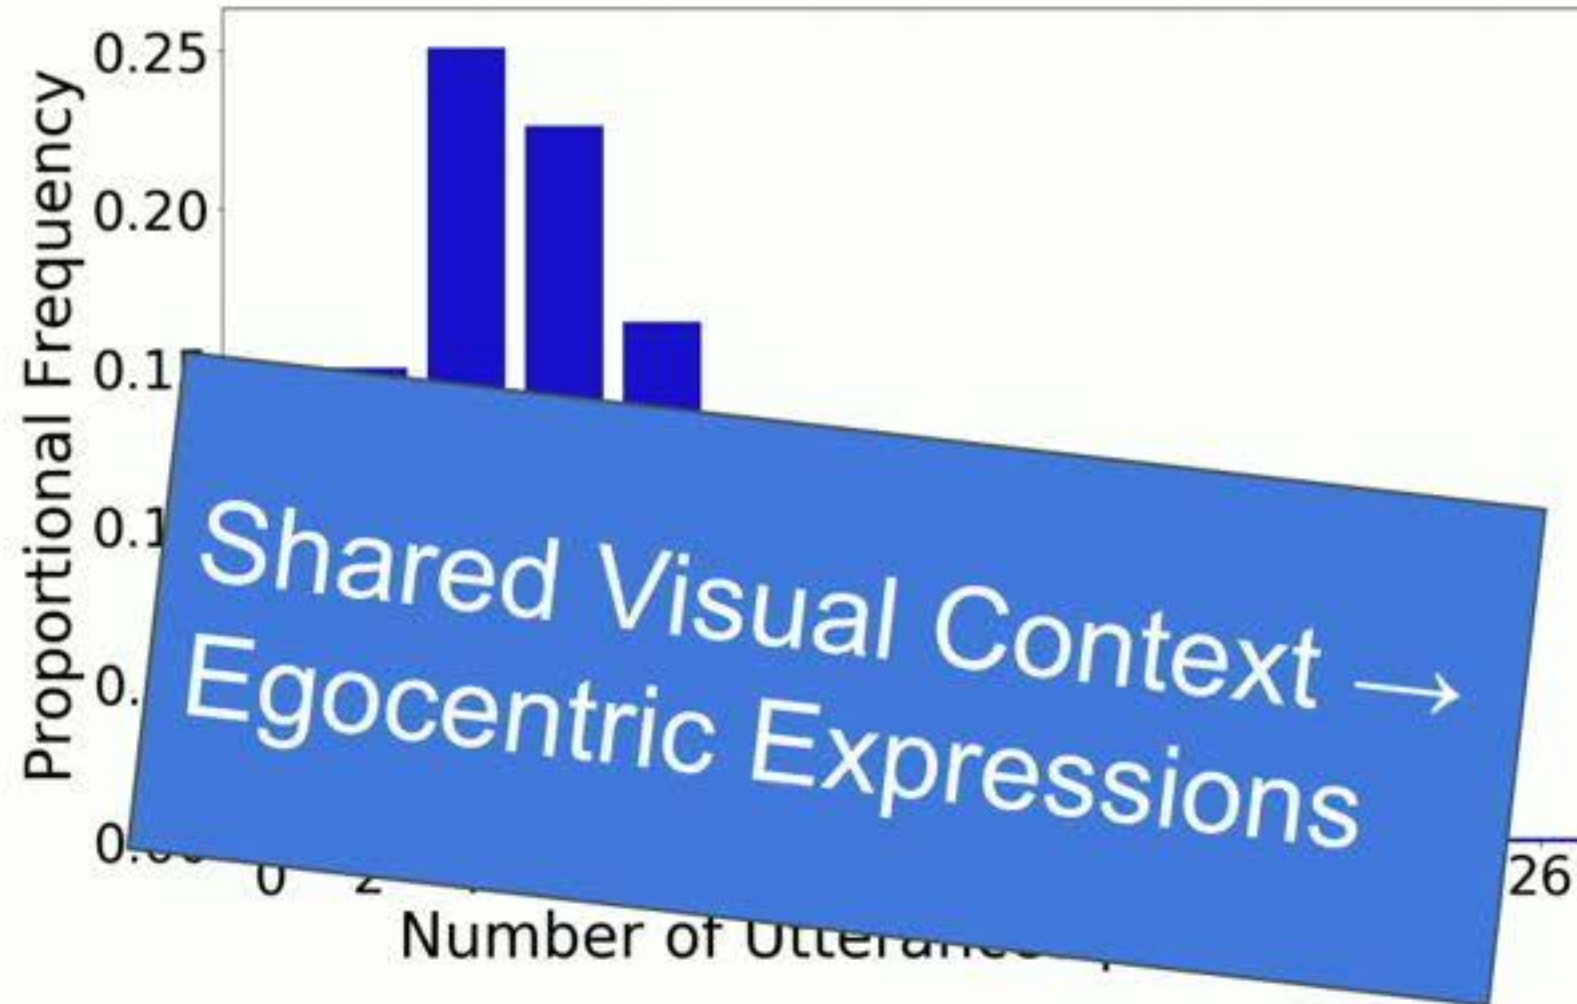
➔ 6 steps


Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 4s  35

Dialog Leads to Long Paths and Rich Language




| | |
|---|---|
| The goal room contains a <i>rug</i>. | ➔ 3 steps |
| Navigator: Should I go to the left or right? | |
| Oracle: Go left and turn right after the bathroom. | ➔ 4 steps |
| Navigator: Do I need to go in the room with the run or keep on going right? | |
| Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right. | ➔ 7 steps |
| Navigator: Should I go into the kitchen or to the right? | |
| Oracle: Turn toward the front door and go up the stairs you see on the right. | ➔ 6 steps |
| Navigator: Do I go left or right? | |
| Oracle: Go along the railing to the right. Stop at the room with a brown chair. | ➔ 4s  36 |

Outline

- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - **Navigation from dialog history**
 - **7k dialog-based navigation inputs**
- Next steps

Navigation from Dialog History

| | |
|---|--|
| The goal room contains a <i>rug</i>. | ➔ 3 steps |
| Navigator: Should I go to the left or right? | |
| Oracle: Go left and turn right after the bathroom. | ➔ 4 steps |
| Navigator: Do I need to go in the room with the run or keep on going right? | |
| Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right. | ➔ 7 steps |
| Navigator: Should I go into the kitchen or to the right? | |
| Oracle: Turn toward the front door and go up the stairs you see on the right. | ➔ 6 steps |
| Navigator: Do I go left or right? | |
| Oracle: Go along the railing to the right. Stop at the room with a brown chair. | ➔ 4s  |

Navigation from Dialog History


The goal room contains a *rug*.

Navigator: Should I go to the left or right?


Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.



Oracle



Navigator



Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the run or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

→ 3 steps

Navigation from Dialog History

The goal room contains a *rug*.

➔ 3 steps

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

Navigation from Dialog History

The goal room contains a *rug*.

➔ 3 steps

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

➔ 4 steps

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 3 steps

➔ 4 steps

- **Input:** History so far + visual frame per timestep.

Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 3 steps

➔ 4 steps

- **Input:** History so far + visual frame per timestep.
- **Output:** Navigation action per timestep.

Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 3 steps

➔ 4 steps

- **Input:** History so far + visual frame per timestep.
- **Output:** Navigation action per timestep.
- **Goal:** Get closer to the target object room.

Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 3 steps

➔ 4 steps

- **Input:** History so far + visual frame per timestep.
- **Output:** Navigation action per timestep.
- **Goal:** Get closer to the target object room.
- 2k dialogs → 7k histories.

Navigation from Dialog History

The goal room contains a *rug*.

Navigator: Should I go to the left or right?

Oracle: Go left and turn right after the bathroom.

Navigator: Do I need to go in the room with the rug or keep on going right?

Oracle: Turn right and take the tiny hallway on the right. You will ascend the stairs you find on the right.

Navigator: Should I go into the kitchen or to the right?

Oracle: Turn toward the front door and go up the stairs you see on the right.

Navigator: Do I go left or right?

Oracle: Go along the railing to the right. Stop at the room with a brown chair.

➔ 3 steps

➔ 4 steps


➔ 7 steps

➔ 6 steps

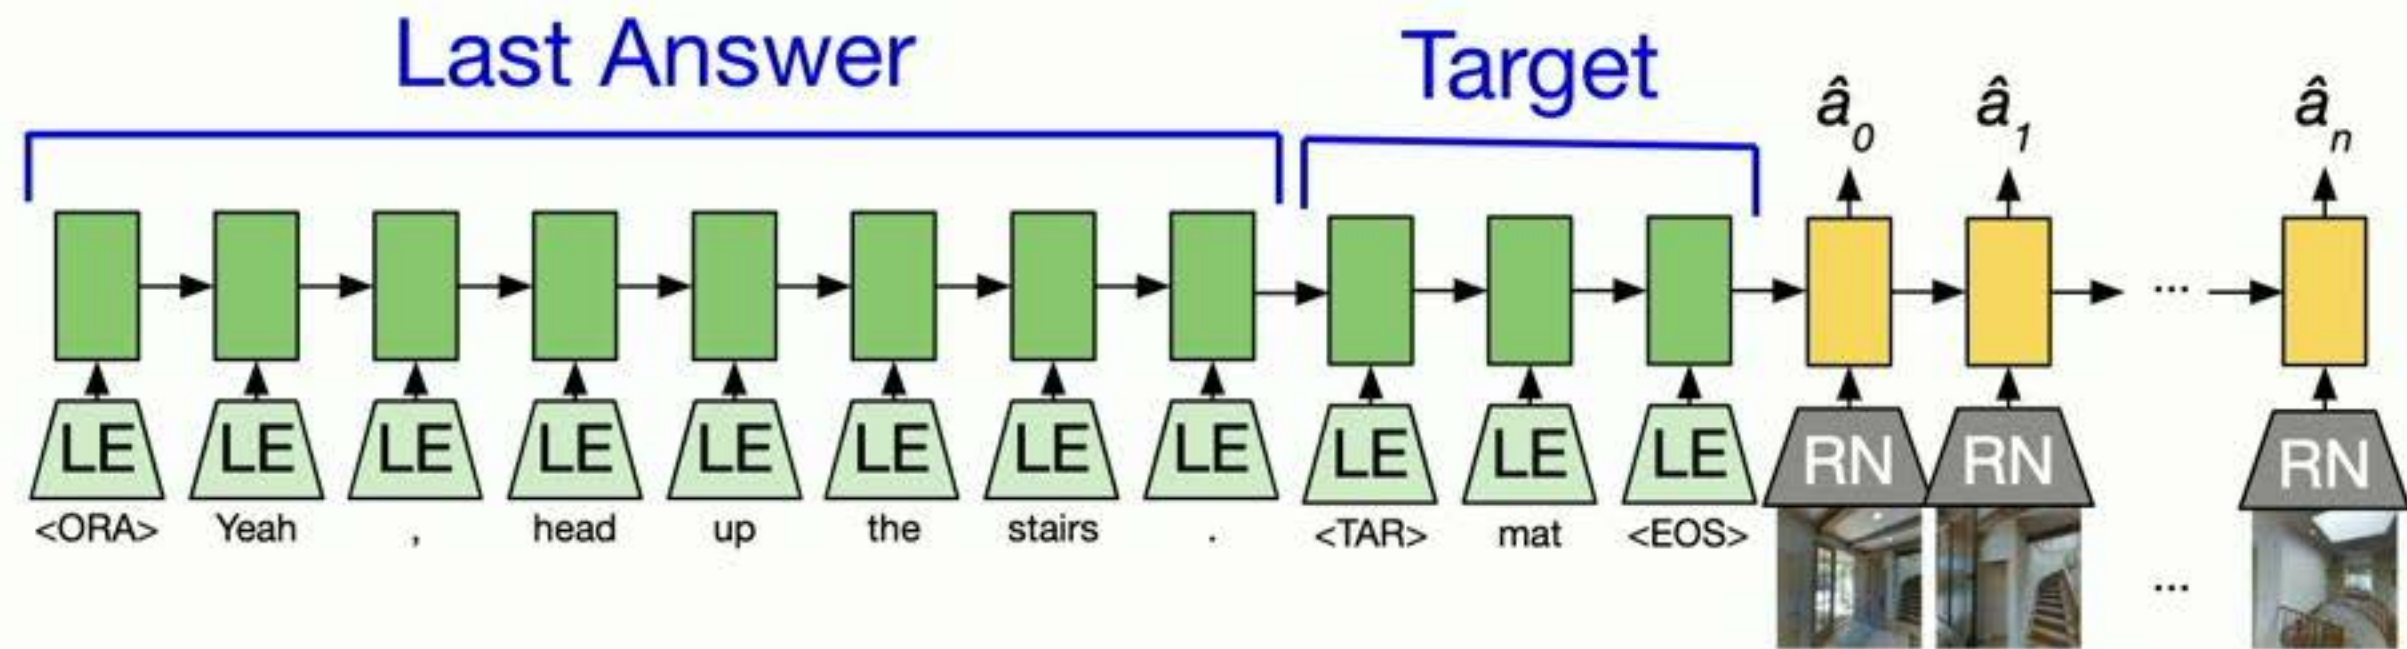
➔ 4 steps

- **Input:** History so far + visual frame per timestep.
- **Output:** Navigation action per timestep.
- **Goal:** Get closer to the target object room.
- 2k dialogs → 7k histories.

Initial, Sequence-to-Sequence Model



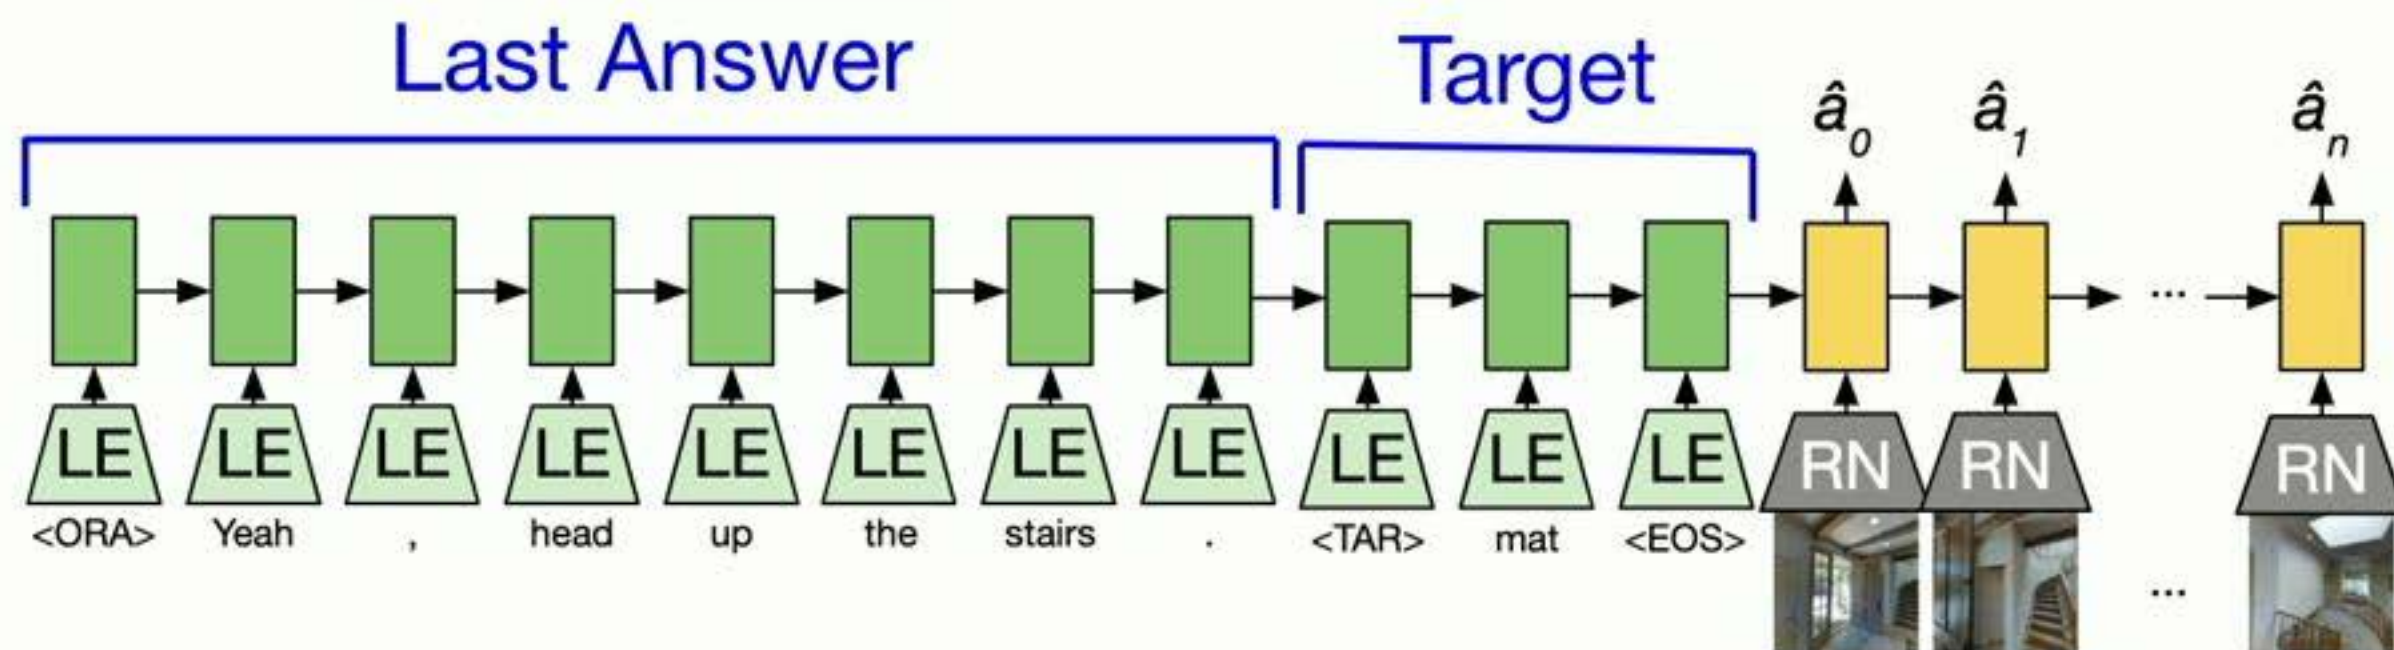
Initial, Sequence-to-Sequence Model



Initial, Sequence-to-Sequence Model

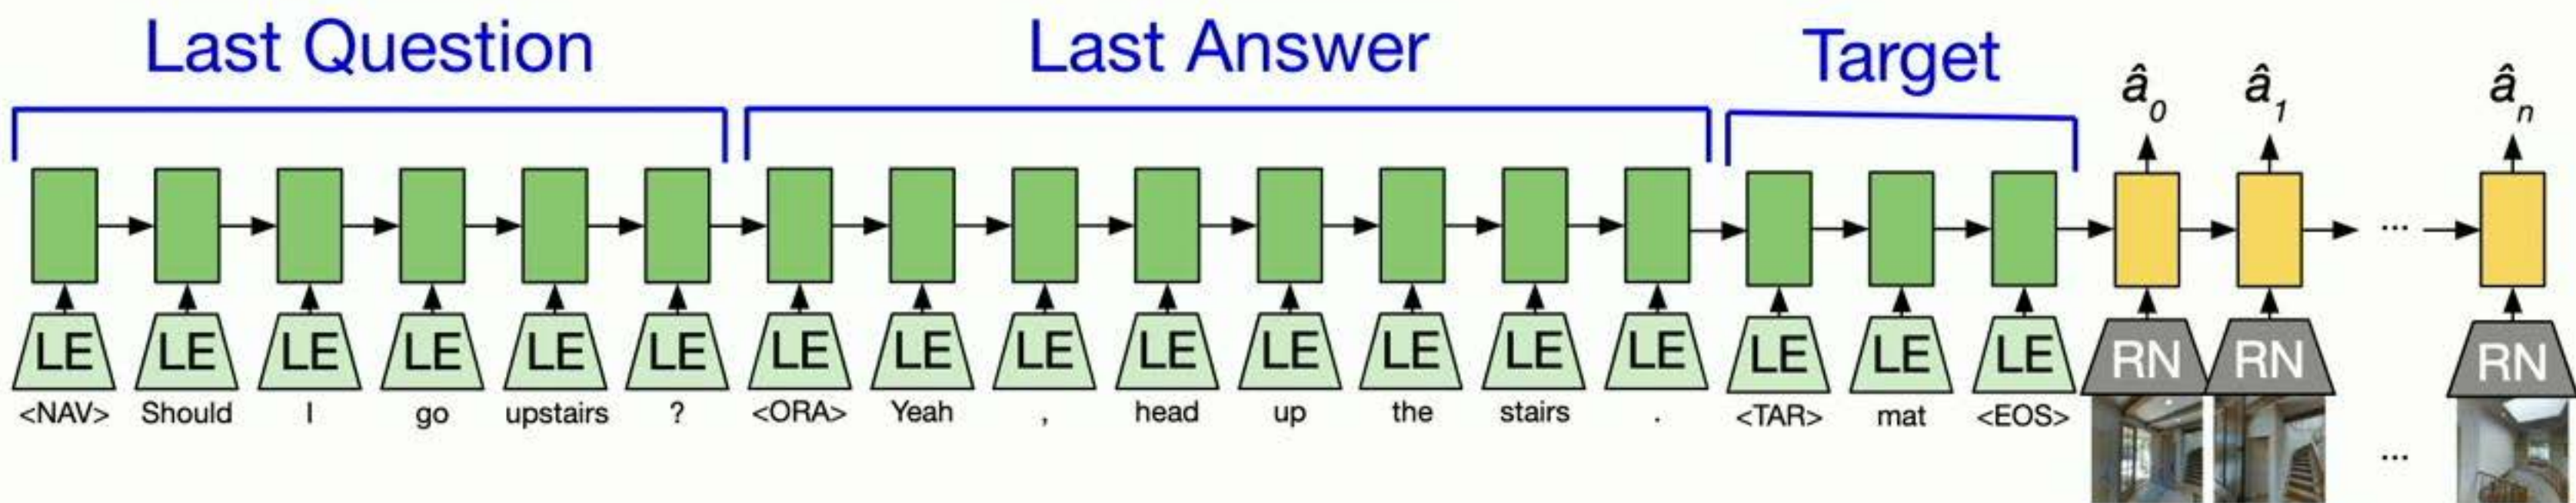
Navigator. Should I turn left down the hallway ahead?

Oracle: *ya*



Initial, Sequence-to-Sequence Model

Navigator. Should I turn left down the hallway ahead?
Oracle: *ya*



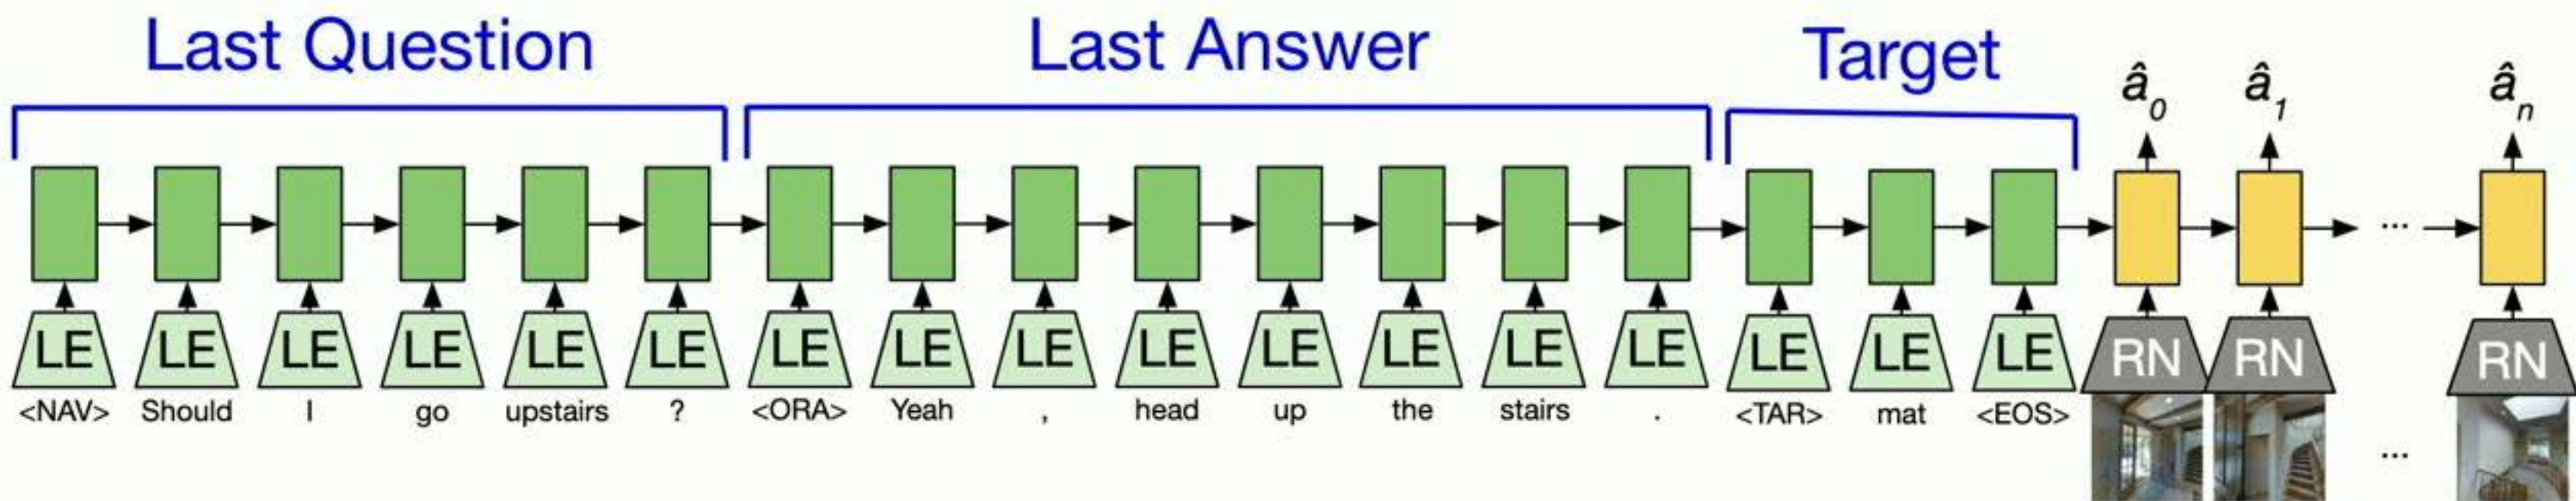
Initial, Sequence-to-Sequence Model

Oracle: Through the lobby. So go through the door next to the green towel. Go to the left door next to **the two yellow lights**.

Walk straight to the end of the hallway and stop


...

Navigator: Are these **the yellow lights** you were talking about?



Initial, Sequence-to-Sequence Model

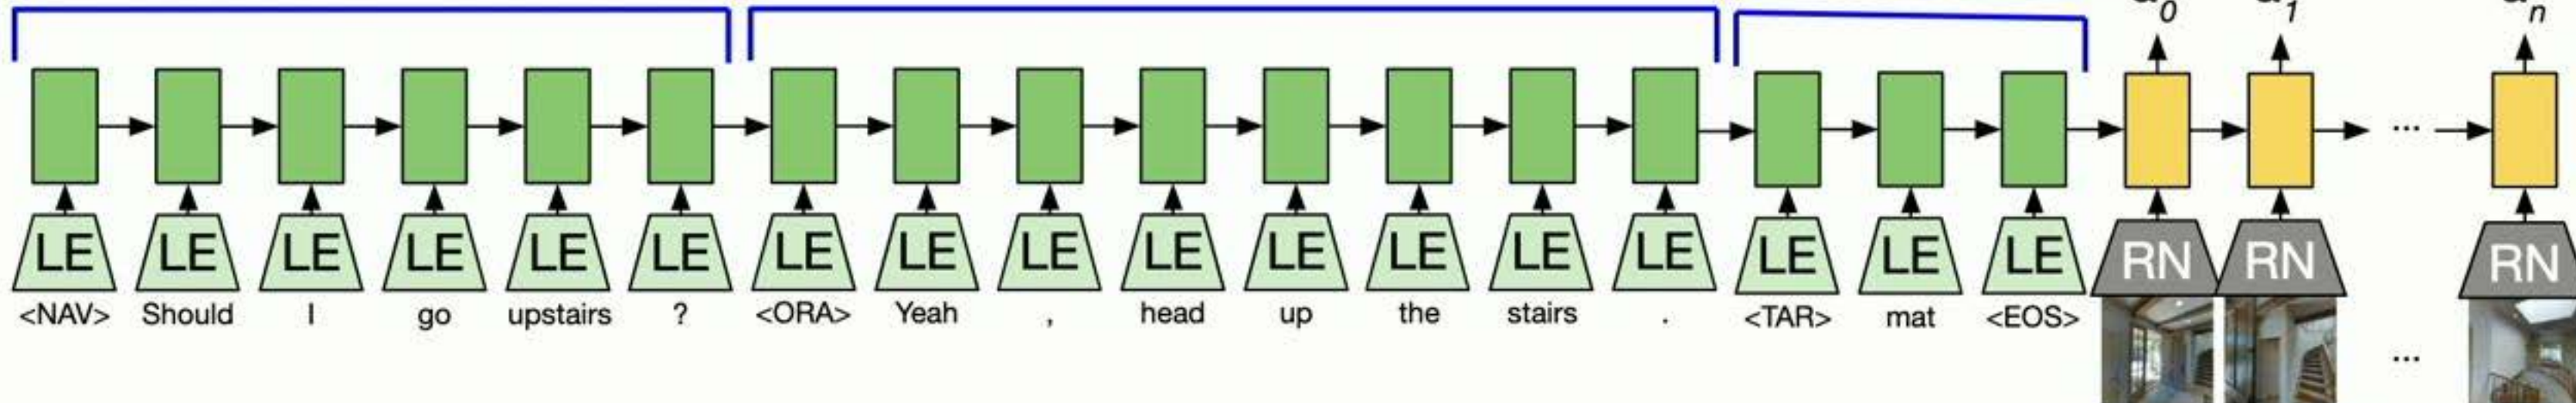
All Previous Questions and Answers




Last Question

Last Answer

Target




Evaluation - Validation (best perf over 200 epochs)



- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Initial, Sequence-to-Sequence Model

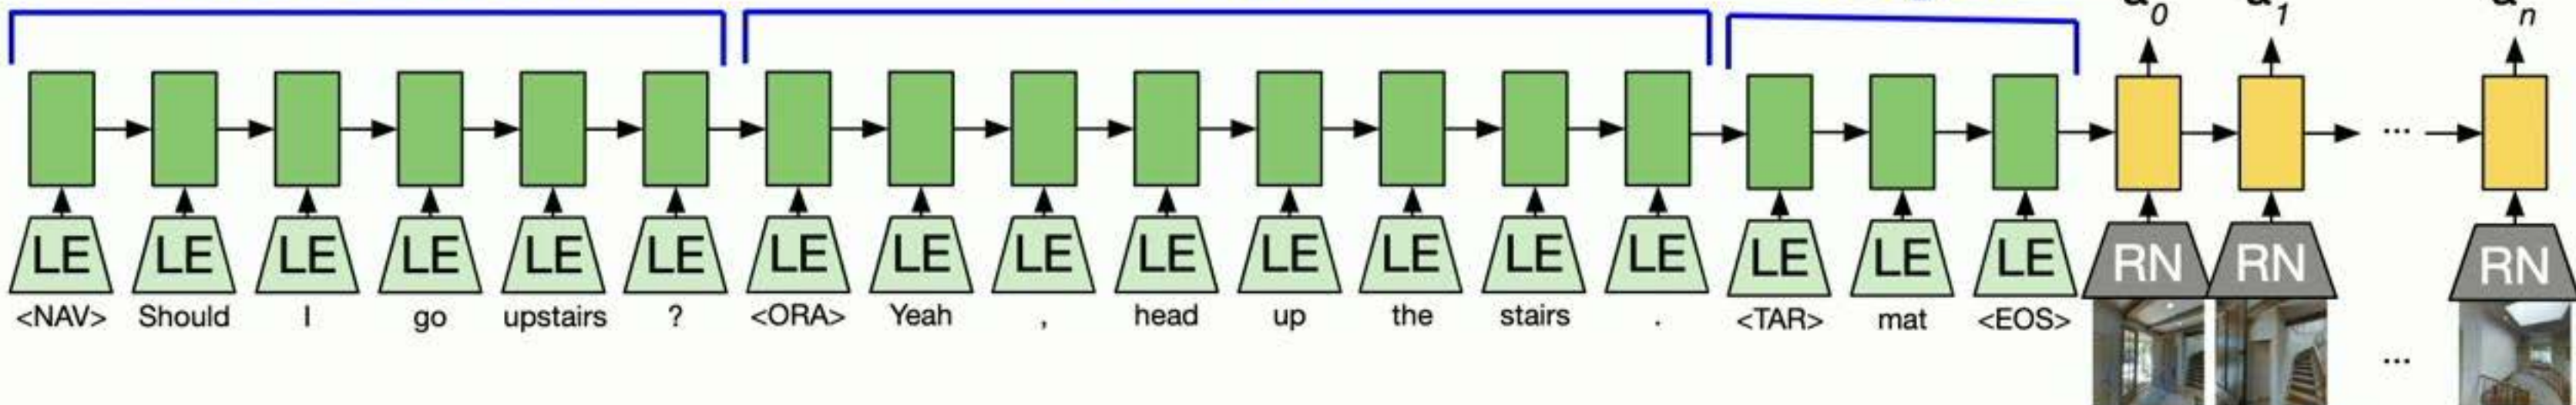
All Previous Questions and Answers




Last Question

Last Answer

Target




Evaluation - Validation (best perf over 200 epochs)



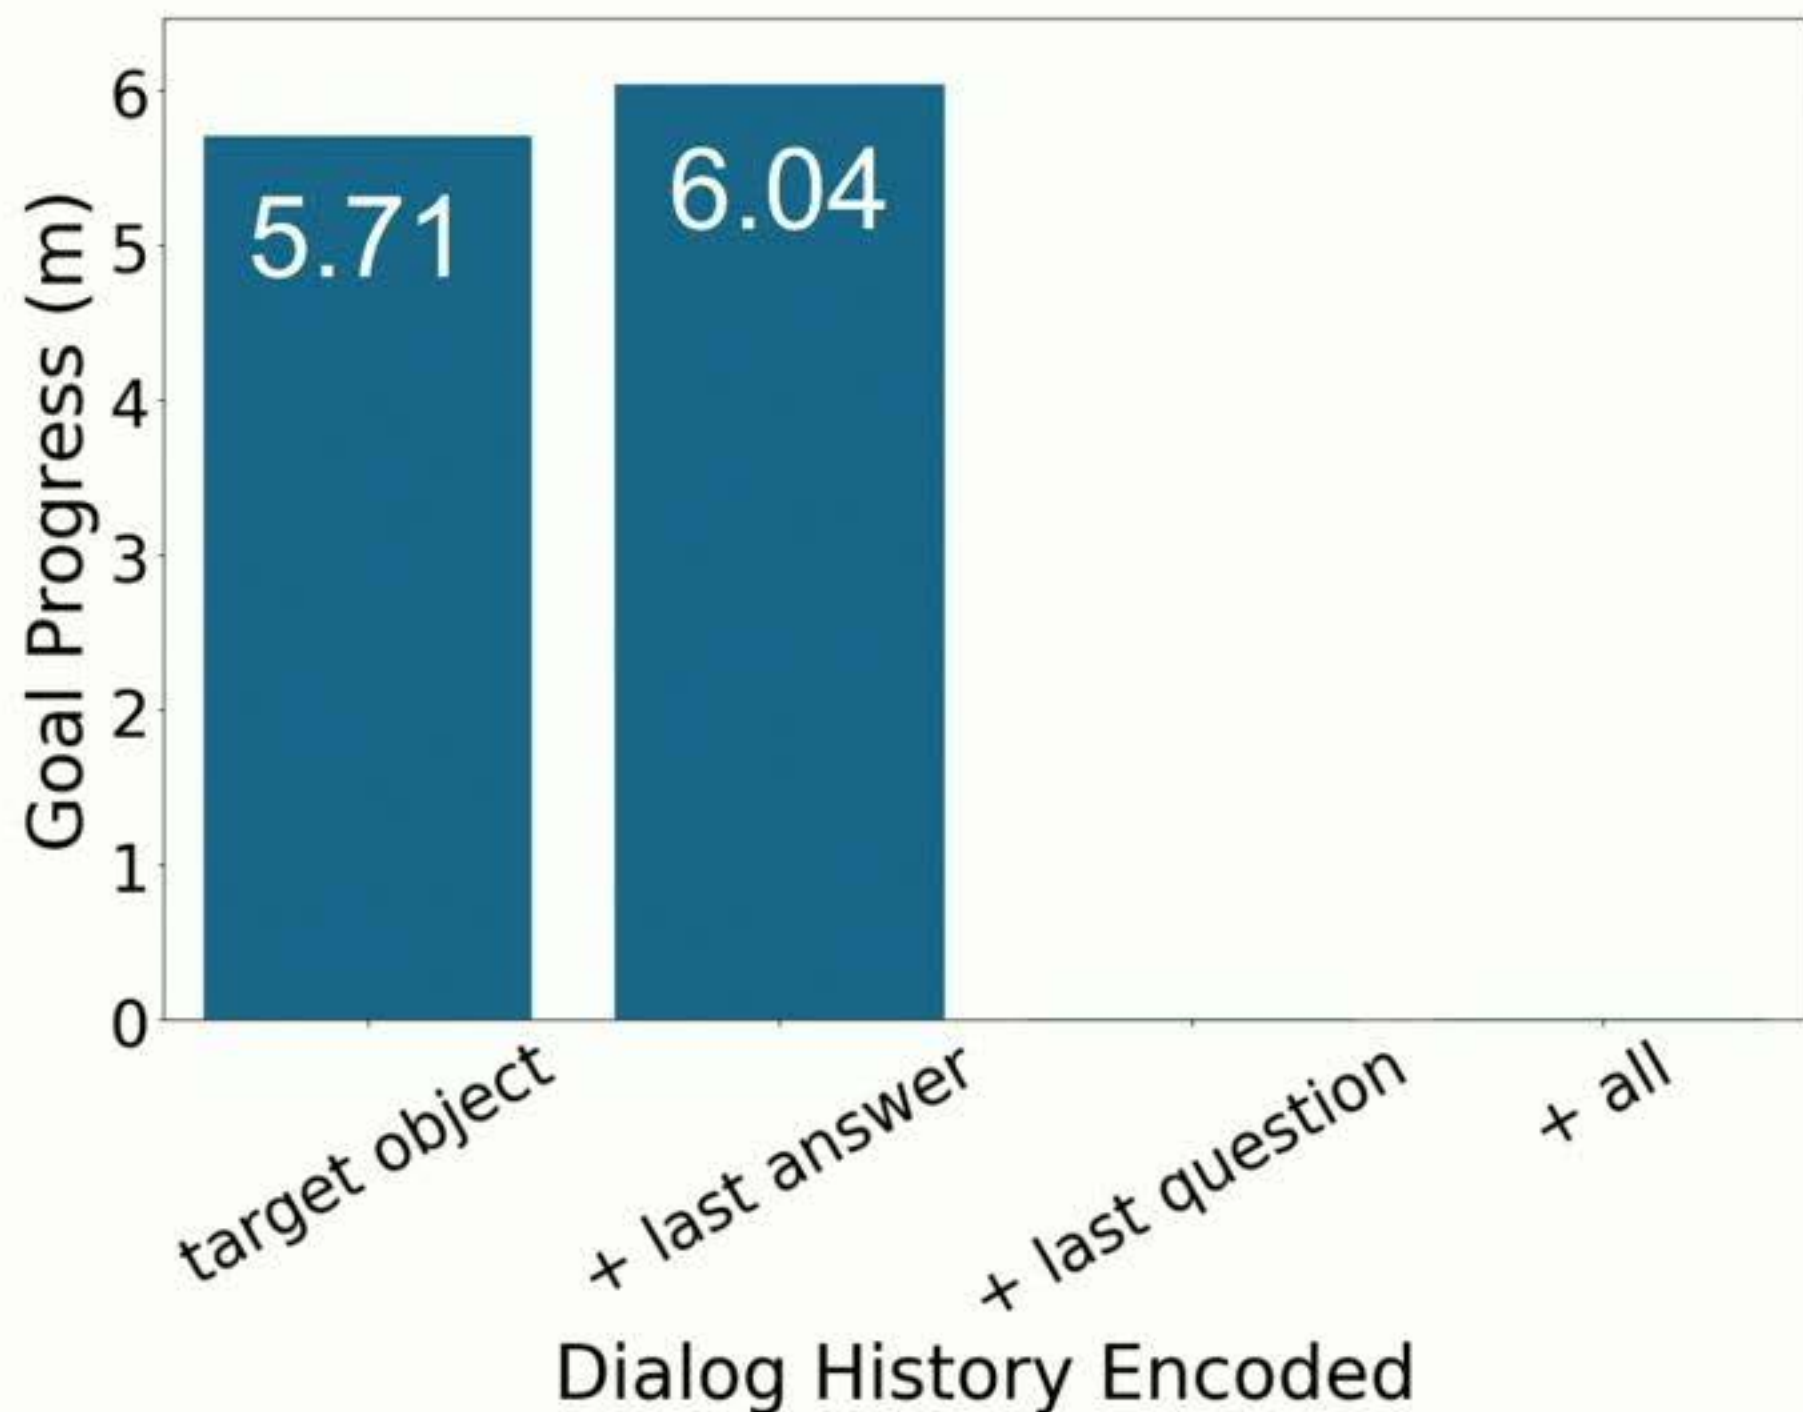
- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Validation (best perf over 200 epochs)




- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Validation (best perf over 200 epochs)




- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Validation (best perf over 200 epochs)



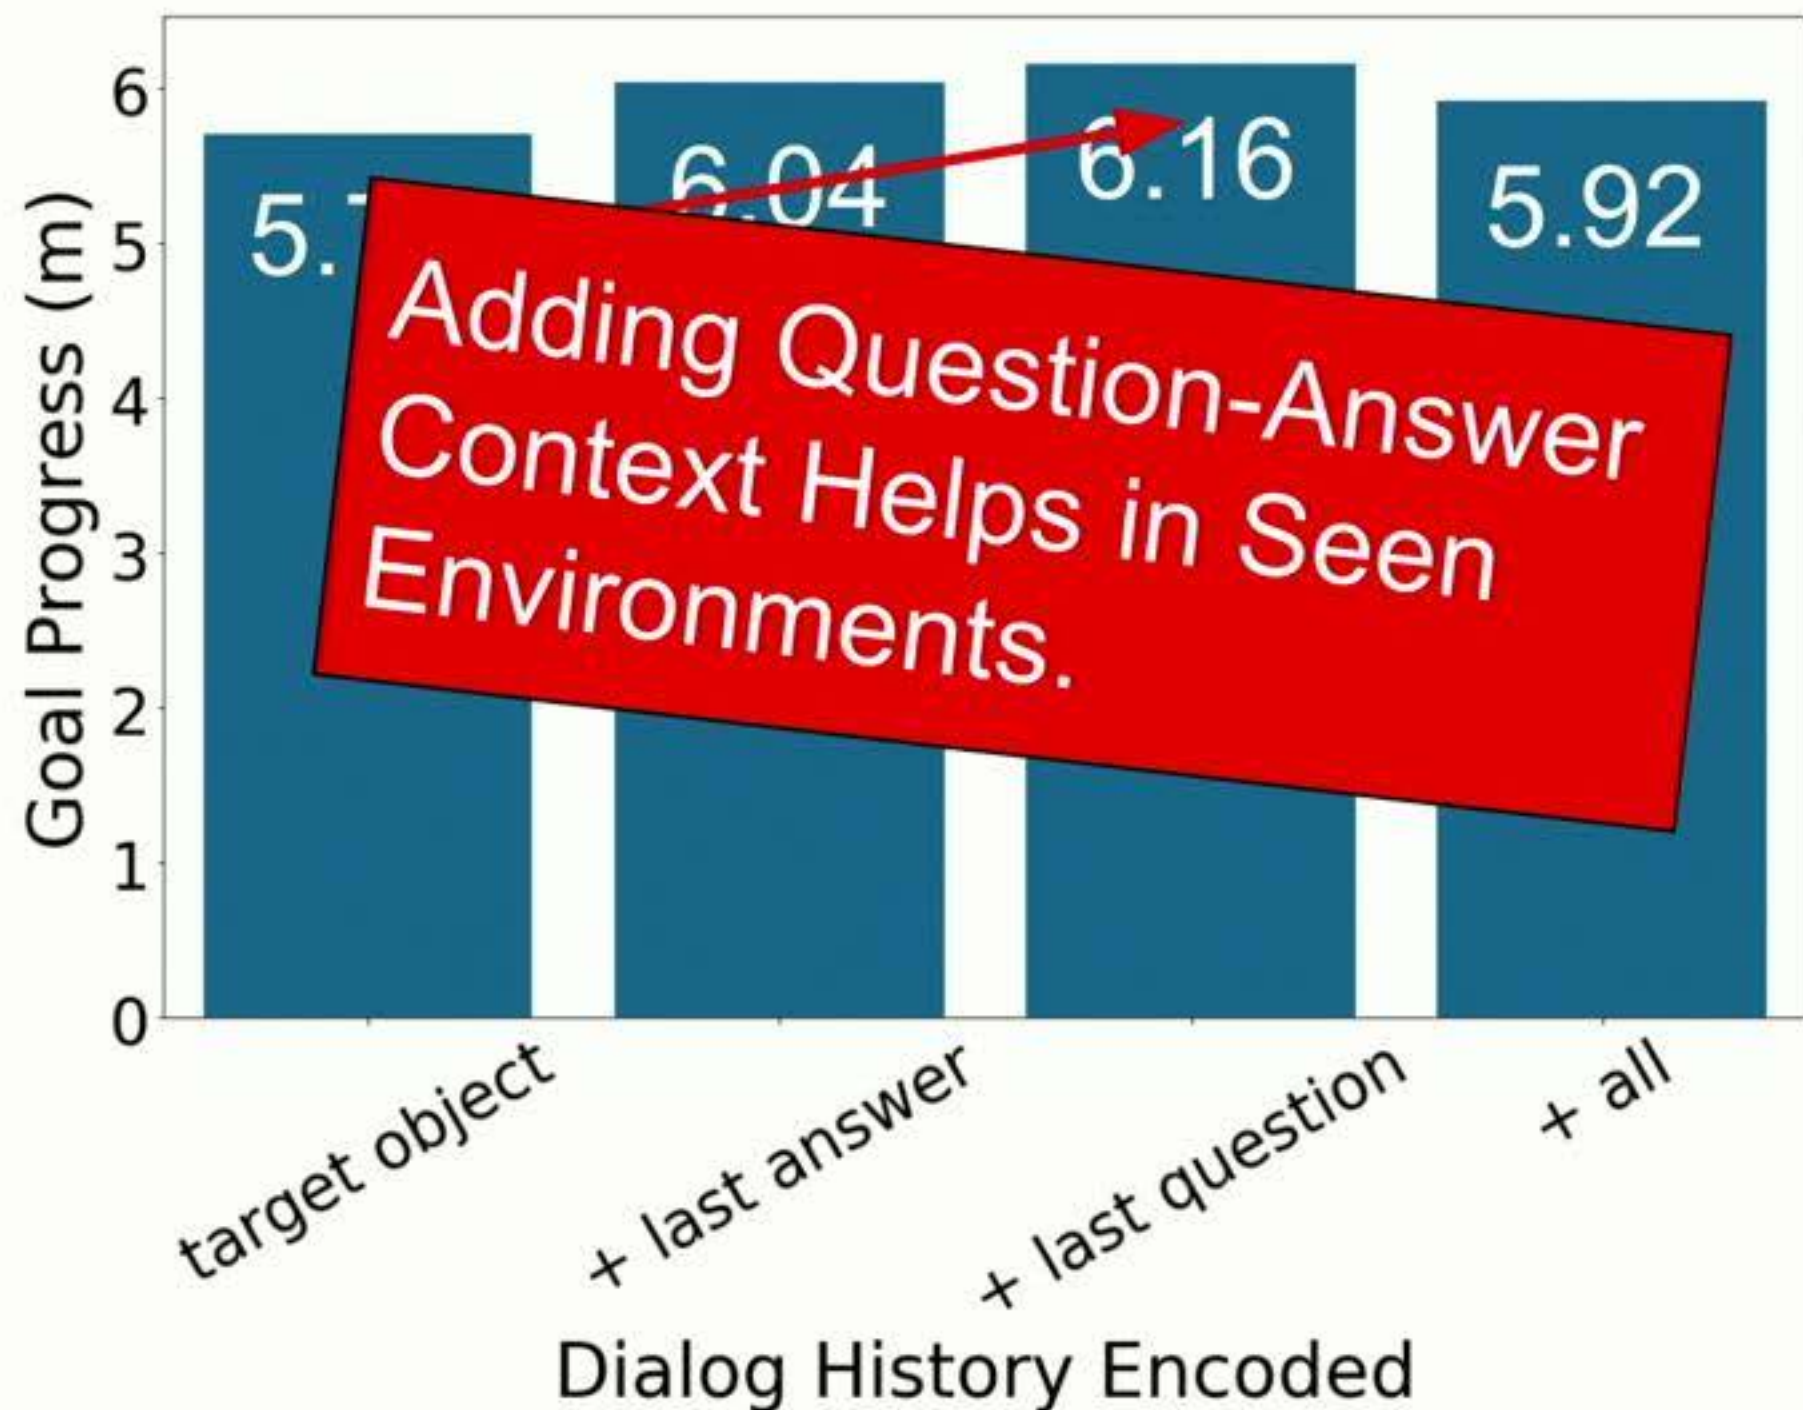
- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Validation (best perf over 200 epochs)




- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Validation (best perf over 200 epochs)




- **Seen Environments:**
 - Novel dialogs.
 - Houses seen at training time.

Evaluation - Test (epoch of best Val Unseen)




- **Unseen Envs:**
 - Novel dialogs.
 - Novel houses not seen during training.

Evaluation - Test (epoch of best Val Unseen)



- **Unseen Envs:**
 - Novel dialogs.
 - Novel houses not seen during training.


Evaluation - Test (epoch of best Val Unseen)



- **Unseen Envs:**

- Novel dialogs.
- Novel houses not seen during training.


Evaluation - Test (epoch of best Val Unseen)



- **Unseen Envs:**

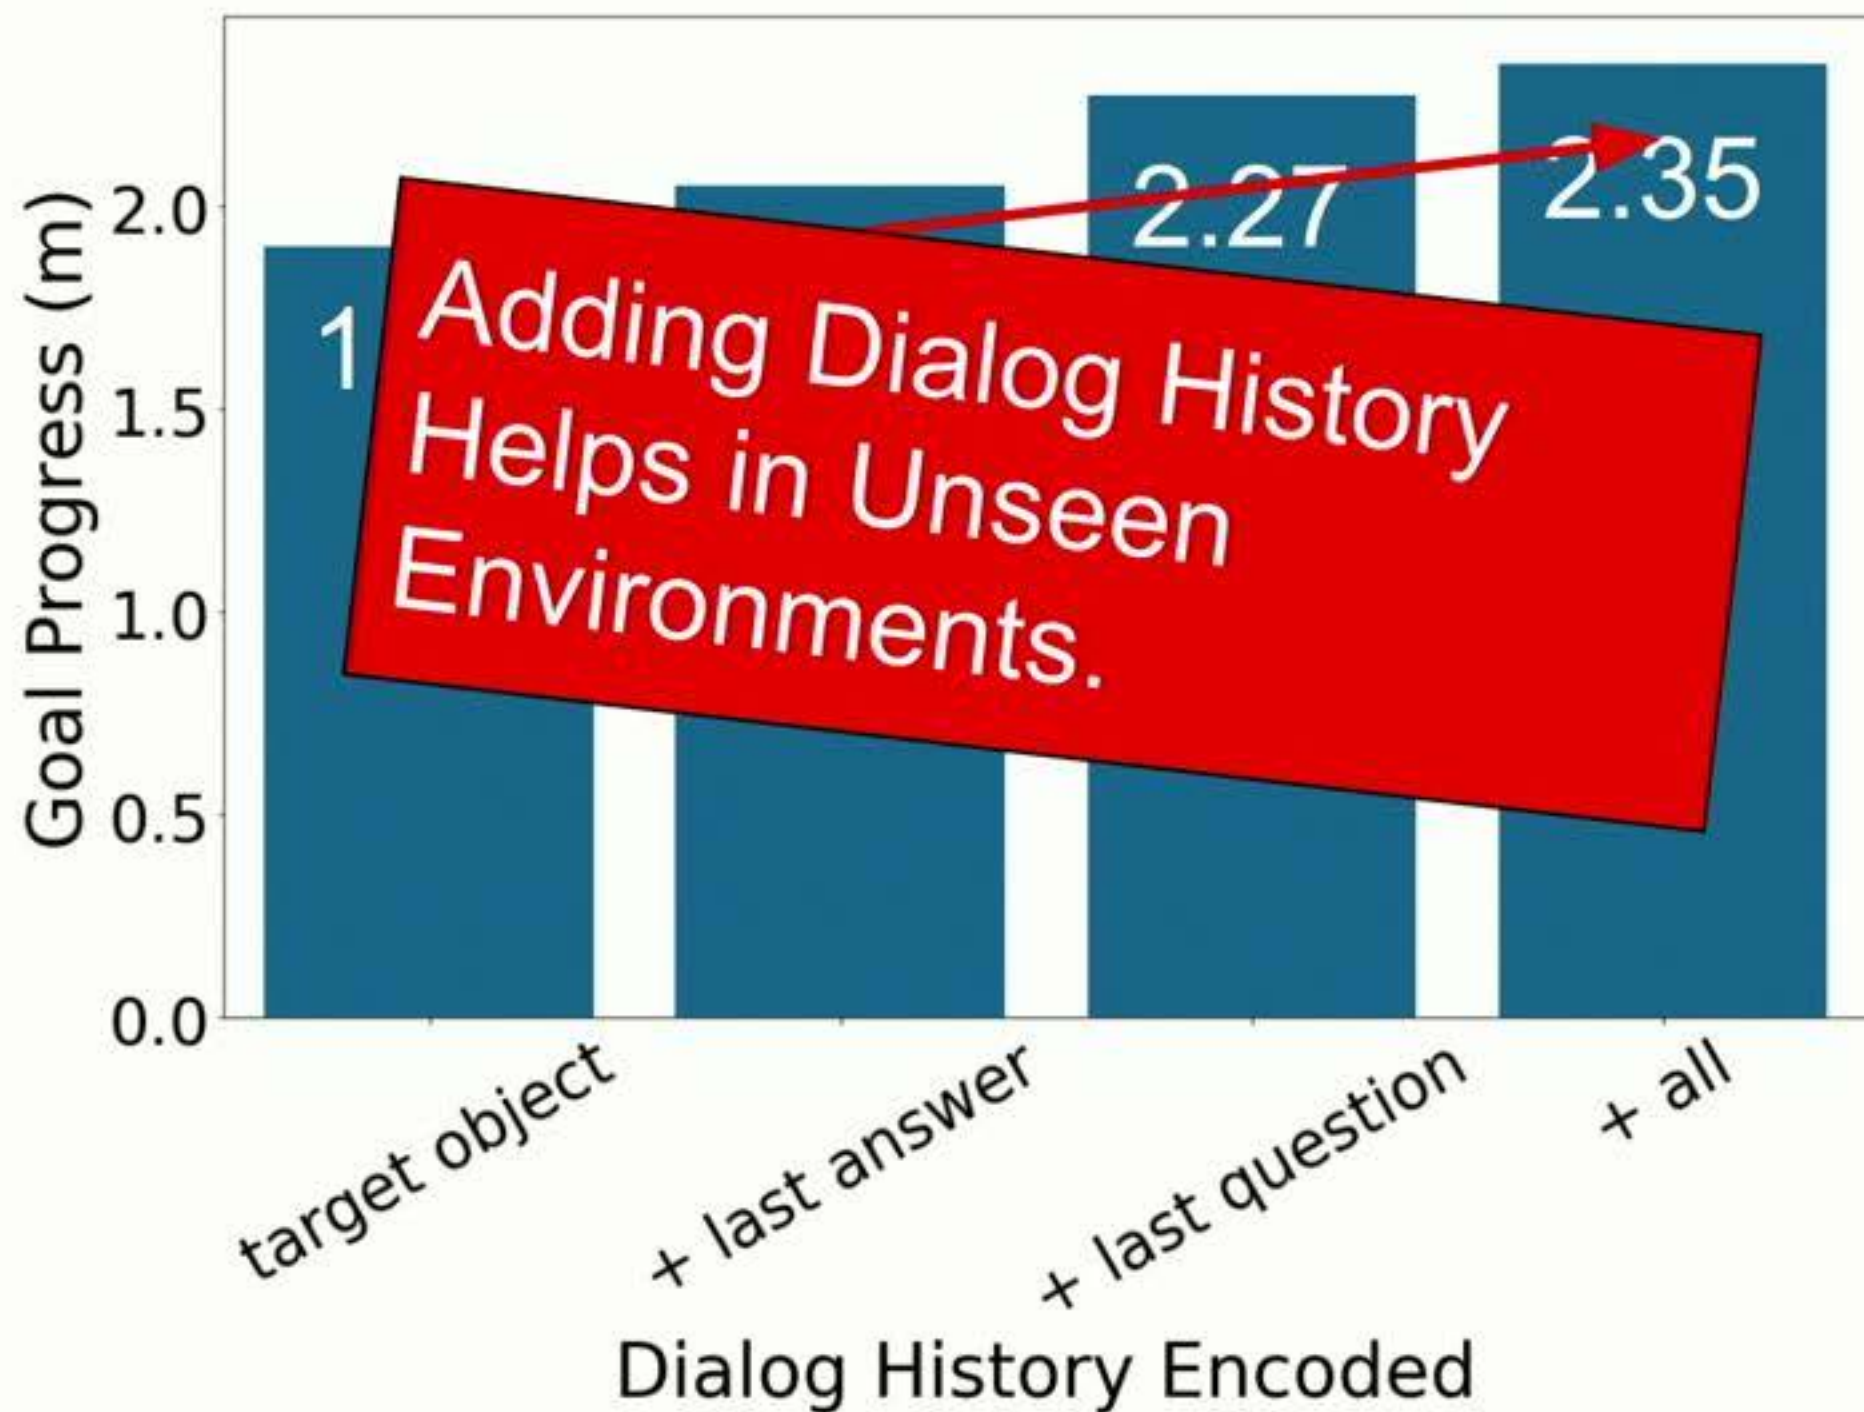
- Novel dialogs.
- Novel houses not seen during training.

Evaluation - Test (epoch of best Val Unseen)




- **Unseen Envs:**
 - Novel dialogs.
 - Novel houses not seen during training.

Evaluation - Test (epoch of best Val Unseen)




- **Unseen Envs:**
 - Novel dialogs.
 - Novel houses not seen during training.


Evaluation - Unimodal Baselines



Evaluation - Unimodal Baselines





Evaluation - Unimodal Baselines




Evaluation - Unimodal Baselines

Action-
only





Evaluation - Unimodal Baselines

Action-
only




Vis-
only





Evaluation - Unimodal Baselines

Action-
only




Vis-
only





Evaluation - Unimodal Baselines

Action-
only




Vis-
only




Evaluation - Unimodal Baselines



Action-
only



Vis-
only

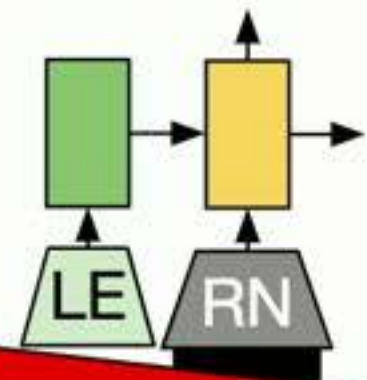


Lang-
only



Evaluation - Unimodal Baselines

Action-
only





Initial Model Uses
Multimodal Input in
Unseen Environments.

only




Lang-
only




Evaluation - Unimodal Baselines



Action-
only



Vis-
only




Lang-
only




Evaluation - Unimodal Baselines



Action-
only



Vis-
only

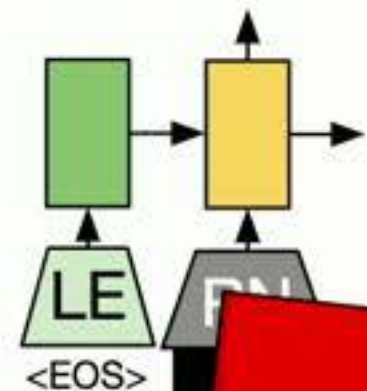


Lang-
only




Evaluation - Unimodal Baselines


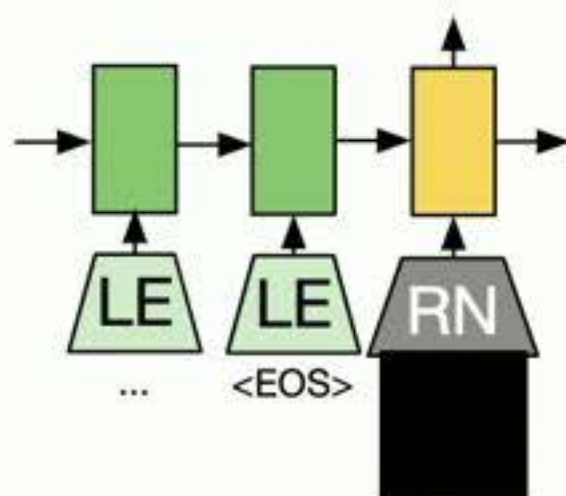
Action-
only



Vis-
only



Lang-
only



Lots of Headroom for
Future Models

Evaluation Lessons

Evaluation Lessons

- Dialog history:

Evaluation Lessons

- Dialog history:
 - Longer context leads to better performance.
 - Particularly helpful in unseen environments.

Evaluation Lessons

- Dialog history:
 - Longer context leads to better performance.
 - Particularly helpful in unseen environments.
- Unimodal baselines:
 - Initial model makes use of multimodal information

Evaluation Lessons

- Dialog history:
 - Longer context leads to better performance.
 - Particularly helpful in unseen environments.
- Unimodal baselines:
 - Initial model makes use of multimodal information

Evaluation Lessons

- Dialog history:
 - Longer context leads to better performance.
 - Particularly helpful in unseen environments.
- Unimodal baselines:
 - Initial model makes use of multimodal information
 - Multimodal most helpful in unseen environments.


Evaluation Lessons

- Dialog history:
 - Longer context leads to better performance.
 - Particularly helpful in unseen environments.
- Unimodal baselines:
 - Initial model makes use of multimodal information
 - Multimodal most helpful in unseen environments.
- Headroom remains for more nuanced models.

Outline


- Language grounding in visual environments
 - For navigation
 - Unimodal bias [Thomason et al., NAACL'19]
- Vision-and-Dialog Navigation [Thomason et al., *in sub*]
 - New dataset - CVDN
 - Navigation from dialog history
- **Next steps**

Incorporating Navigation History

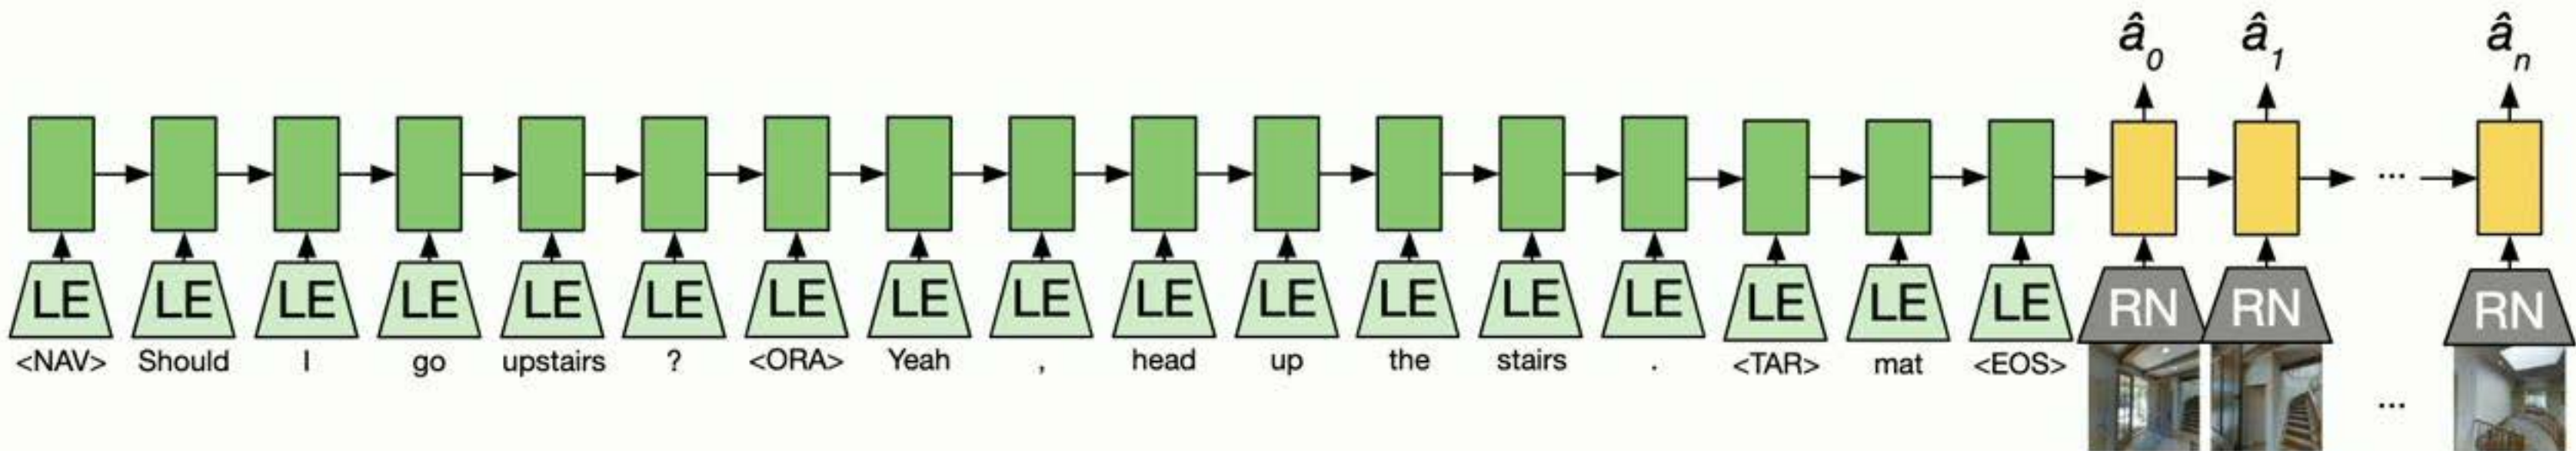



Incorporating Navigation History

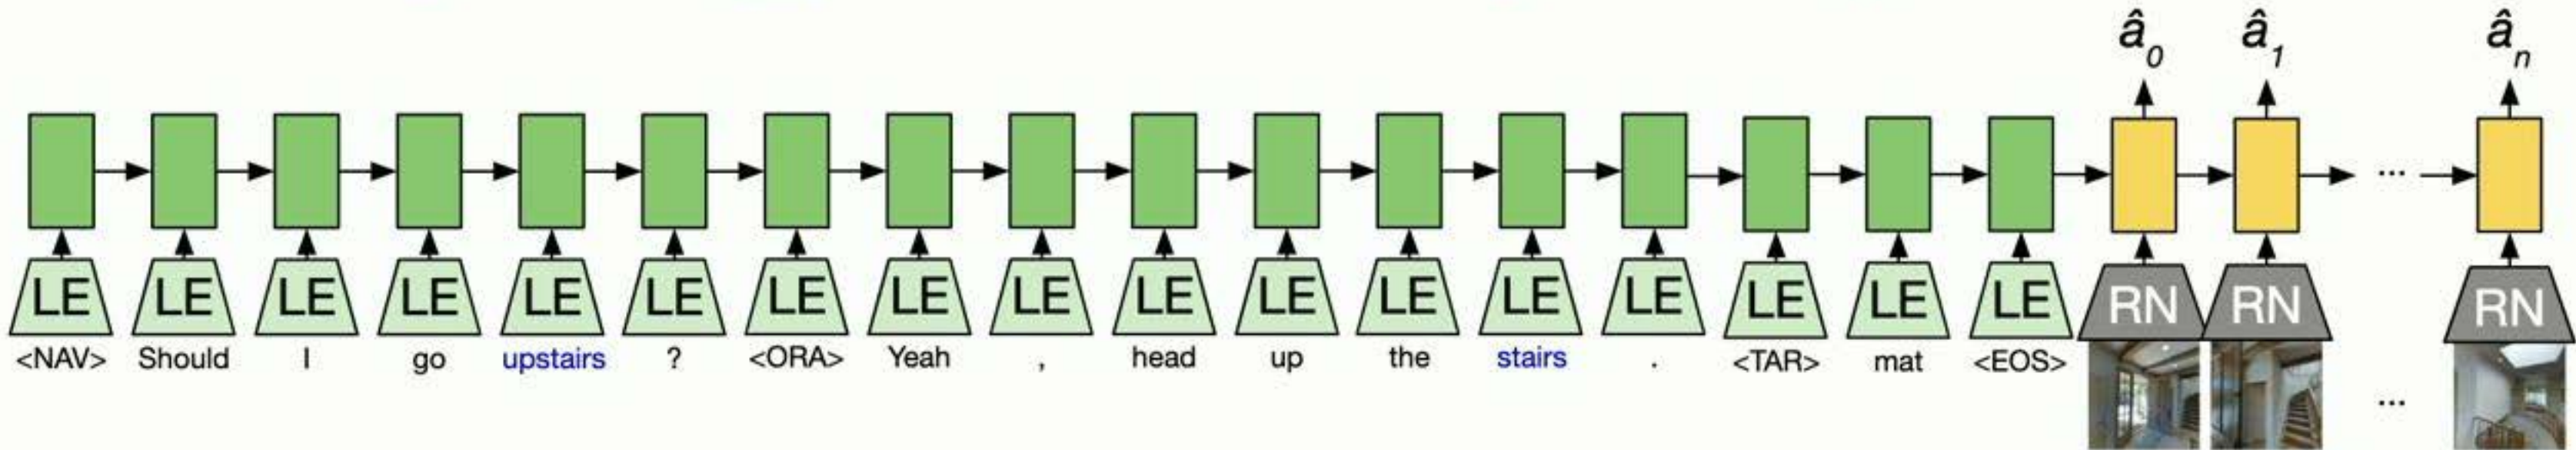

Oracle: You were there briefly but left. There is a turntable behind you a bit. Enter the bedroom next to it.



Incorporating Navigation History

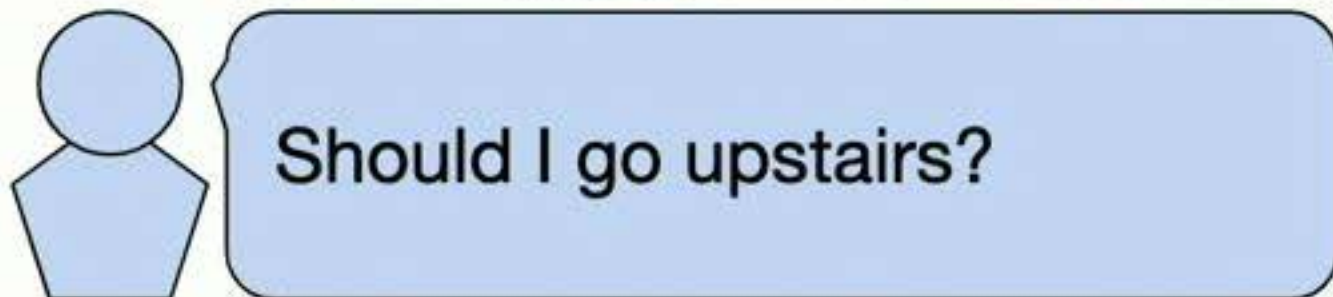
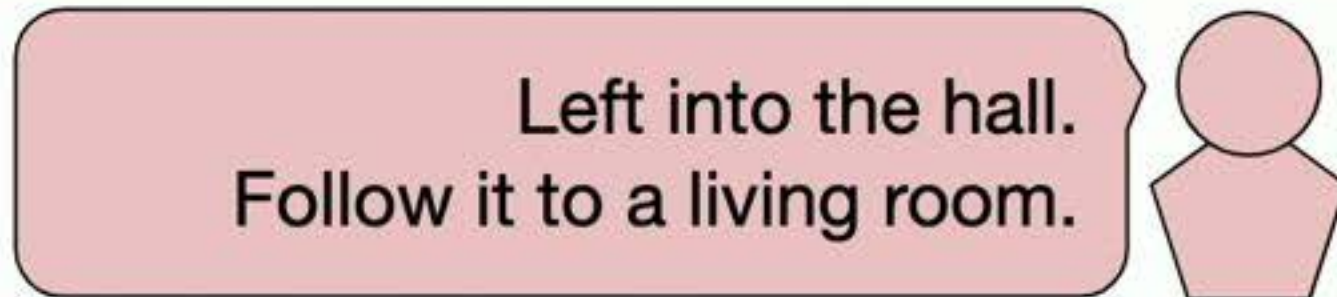
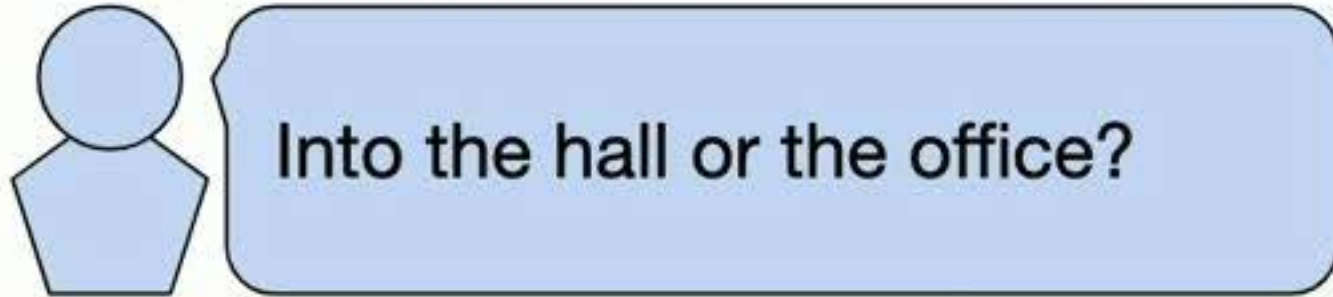


Incorporating Navigation History



Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Visible Only to the **Oracle**




Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Into the hall or the office?

Left into the hall.
Follow it to a living room.



Should I go upstairs?

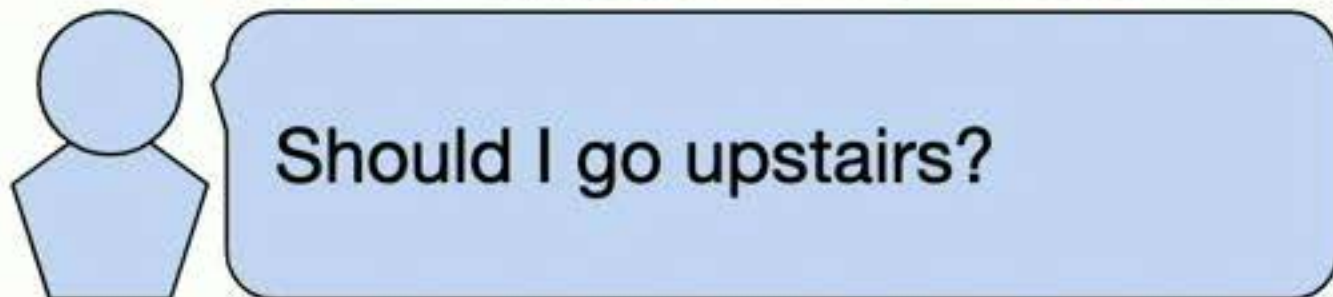
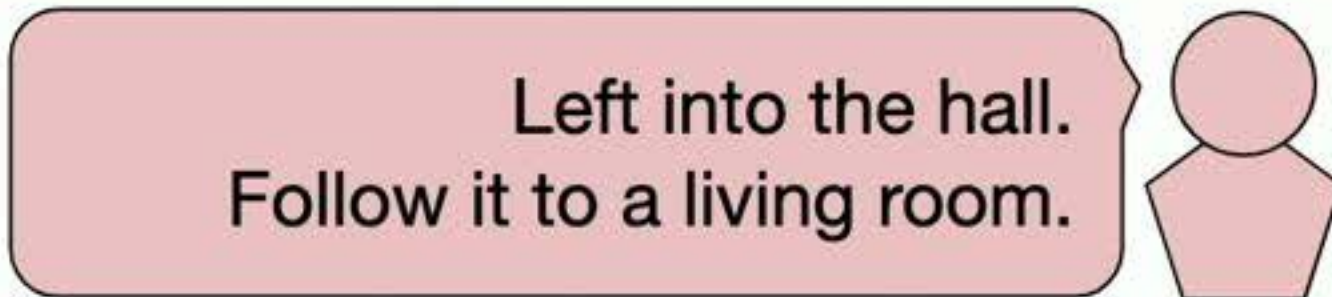
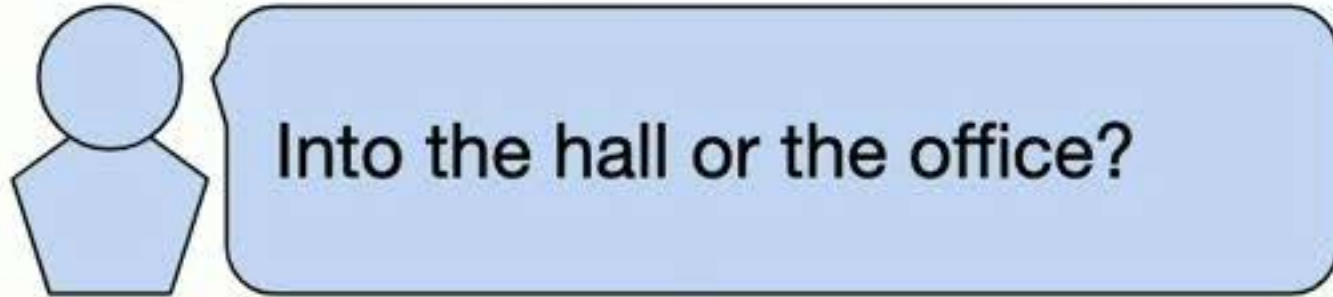
Navigation

Visible Only to the **Oracle**



Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Visible Only to the **Oracle**




Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Into the hall or the office?

Left into the hall.
Follow it to a living room.

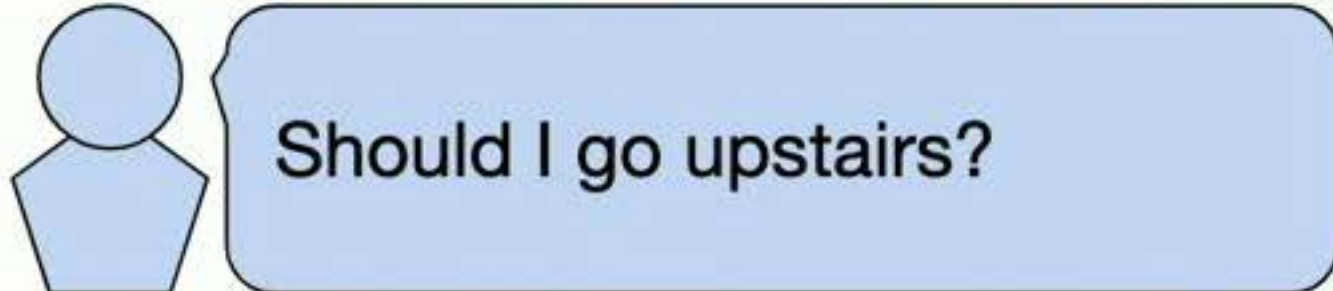
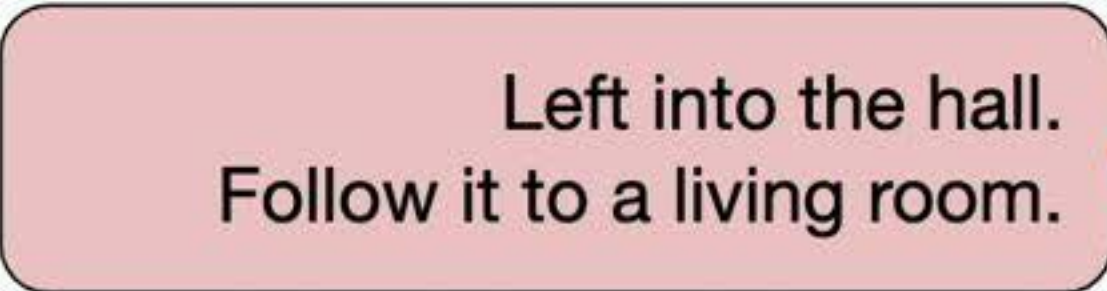
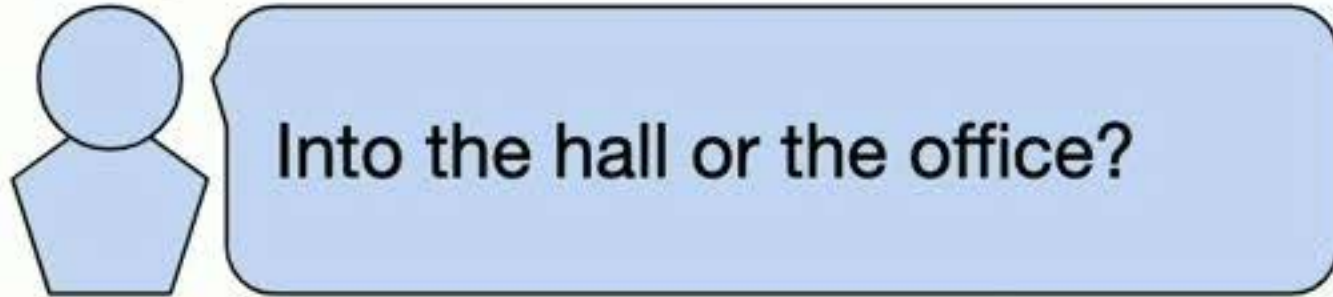


Should I go upstairs?

Question
Generation

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Visible Only to the **Oracle**



Question Answering

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Into the hall or the office?

Left into the hall.
Follow it to a living room.



Should I go upstairs?



Question Answering

Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Into the hall or the office?

Left into the hall.
Follow it to a living room.




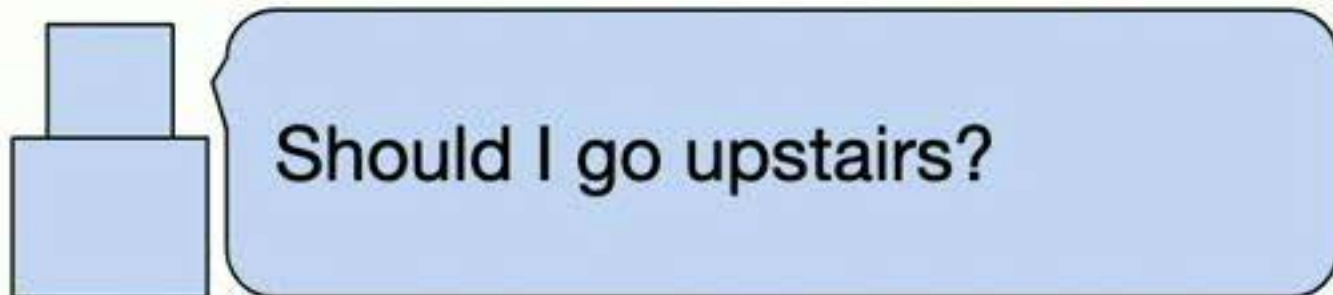
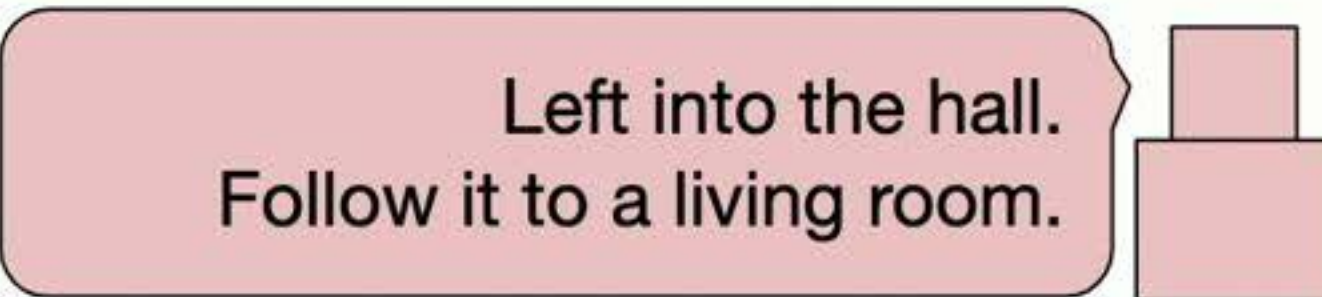
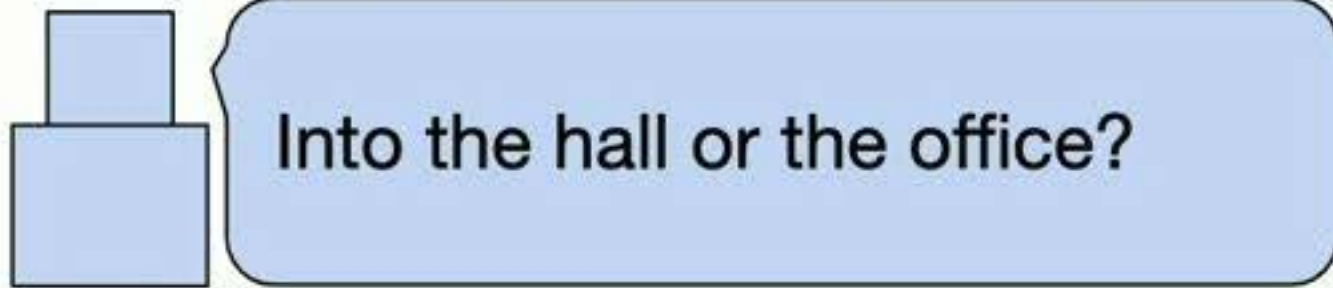
Should I go upstairs?

Visible Only to the **Oracle**



Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.



Navigation

Question
Generation

Visible Only to the **Oracle**



Visible to both **Navigator** and **Oracle**

Hint: The goal room contains a *mat*.




Into the hall or the office?

Left into the hall.
Follow it to a living room.




Should I go upstairs?



Navigation

Question
Generation


Visible Only to the **Oracle**



Question
Answering


Hint: The goal room contains a *mat*.

- Environment exploration



Navigation


Question
Generation



Question
Answering


Hint: The goal room contains a *mat*.

- Environment exploration



Navigation


Question
Generation



Question
Answering


Hint: The goal room contains a *mat*.

- Environment exploration
- Self-play



Navigation


Question
Generation



Question
Answering


Hint: The goal room contains a *mat*.

- Environment exploration
- Self-play
- “Language” evolution



Navigation


Question
Generation



Question
Answering


Hint: The goal room contains a *mat*.

- Environment exploration
- Self-play
- “Language” evolution
- Reinforcement Learning



Navigation

Question
Generation



Question
Answering

Bringing Robots from Industrial to Human Spaces

Industrial



Bringing Robots from Industrial to Human Spaces

Industrial



Human



Bringing Robots from Industrial to Human Spaces

Industrial



Natural Language



Navigation



Additional Safety



Robust Perception



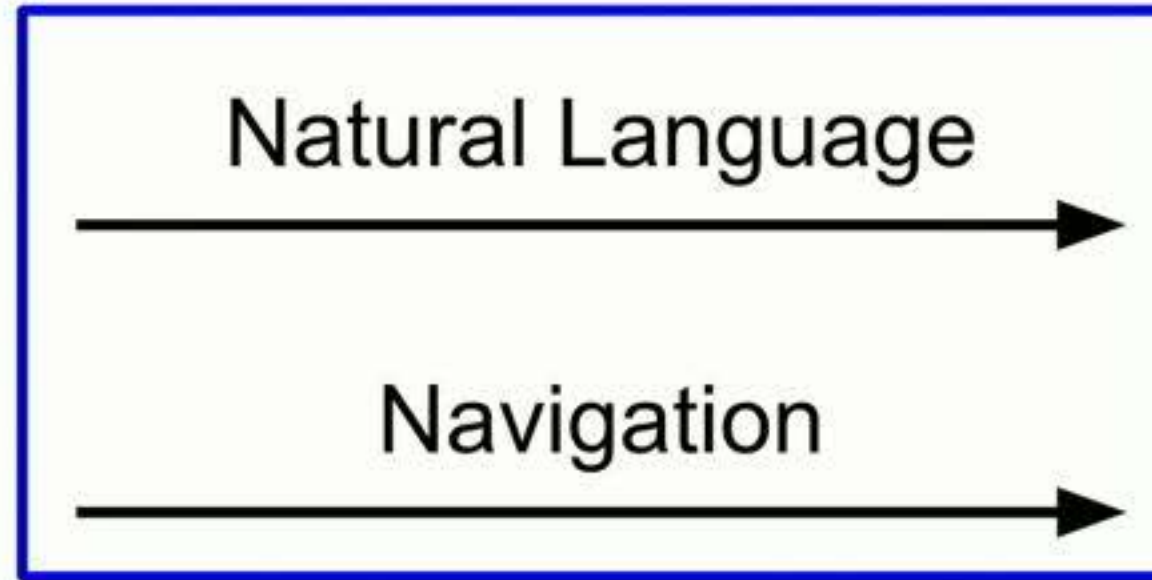
⋮

Human




Bringing Robots from Industrial to Human Spaces

Industrial



Human



⋮



Takeaways

Takeaways

- For vision-and-language navigation:
 - Unimodal ablations expose dataset bias.

Takeaways

- For vision-and-language navigation:
 - Unimodal ablations expose dataset bias.

Takeaways

- For vision-and-language navigation:
 - Unimodal ablations expose dataset bias.
- For vision-and-dialog navigation:

Takeaways

- For vision-and-language navigation:
 - Unimodal ablations expose dataset bias.
- For vision-and-dialog navigation:
 - Cooperative dialog facilitates a mix of high- and low-level language.

Takeaways

- For vision-and-language navigation:
 - Unimodal ablations expose dataset bias.
- For vision-and-dialog navigation:
 - Cooperative dialog facilitates a mix of high- and low-level language.
 - Dialog context helps agents infer better navigation actions.

Vision-and-Dialog Navigation

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]
- Cooperative Vision-and-Dialog Navigation dataset

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]
- Cooperative Vision-and-Dialog Navigation dataset
 - Data+code: <https://github.com/mmurray/cvdn/>

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]
- Cooperative Vision-and-Dialog Navigation dataset
 - Data+code: <https://github.com/mmurray/cvdn/>
 - Live demo: <https://cvdn.dev/>

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]
- Cooperative Vision-and-Dialog Navigation dataset
 - Data+code: <https://github.com/mmurray/cvdn/>
 - Live demo: <https://cvdn.dev/>
- Navigation from Dialog History task

Vision-and-Dialog Navigation

- Jesse Thomason, Daniel Gordon, and Yonatan Bisk.
“Shifting the Baseline: Single Modality Performance on Visual Navigation and QA”,
NAACL’19 [<https://arxiv.org/abs/1811.00613>]
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
“Vision-and-Dialog Navigation”, *in submission* [<https://arxiv.org/abs/1907.04957>]
- Cooperative Vision-and-Dialog Navigation dataset
 - Data+code: <https://github.com/mmurray/cvdn/>
 - Live demo: <https://cvdn.dev/>
- Navigation from Dialog History task
 - Leaderboard coming soon!

Future Work: More Expressive Simulator



- High-level instructions.
 - With optional, accompanying low-level.
- Navigation
+ manipulation
+ interaction.
- No dialog to support clarifications.

“Put a slice of bread in the microwave.”