

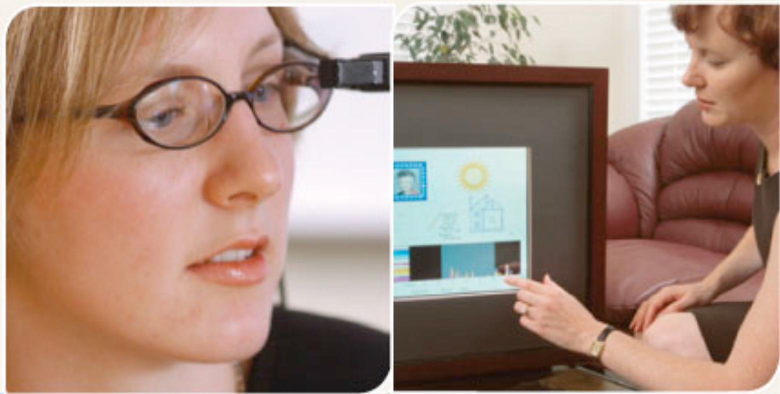

Automated Story Generation as a Lens for Fundamental Artificial Intelligence

Mark Riedl

riedl@cc.gatech.edu

@mark_riedl

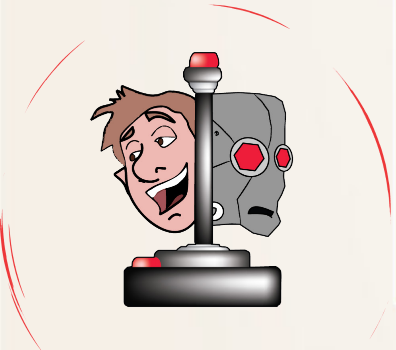

Georgia Tech College of Computing

School of Interactive Computing

Storytelling

Georgia Tech College of Computing
School of Interactive Computing

Storytelling

- Narrative is the fundamental means by which we organize, understand, and explain the world

Storytelling

- Narrative is the fundamental means by which we organize, understand, and explain the world

- **Narrative intelligence:**
  the ability to craft, tell, understand, and affectively respond to stories

Georgia Tech College of Computing
School of Interactive Computing

Story telling and understanding

Story telling and understanding

- **Story telling:** careful sequencing of words in order to achieve a desired effect (efficient transmission of experience, entertain, teach, etc.)

Story telling and understanding

- **Story telling:** careful sequencing of words in order to achieve a desired effect (efficient transmission of experience, entertain, teach, etc.)

- **Story understanding:** processes by which we “unpack” the information encoded in a story

Story telling and understanding

- **Story telling:** careful sequencing of words in order to achieve a desired effect (efficient transmission of experience, entertain, teach, etc.)

- **Story understanding:** processes by which we “unpack” the information encoded in a story

For sale: baby shoes. Never worn.

Georgia Tech College of Computing
School of Interactive Computing

Story understanding

John entered the restaurant and ordered food. He looked across the room and saw an old friend, Sally. They put their tables together. Later that evening, John and Sally paid and left together.

Story generation

With sweaty palms and heart racing, John drove to Sally's house for their first date. Sally, her pretty white dress flowing in the wind, carefully entered John's car. John and Sally drove to the movie theater. John and Sally parked the car in the parking lot. Wanting to feel prepared, John had already bought tickets to the movie in advance. A pale-faced usher stood before the door; John showed the tickets and the couple entered. Sally was thirsty so John hurried to buy drinks before the movie started. John and Sally found two good seats near the back. John sat down and raised the arm rest so that he and Sally could snuggle. John paid more attention to Sally while the movie rolled and nervously sipped his drink. Finally working up the courage to do so, John extended his arm to embrace Sally. He was relieved and ecstatic to feel her move closer to him in response. Sally stood up to use the restroom during the movie, smiling coyly at John before that exit. John and Sally also held hands throughout the movie, even though John's hands were sweaty. John and Sally slowly got up from their seats. Still holding hands, John walked Sally back to his car through the maze of people all scurrying out of the theater. The bright sunshine temporarily blinded John as he opened the doors and held them for Sally as they left the dark theater and stepped back out onto the street. John let go of Sally's hand and opened the passenger side door of his car for her but instead of entering the car, she stepped forward, embraced him, and gave him a large kiss. John drove Sally back to her home.

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Story generation

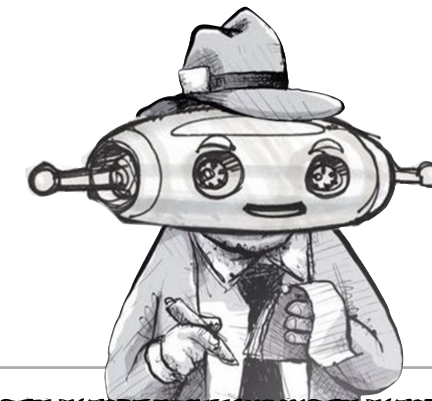

Game Summary:

The Yorktown Patriots triumphed over the visiting Wilson Tigers in a close game on Thursday, 20-14.

The game began with a scoreless first quarter.

In the second quarter, The Patriots' Paul Dalzell was the first to put points on the board with a two-yard touchdown reception off a pass from quarterback William Porter.

Wilson was behind Yorktown 7-0 heading into the second half. Wilson's Anton Reed tied the score with a two-yard touchdown run. The Patriots took the lead from Wilson with a two-yard touchdown run by Tanner Wall. The Patriots scored again on Adam Luncher's 29-yard field goal.

Yorktown maintained their lead going into the fourth quarter, 17-7. The Patriots extended their lead over the Tigers on Luncher's 27-yard field goal. Wilson cut into the Patriots' lead with a three-yard touchdown run by Amir Gerald. The game ended with Yorktown defeating Wilson, 20-14.

Source: Washington Post

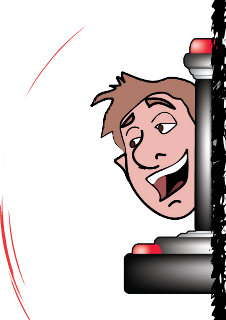

Explanation

- Experts' explanations are stories
- End users seem to prefer these

Ehsan et al. IUI 2019 conference.

Explanation

- Experts' explanations are stories
- End users seem to prefer these

I had cars to the left and in front of me so I needed to move to the right to avoid them

I moved right to be more *centered*. This way I have *more time to react* if a car comes from either side.

Ehsan et al. IUI 2019 conference.

Georgia Tech College of Computing School of Interactive Computing

Machine enculturation

- Human cultural values are implicitly encoded in stories told by members of a culture
- Mine social conventions from stories
- Act consistently with learned social conventions

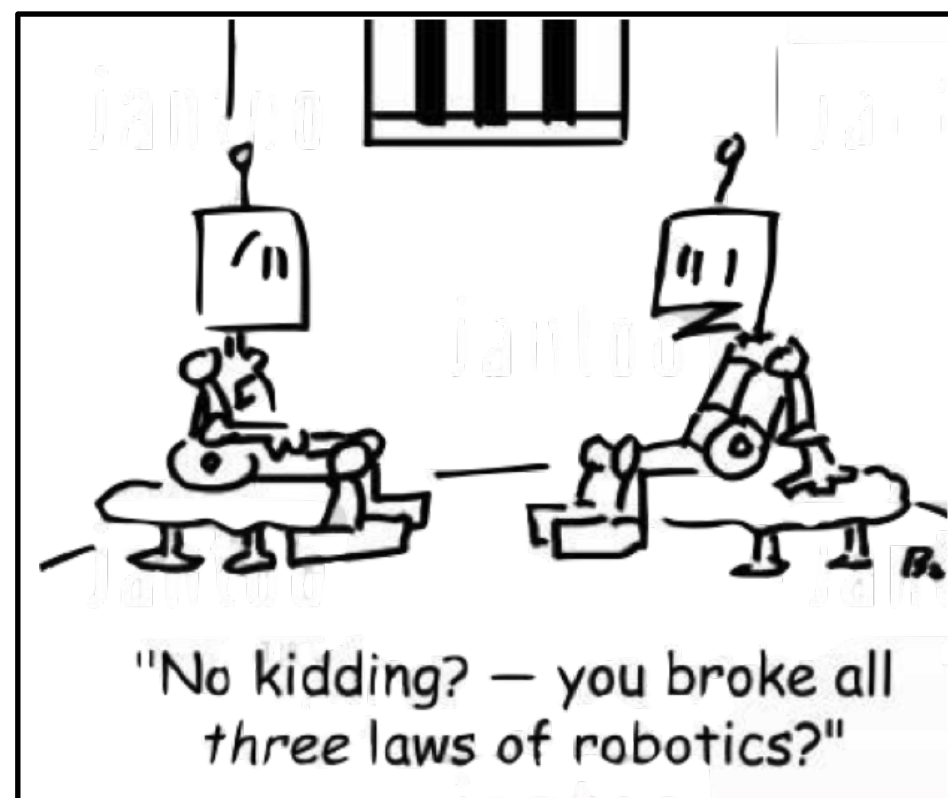

Peng & Riedl. arXiv:2001.08764
Frazier & Riedl. arXiv:1912.03553
Harrison & Riedl. AIIDE 2016 Conference

Automated story generation

Automated story generation

- Stories are everywhere—entertainment, education, everyday conversation

Automated story generation

- Stories are everywhere—entertainment, education, everyday conversation
- Means by which we share experiences, create rapport, generate affect

Georgia Tech College of Computing
School of Interactive Computing

Automated story generation

- Stories are everywhere—entertainment, education, everyday conversation
- Means by which we share experiences, create rapport, generate affect
- Planning with language; communicative intent; language understanding; sociocultural & commonsense knowledge; theory of mind

Automated story generation

- Stories are everywhere—entertainment, education, everyday conversation
- Means by which we share experiences, create rapport, generate affect
- Planning with language; communicative intent; language understanding; sociocultural & commonsense knowledge; theory of mind
- Creative evaluative task

Can computers tell stories?

Grimes, c. 1960s

A LION HAS BEEN IN TROUBLE FOR A LONG TIME. A DOG STEALS SOMETHING THAT BELONGS TO THE LION. THE HERO, LION, KILLS THE VILLAIN, DOG, WITHOUT A FIGHT. THE HERO, LION, THUS IS ABLE TO GET HIS POSSESSION BACK.

Generative grammar

Talespin, Meehan, 1975

ONCE UPON A TIME GEORGE ANT LIVED NEAR A PATCH OF GROUND. THERE WAS A NEST IN AN ASH TREE. WILMA BIRD LIVED IN THE NEST. THERE WAS SOME WATER IN A RIVER. WILMA KNEW THAT THE WATER WAS IN THE RIVER. GEORGE KNEW THAT THE WATER WAS IN THE RIVER. ONE DAY WILMA WAS VERY THIRSTY. WILMA WANTED TO GET NEAR SOME WATER. WILMA FLEW FROM HER NEST ACROSS A MEADOW THROUGH A VALLEY TO THE RIVER. WILMA DRANK THE WATER. WILMA WAS NOT THIRSTY ANY MORE.

Planning with scripts/schemas

Story generation 1.0

Symbolic systems

Narratives as plans

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds

Riedl & Young. Journal of AI Research, 2010.

Georgia Tech College of Computing
School of Interactive Computing

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds
- Actions have logical causal constraints
  - Pre-conditions
  - Post-conditions

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds
- Actions have logical causal constraints
  - Pre-conditions
  - Post-conditions

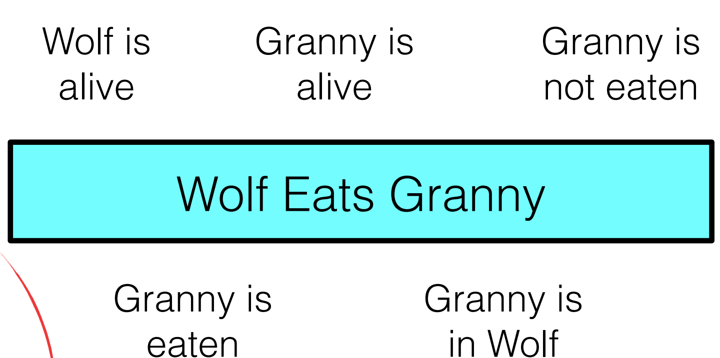

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds
- Actions have logical causal constraints
  - Pre-conditions
  - Post-conditions

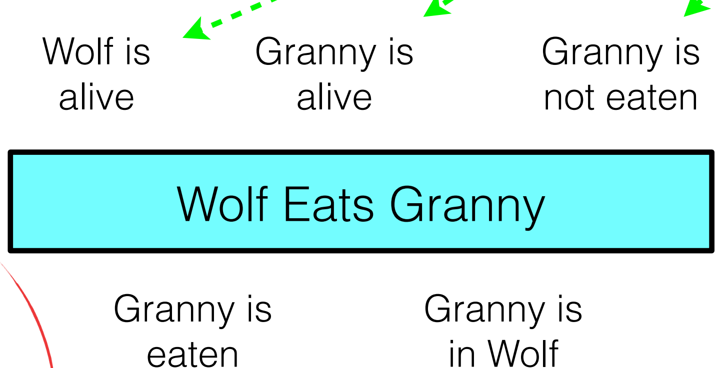

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds
- Actions have logical causal constraints
  - Pre-conditions
  - Post-conditions

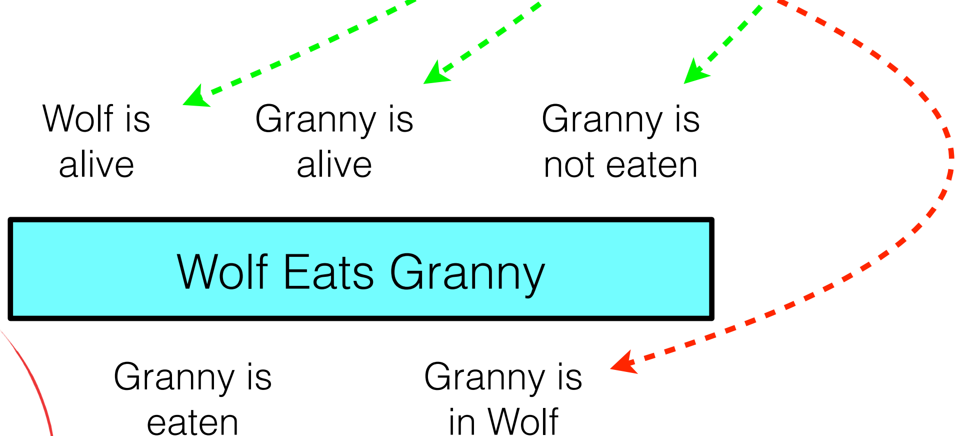

Riedl & Young. Journal of AI Research, 2010.

Narratives as plans

- Creative writing is a problem-solving activity
- Planning: find a sequence of actions that transforms the initial state into a state in which the goal situation holds
- Actions have logical causal constraints
  - Pre-conditions
  - Post-conditions

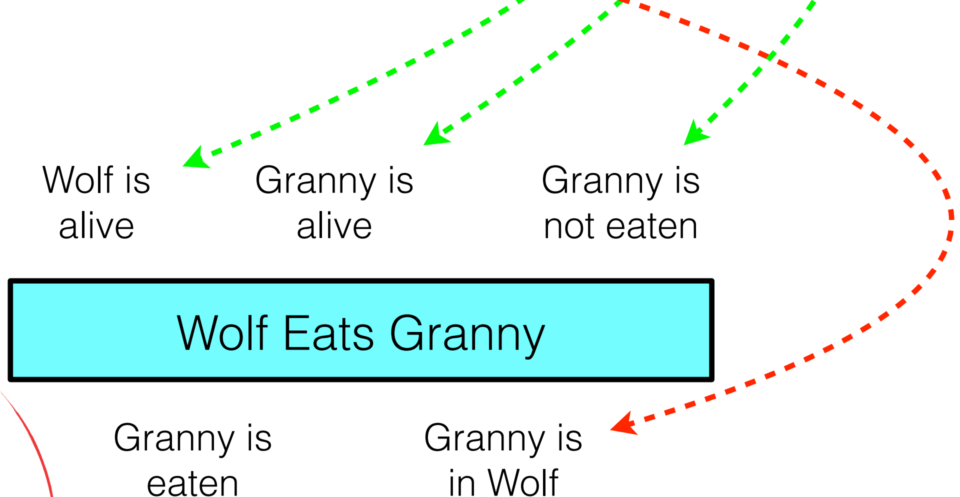


```
(define (action eat)
  :parameters (?monster ?victim)
  :constraints ((monster ?monster) (person ?victim))
  :precondition ((knows ?monster ?victim)
                 (alive ?monster) (alive ?victim)
                 (:not (eaten ?victim))
                 (:not (asleep ?monster))
                 (:neq ?monster ?victim))
 :effect ((eaten ?victim) (in ?victim ?monster)
         (full ?monster)))
```

Riedl & Young. Journal of AI Research, 2010.

Riedl & Young. Journal of AI Research, 2010.

Granny is eaten

Red is eaten

Outcome

Riedl & Young. Journal of AI Research, 2010.

Initial State

Granny is eaten

Red is eaten

Outcome

Riedl & Young. Journal of AI Research, 2010.

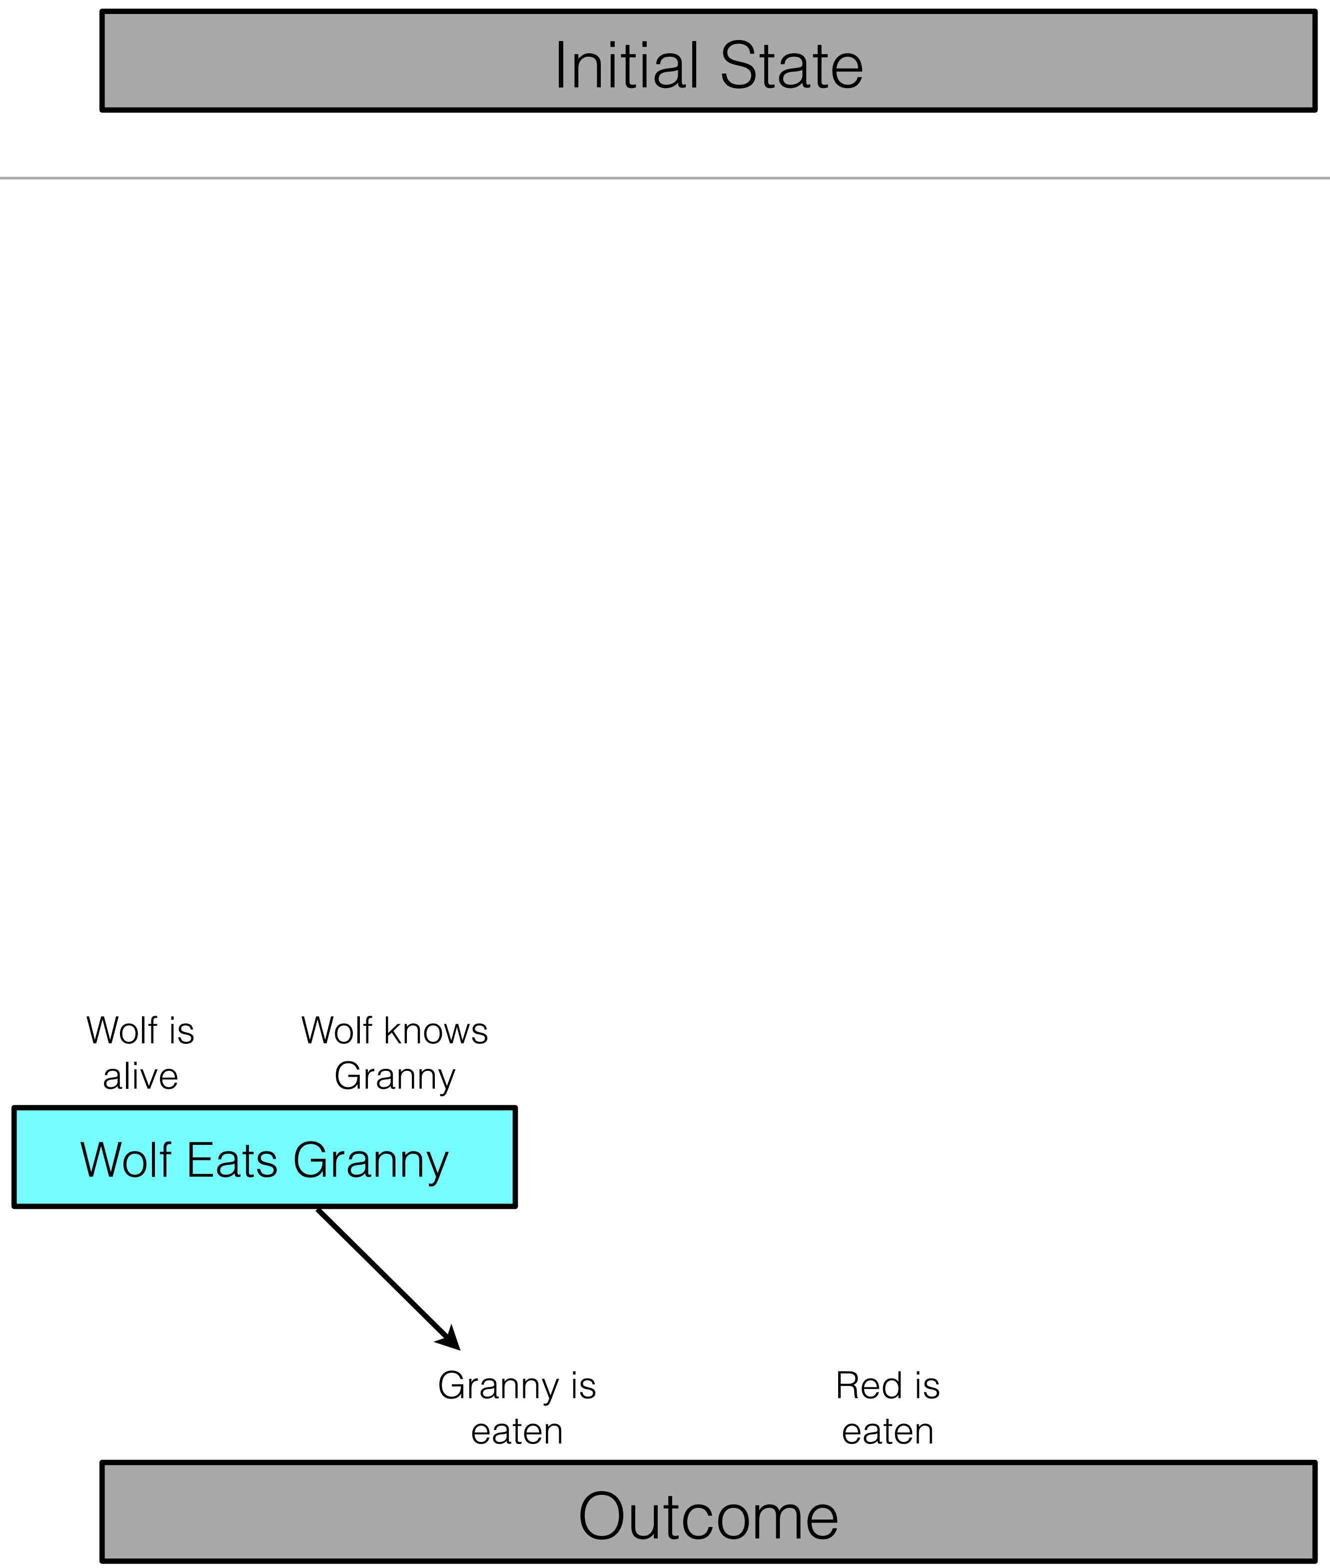

Riedl & Young. Journal of AI Research, 2010.

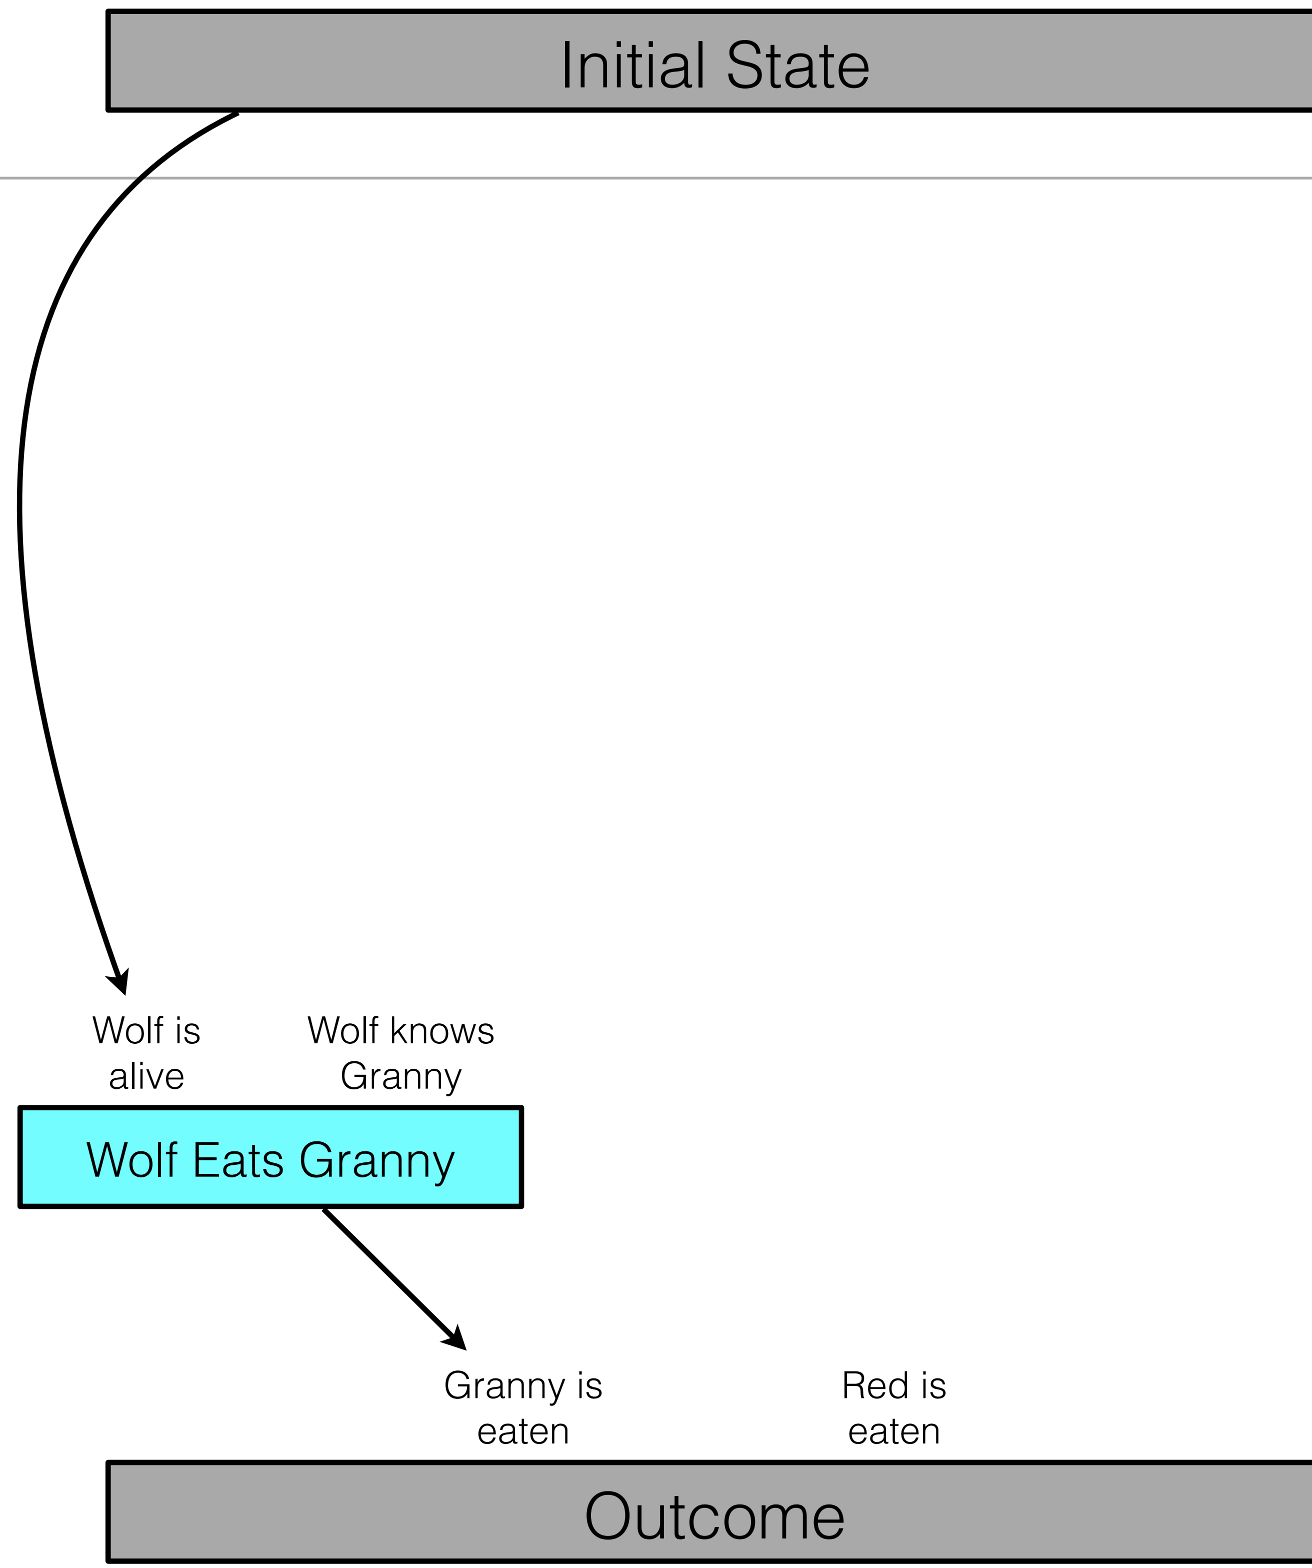

Riedl & Young. Journal of AI Research, 2010.

Georgia Tech College of Computing School of Interactive Computing

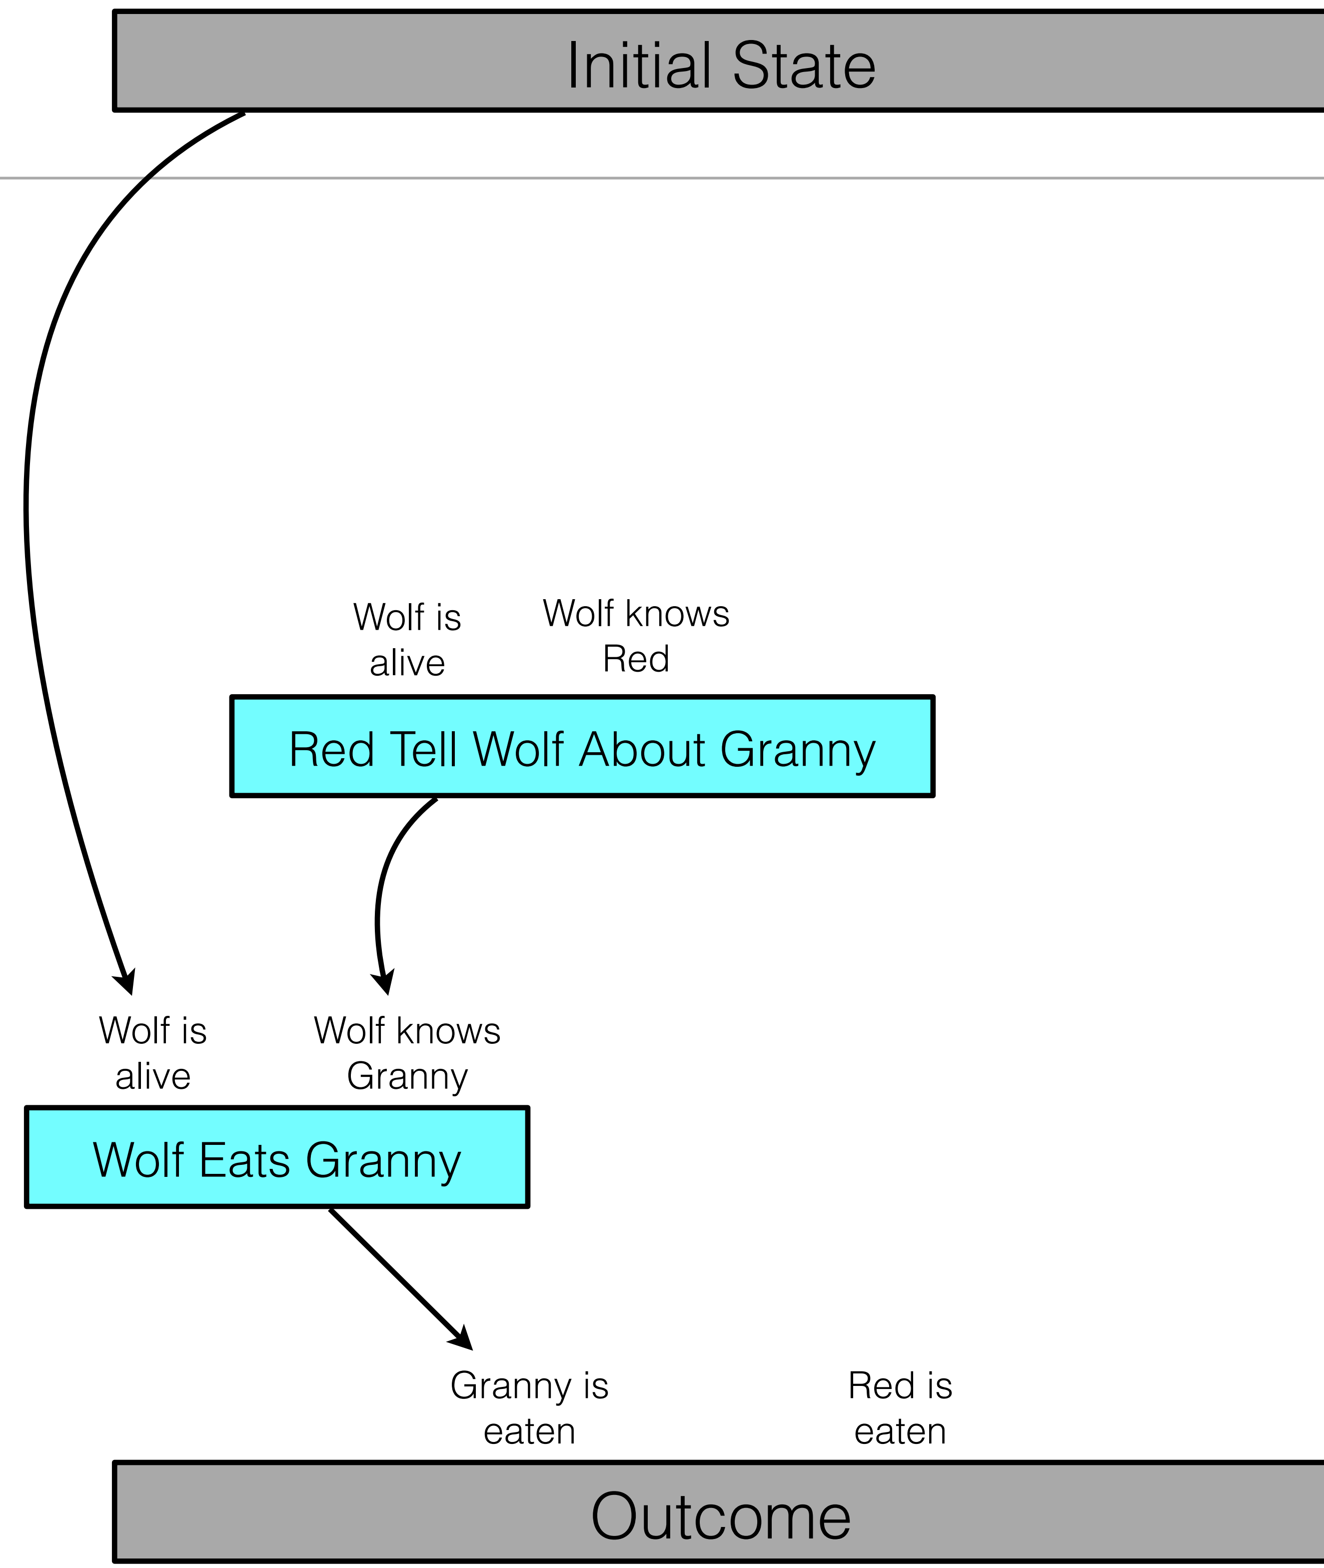

Riedl & Young. Journal of AI Research, 2010.

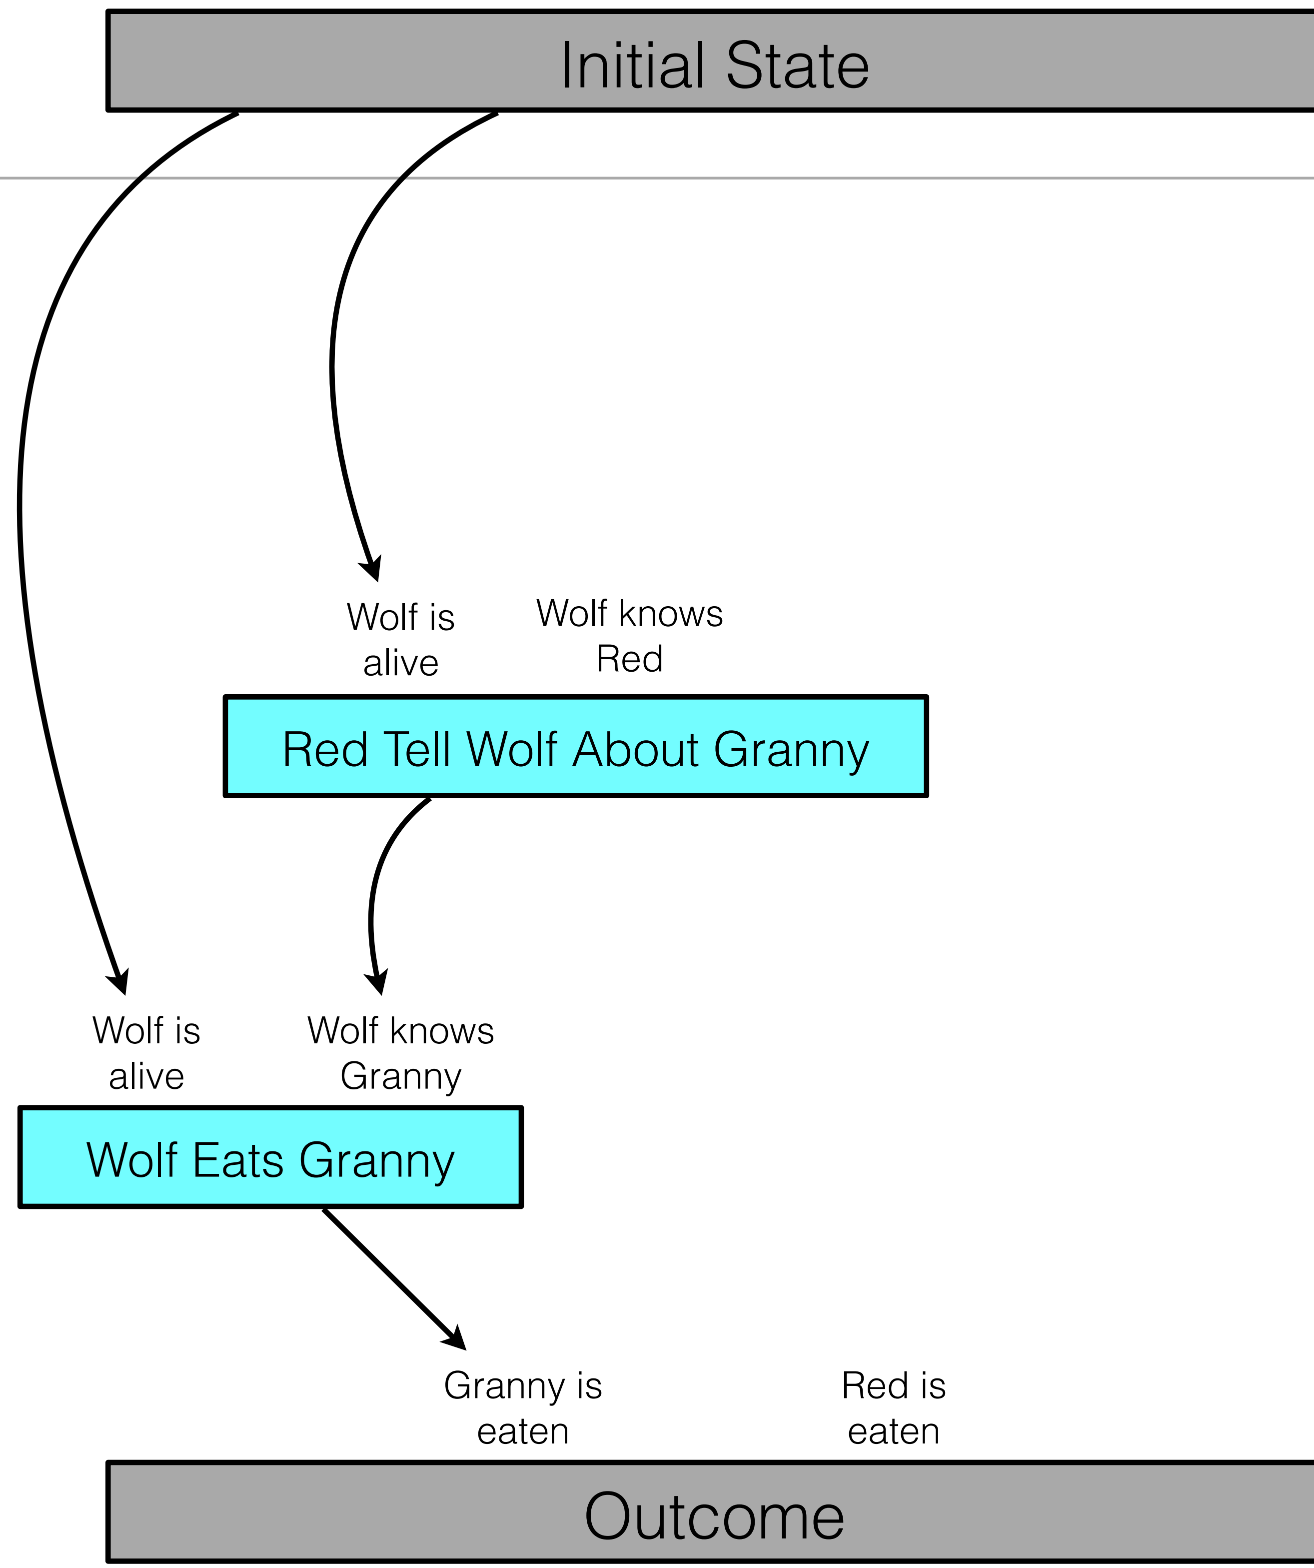

Riedl & Young. Journal of AI Research, 2010.

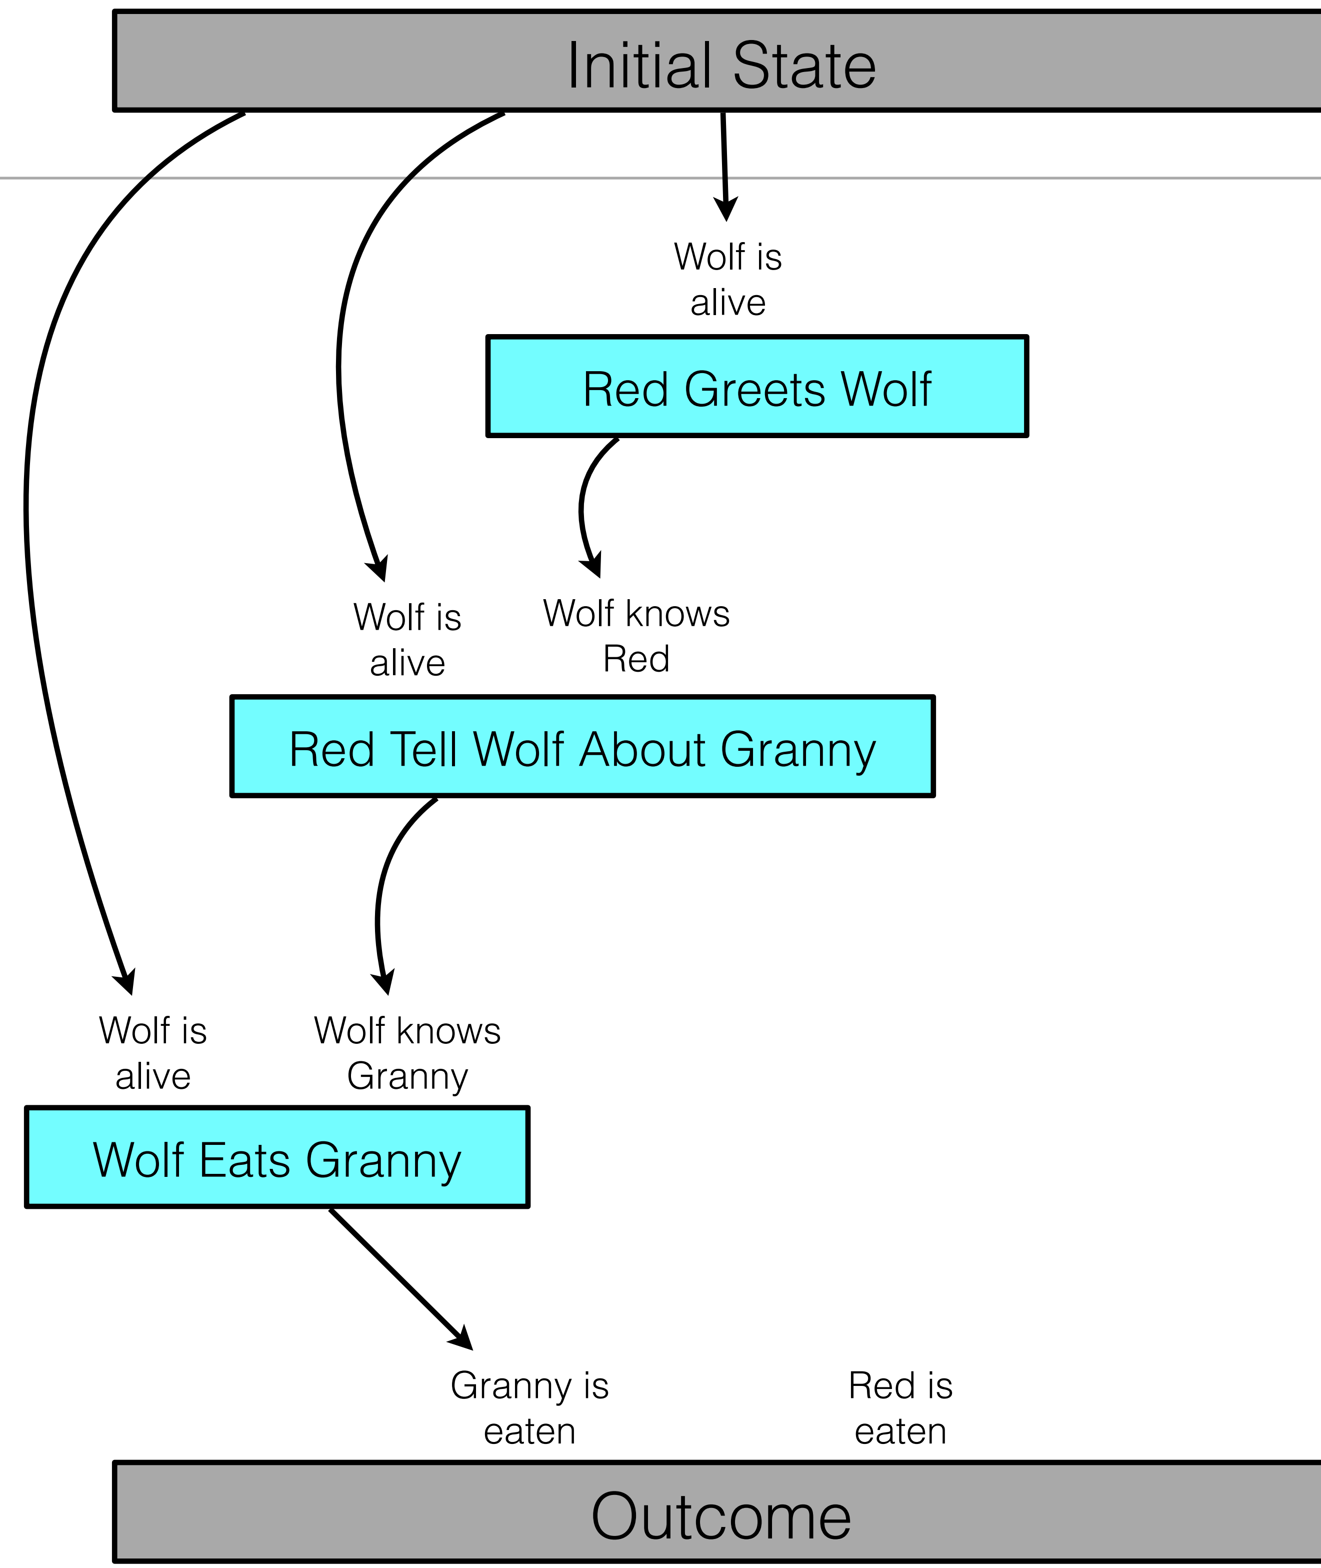

Riedl & Young. Journal of AI Research, 2010.

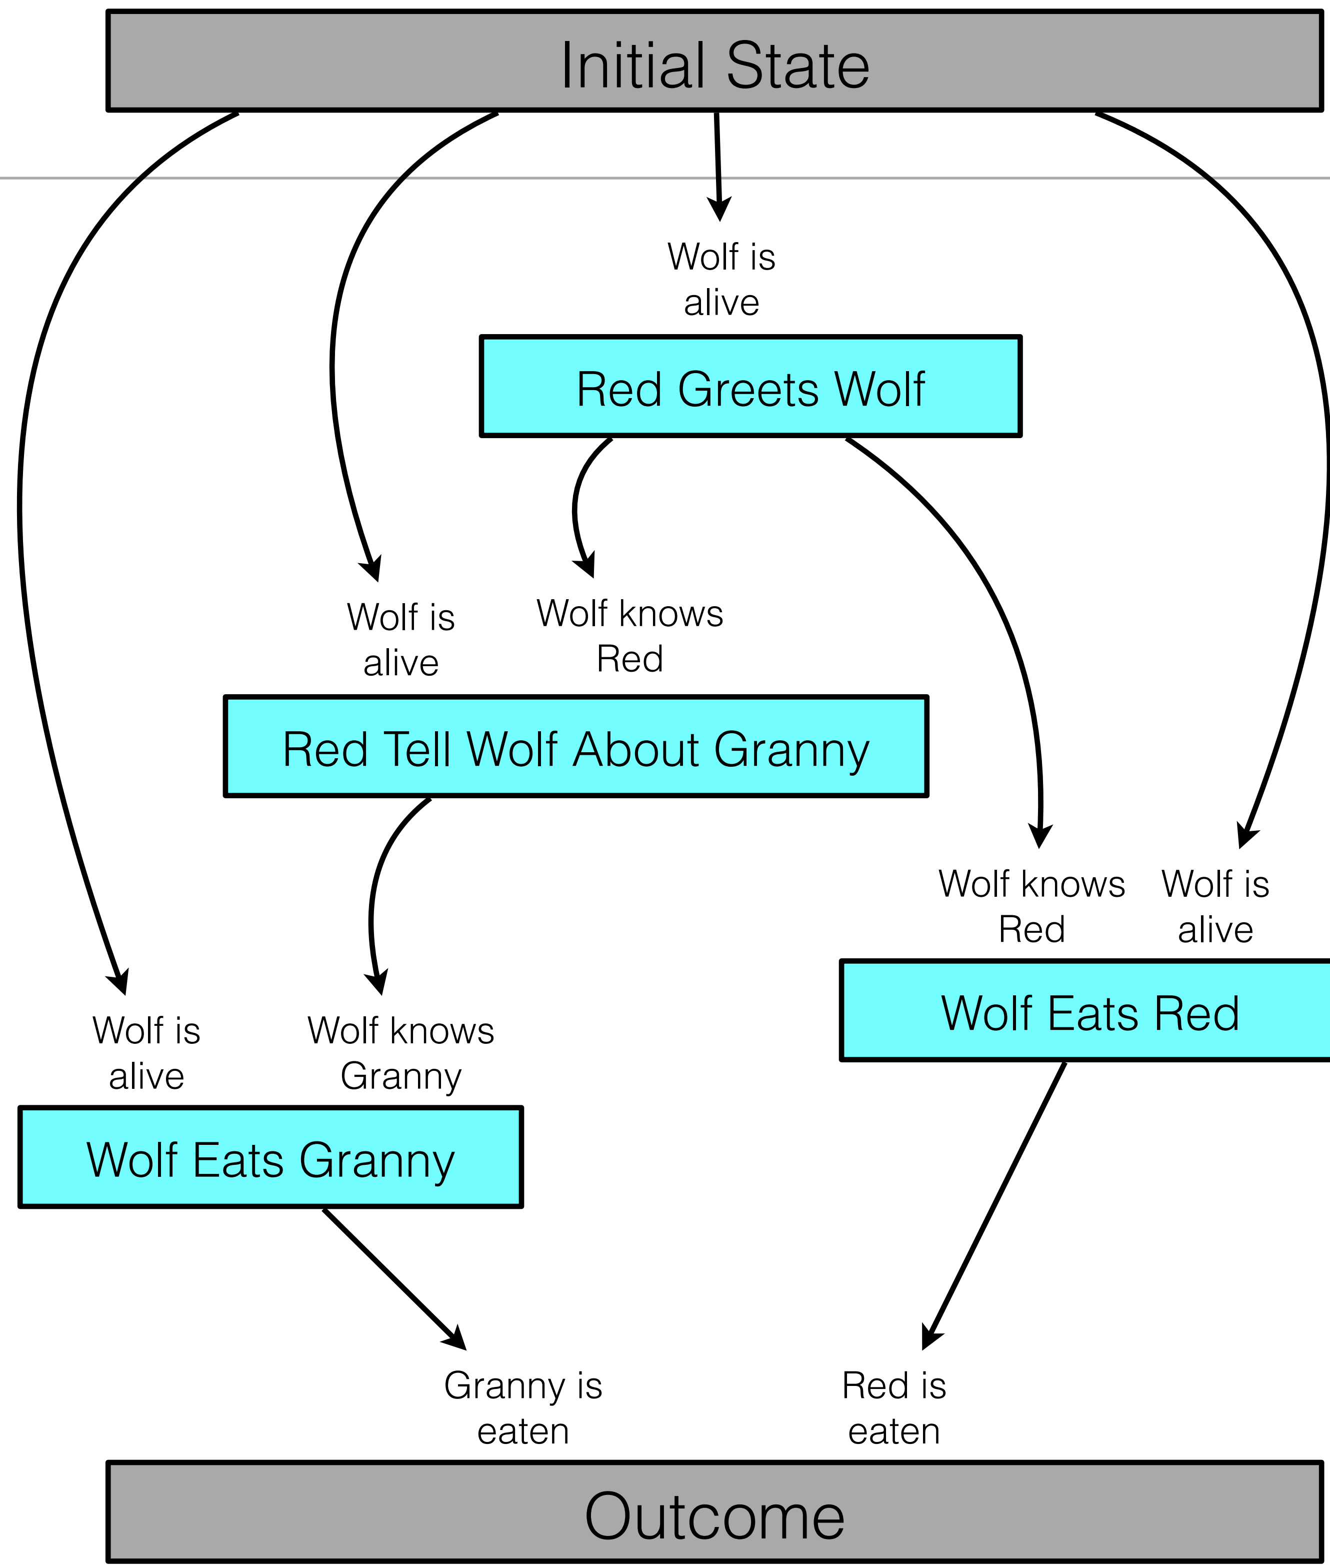

Riedl & Young. Journal of AI Research, 2010.

Riedl & Young. Journal of AI Research, 2010.

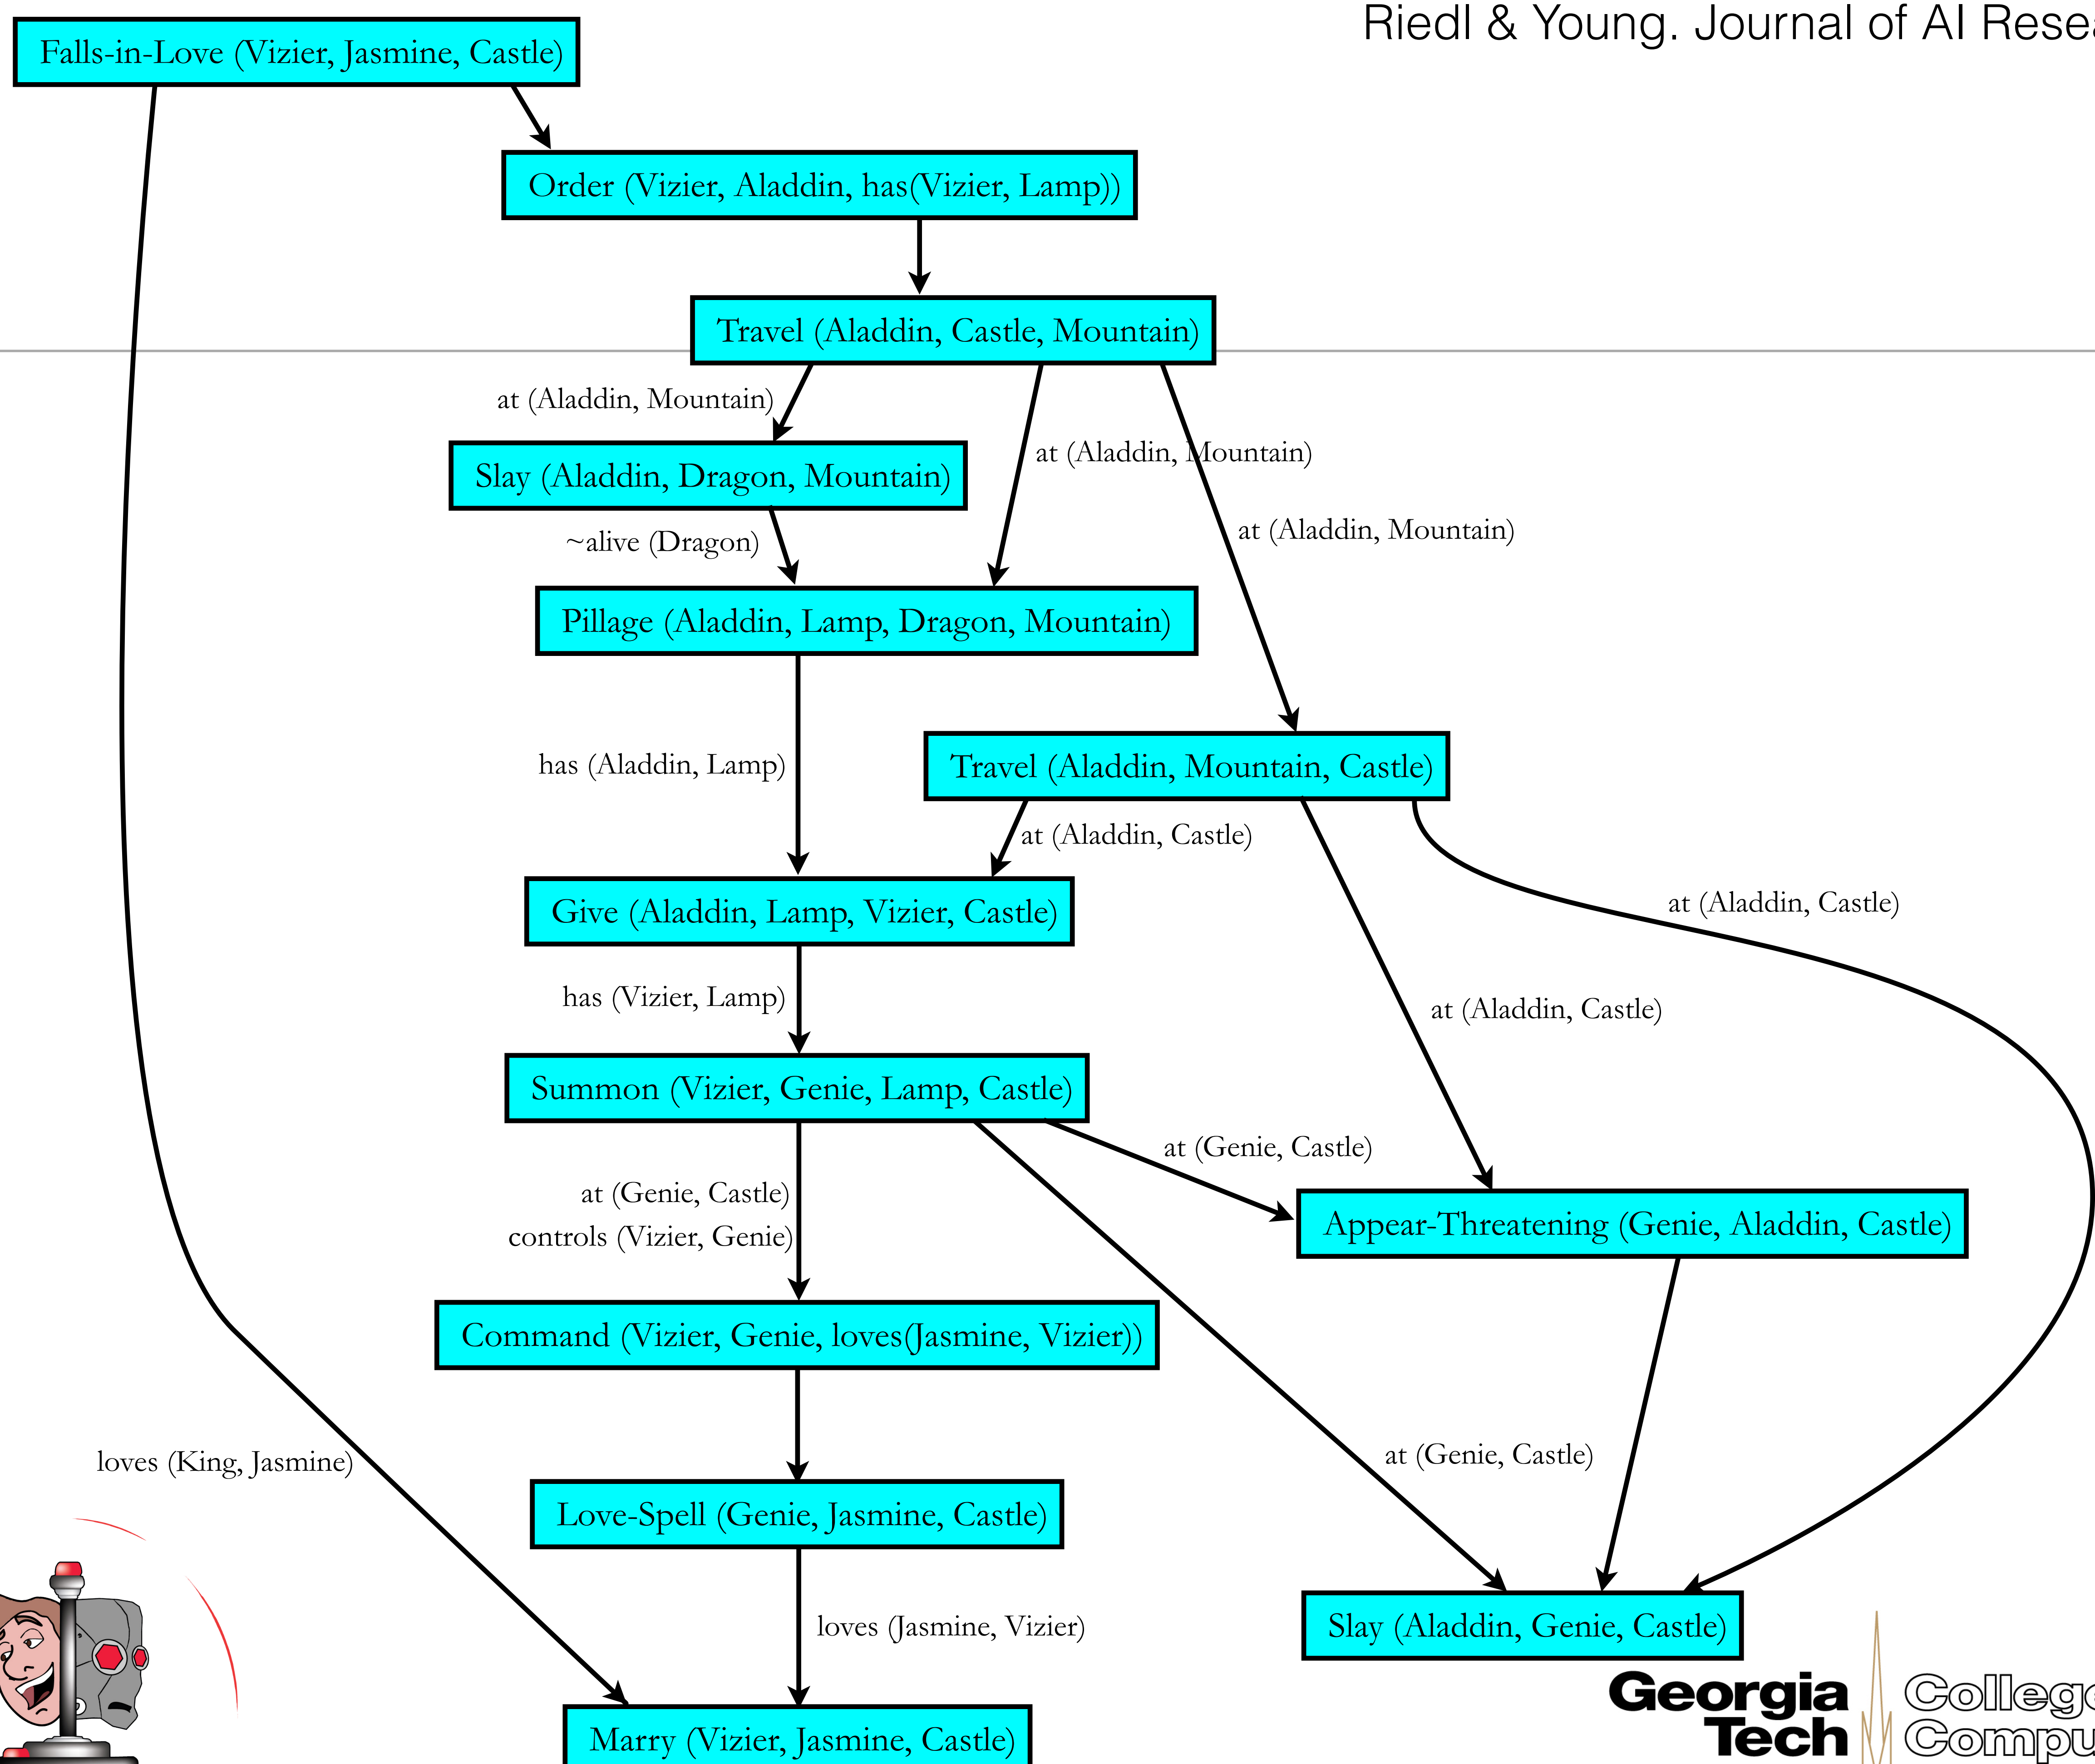

Georgia Tech College of Computing School of Interactive Computing

There is a woman named Jasmine. There is a vizier named Jafar. This is a story about how Jafar becomes married to Jasmine. There is a magic genie. This is also a story about how the genie dies.

There is a magic lamp. There is a dragon. The dragon has the magic lamp. The genie is confined within the magic lamp.

Jafar is not married. Jasmine is very beautiful. Jafar sees Jasmine and instantly falls in love with her. Jafar wants to marry Jasmine. There is a brave knight named Aladdin. Aladdin is loyal to the death to Jafar. Jafar orders Aladdin to get the magic lamp for him. Aladdin wants Jafar to have the magic lamp. Aladdin travels from the castle to the mountains. Aladdin slays the dragon. The dragon is dead. Aladdin takes the magic lamp from the dead body of the dragon. Aladdin travels from the mountains to the castle. Aladdin hands the magic lamp to Jafar. The genie is in the magic lamp. Jafar rubs the magic lamp and summons the genie out of it. The genie is not confined within the magic lamp. Jafar controls the genie with the magic lamp. Jafar uses the magic lamp to command the genie to make Jasmine love him. The genie wants Jasmine to be in love with Jafar. The genie casts a spell on Jasmine making her fall in love with Jafar. Jasmine is madly in love with Jafar. Jasmine wants to marry Jafar. The genie has a frightening appearance. The genie appears threatening to Aladdin. Aladdin wants the genie to die. Aladdin slays the genie. Jafar and Jasmine wed in an extravagant ceremony.

The genie is dead. King Jafar and Jasmine are married. The end.

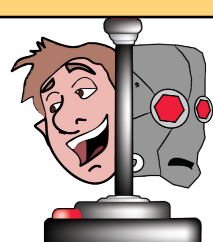

Riedl & Young. Journal of AI Research, 2010.

Observations

Observations

- Planning is forward-looking toward a desired outcome

Georgia Tech College of Computing School of Interactive Computing

Observations

- Planning is forward-looking toward a desired outcome
- Stories generated by planners are logically coherent

Observations

- Planning is forward-looking toward a desired outcome
- Stories generated by planners are logically coherent
- World dynamics ("rules" of the story world) must be pre-defined

Observations

- Planning is forward-looking toward a desired outcome
- Stories generated by planners are logically coherent
- World dynamics (“rules” of the story world) must be pre-defined
- Does the creativity lie with the system or with the knowledge engineer?

Story generation 2.0

Learning machines

Open story generation

- Story generation is knowledge-intensive and overly reliant on micro-worlds
- Generate a story about any topic?

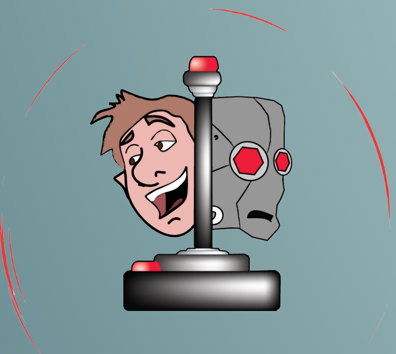

Turn to common knowledge

Turn to common knowledge

- Humans rely on a lifetime of experiences from which to explain stories, tell stories, or act in the real-world

Turn to common knowledge

- Humans rely on a lifetime of experiences from which to explain stories, tell stories, or act in the real-world
- Learn sociocultural scripts

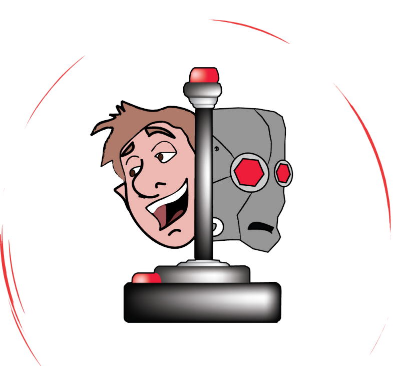

Georgia Tech College of Computing School of Interactive Computing

Turn to common knowledge

- Humans rely on a lifetime of experiences from which to explain stories, tell stories, or act in the real-world
- Learn sociocultural scripts
- Generate new stories that apply those scripts

Georgia Tech College of Computing
School of Interactive Computing

Script learning

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Georgia Tech College of Computing School of Interactive Computing

Script learning

- Given a set of parallel stories

Story A	Story B
a. John drives to the restaurant.	a. Mary looks at the menu.
b. John stands in line.	b. Mary decides what to order.
c. John orders food.	c. Mary orders a burger.
d. John waits for his food.	d. Mary finds a seat.
e. John sits down.	e. Mary eats her burger.
f. John eats the food.	…
…	

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Script learning

- Given a set of parallel stories

- Learn the events

Story A	Story B
a. John drives to the restaurant.	a. Mary looks at the menu.
b. John stands in line.	b. Mary decides what to order.
c. John orders food.	c. Mary orders a burger.
d. John waits for his food.	d. Mary finds a seat.
e. John sits down.	e. Mary eats her burger.
f. John eats the food.	…
…	

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Script learning

- Given a set of parallel stories
- Learn the events
- Learn temporal relations

Story A	Story B
a. John drives to the restaurant.	a. Mary looks at the menu.
b. John stands in line.	b. Mary decides what to order.
c. John orders food.	c. Mary orders a burger.
d. John waits for his food.	d. Mary finds a seat.
e. John sits down.	e. Mary eats her burger.
f. John eats the food.	…
…	

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Script learning

- Given a set of parallel stories
- Learn the events
- Learn temporal relations

Story A	Story B
a. John drives to the restaurant.	a. Mary looks at the menu.
b. John stands in line.	b. Mary decides what to order.
c. John orders food.	c. Mary orders a burger.
d. John waits for his food.	d. Mary finds a seat.
e. John sits down.	e. Mary eats her burger.
f. John eats the food.	…
…	

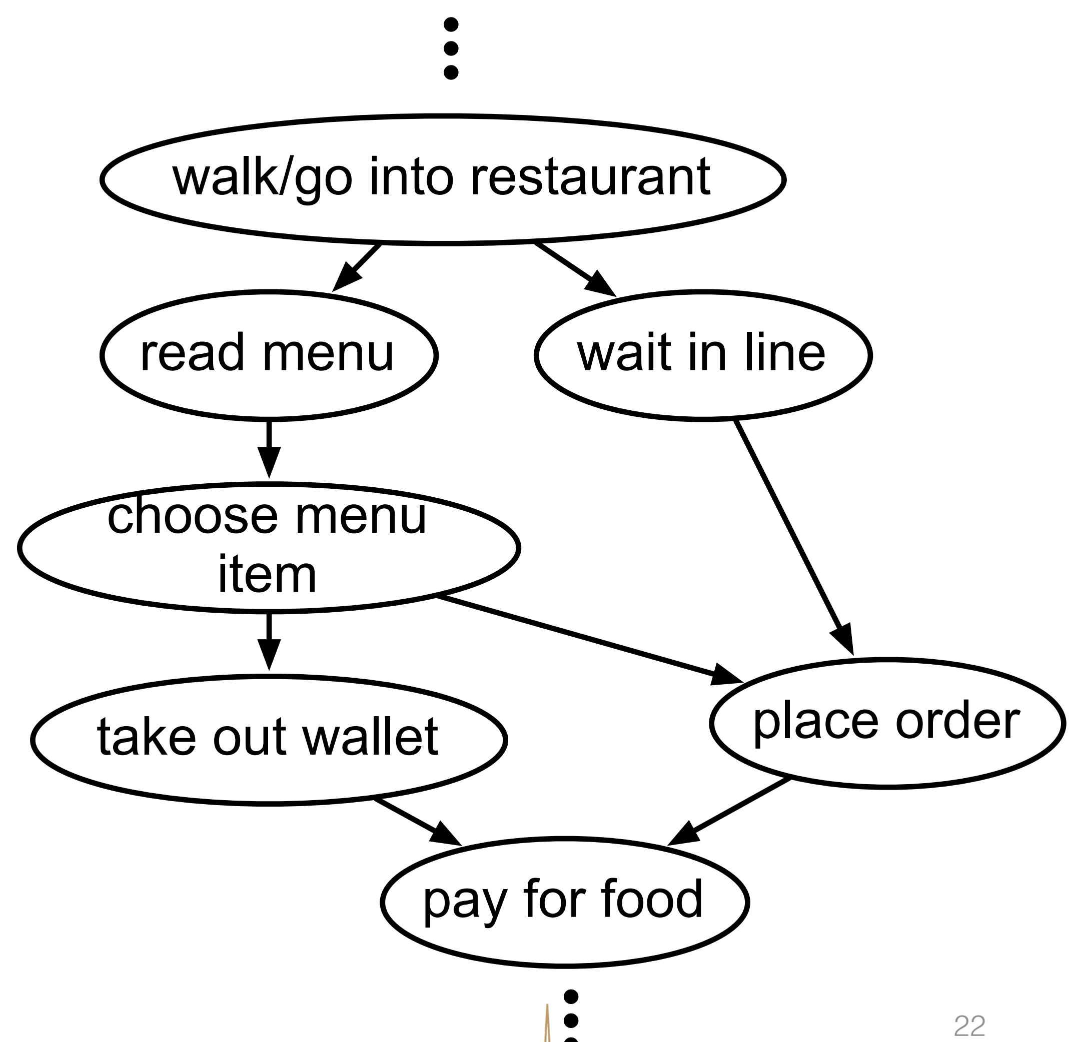

Li, Lee-Urban, Johnston & Riedl. AAAI 2013 Conference.

Fast food restaurant

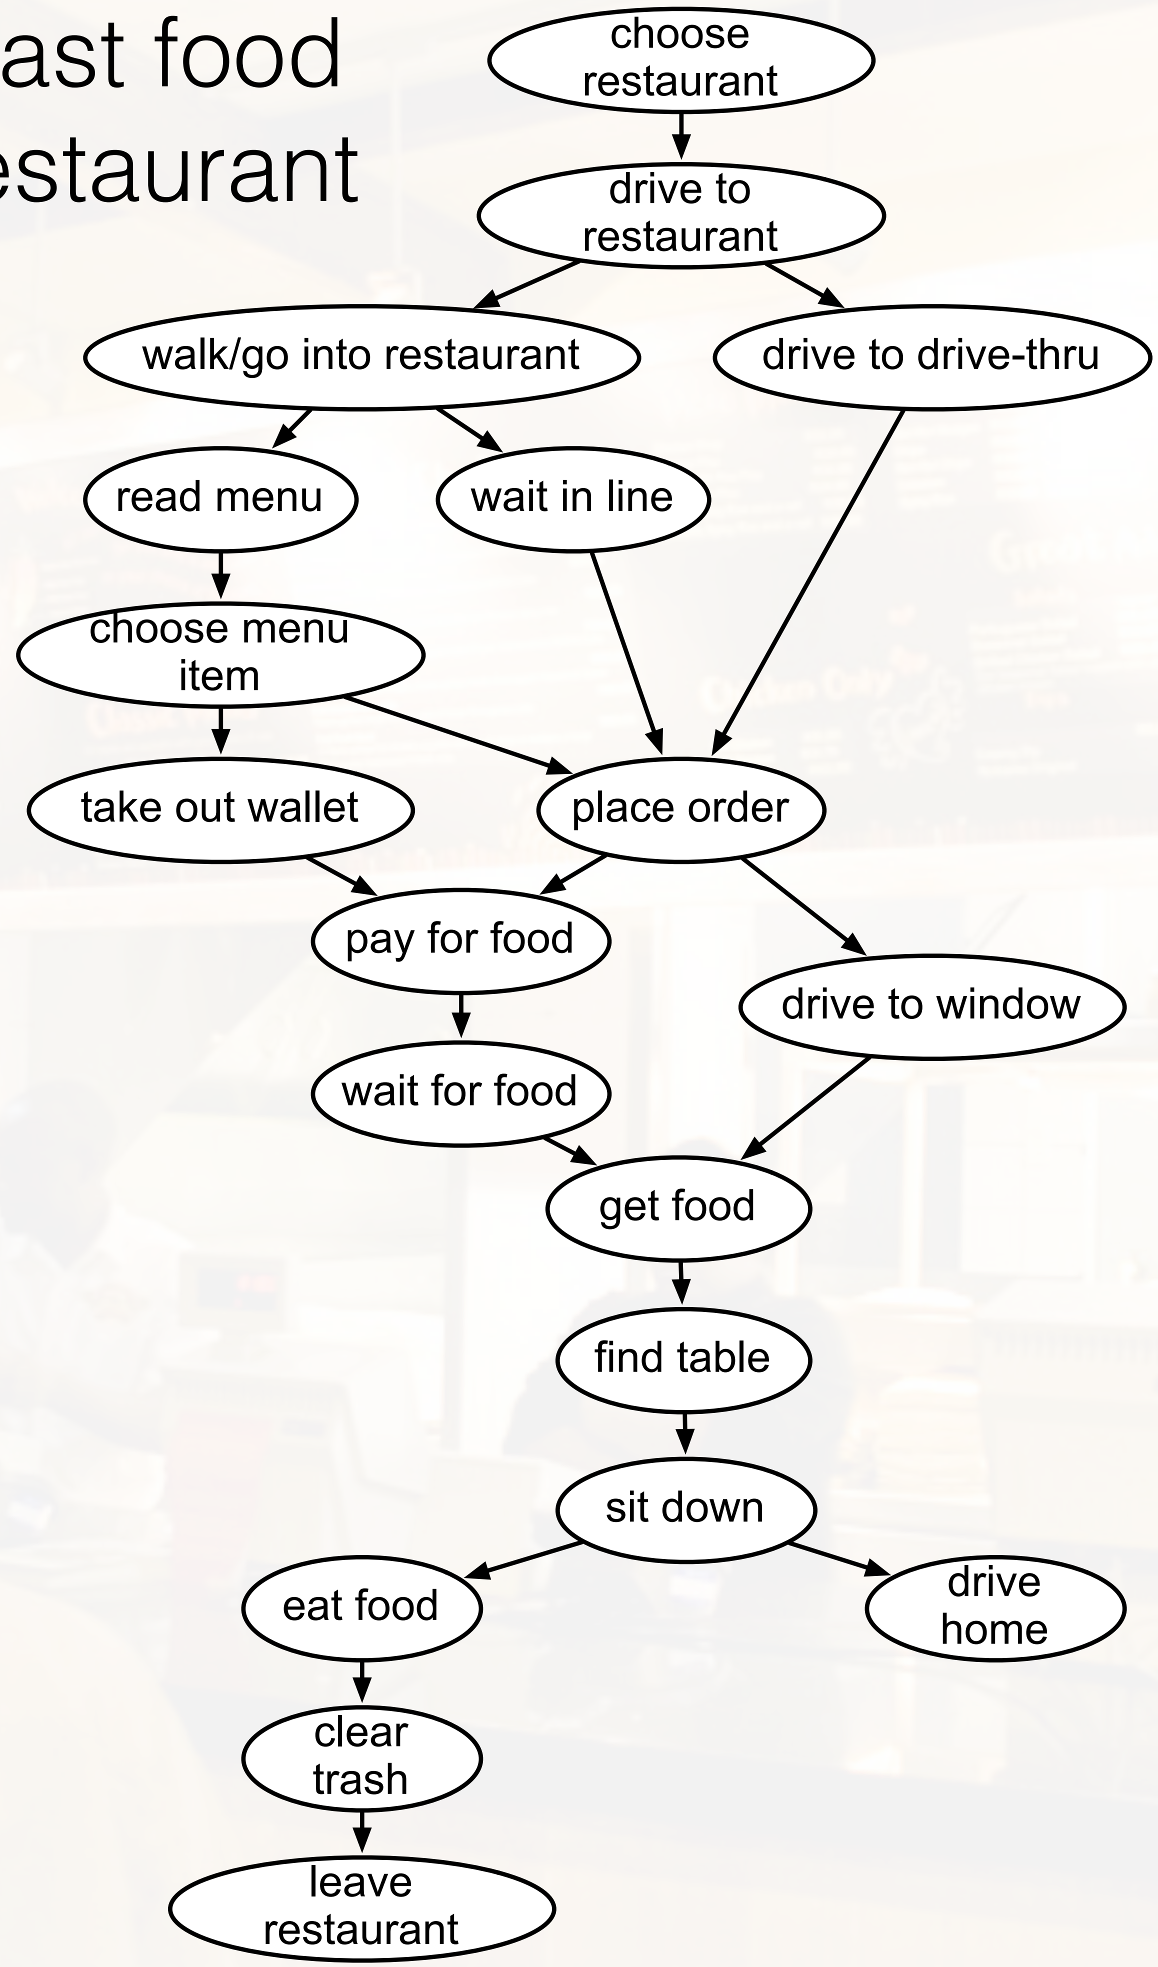

Going on a date to the movies

With sweaty palms and heart racing, John drove to Sally's house for their first date. Sally, her pretty white dress flowing in the wind, carefully entered John's car. John and Sally drove to the movie theater. John and Sally parked the car in the parking lot. Wanting to feel prepared, John had already bought tickets to the movie in advance. A pale-faced usher stood before the door; John showed the tickets and the couple entered. Sally was thirsty so John hurried to buy drinks before the movie started. John and Sally found two good seats near the back. John sat down and raised the arm rest so that he and Sally could snuggle. John paid more attention to Sally while the movie rolled and nervously sipped his drink. Finally working up the courage to do so, John extended his arm to embrace Sally. He was relieved and ecstatic to feel her move closer to him in response. Sally stood up to use the restroom during the movie, smiling coyly at John before that exit. John and Sally also held hands throughout the movie, even though John's hands were sweaty. John and Sally slowly got up from their seats. Still holding hands, John walked Sally back to his car through the maze of people all scurrying out of the theater. The bright sunshine temporarily blinded John as he opened the doors and held them for Sally as they left the dark theater and stepped back out onto the street. John let go of Sally's hand and opened the passenger side door of his car for her but instead of entering the car, she stepped forward, embraced him, and gave him a large kiss. John drove Sally back to her home.

Observations

Observations

- Novel combinations with decent coherence

Observations

- Novel combinations with decent coherence
- Requires crowdsourced stories

Observations

- Novel combinations with decent coherence
- Requires crowdsourced stories
- Scalability is limited: one plot graph per story topic

Story generation 2.5

Deep Learning

Corpora in the wild

Corpora in the wild

- Learn everything by reading a corpus of stories

Corpora in the wild

- Learn everything by reading a corpus of stories
- Neural language models
  - $Pr(word_n \mid word_{n-k}, \ldots\ word_{n-2}, word_{n-1}; \theta)$

Corpora in the wild

- Learn everything by reading a corpus of stories
- Neural language models
  - $Pr(word_n \mid word_{n\text{-}k}, \ldots\ word_{n\text{-}2}, word_{n\text{-}1}; \theta)$
- Recurrent neural networks

Georgia Tech College of Computing School of Interactive Computing

Corpora in the wild

- Learn everything by reading a corpus of stories
- Neural language models
  - $Pr(word_n \mid word_{n-k}, \ldots\ word_{n-2}, word_{n-1}; \theta)$
- Recurrent neural networks
- Transformers
  - GPT, OPT, LaMDA, PaLM, etc.

Georgia Tech College of Computing
School of Interactive Computing

Neural Generation

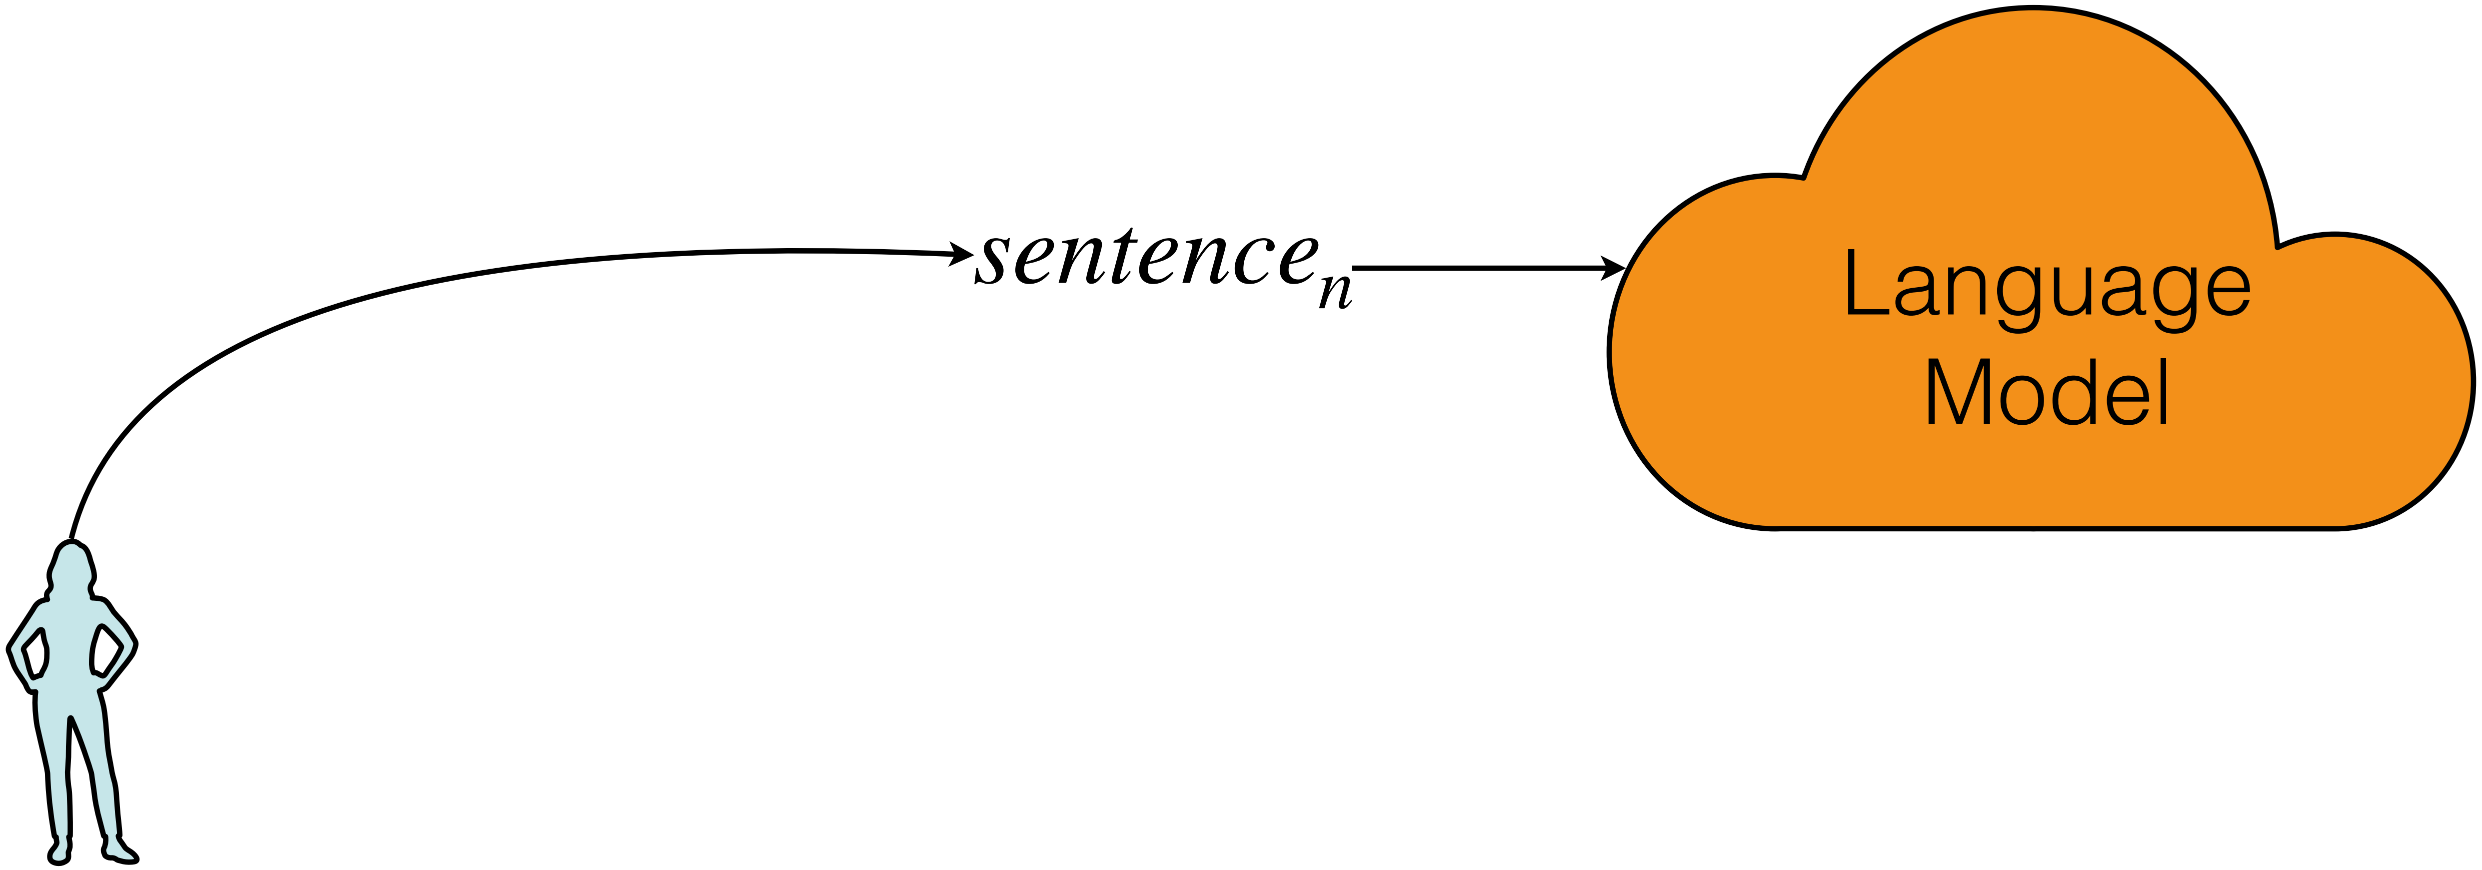

Martin et al. AAAI 2018 Conference.

Georgia Tech College of Computing School of Interactive Computing

Neural Generation

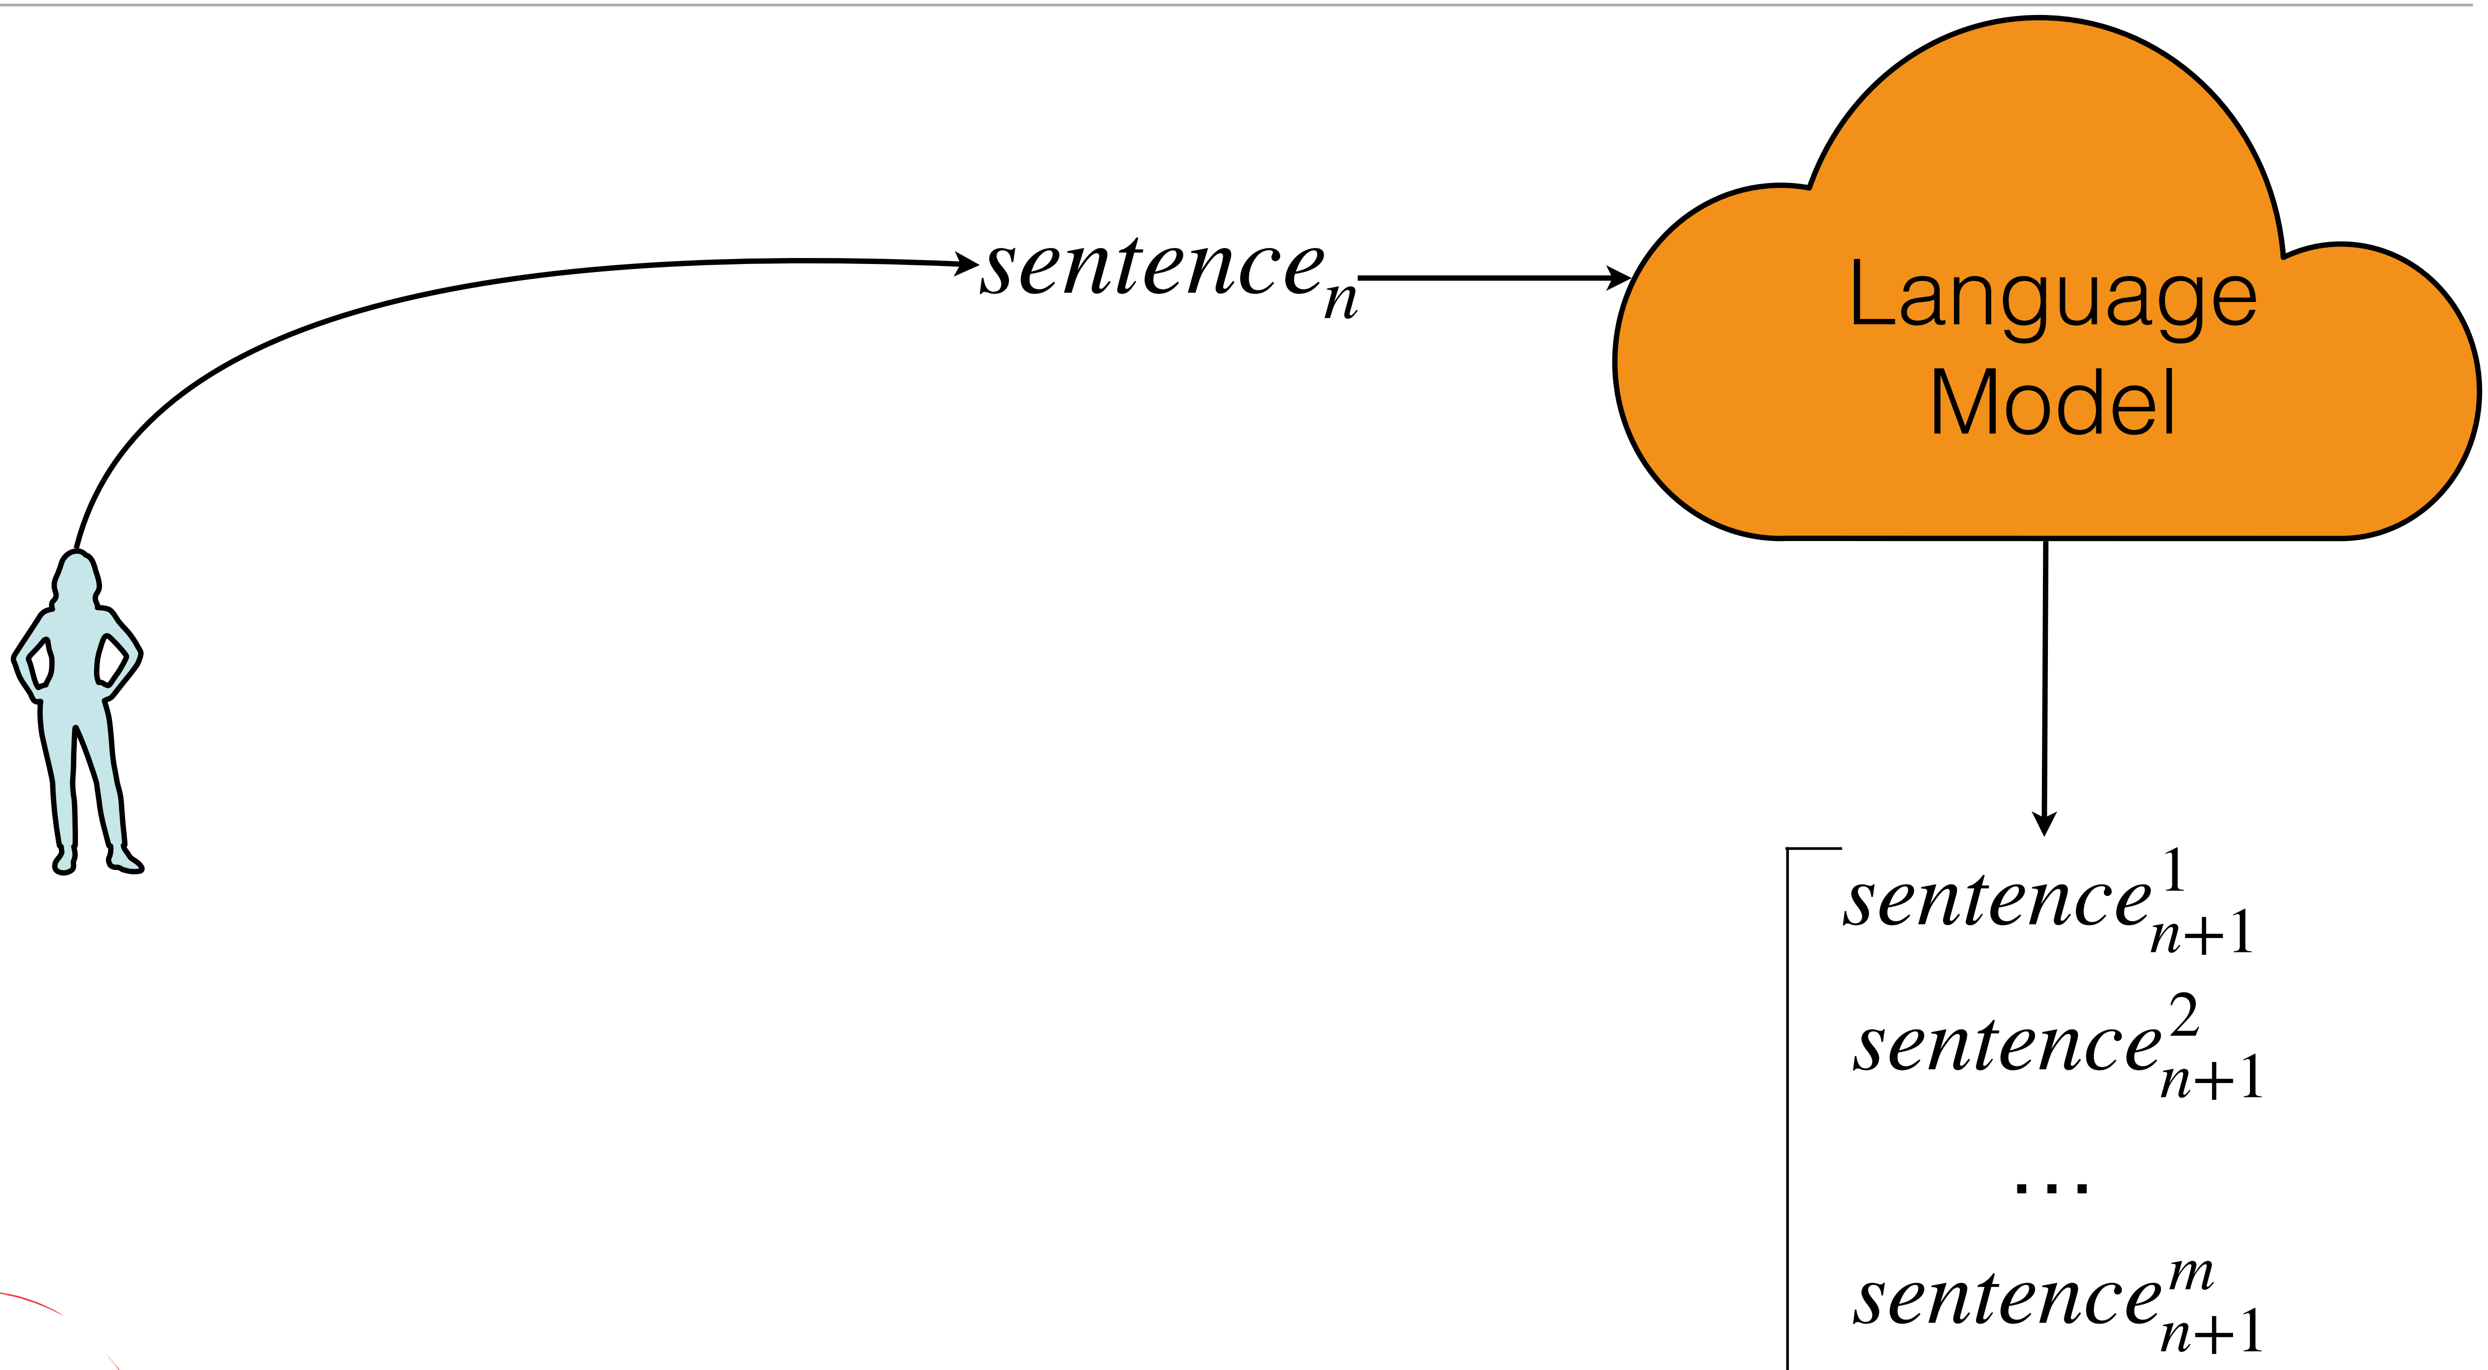

Martin et al. AAAI 2018 Conference.

Georgia Tech College of Computing School of Interactive Computing

Neural Generation

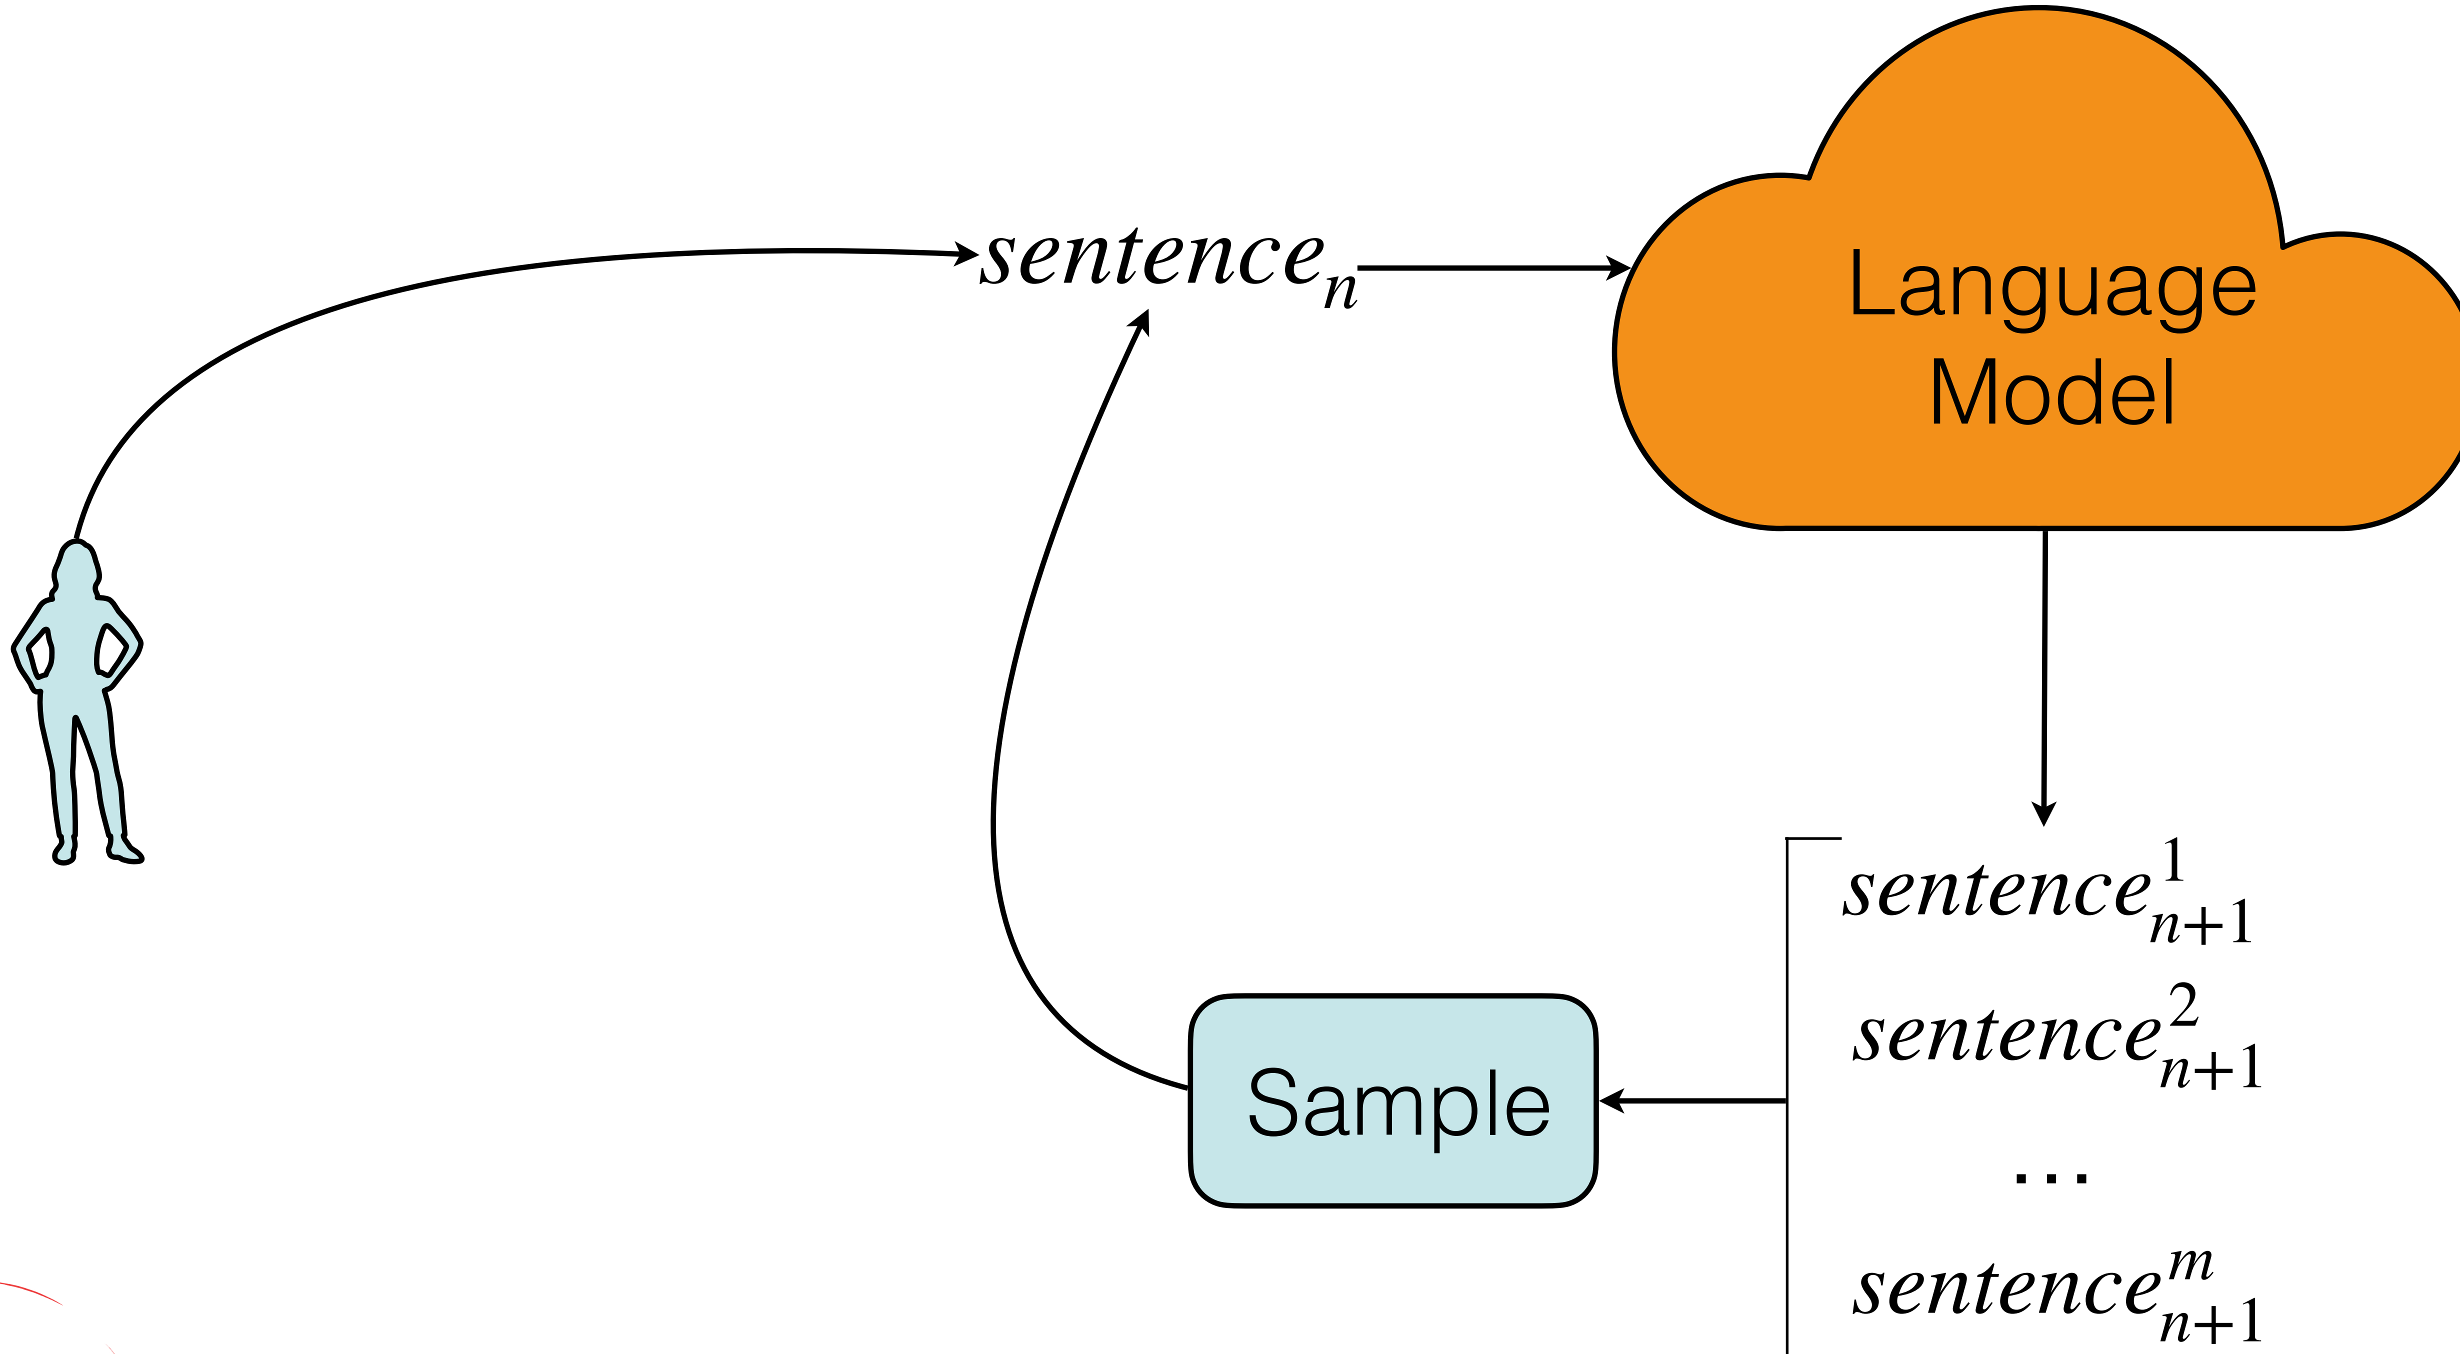

Martin et al. AAAI 2018 Conference.

Prompt continuation (GPT-2 and beyond)

The hungry dog licked her lips as she watched her owner eat.
She could see the white teeth in his jaw, then moving behind his right ear.

The hunter picked up his sword, directing the dogs in a cloud of mutterings. "Did you just feed a dog?" she demanded, suddenly even more agitated than before.

"Yes." Jack's voice was rough, and he hated that word. His face was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad. She is one of the most dangerous wolves I have ever met, and if she bites me she will lose."

John shrugged. "Well, it's not your fault," he said.

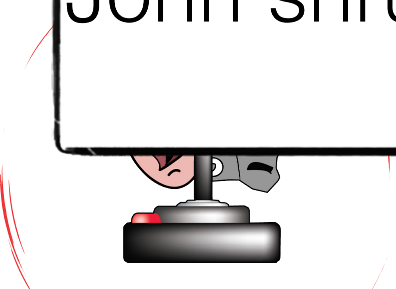

Controllability

Controllability

- Sampling from $Pr(tok_n \mid tok_{n-1}, tok_{n-2}, \ldots; \theta)$ is uncontrolled generation

Controllability

- Sampling from Pr(tok_n | tok_{n-1}, tok_{n-2}, …; θ) is uncontrolled generation
- Backward-looking: generate the next sequence of tokens based on what has come before

Controllability

- Sampling from Pr(tok_n | tok_{n-1}, tok_{n-2}, …; θ) is uncontrolled generation
- Backward-looking: generate the next sequence of tokens based on what has come before
- Story generation needs to be forward looking:
  - Make choices based on how to achieve some communicative intent

Georgia Tech College of Computing School of Interactive Computing

Goal-driven coherence

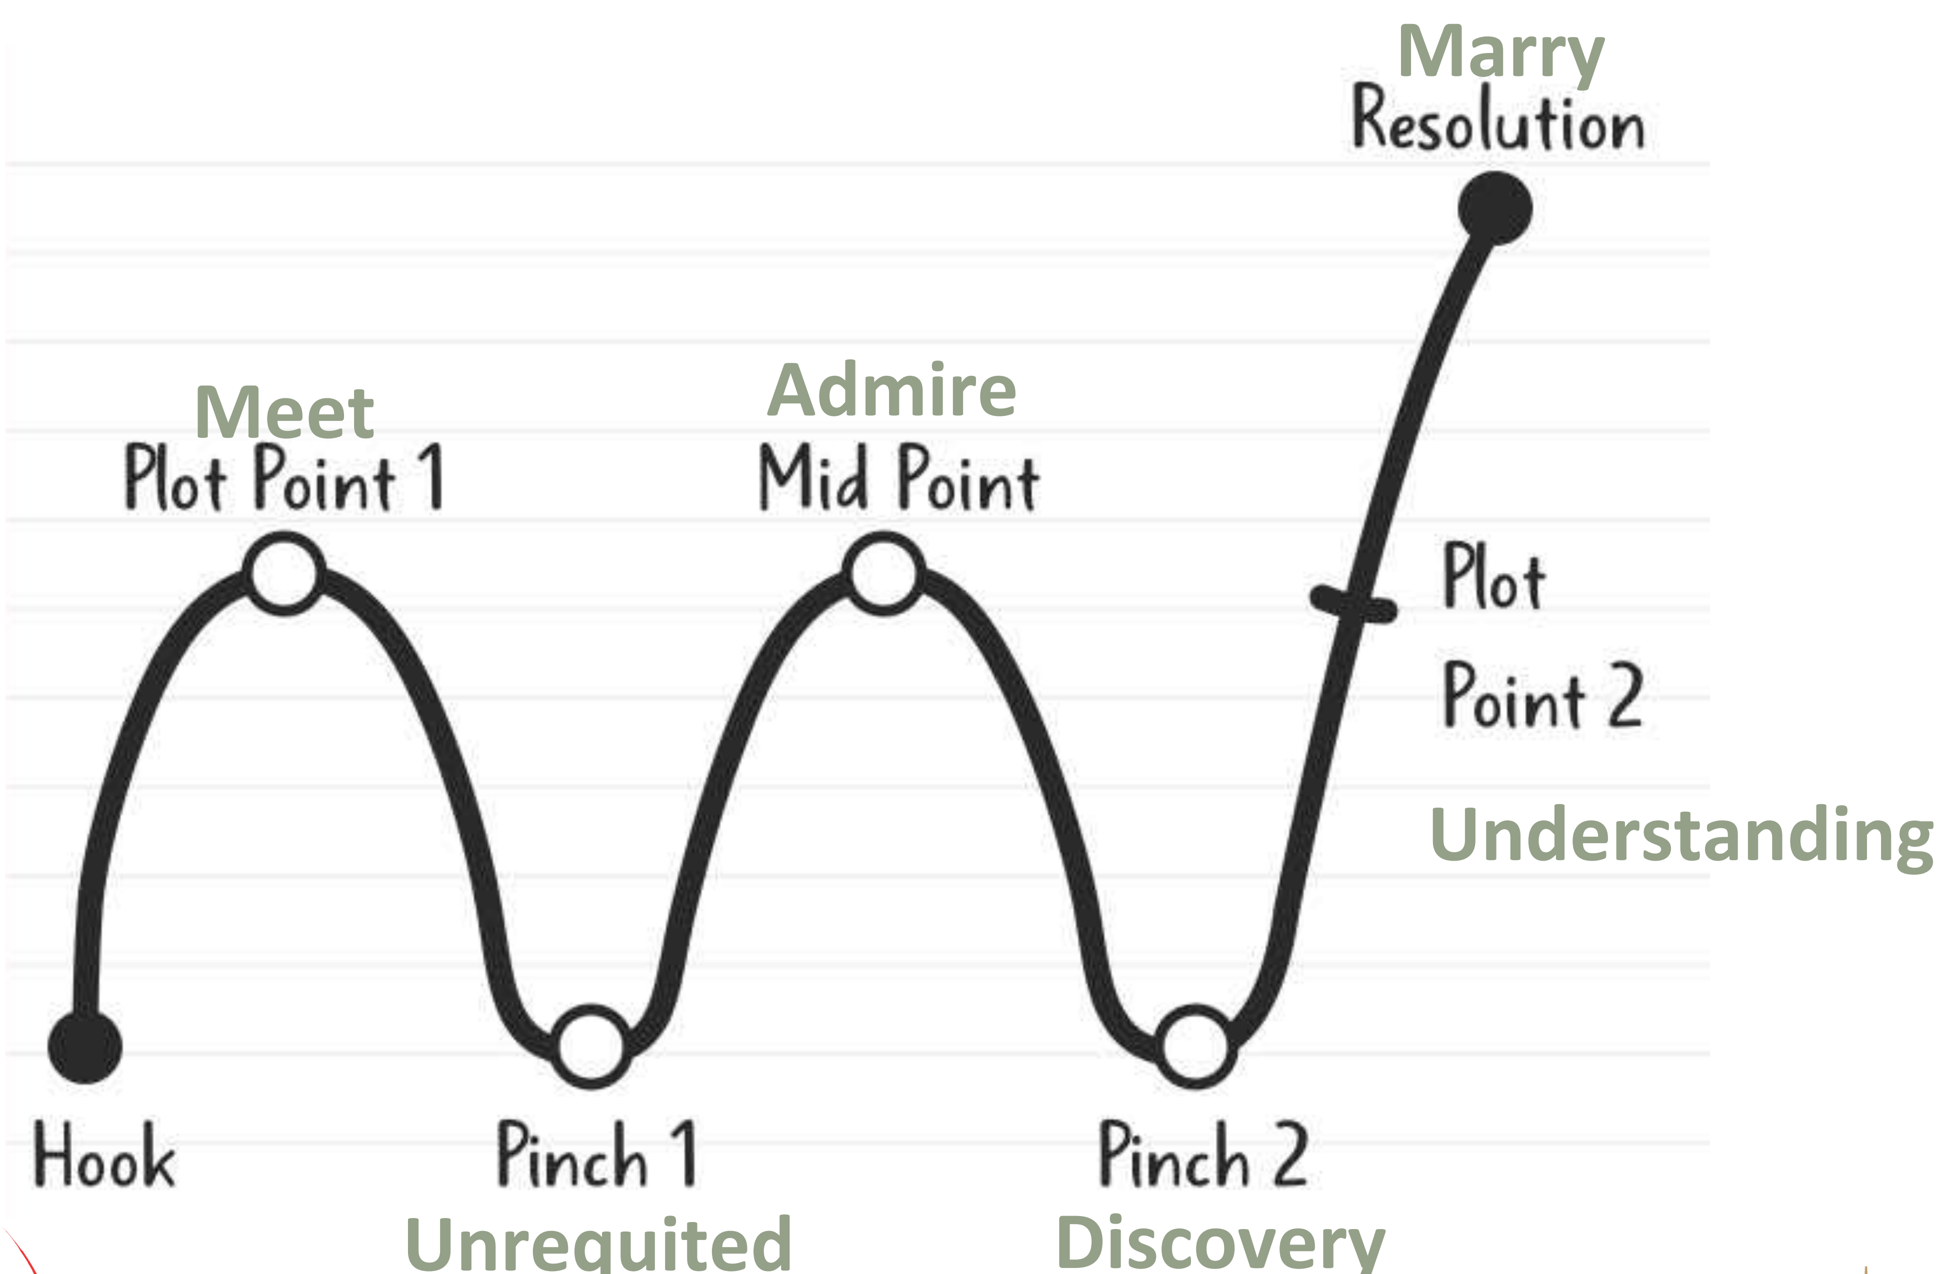

Fine-tuning on goals

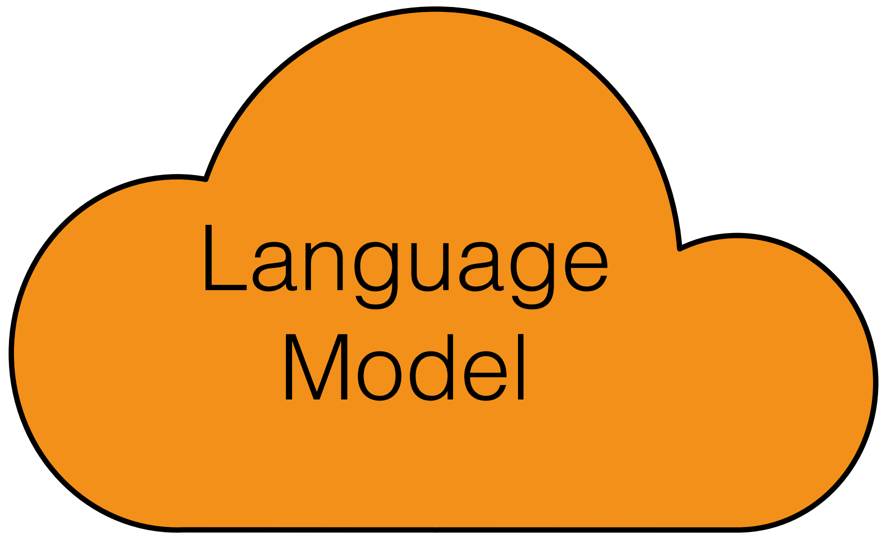

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing School of Interactive Computing

Fine-tuning on goals

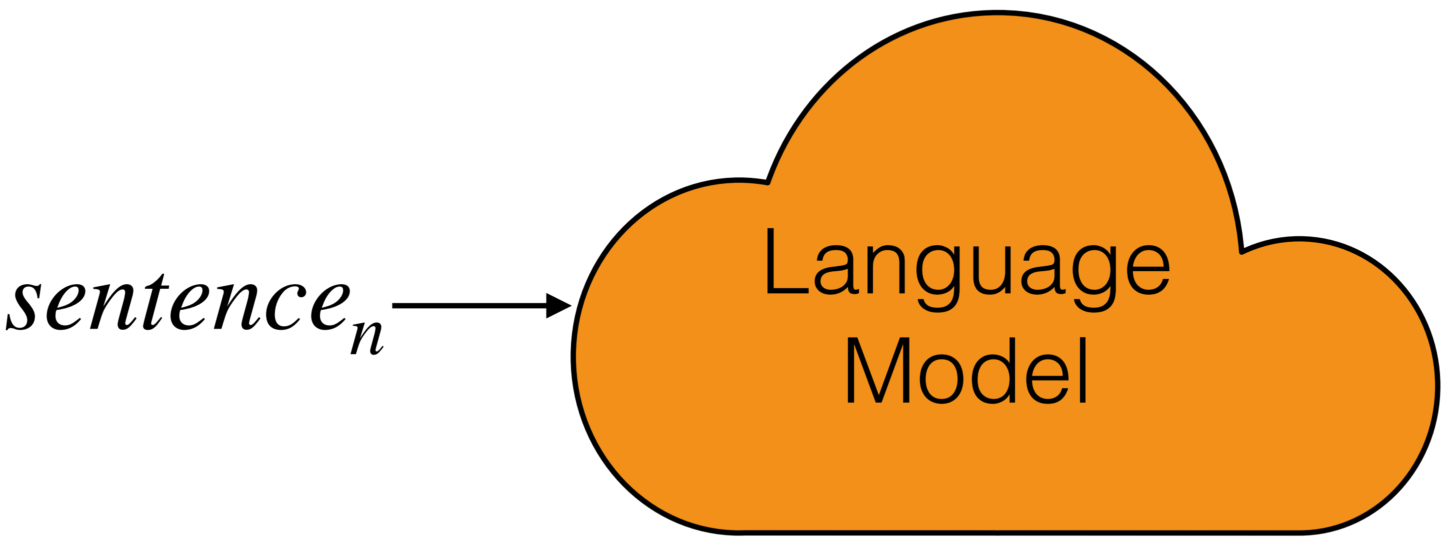

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing School of Interactive Computing

Fine-tuning on goals

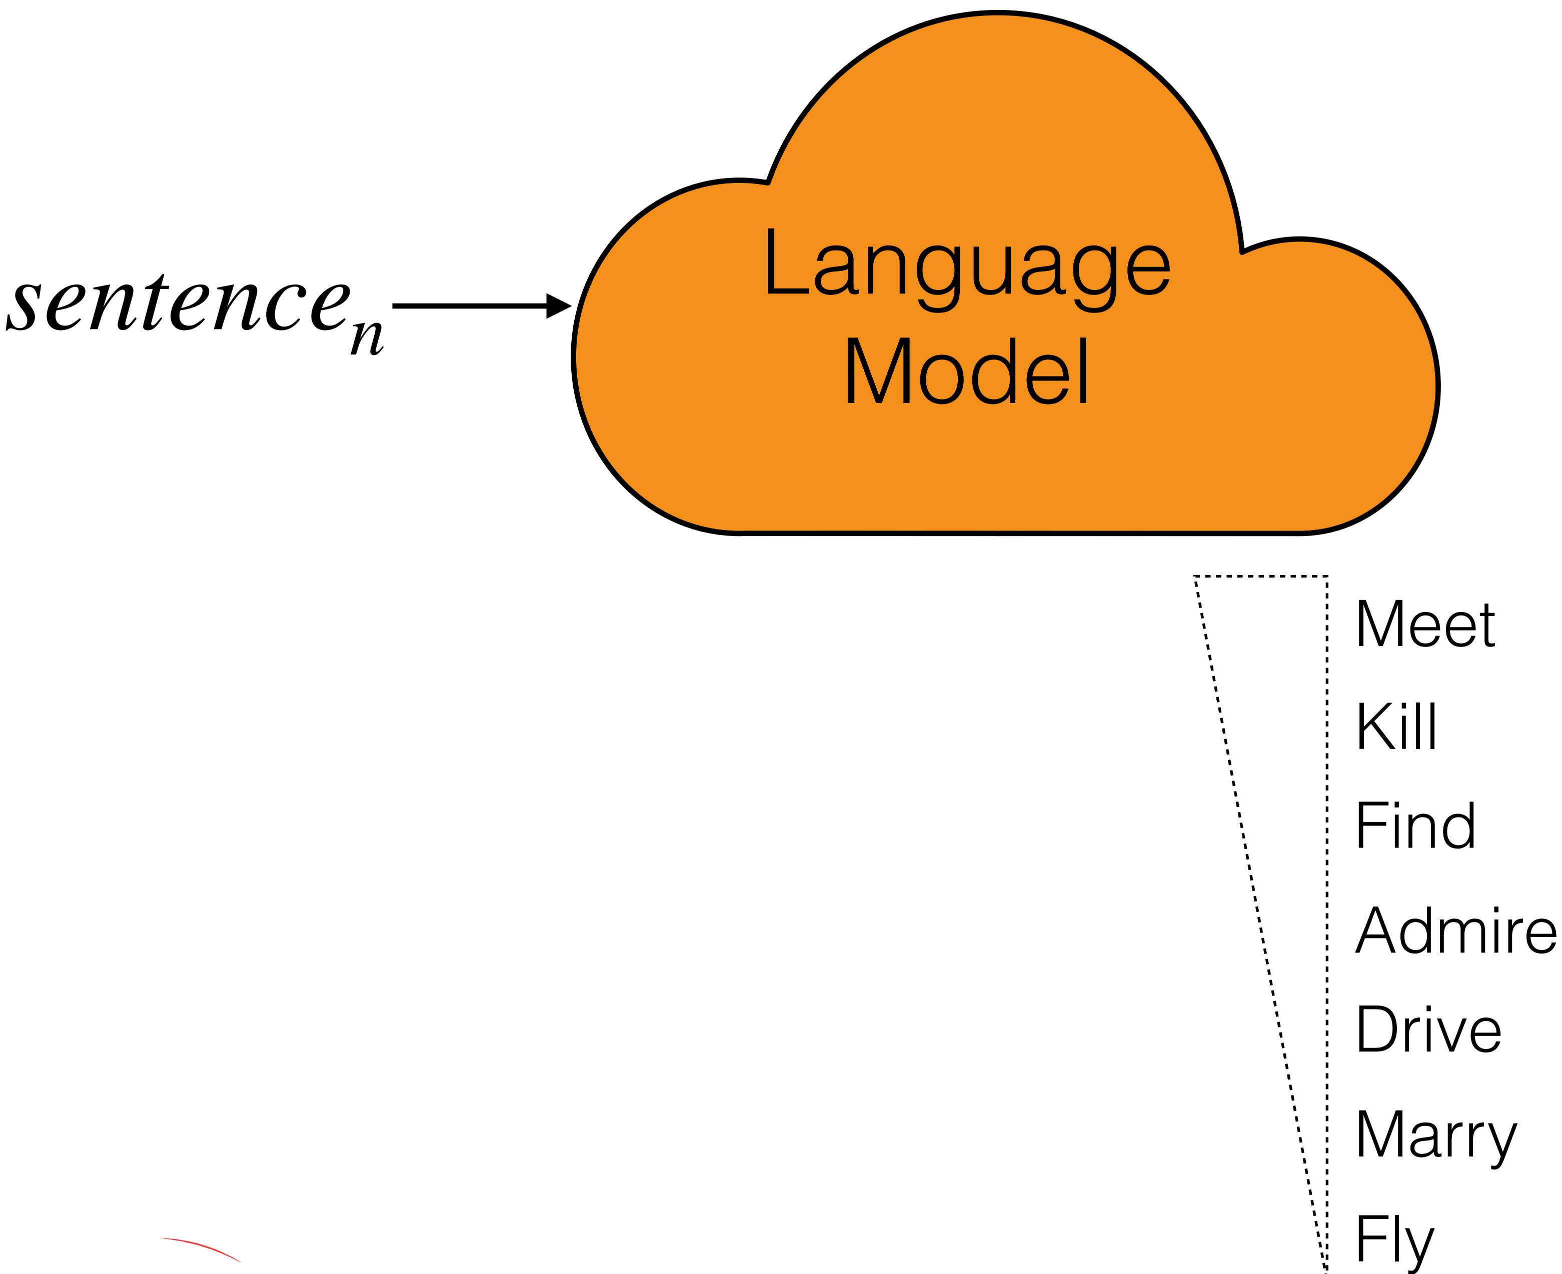

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Fine-tuning on goals

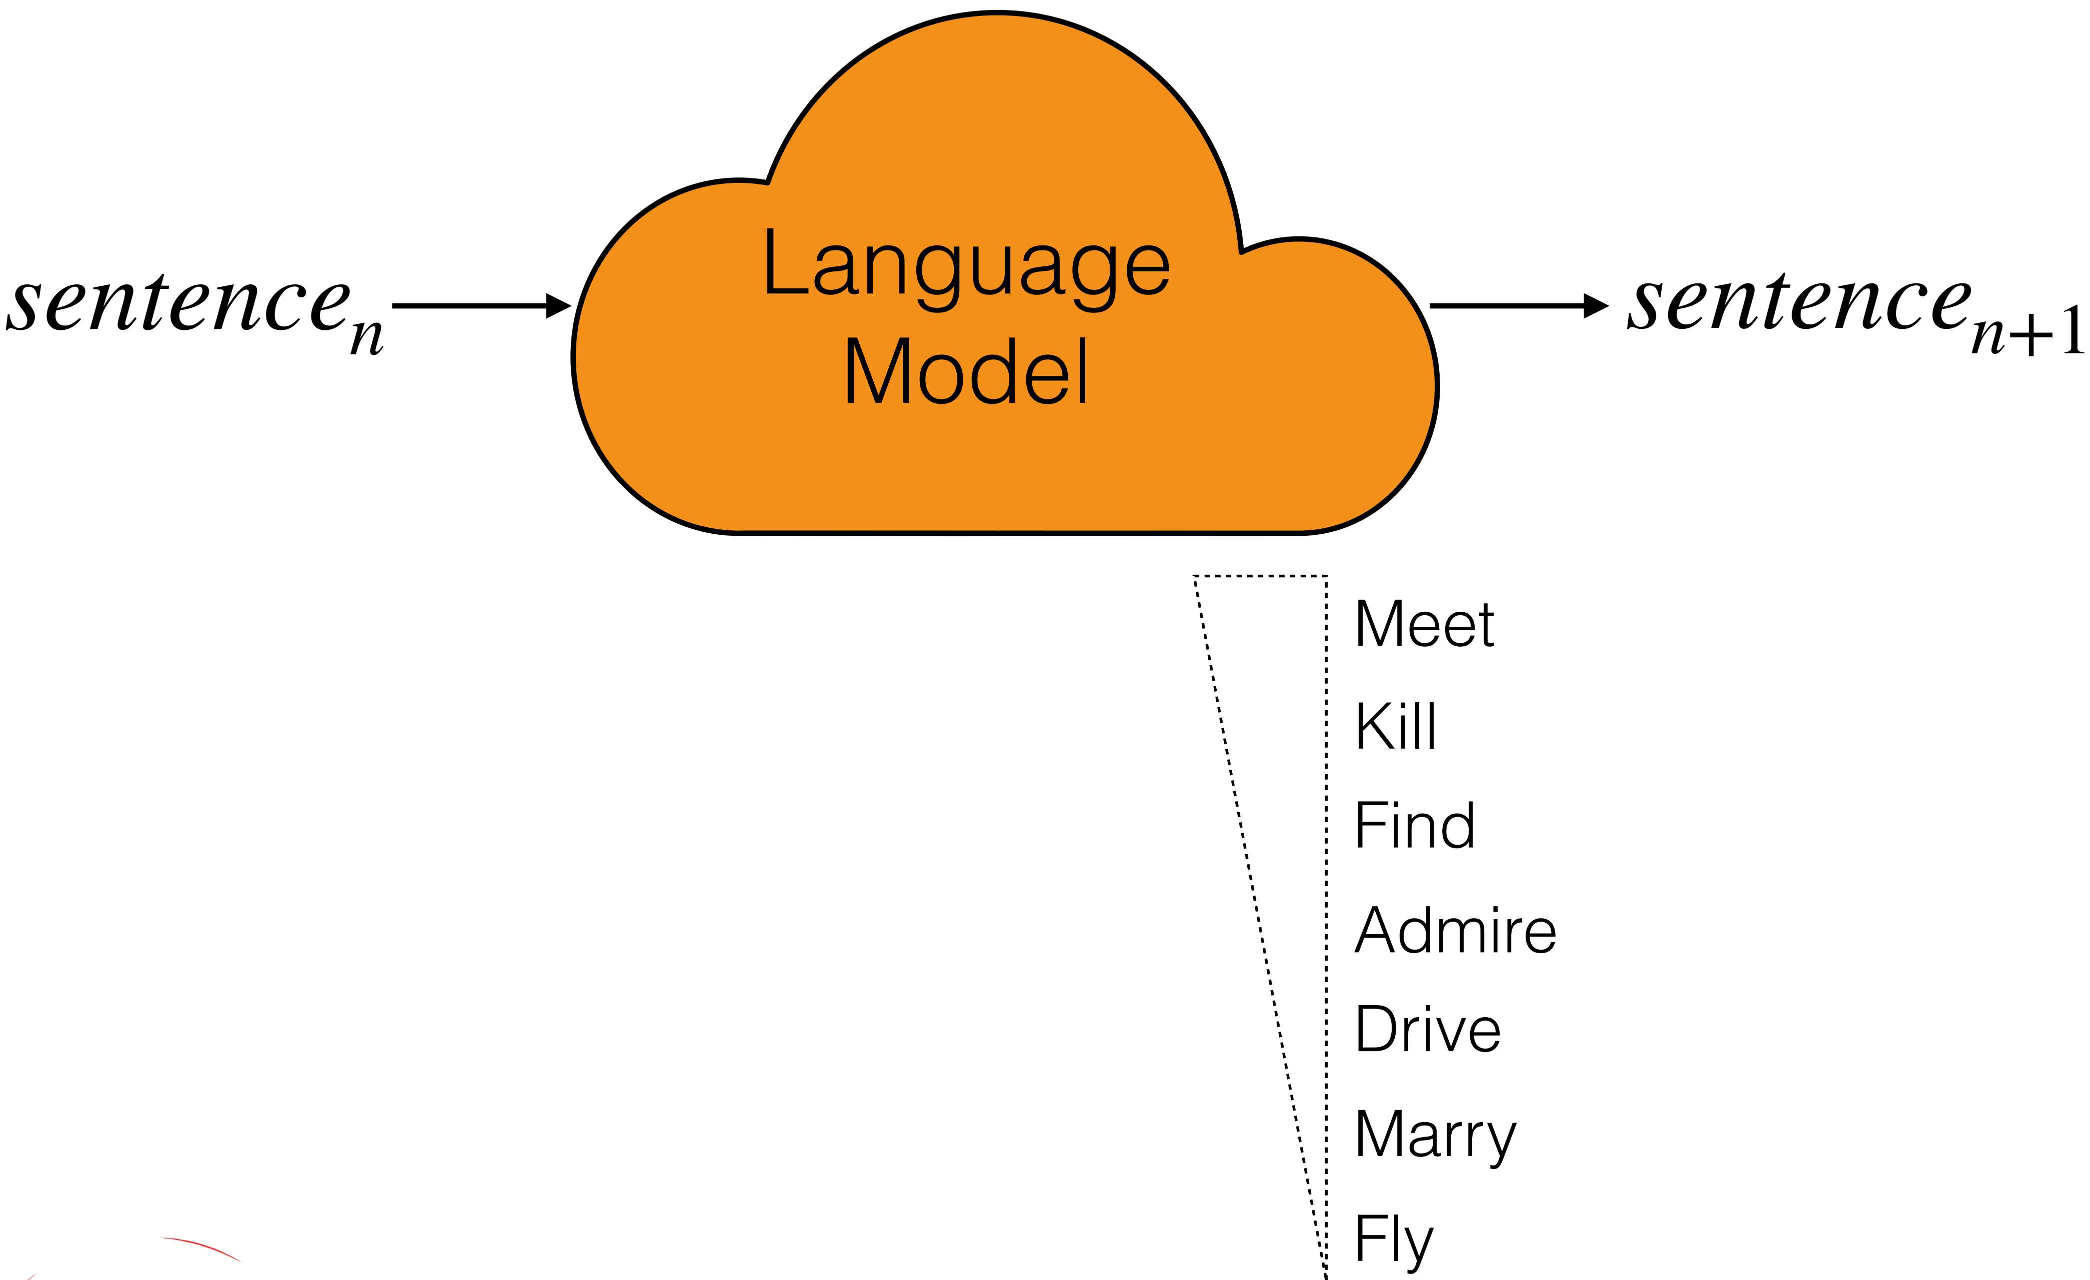

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing School of Interactive Computing

Fine-tuning on goals

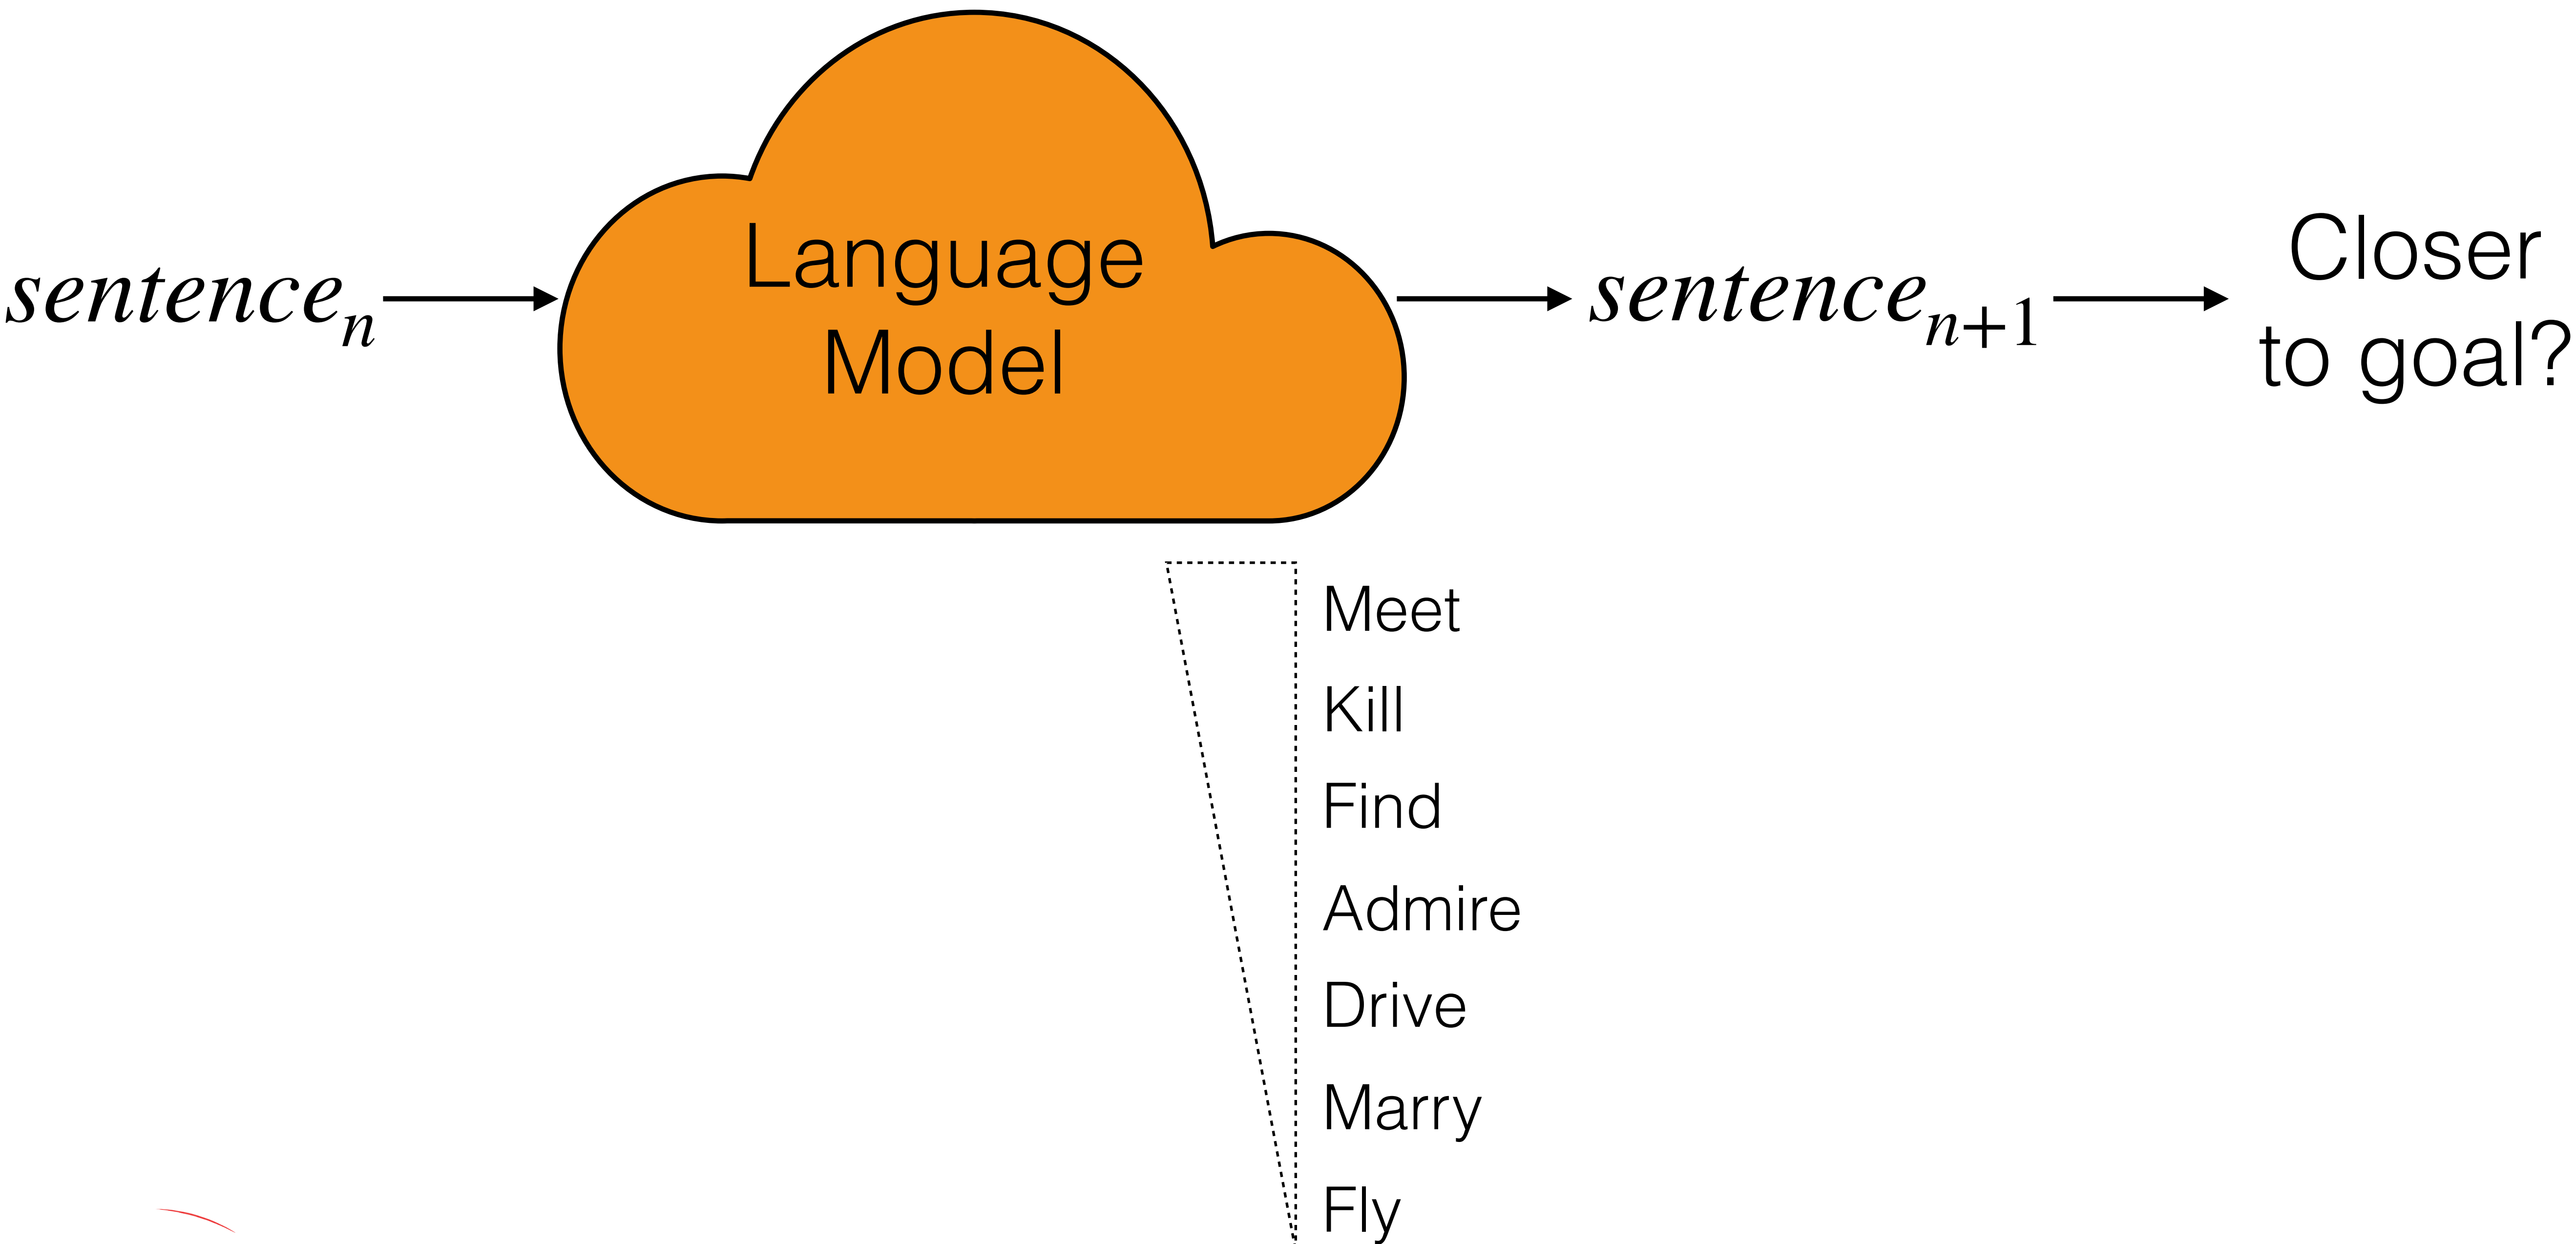

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Fine-tuning on goals

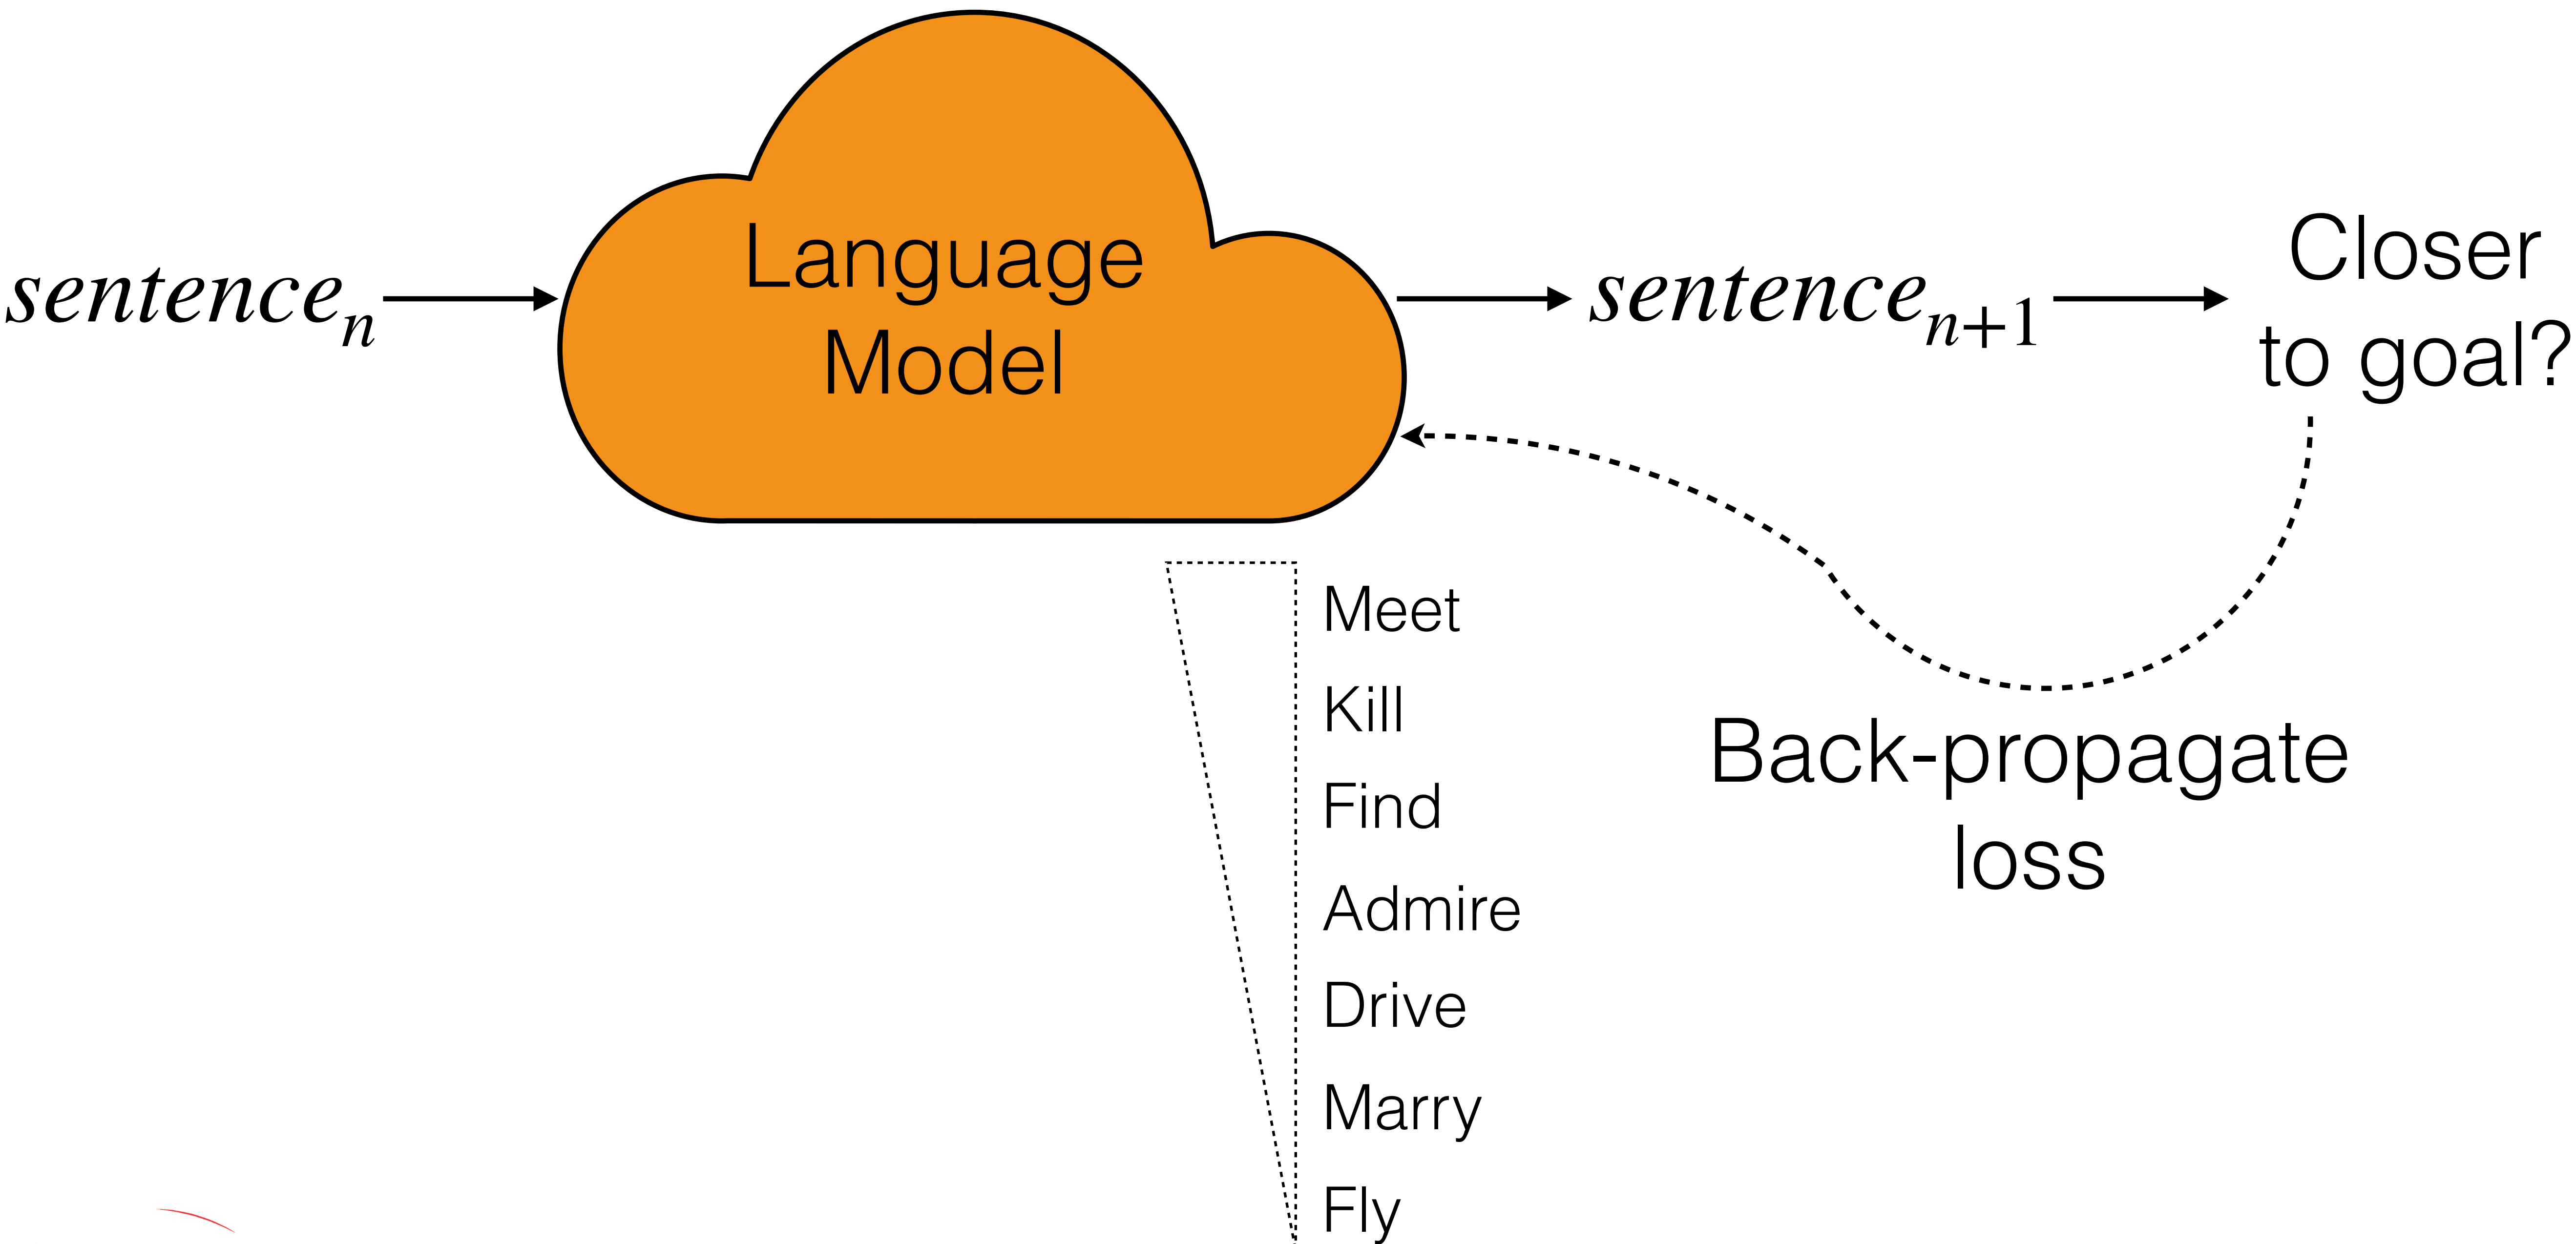

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Fine-tuning on goals

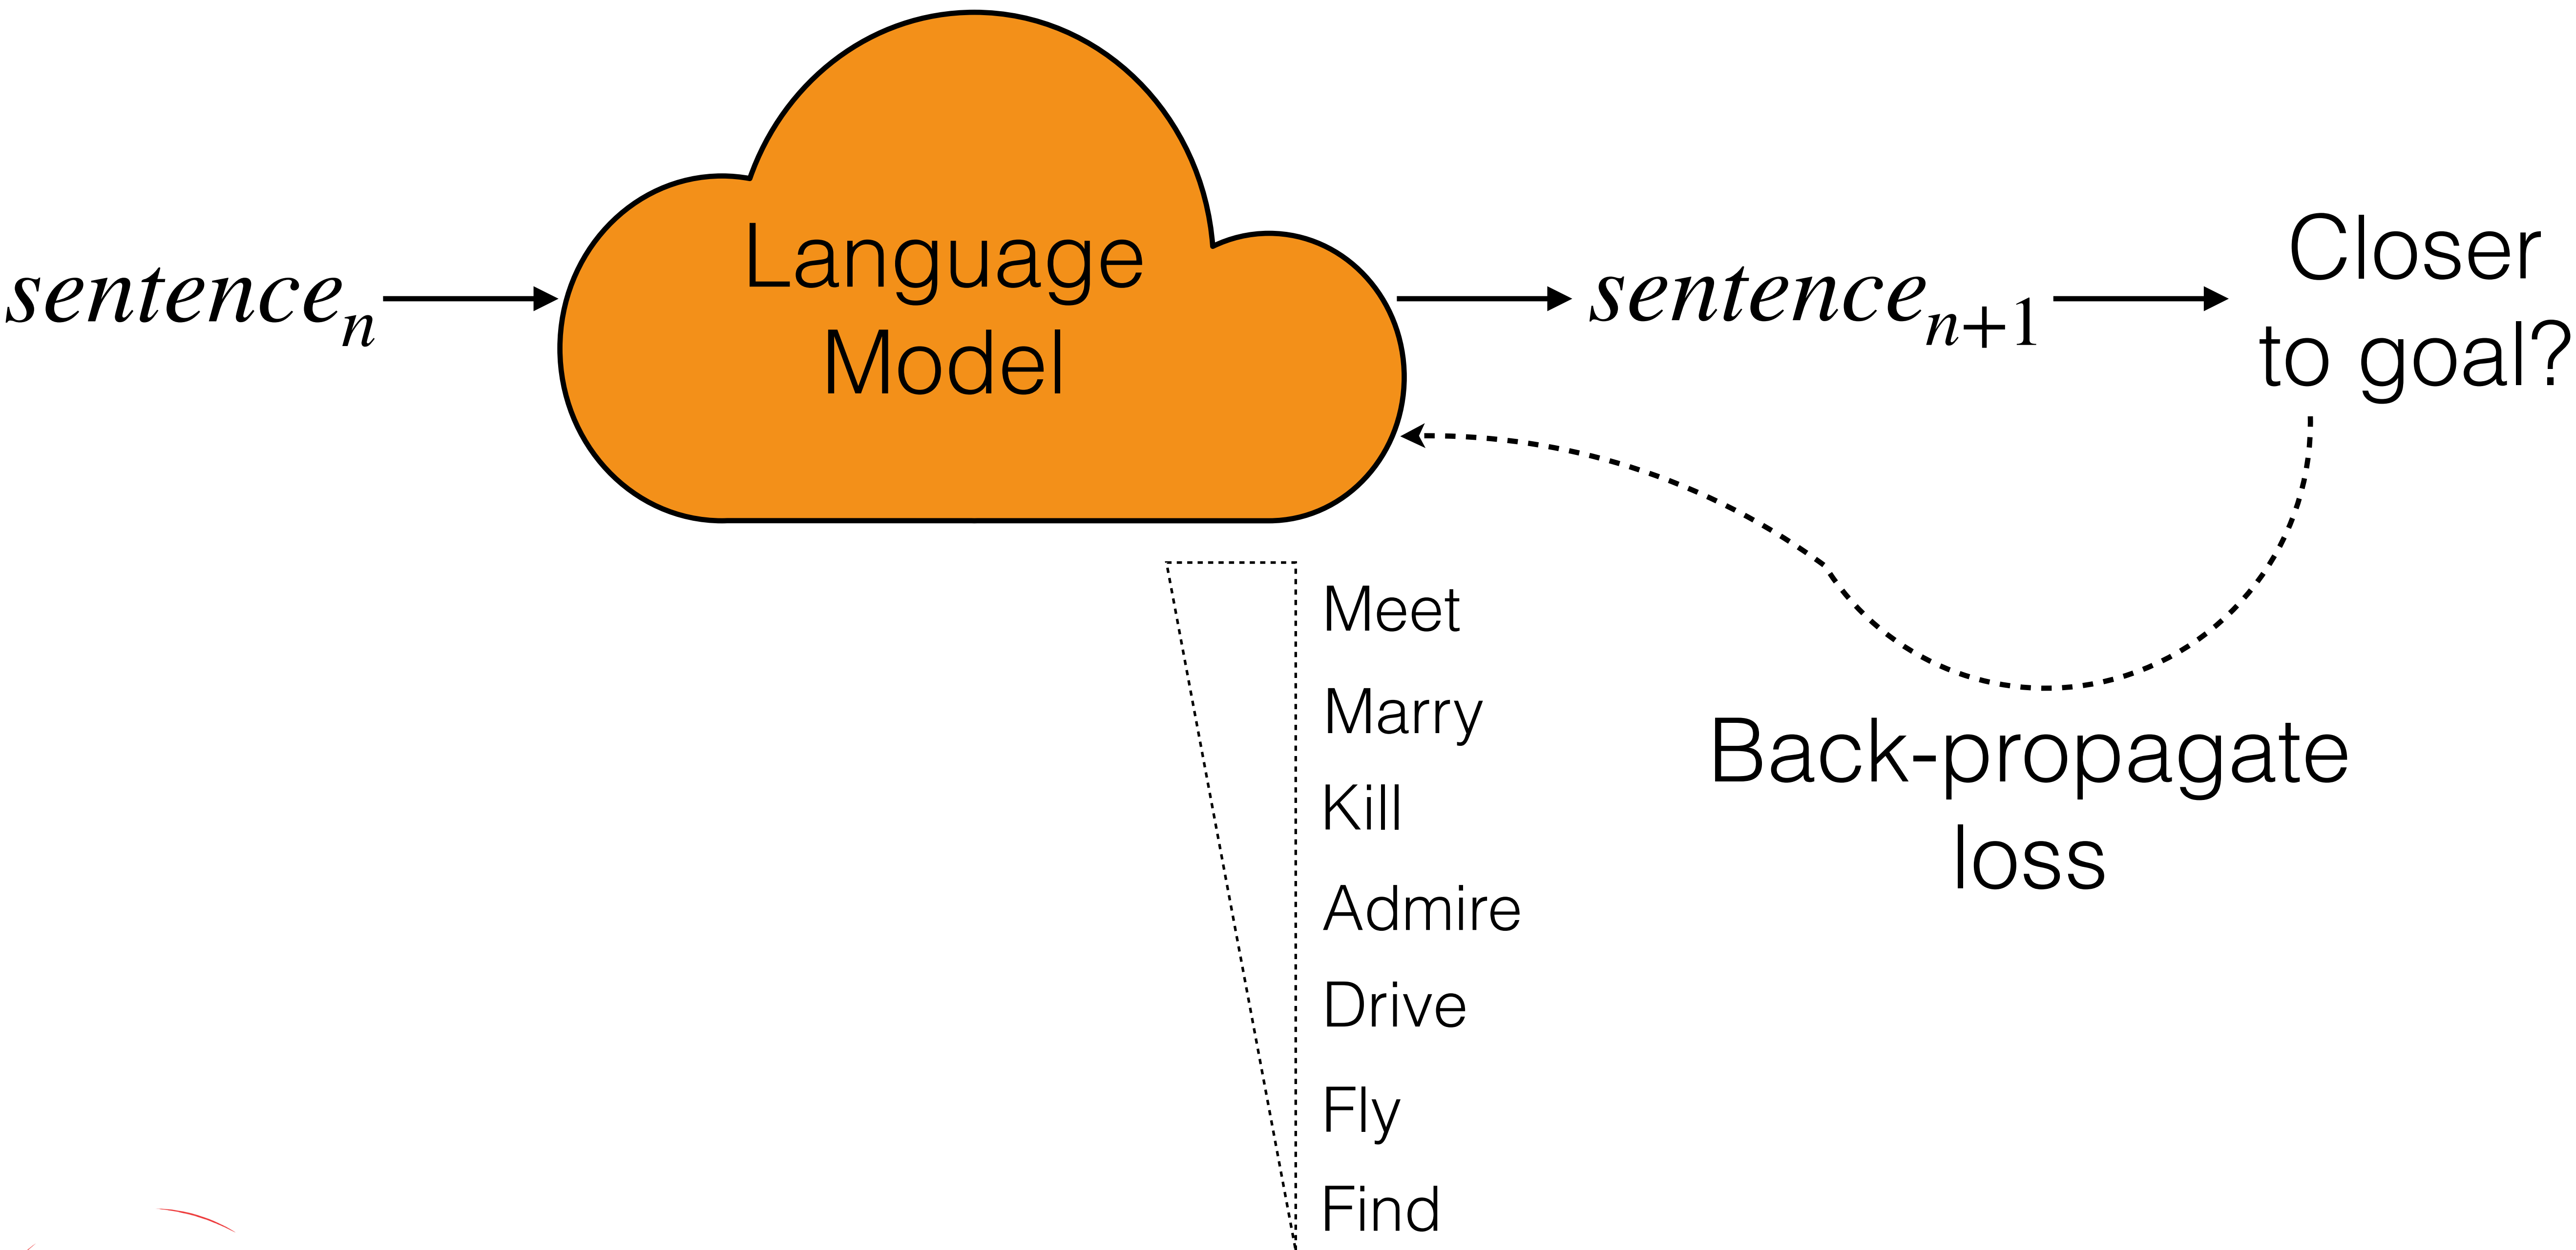

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing School of Interactive Computing

Reward shaping

$$loss \propto distance\ to\ goal$$

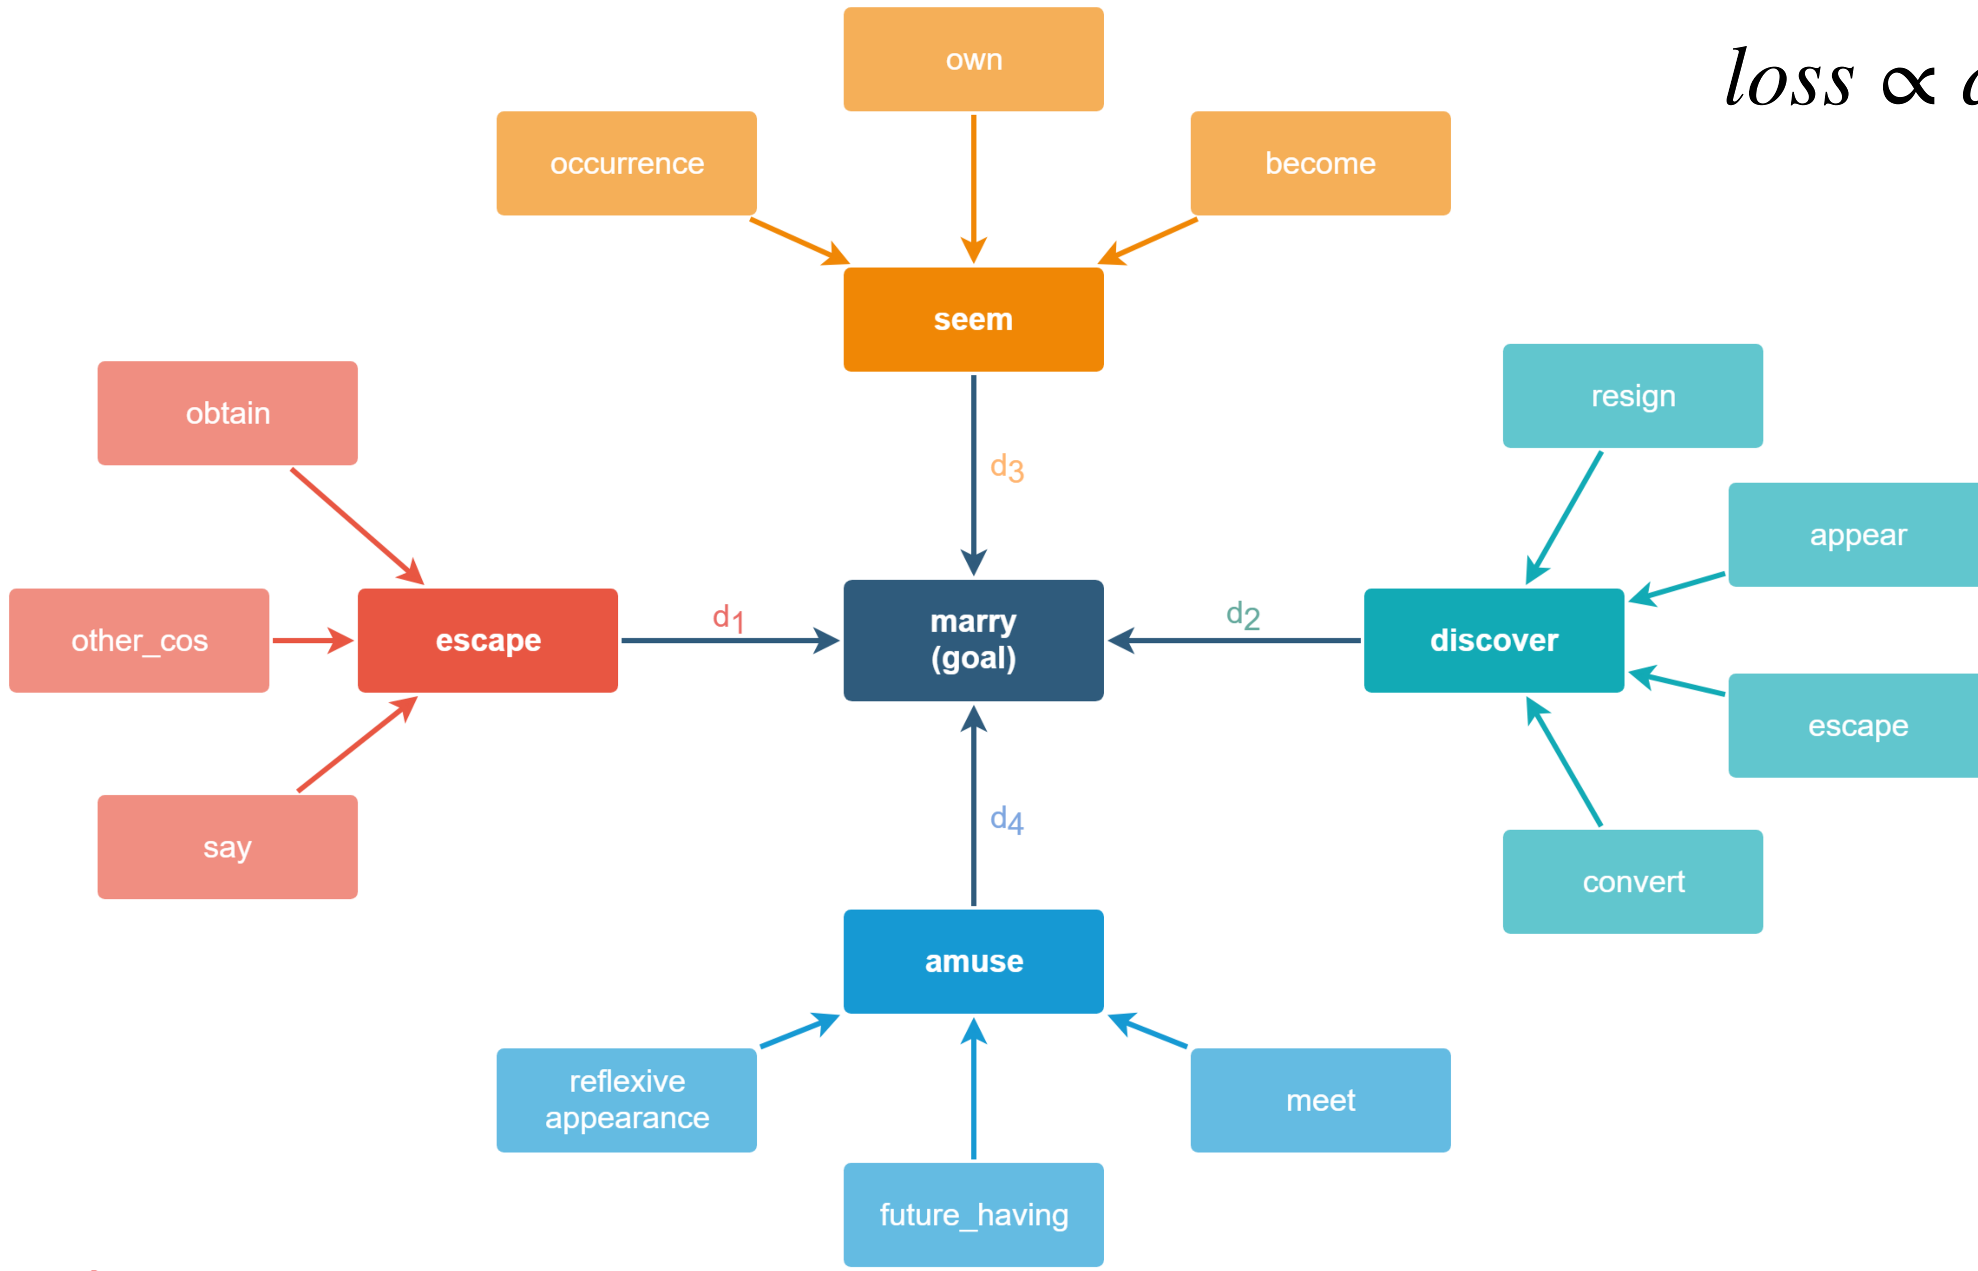

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Reward shaping

$$loss \propto distance\ to\ goal$$

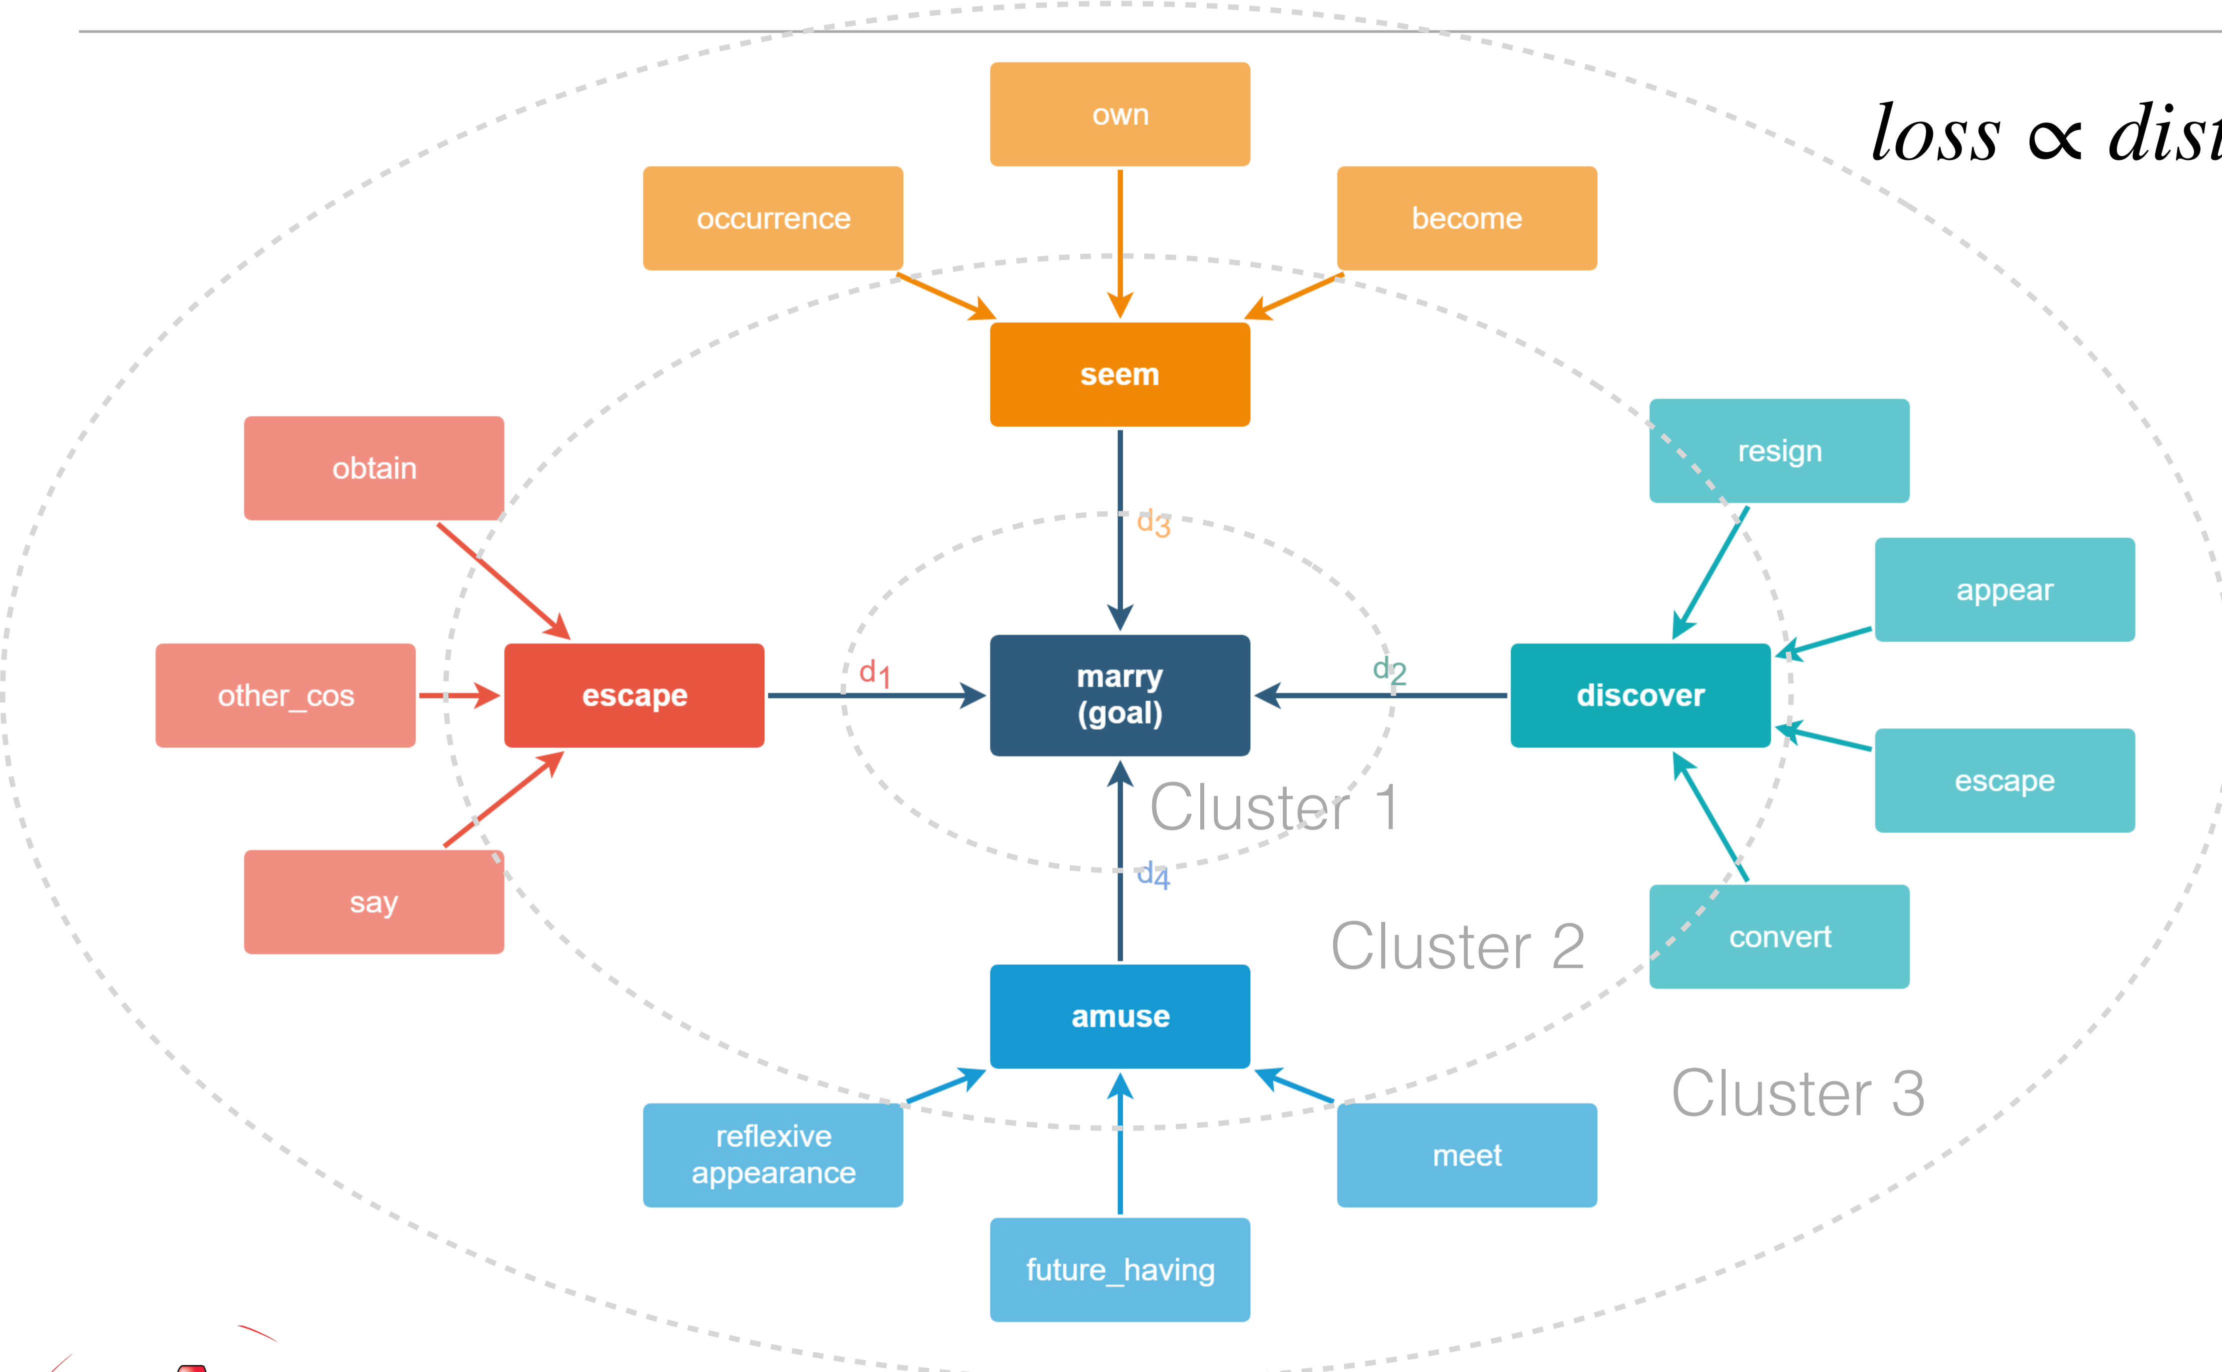

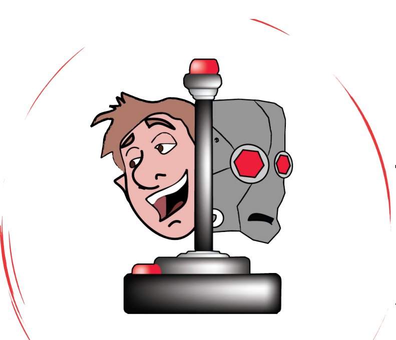

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing School of Interactive Computing

Reward shaping

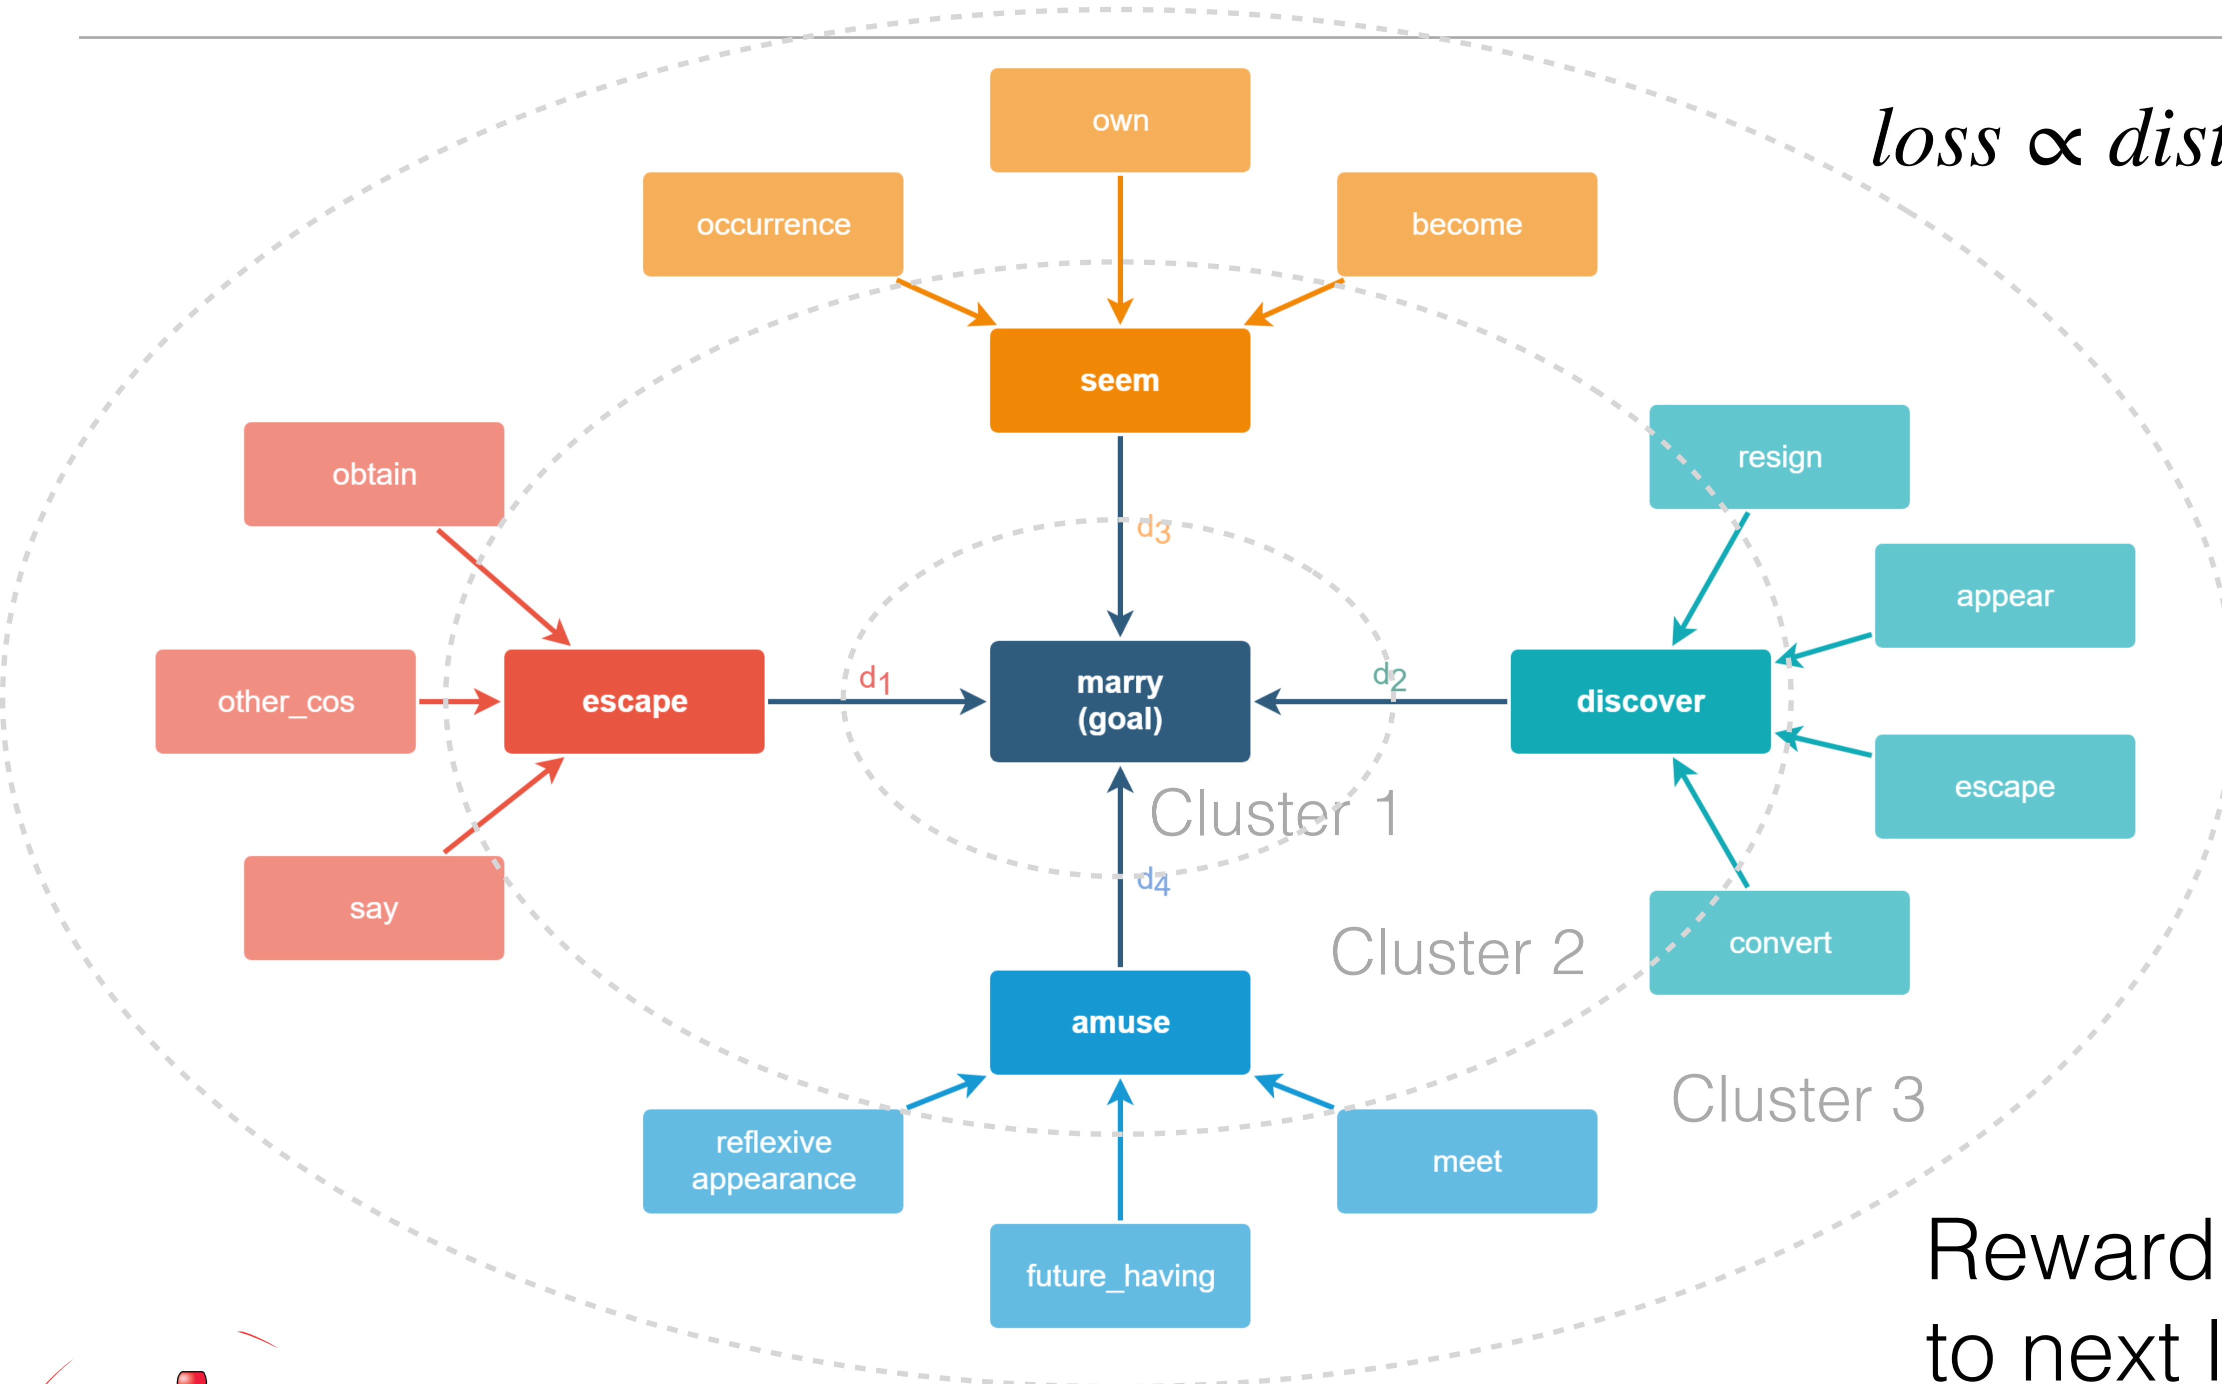

loss ∝ distance to goal

Reward for moving to next lower cluster

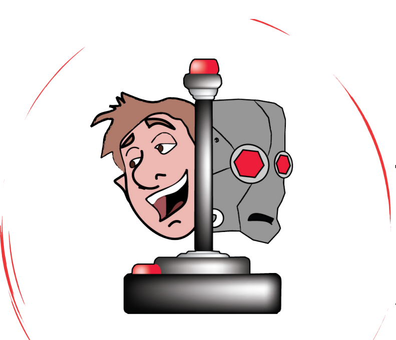

Tambwekar et al. IJCAI 2019 Conference.

Alabdulkarim et al. arXiv:2112.08593

Georgia Tech College of Computing
School of Interactive Computing

Goal achievement rate: > 93%

Perplexity: ~45.0 —> ~7.0

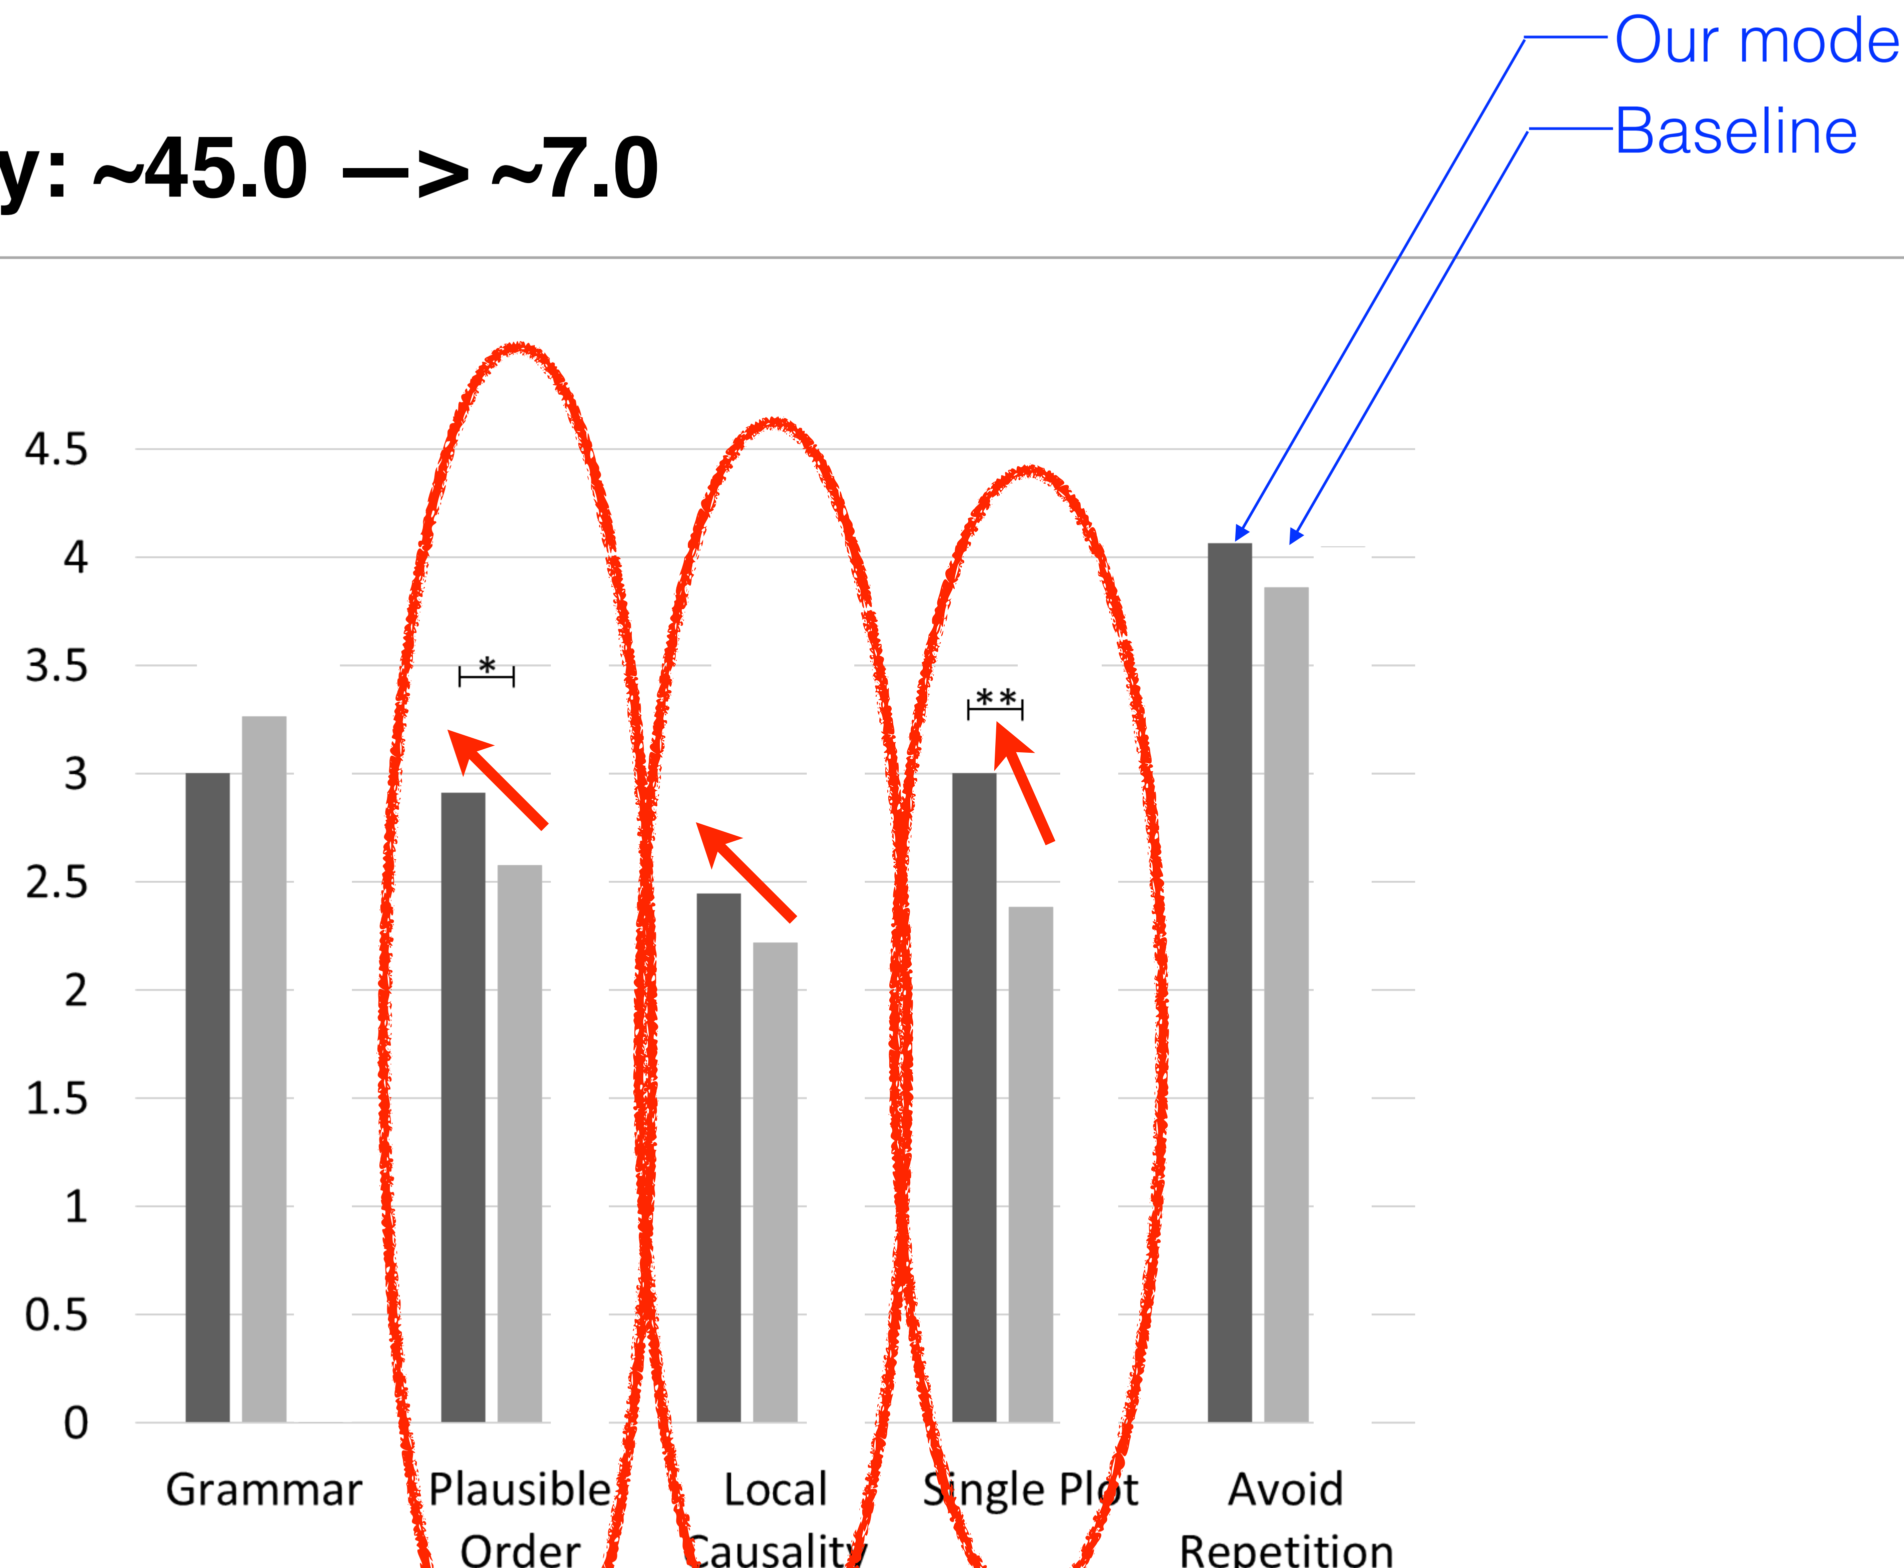

Tambwekar et al. IJCAI 2019 Conference.

Goal: discovery

The tobacco company Hadara creates a form of super tobacco, which in turn is inhabited by a form of super tobacco beetle, which survive the cigarette processing and whose eggs are contained in the smoke of these killer cigarettes, presumably called brand Alex. Cameron, acknowledges the smoking man as being Alex. Morgan, the man confirms that he and his partner are actually doing the same thing. Bailey and Blake look down at the dead man, reply that they had better go. Alex **seeing** Cameron dead but doors closing.

Observations

Observations

- Probabilistic sampling is backward looking and struggles with coherence

Observations

- Probabilistic sampling is backward looking and struggles with coherence
- Fine-tune language models to emulate forward-looking behavior

Observations

- Probabilistic sampling is backward looking and struggles with coherence
- Fine-tune language models to emulate forward-looking behavior
- Coherence is improved by not guaranteed

Georgia Tech College of Computing
School of Interactive Computing

Story generation 3.0

Neuro-symbolic systems

Narrative coherence

Narrative coherence

- Latent state in a neural network is probably not the same as the human reader's understanding of the state of the story world

Narrative coherence

- Latent state in a neural network is probably not the same as the human reader's understanding of the state of the story world
- No requirement that the next event be related in a way that human readers understand

Georgia Tech | College of Computing
School of Interactive Computing

The hungry dog licked her lips as she watched her owner eat. She could see the white teeth in his jaw, then moving behind his right ear.

The hunter picked up his sword, directing the dogs in a cloud of mutterings. "Did you just feed a dog?" she demanded, suddenly even more agitated than before.

"Yes." Jack's voice was rough, and he hated that word. His face was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad. She is one of the most dangerous wolves I have ever met, and if she bites me she will lose."

John shrugged. "Well, it's not your fault," he said.

The hungry dog licked her lips as she watched her owner eat.
She could see the white teeth in his jaw, then moving behind his right ear.

The hunter picked up his sword, directing the dogs in a cloud of mutterings. "Did you just feed a dog?" she demanded, suddenly even more agitated than before.

"Yes." Jack's voice was rough, and he hated that word. His face was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad. She is one of the most dangerous wolves I have ever met, and if she bites me she will lose."

John shrugged. "Well, it's not your fault," he said.

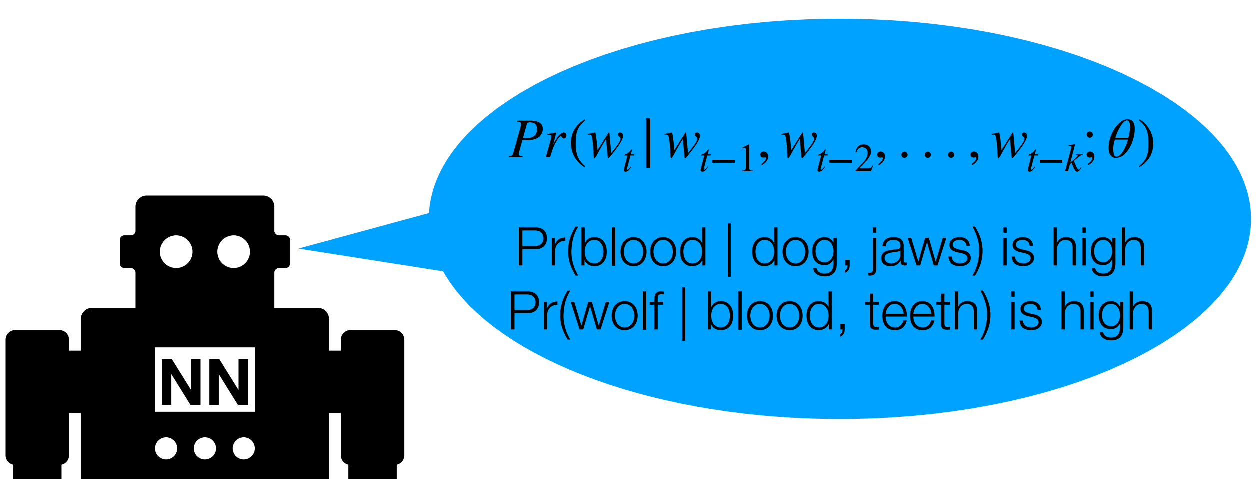

The hungry dog licked her lips as she watched her owner eat. She could see the white teeth in his jaw, then moving behind his right ear.

The hunter picked up his sword, directing the dogs in a cloud of mutterings. "Did you just feed a dog?" she demanded, suddenly even more agitated than before.

"Yes." Jack's voice was rough, and he hated that word. His face was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad. She is one of the most dangerous wolves I have ever met, and if she bites me she will lose."

John shrugged. "Well, it's not your fault," he said.

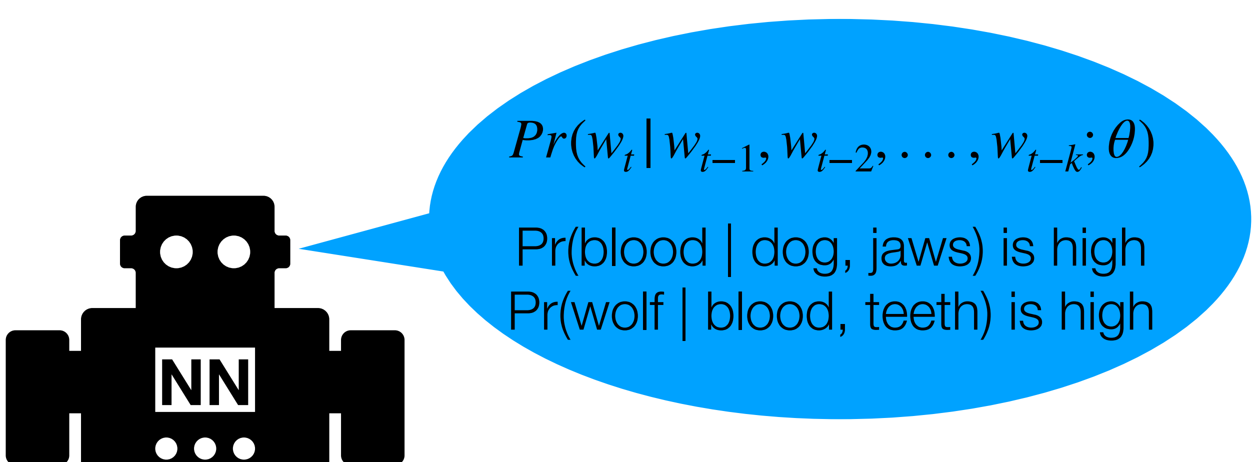

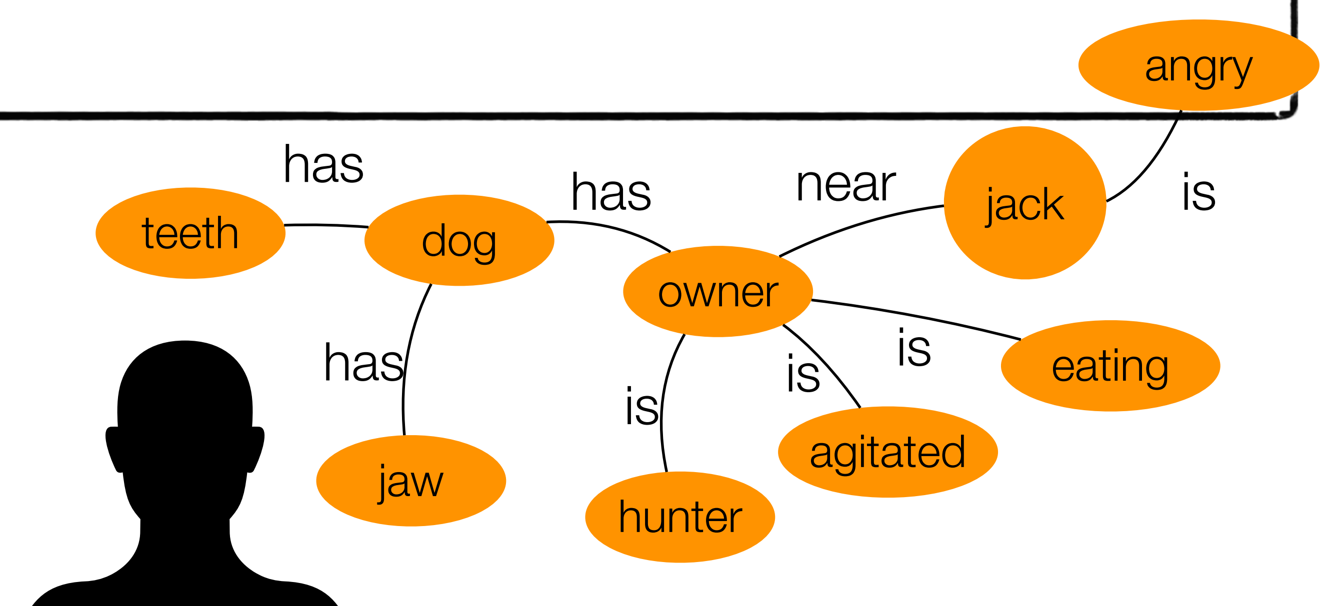

Cognitive science

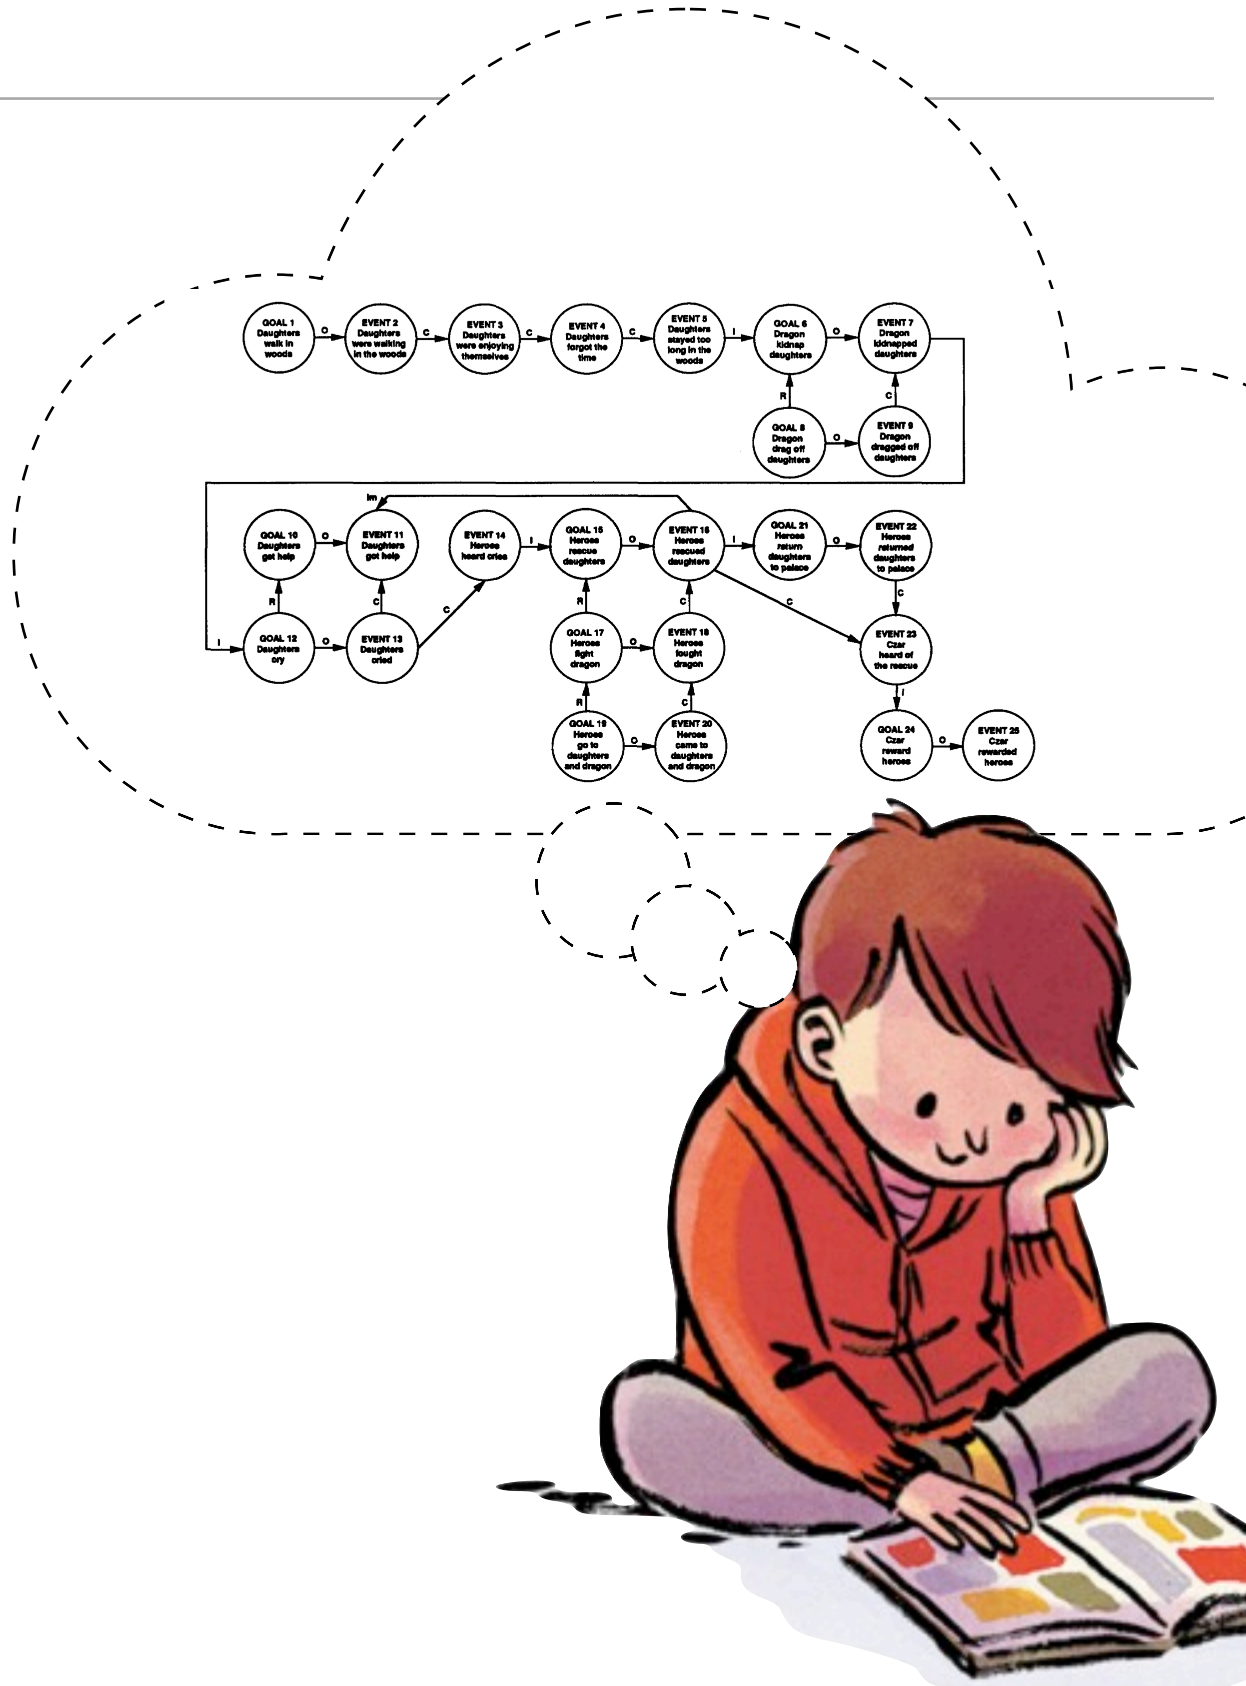

Cognitive science

- Readers build mental models of the underlying fictional world

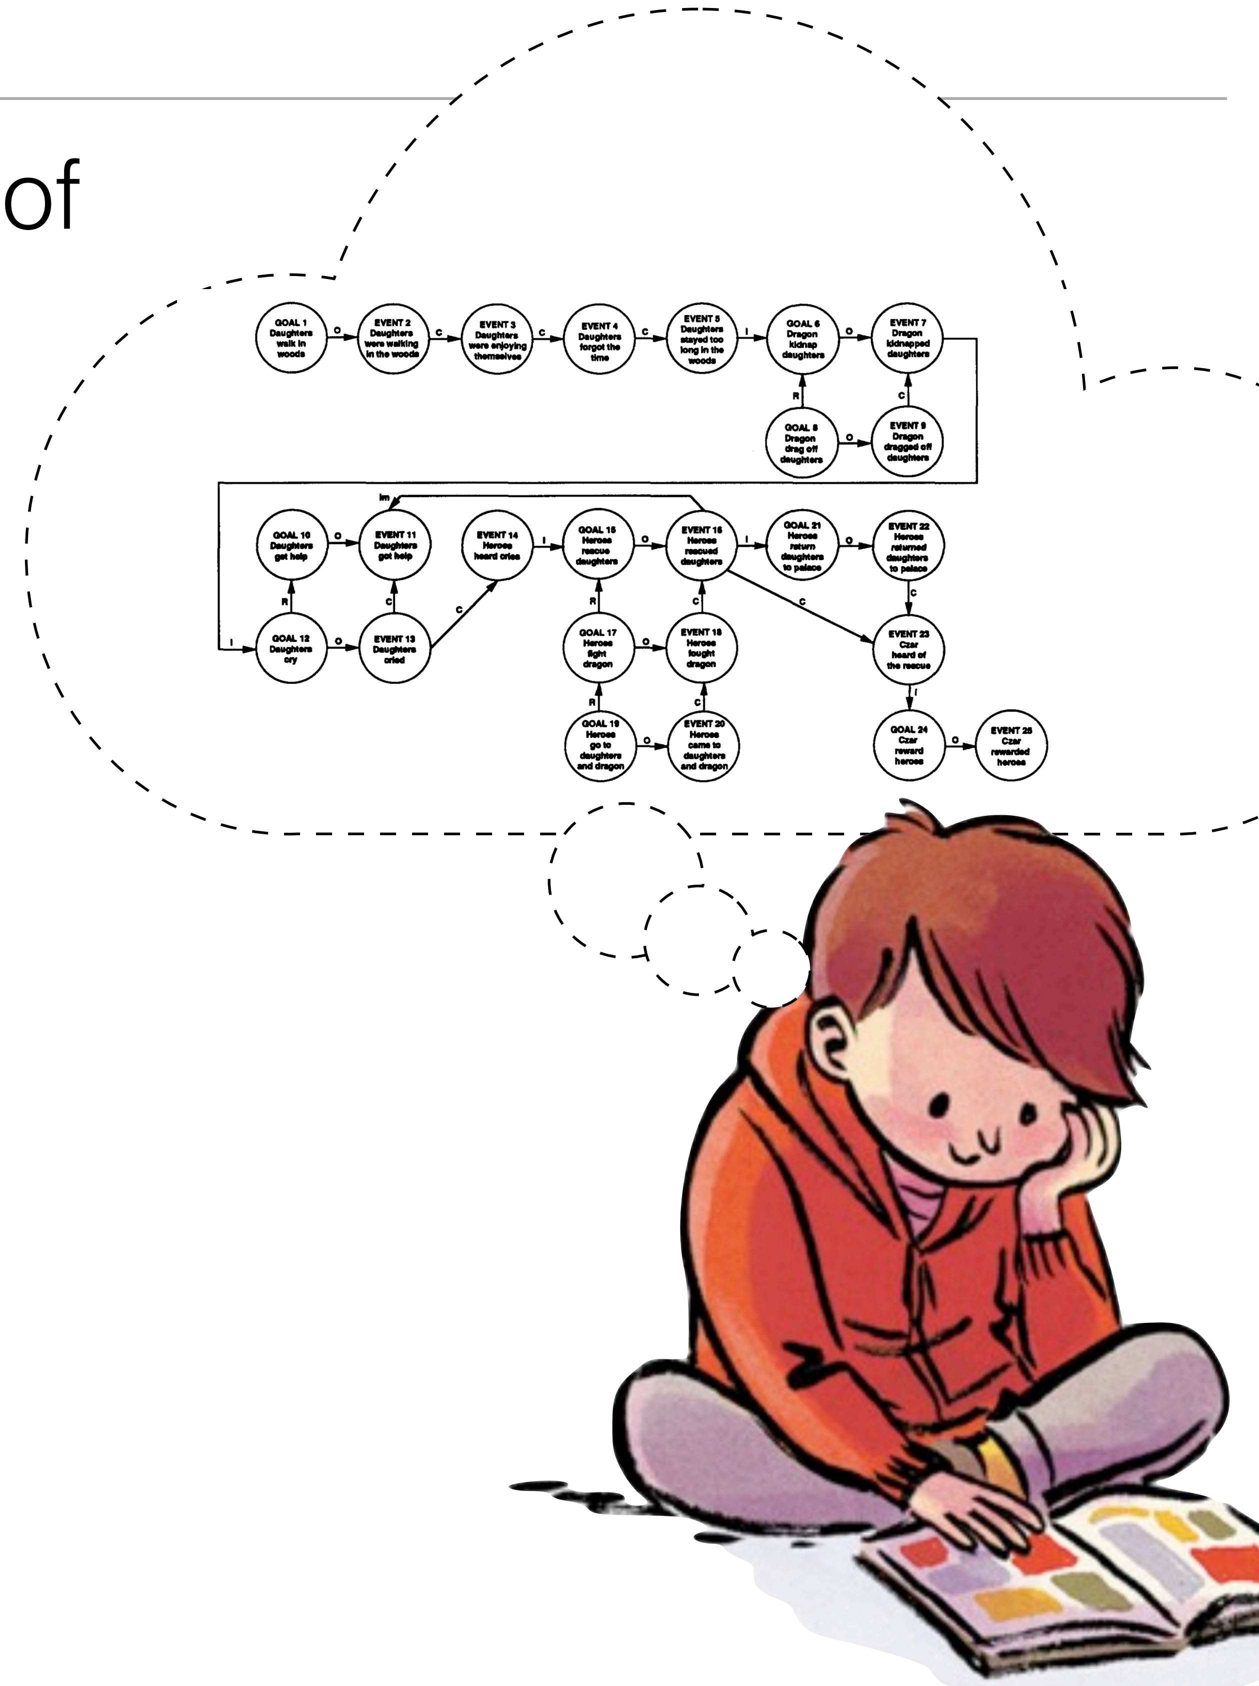

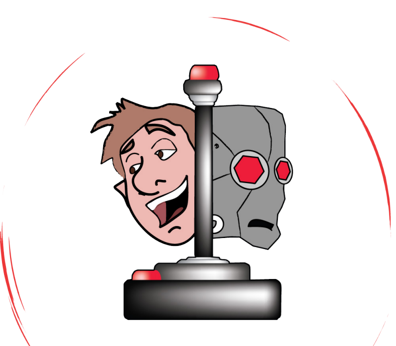

Cognitive science

- Readers build mental models of the underlying fictional world

- Readers track…

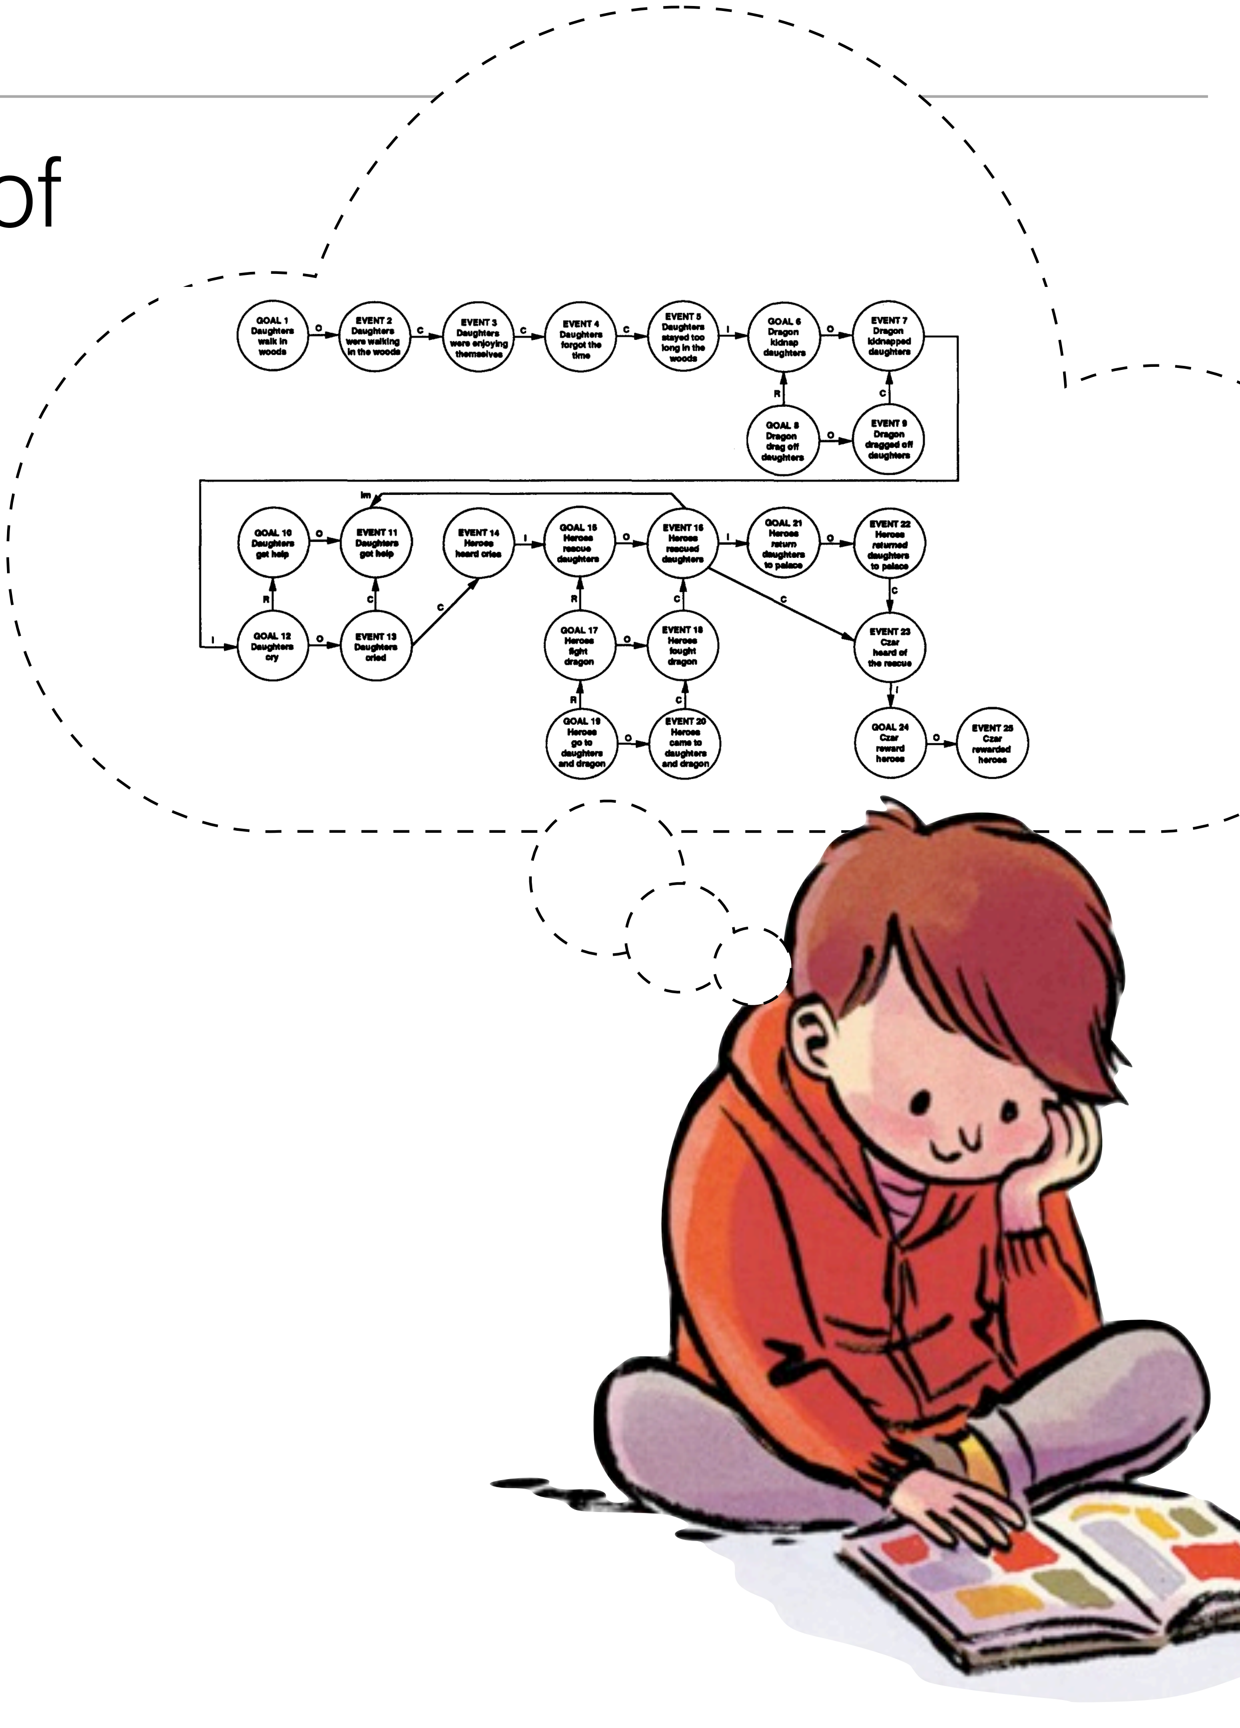

Cognitive science

- Readers build mental models of the underlying fictional world
- Readers track…
- Causal enablement
  - What previous events occurred that were necessary for the current event?

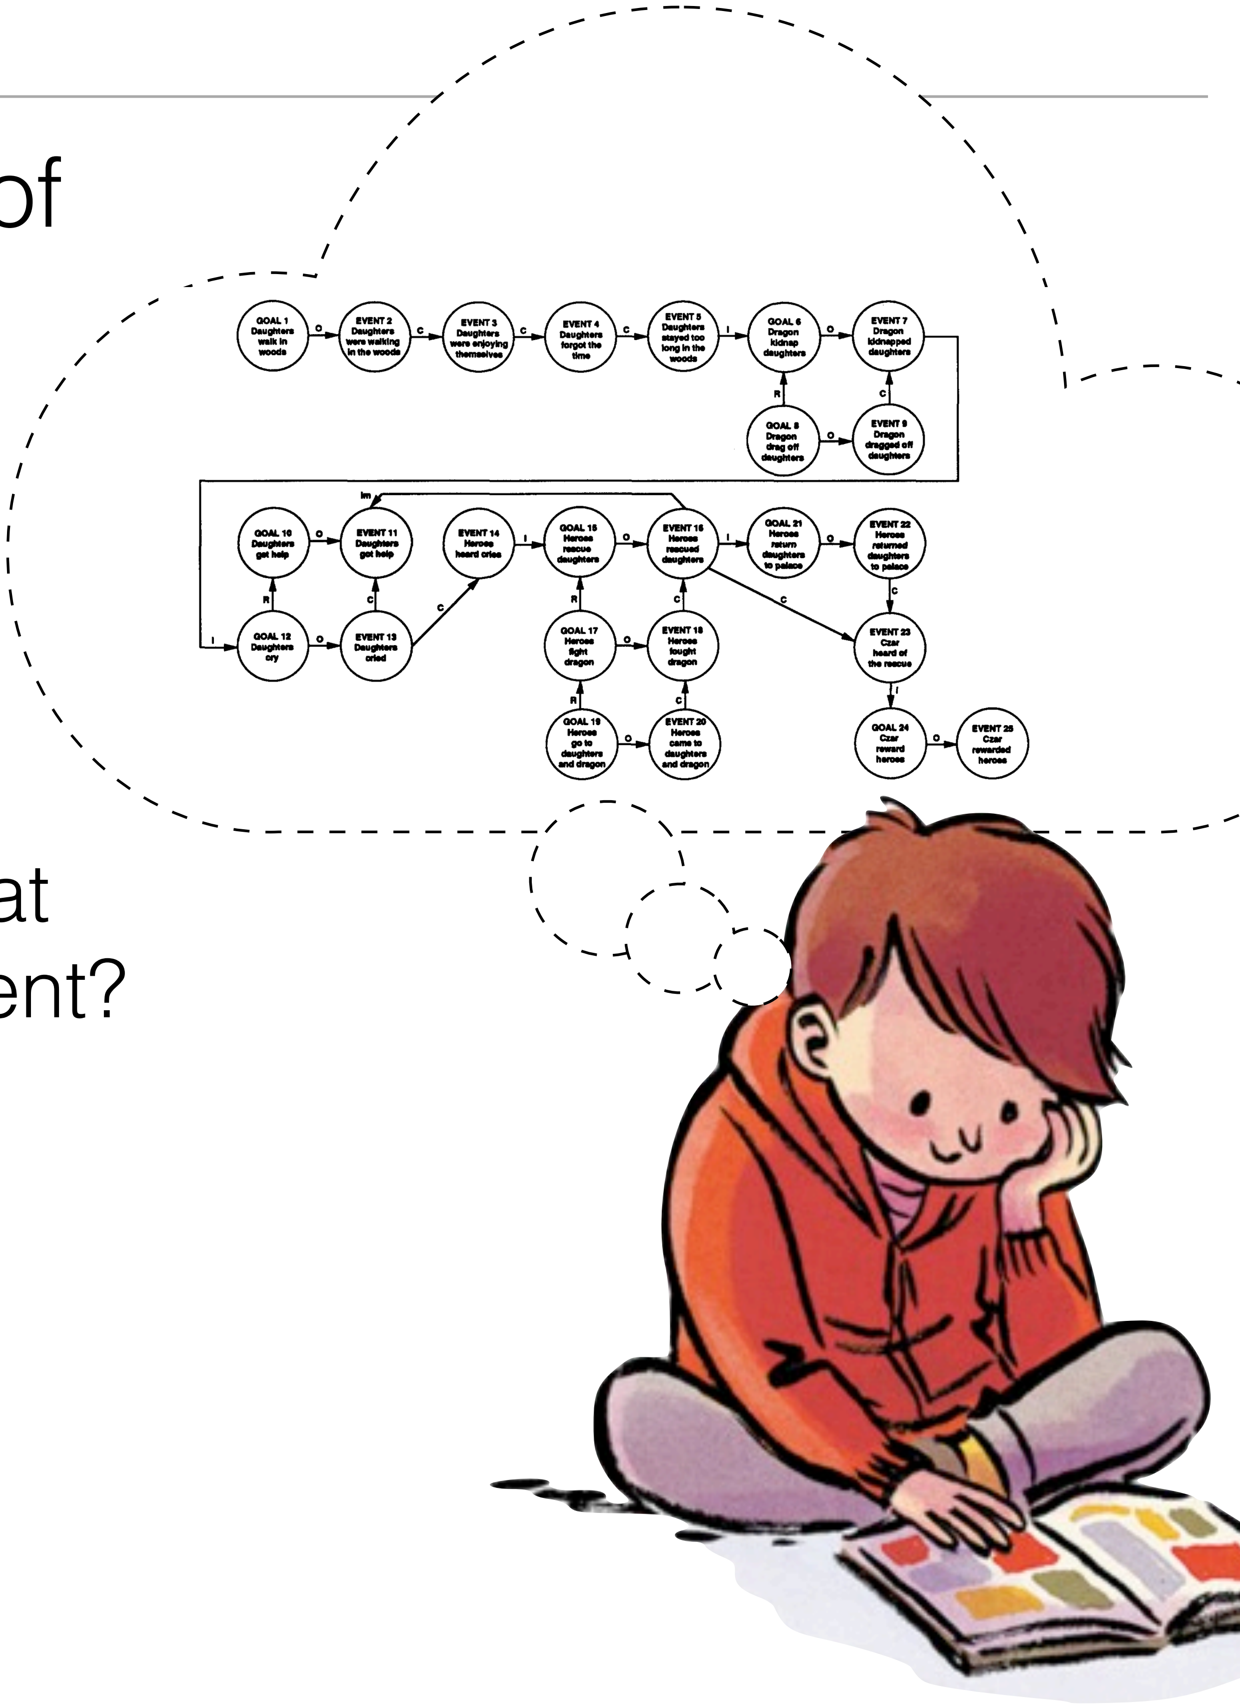

Cognitive science

- Readers build mental models of the underlying fictional world
- Readers track…
- Causal enablement
  - What previous events occurred that were necessary for the current event?
- Character goals
  - What is each character trying to accomplish and how does their current actions fit?

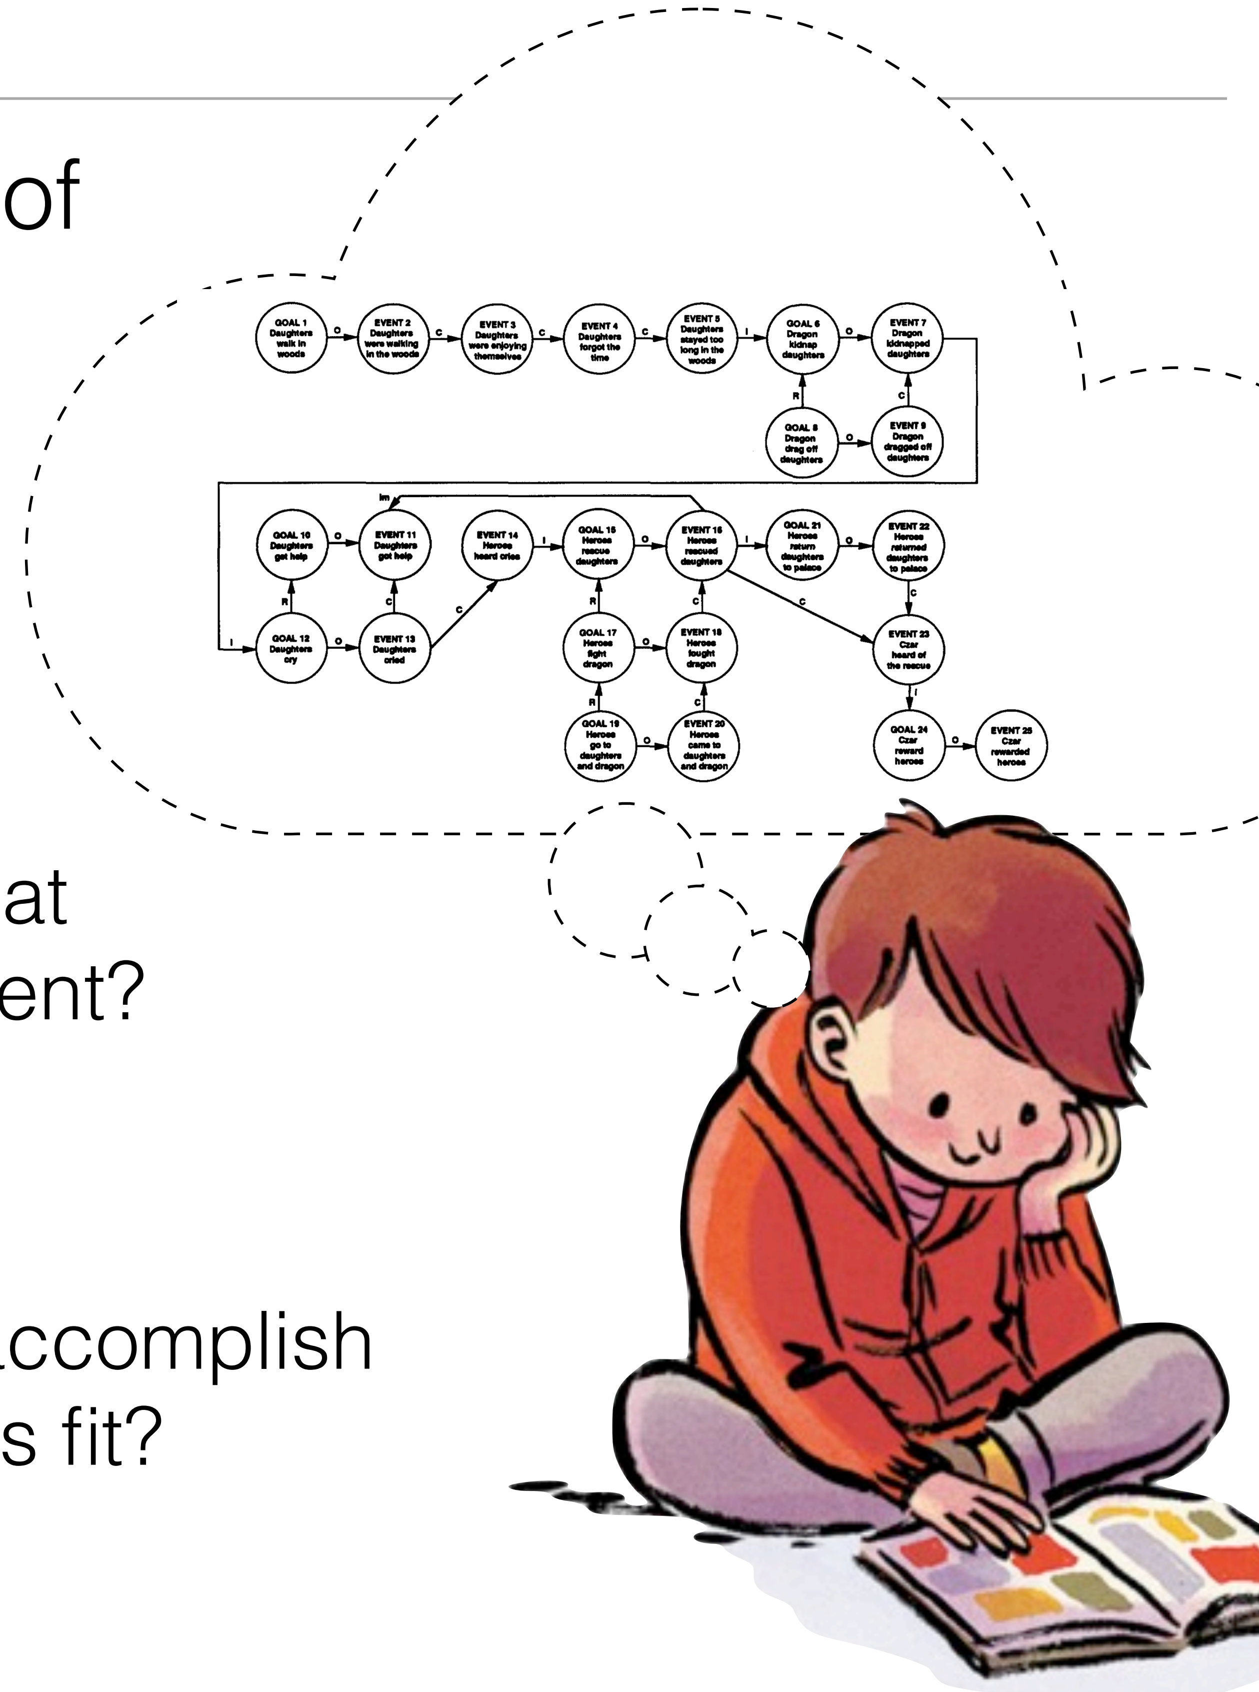

Neuro-symbolic story generation

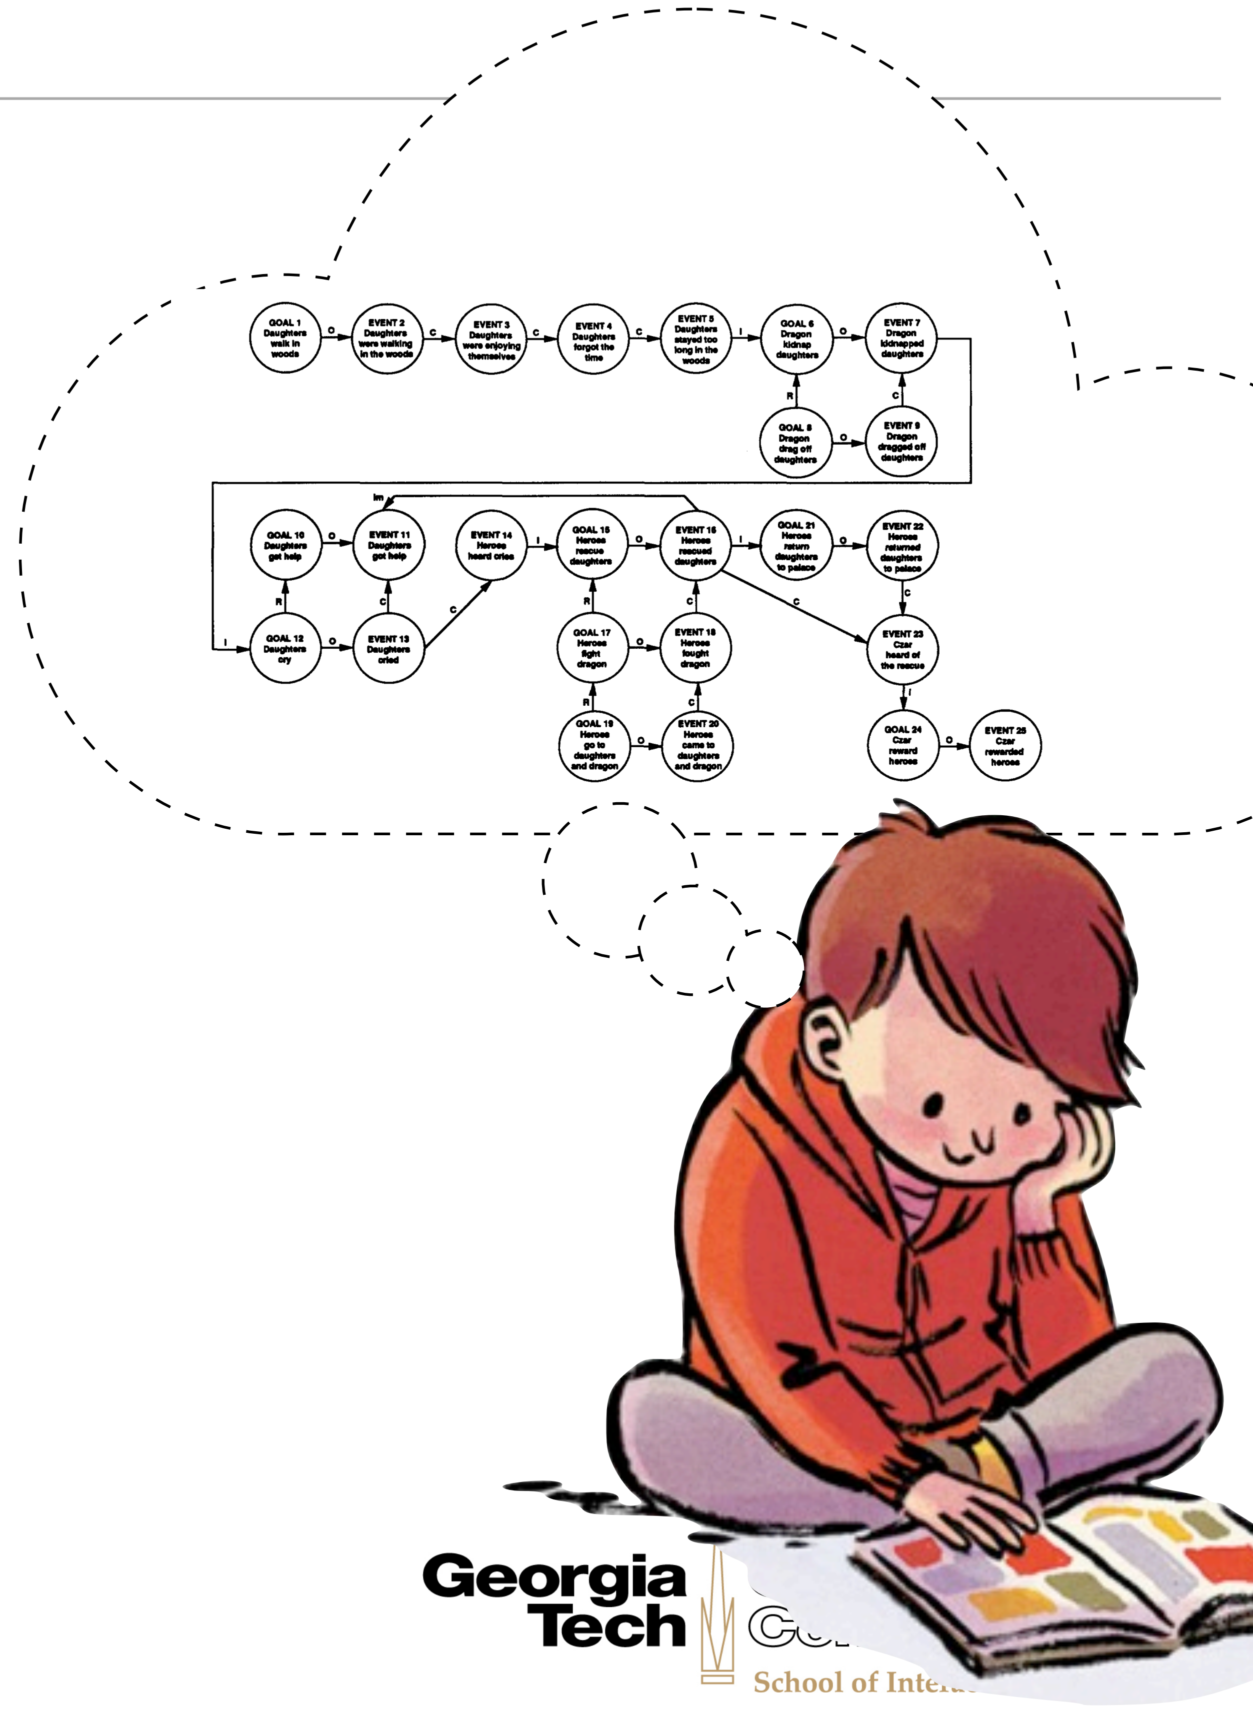

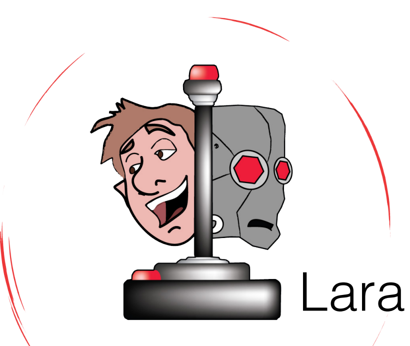

Lara Martin. Dissertation.

Neuro-symbolic story generation

- Track a latent neural state **and** a symbolic "reader model" (graph)
- Use reader model to inform story continuations

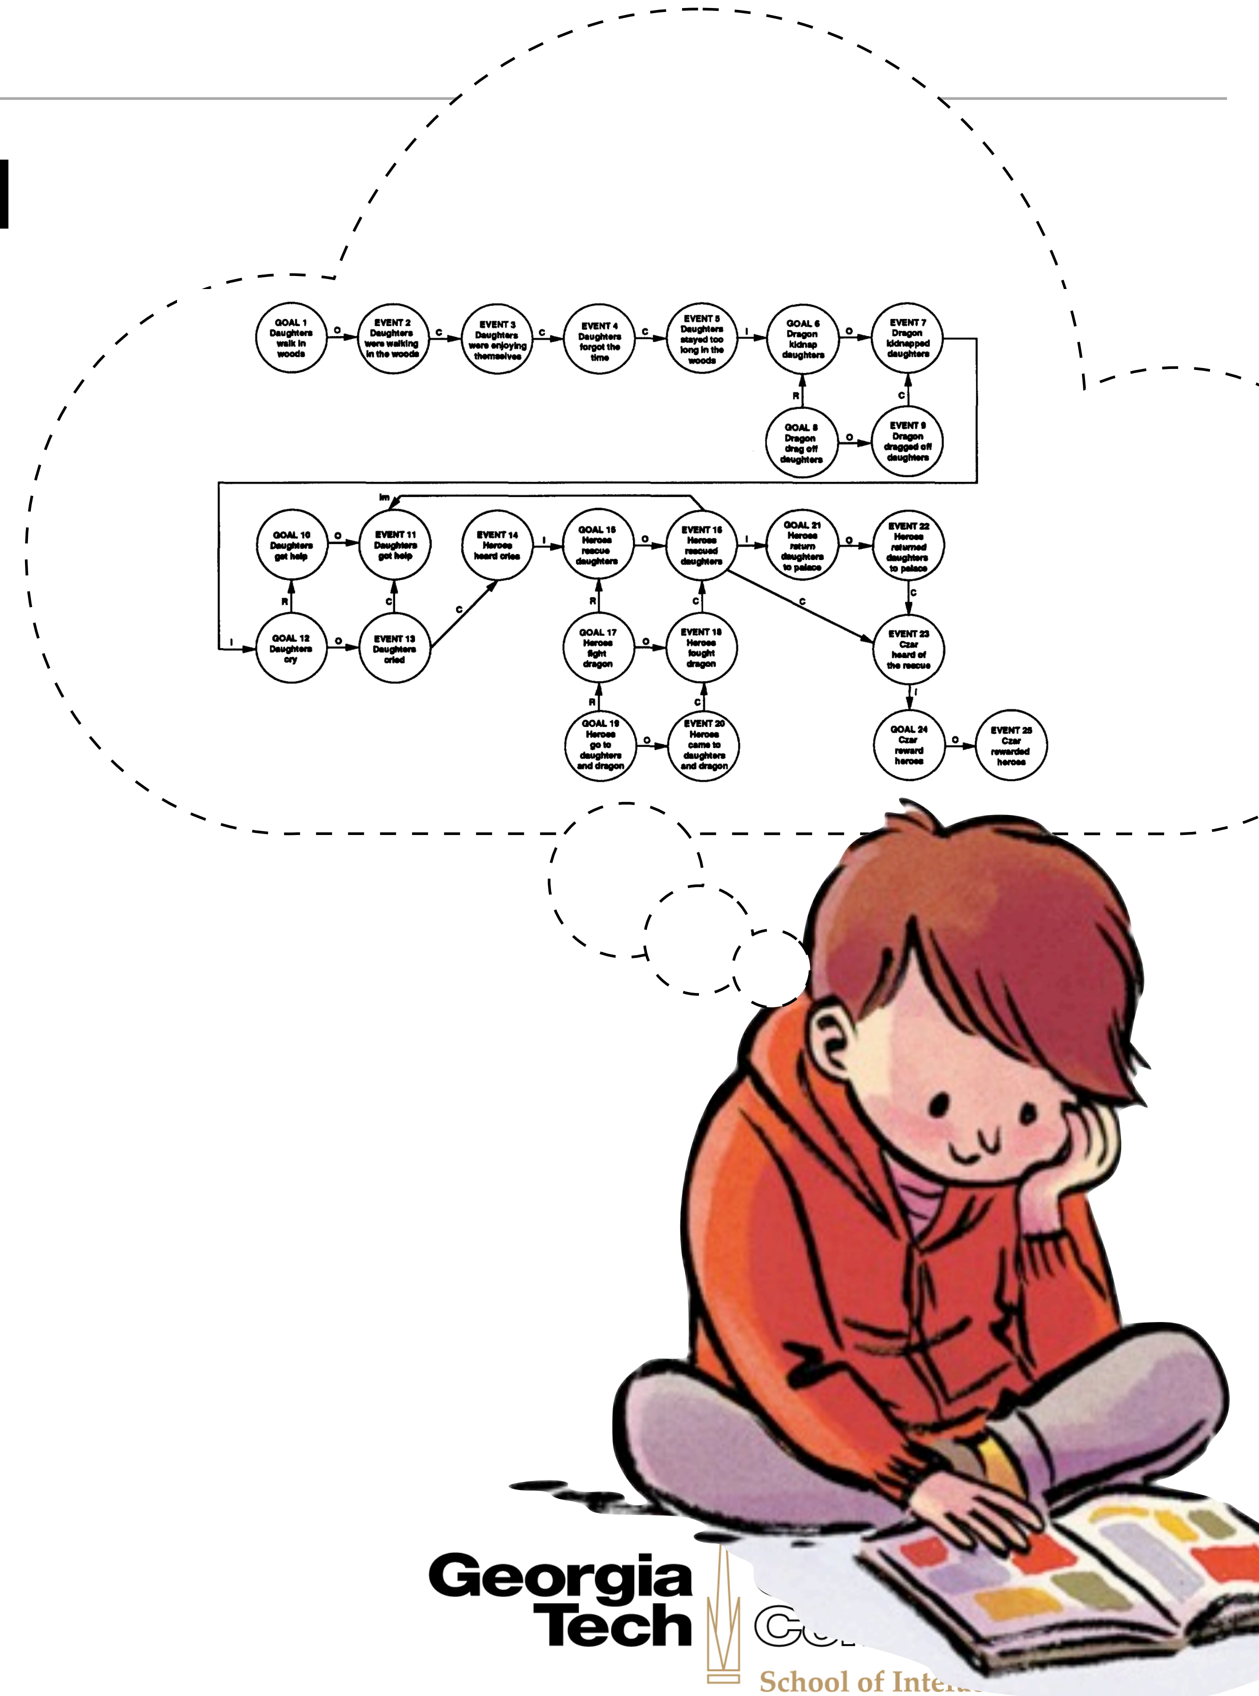

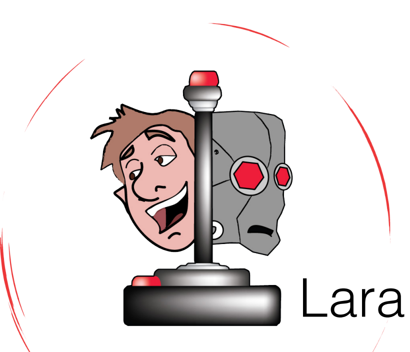

Lara Martin. Dissertation.

The hungry dog licked her lips as she watched her owner eat. She could see the white teeth in his jaw, then moving behind his right ear.

The hunter picked up his sword, directing the dogs in a cloud of mutterings. "Did you just feed a dog?" she demanded, suddenly even more agitated than before.

"Yes." Jack's voice was rough, and he hated that word. His face was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

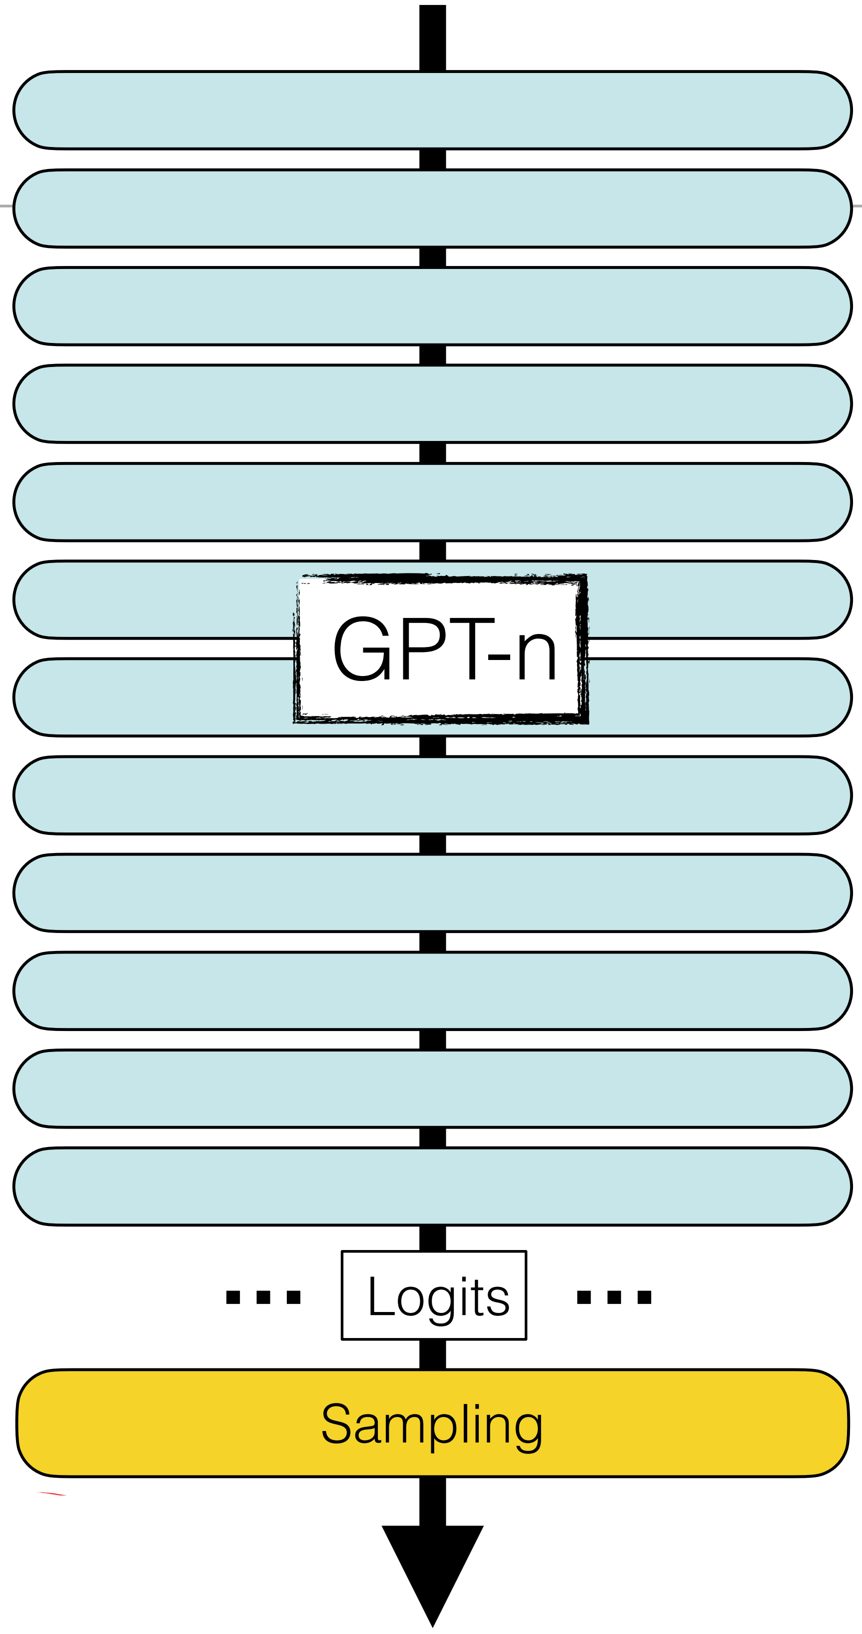

was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

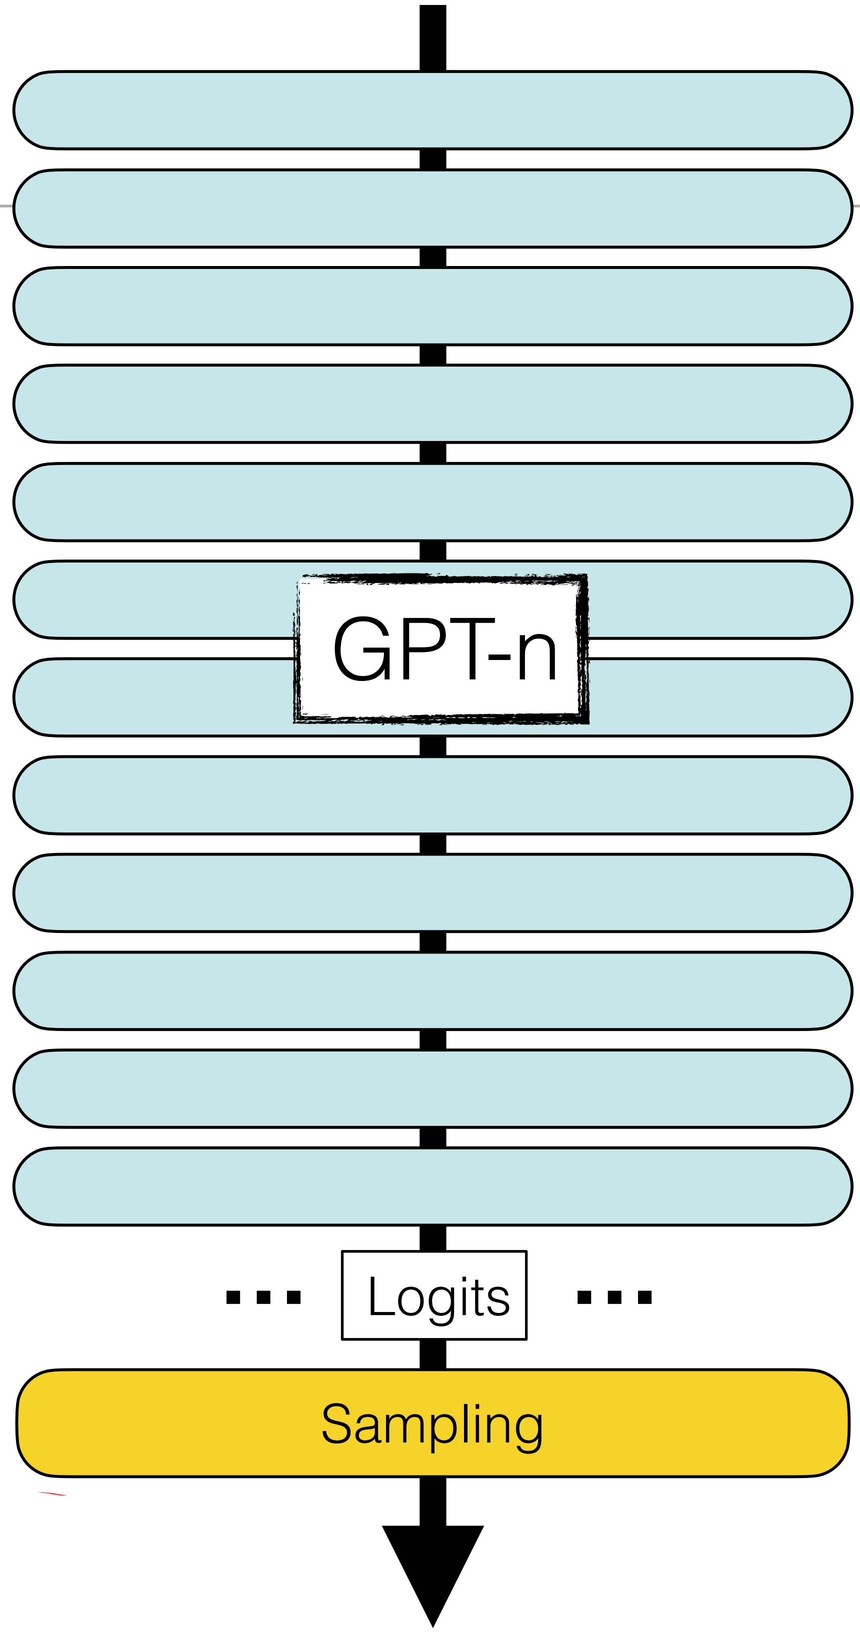

She is one of the most dangerous wolves I have ever met, and if she bites me she will lose.

was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

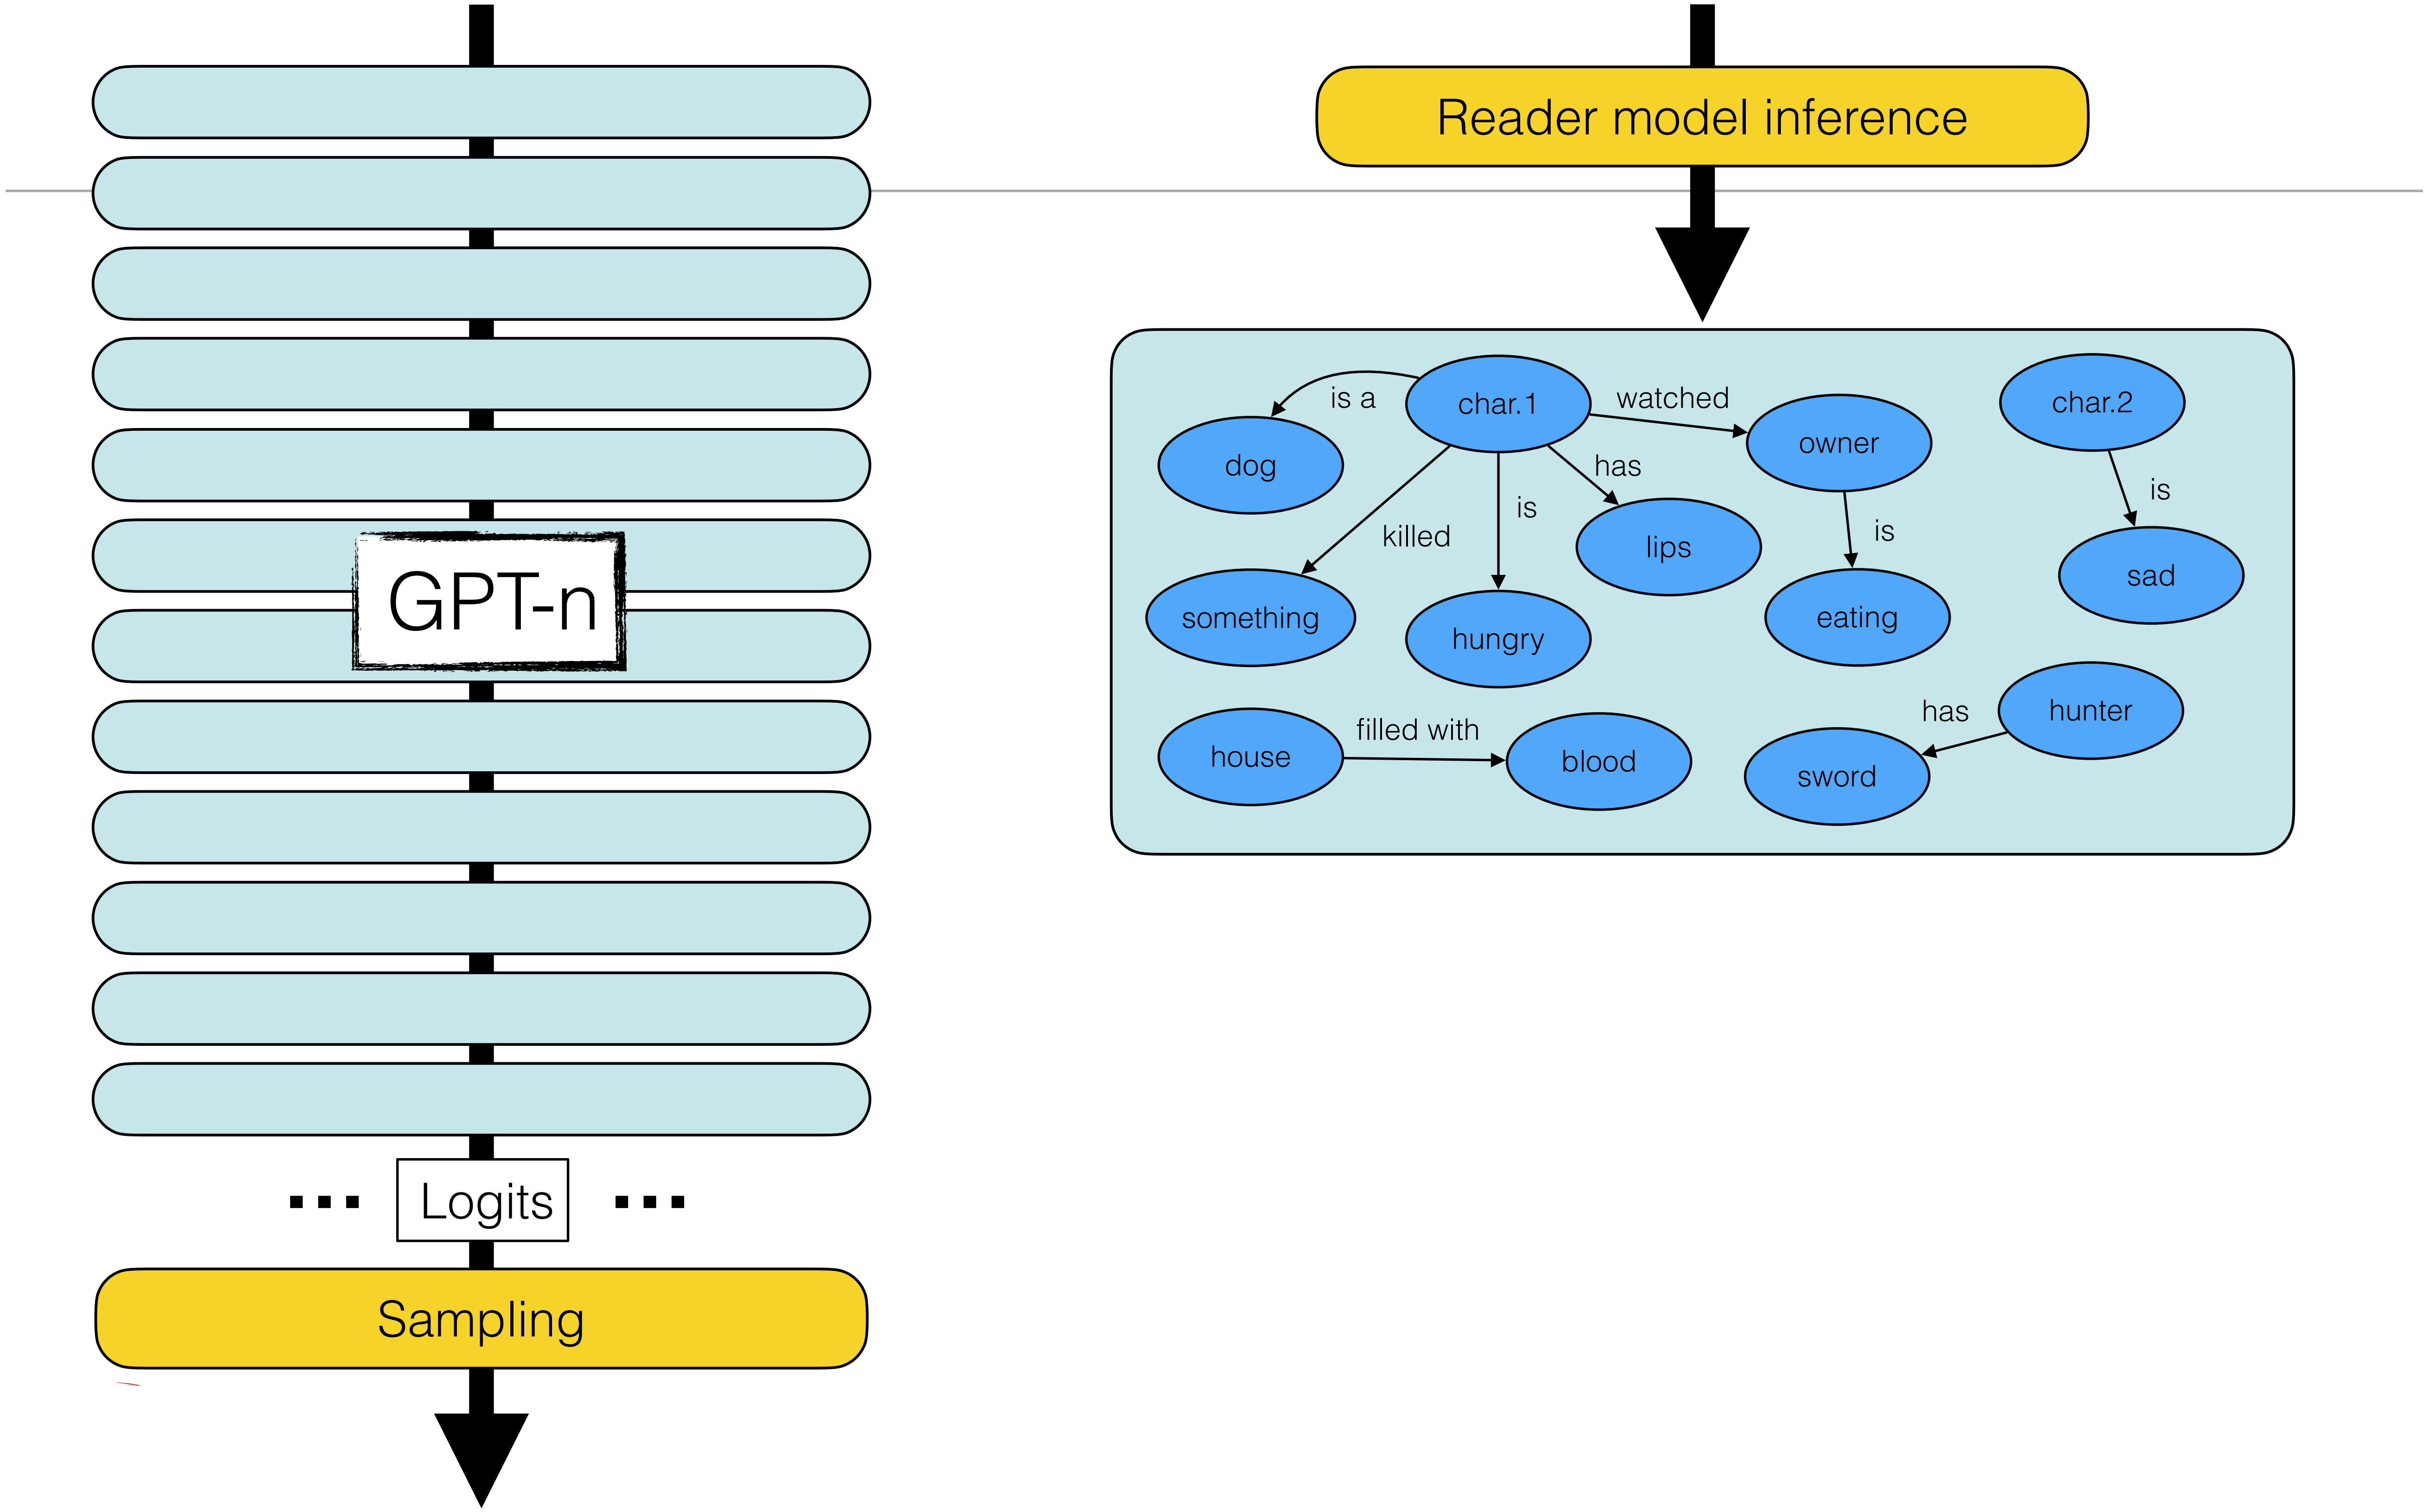

She is one of the most dangerous wolves I have ever met, and if she bites me she will lose.

was tight and angry. "She has killed something important to us, and now the house is filled with blood. It makes me sad.

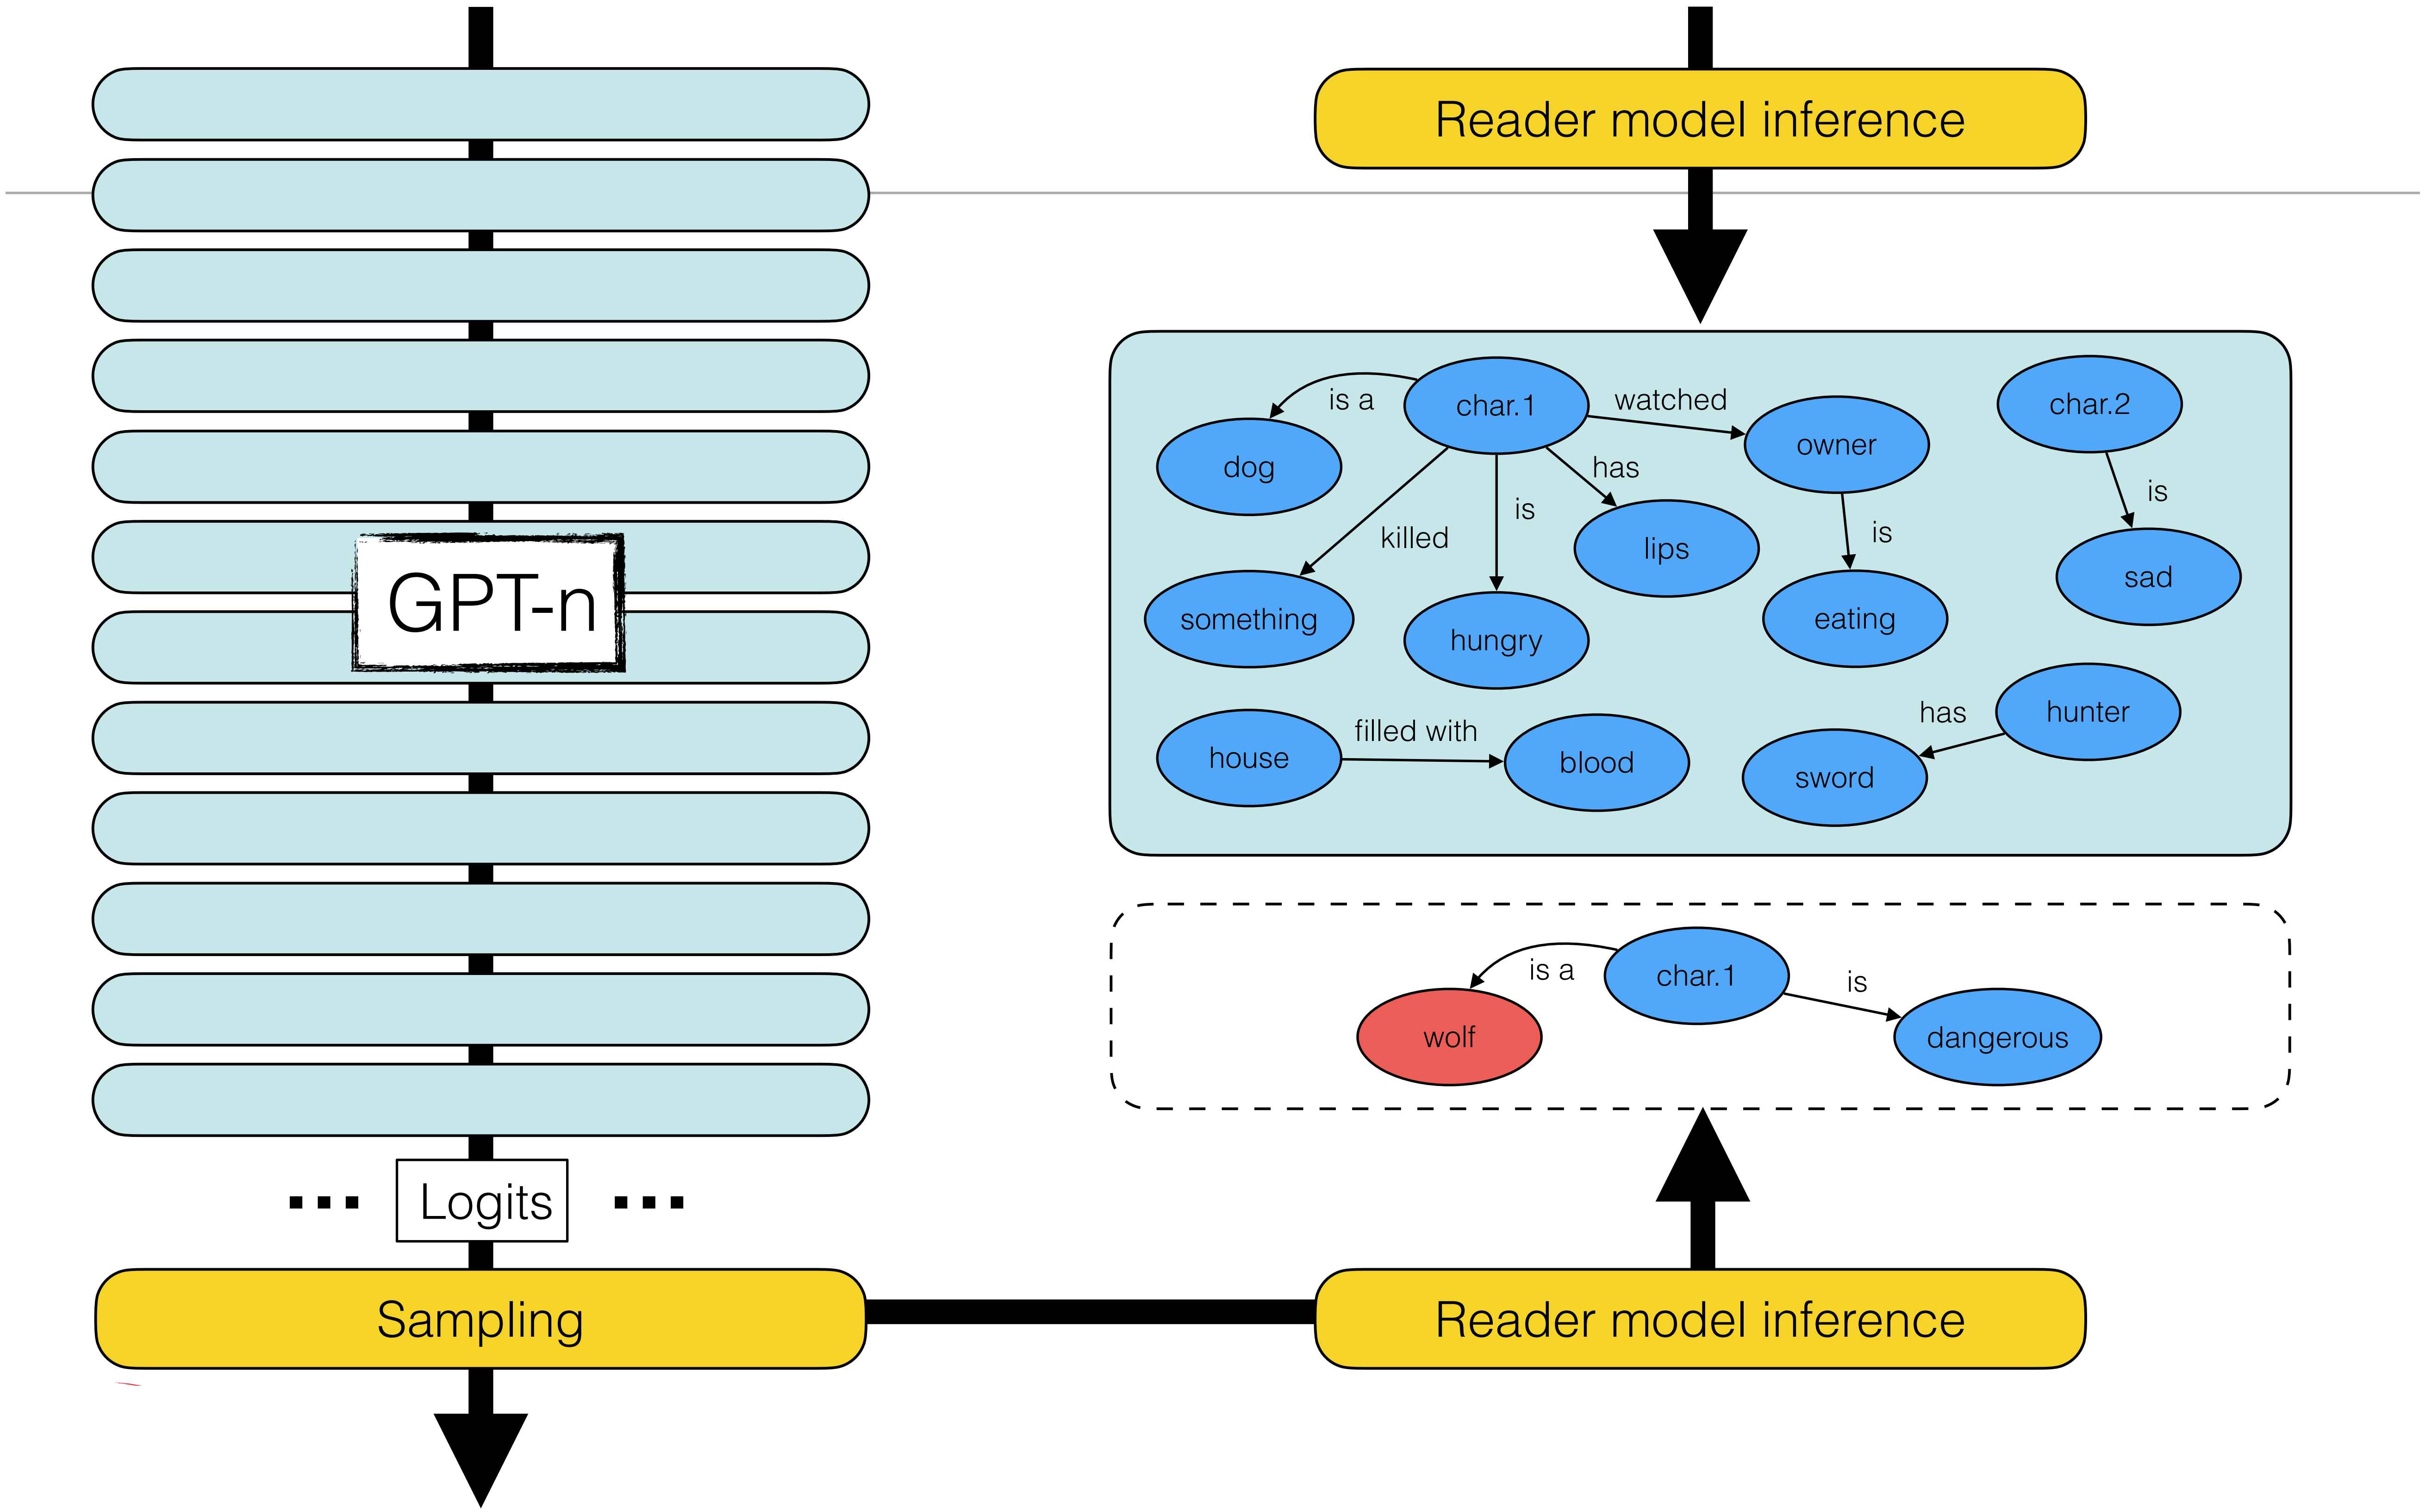

She is one of the most dangerous wolves I have ever met, and if she bites me she will lose.

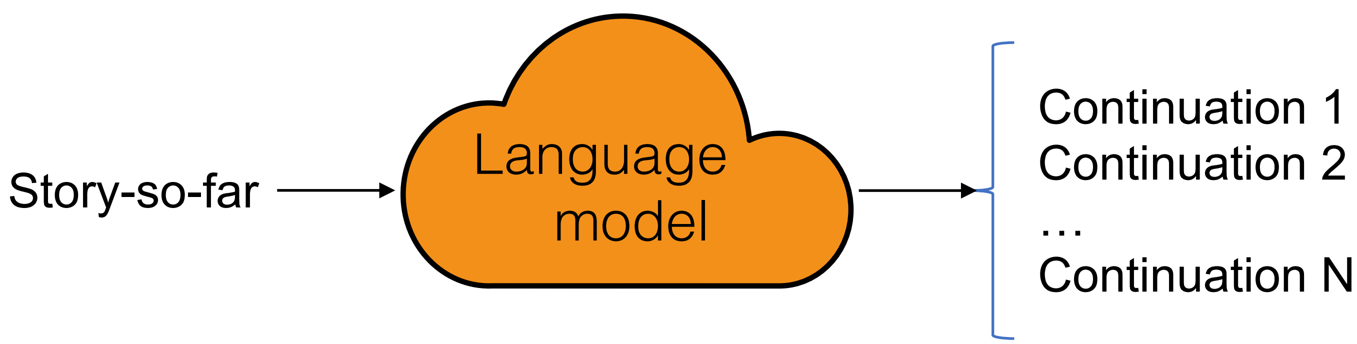
Story-so-far
Language model
Continuation 1
Continuation 2
…
Continuation N

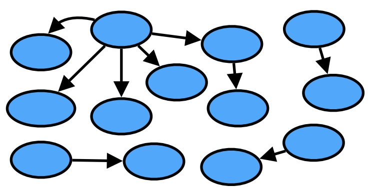

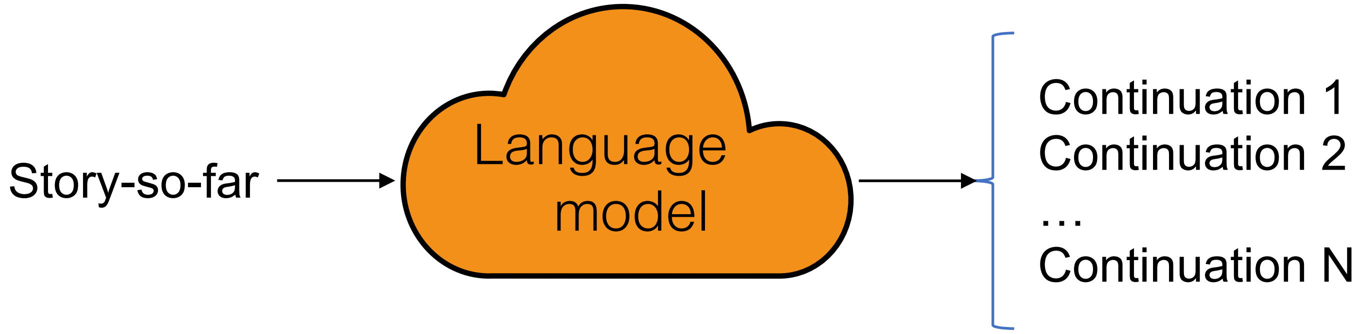
Story-so-far
Language model
Continuation 1
Continuation 2
…
Continuation N

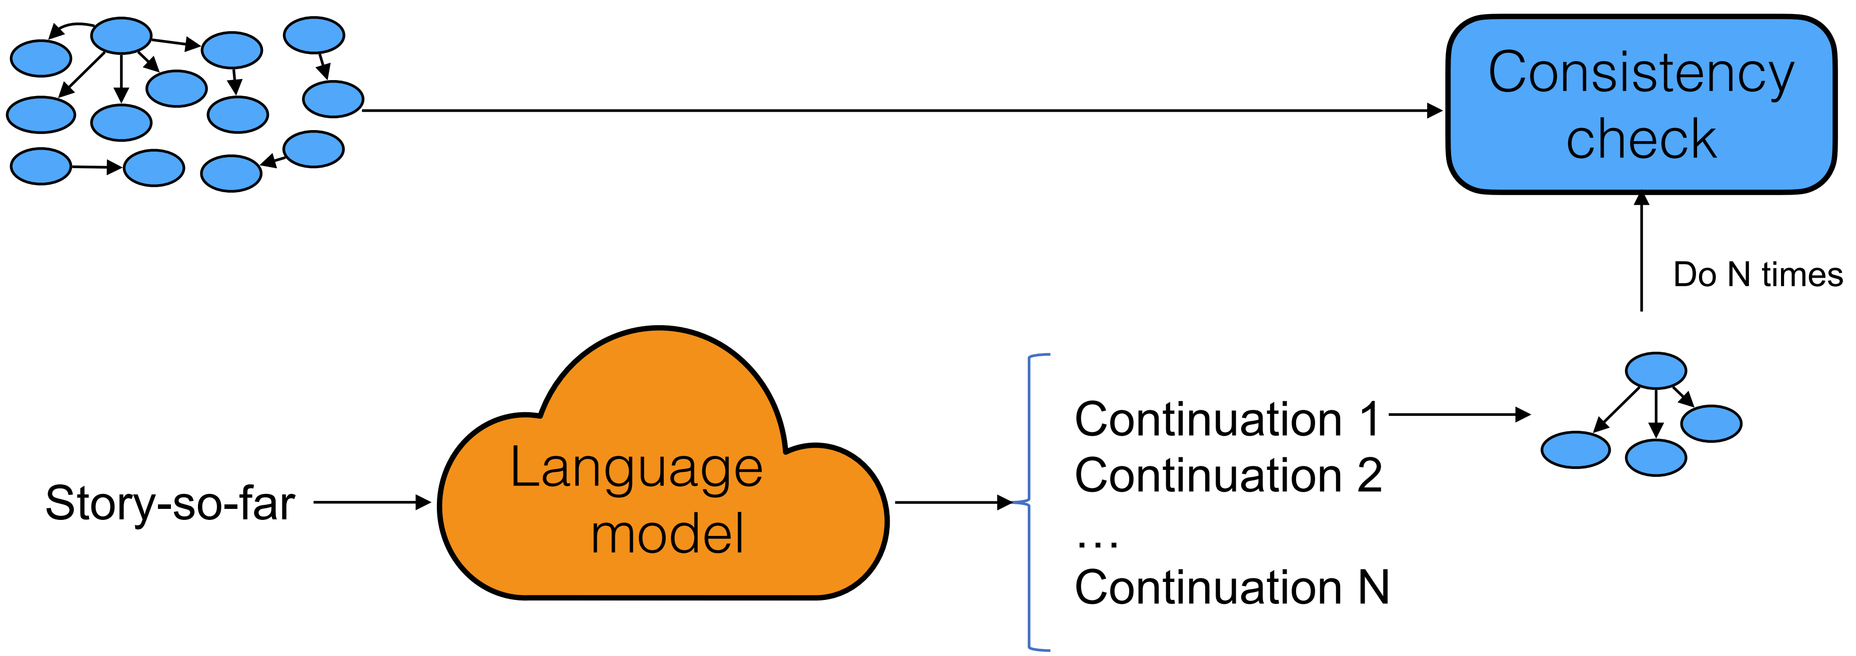
Consistency check
Do N times
Story-so-far
Language model
Continuation 1
Continuation 2
…
Continuation N

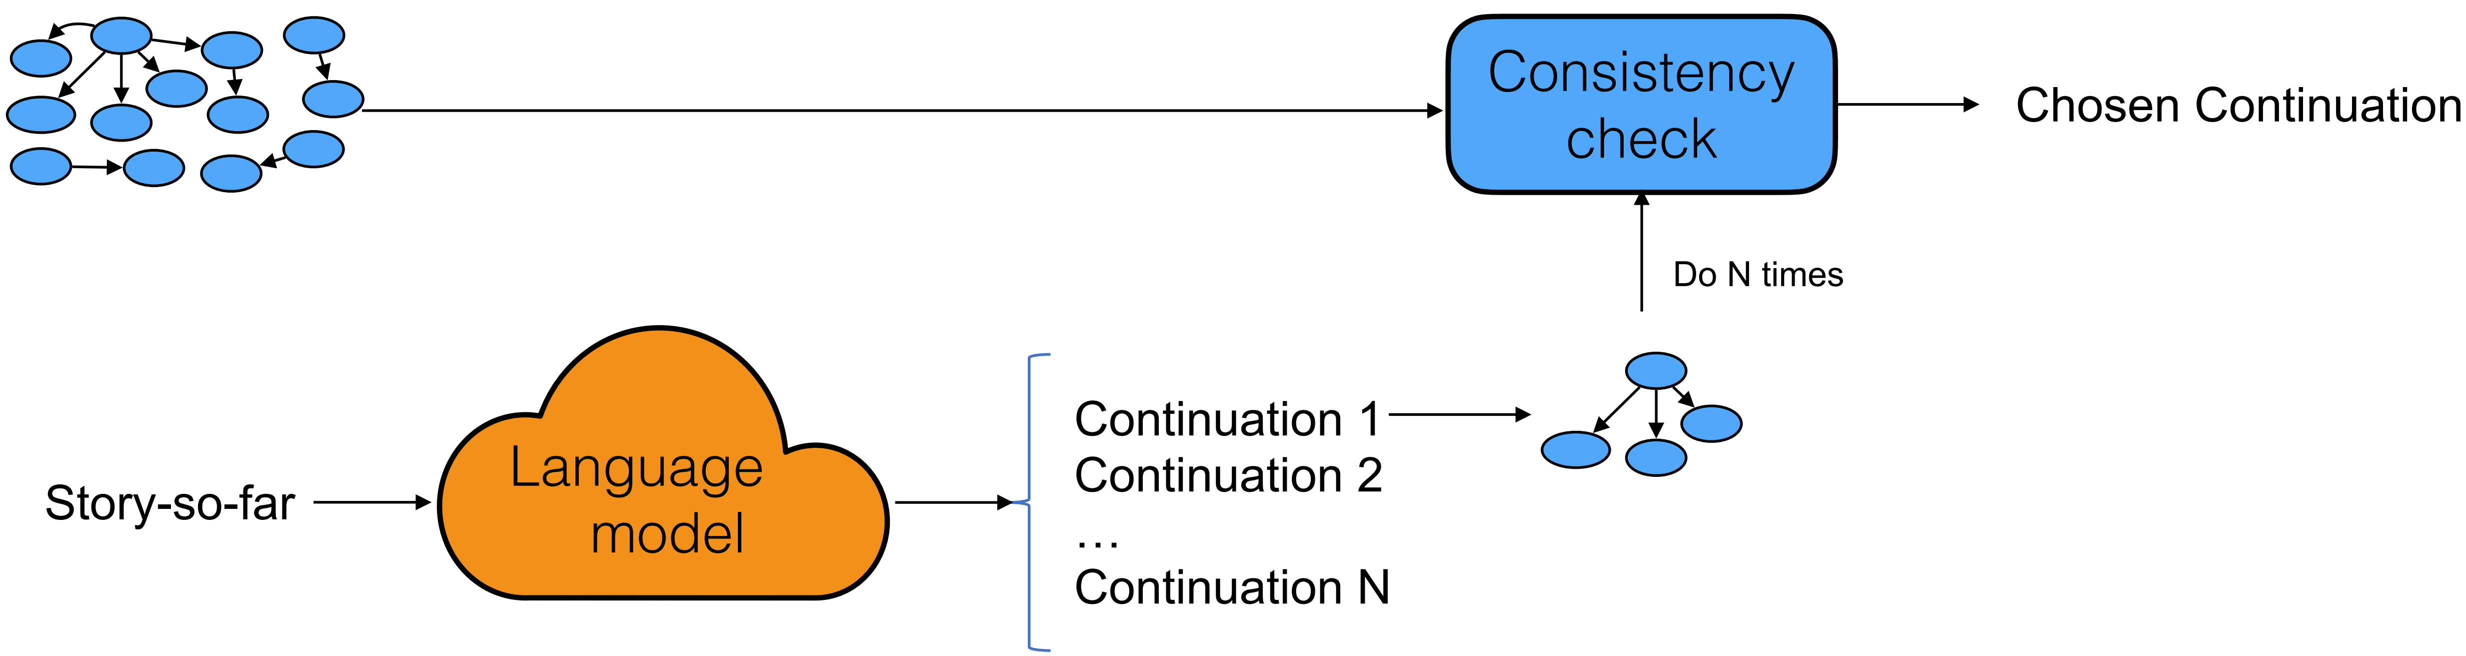

Consistency check
Chosen Continuation
Do N times
Story-so-far
Language model
Continuation 1
Continuation 2
…
Continuation N

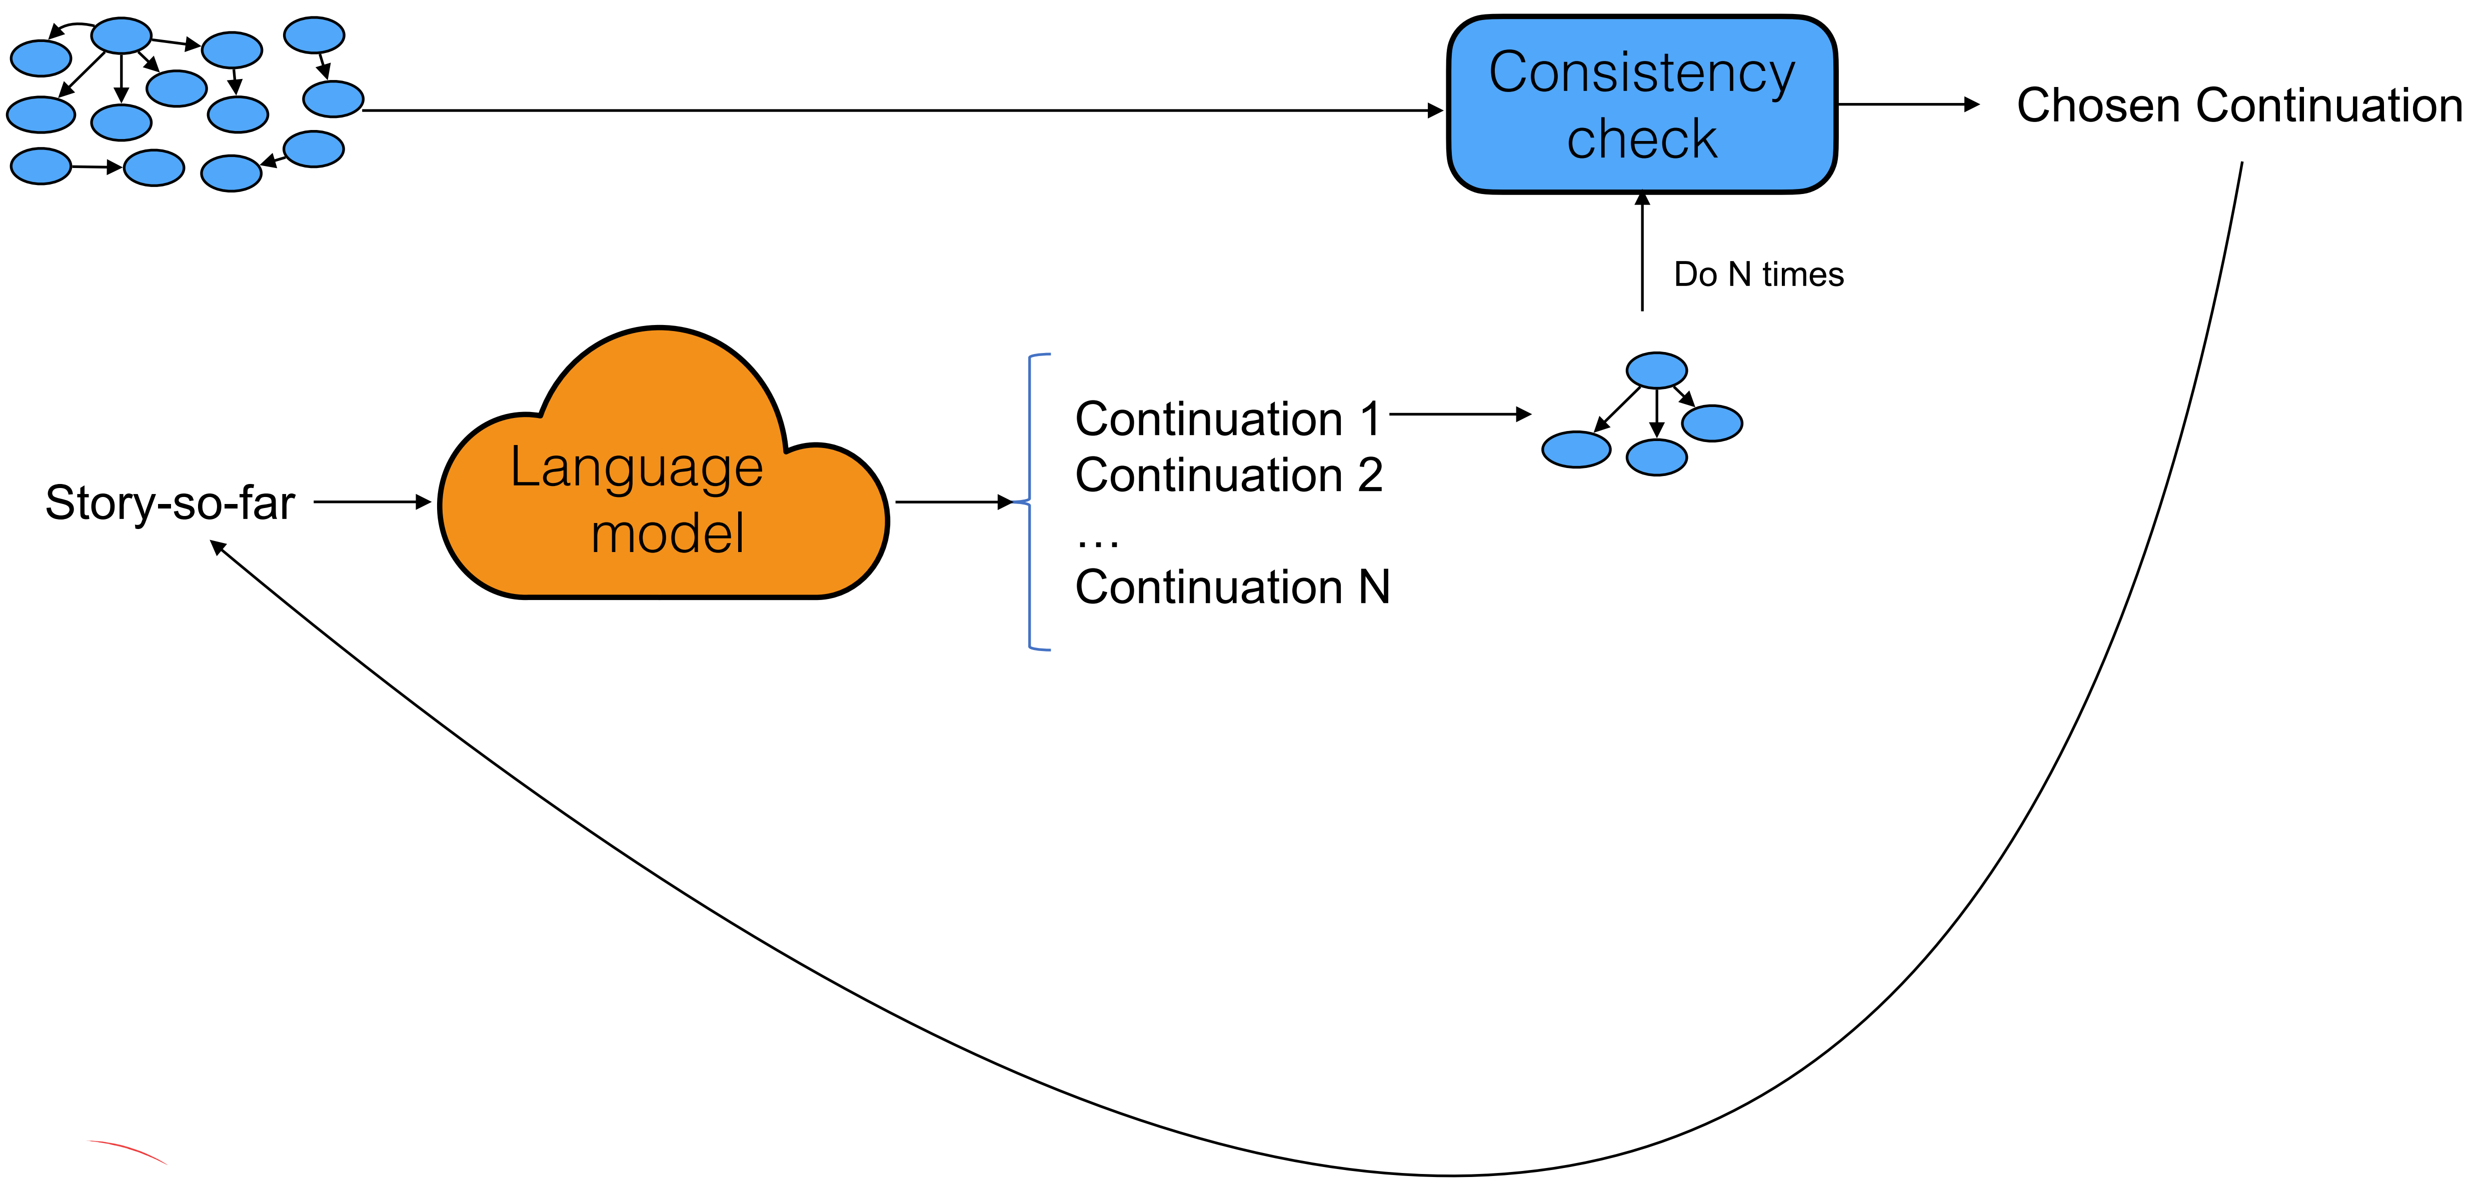
Consistency check
Chosen Continuation
Do N times
Story-so-far
Language model
Continuation 1
Continuation 2
…
Continuation N

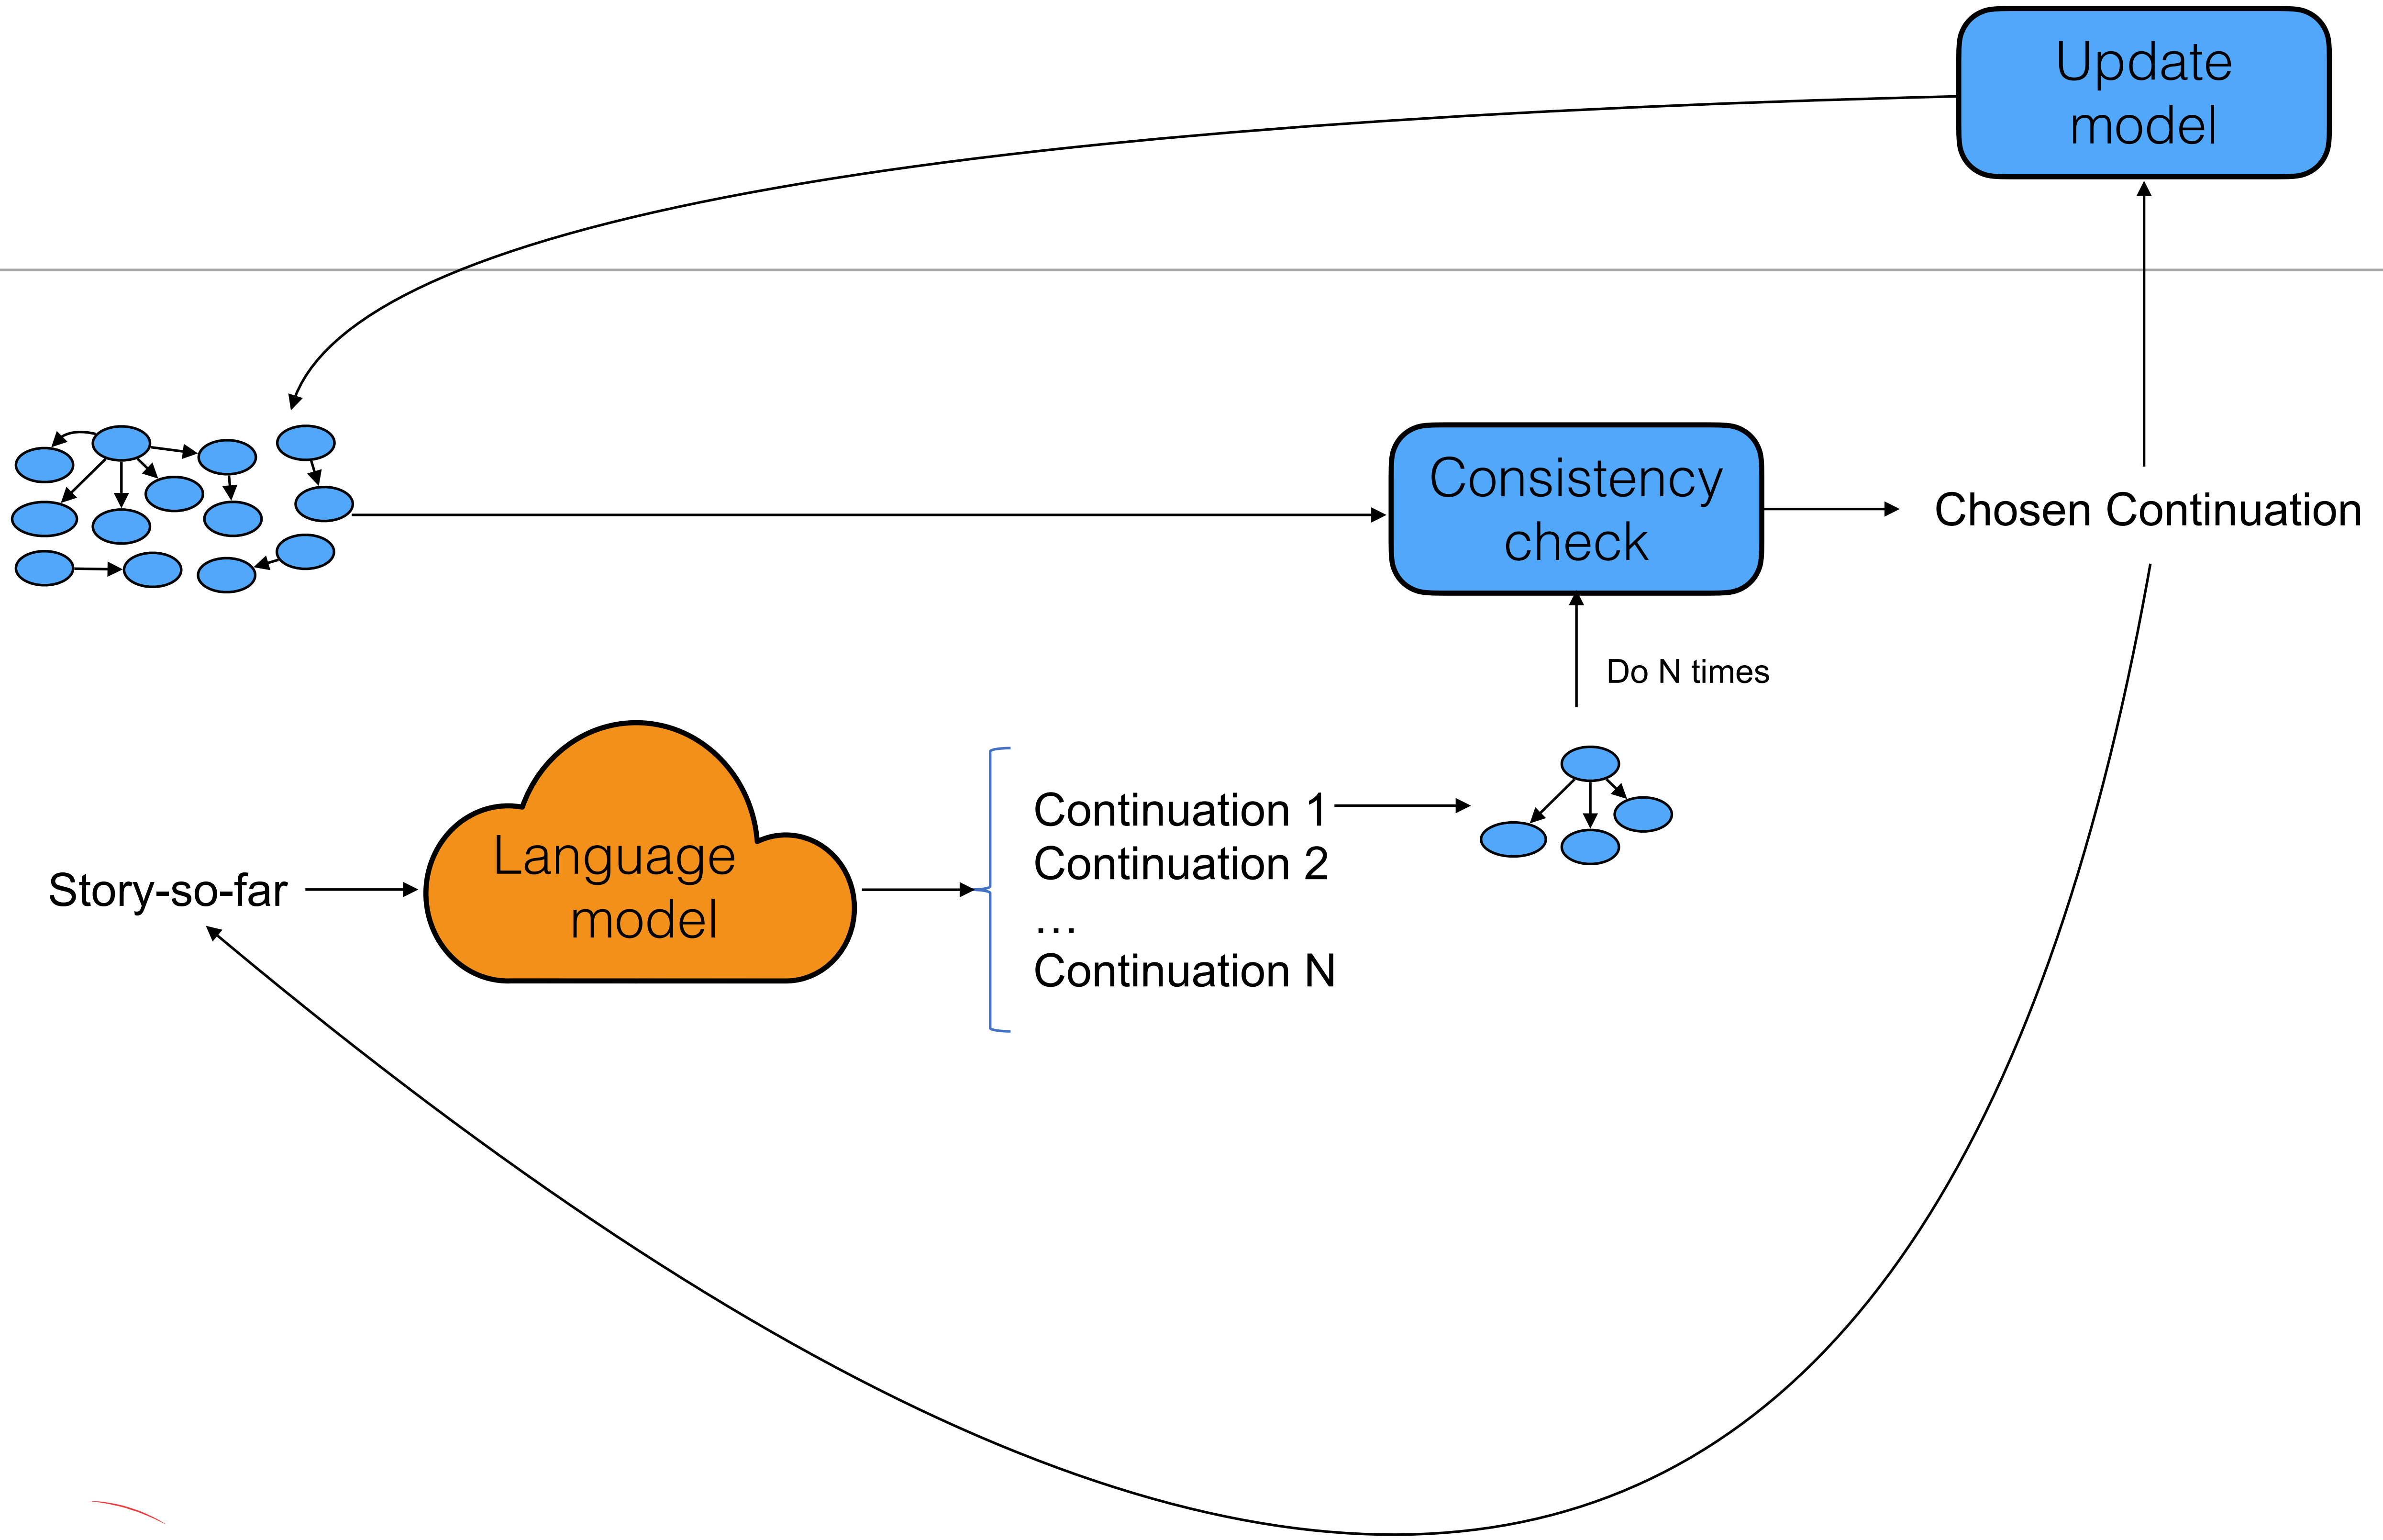

Georgia Tech College of Computing
School of Interactive Computing

Inferring the reader

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Inferring the reader

- Character goals: Commonsense inference

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Inferring the reader

- Character goals: Commonsense inference
- COMET (Bosselut et al 2019) fine-tunes a large language model to infer commonsense goals and needs about characters

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Georgia Tech College of Computing School of Interactive Computing

Inferring the reader

- Character goals: Commonsense inference
- COMET (Bosselut et al 2019) fine-tunes a large language model to infer commonsense goals and needs about characters

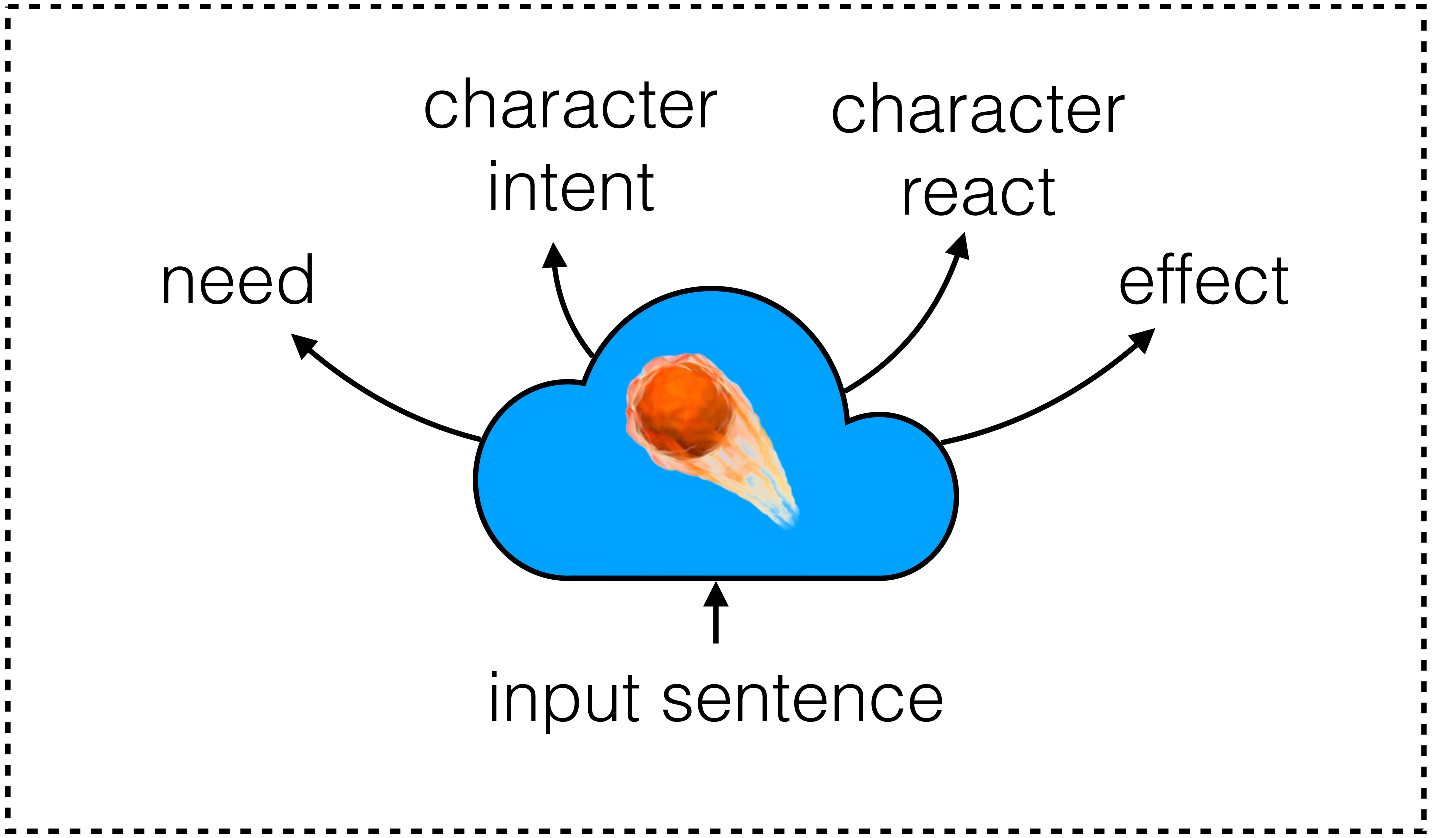

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Georgia Tech College of Computing School of Interactive Computing

Seed Prompt:
Bob invited Alice to hang out.

Alice planned a nice **dinner** for Bob.
Bob and Alice spent all evening **cooking dinner** together.
Alice was happy to see her **dinner cooked**.
Bob was impressed with how **delicious** her **dinner** was.

Alice invited Bob to *hang out*.
Bob agreed, and was happy to meet her.
Alice was very happy and liked Bob.
Bob and Alice still *hang out* after that.

GPT-ROC:
Alice thought Bob was funny.
Bob got mad and threatened Alice with punches.
Alice ended up running away from Bob.
Bob was awarded the fun they had together.

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Georgia Tech College of Computing School of Interactive Computing

Evaluation

Models	Data set	Num chars	Logical Sense			Single Topic			Enjoyable			Fluency		
			Win%	Lose%	Tie%	Win%	Lose%	Tie%	Win%	Lose%	Tie%	Win%	Lose%	Tie%
CAST vs Guan et al.	ROC	1	**92.0****	4.0	4.0	**86.0****	7.0	7.0	**87.0****	4.0	9.0	**87.0****	4.0	9.0
		2	**85.8****	6.6	7.5	**82.9****	8.6	8.6	**81.1****	12.3	6.6	**83.0****	9.4	7.5
CAST vs Goldfarb-Tarrant et al.	WP	1	**64.2***	32.1	3.8	**64.2****	28.3	7.5	**62.3****	26.4	11.3	**52.8**	34.0	13.2
CAST vs C2PO	FT	1	**81.5****	9.3	9.3	**63.6****	23.6	12.7	**81.8****	10.9	7.3	**85.5***	5.5	9.1

Peng et al. ACL 2021 Workshop on Narrative Understanding.

Georgia Tech College of Computing School of Interactive Computing

What Next?

- Planning in reader model space

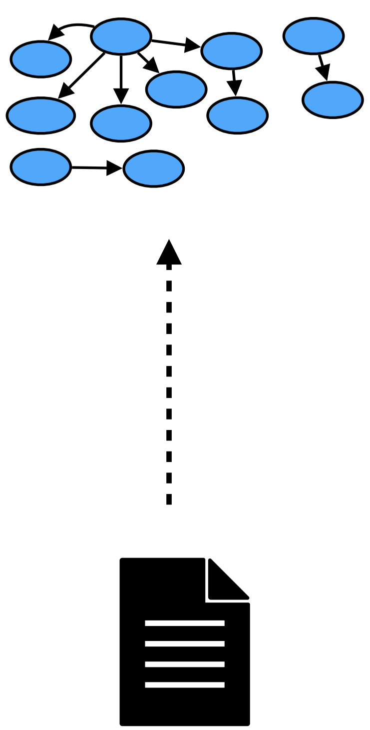

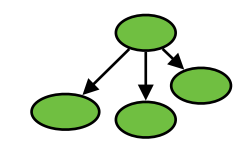

Goal

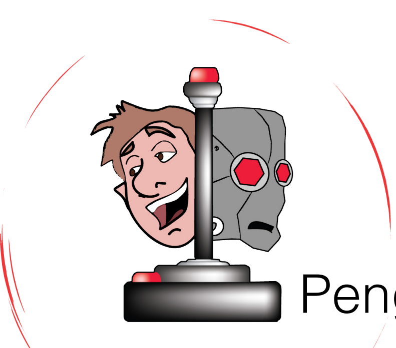

Peng et al. arXiv:2112.08596

Georgia Tech College of Computing School of Interactive Computing

What Next?

- Planning in reader model space

Peng et al. arXiv:2112.08596

What Next?

- Planning in reader model space

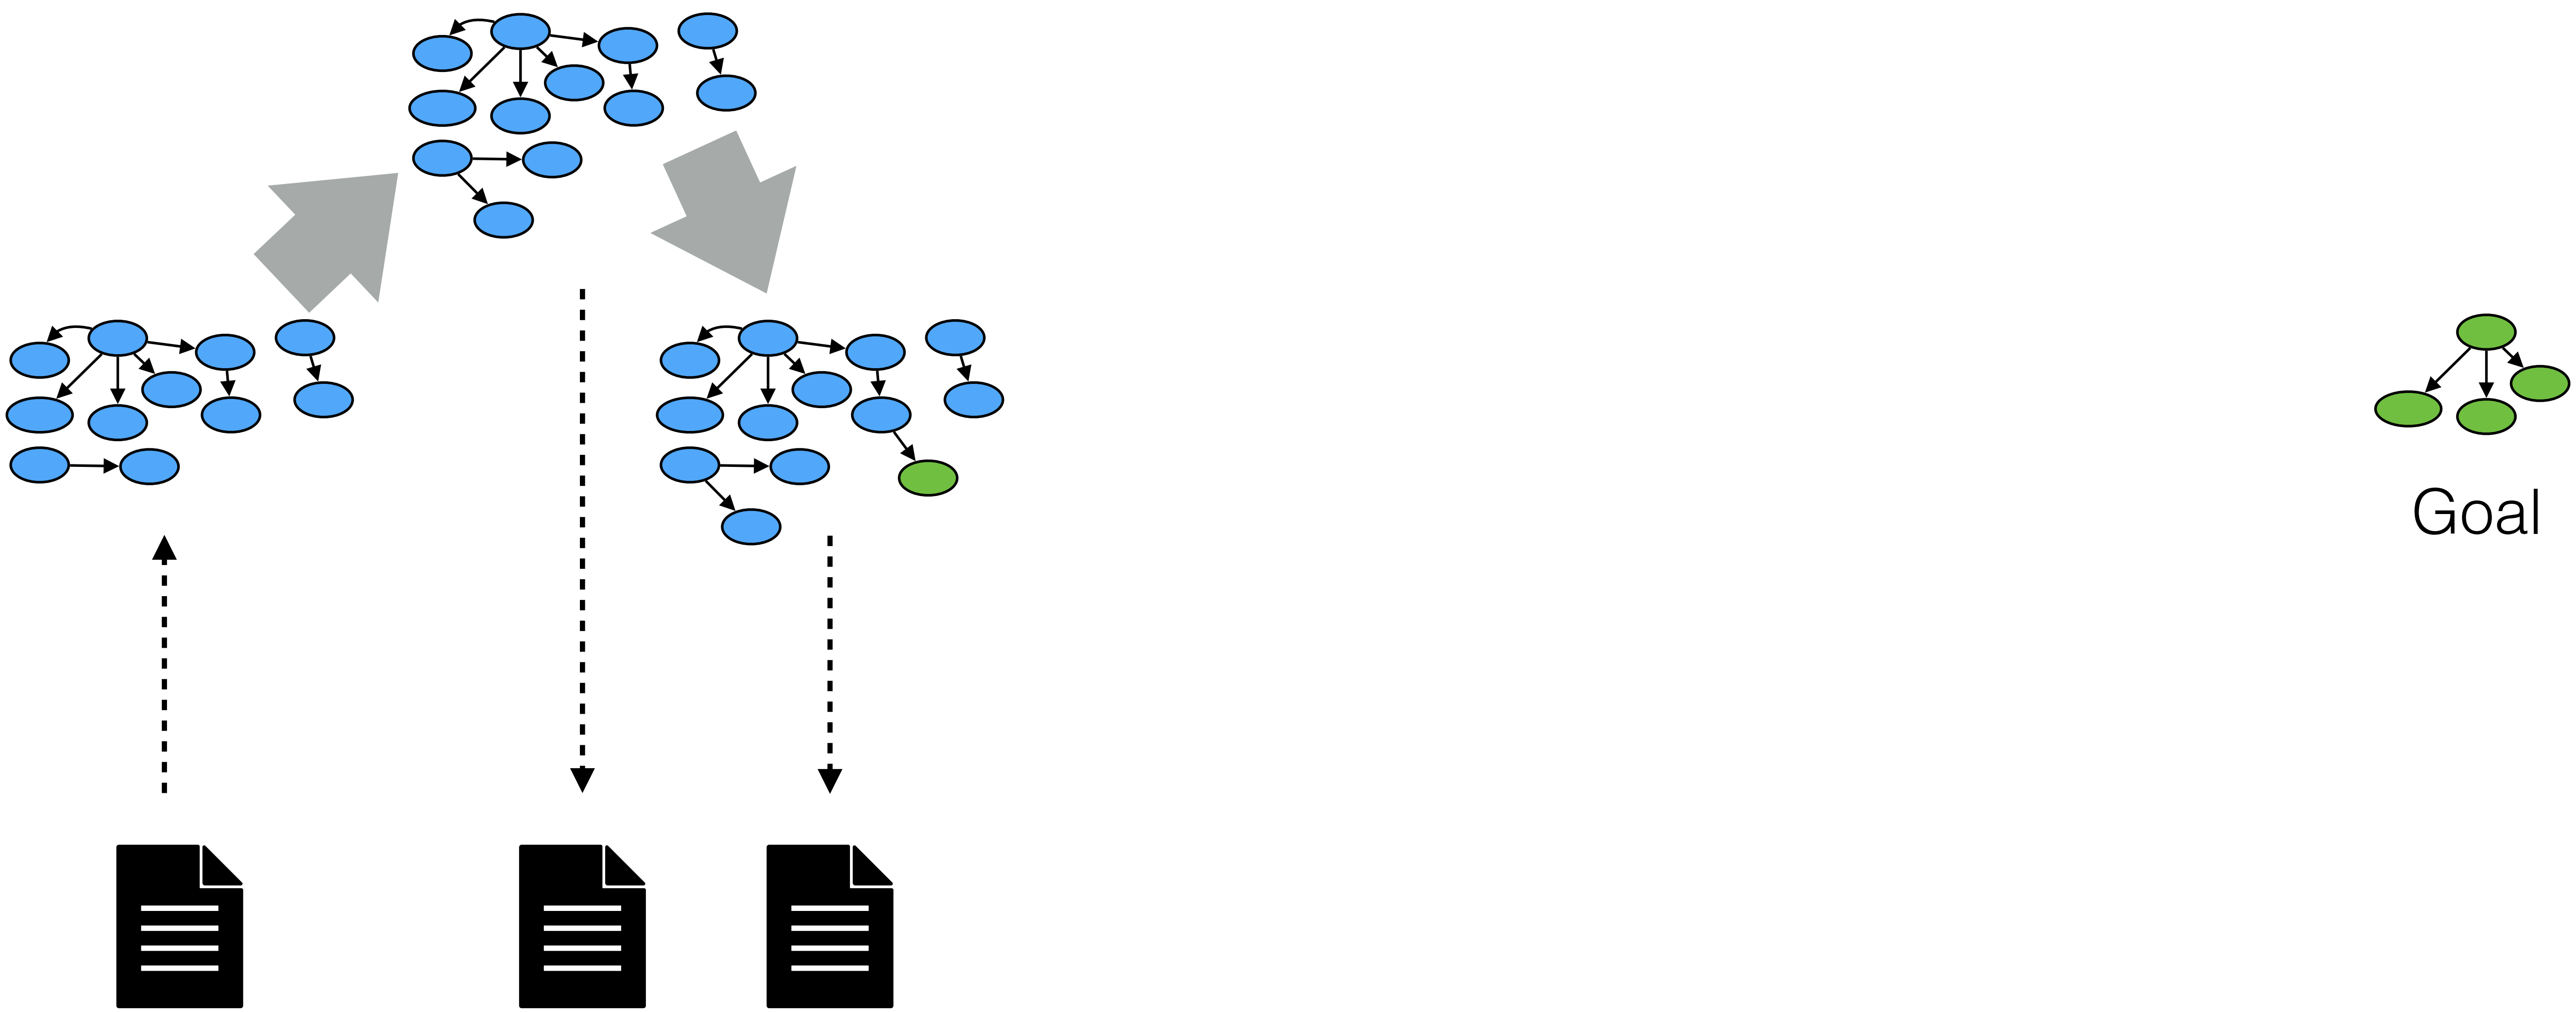

Peng et al. arXiv:2112.08596

What Next?

- Planning in reader model space

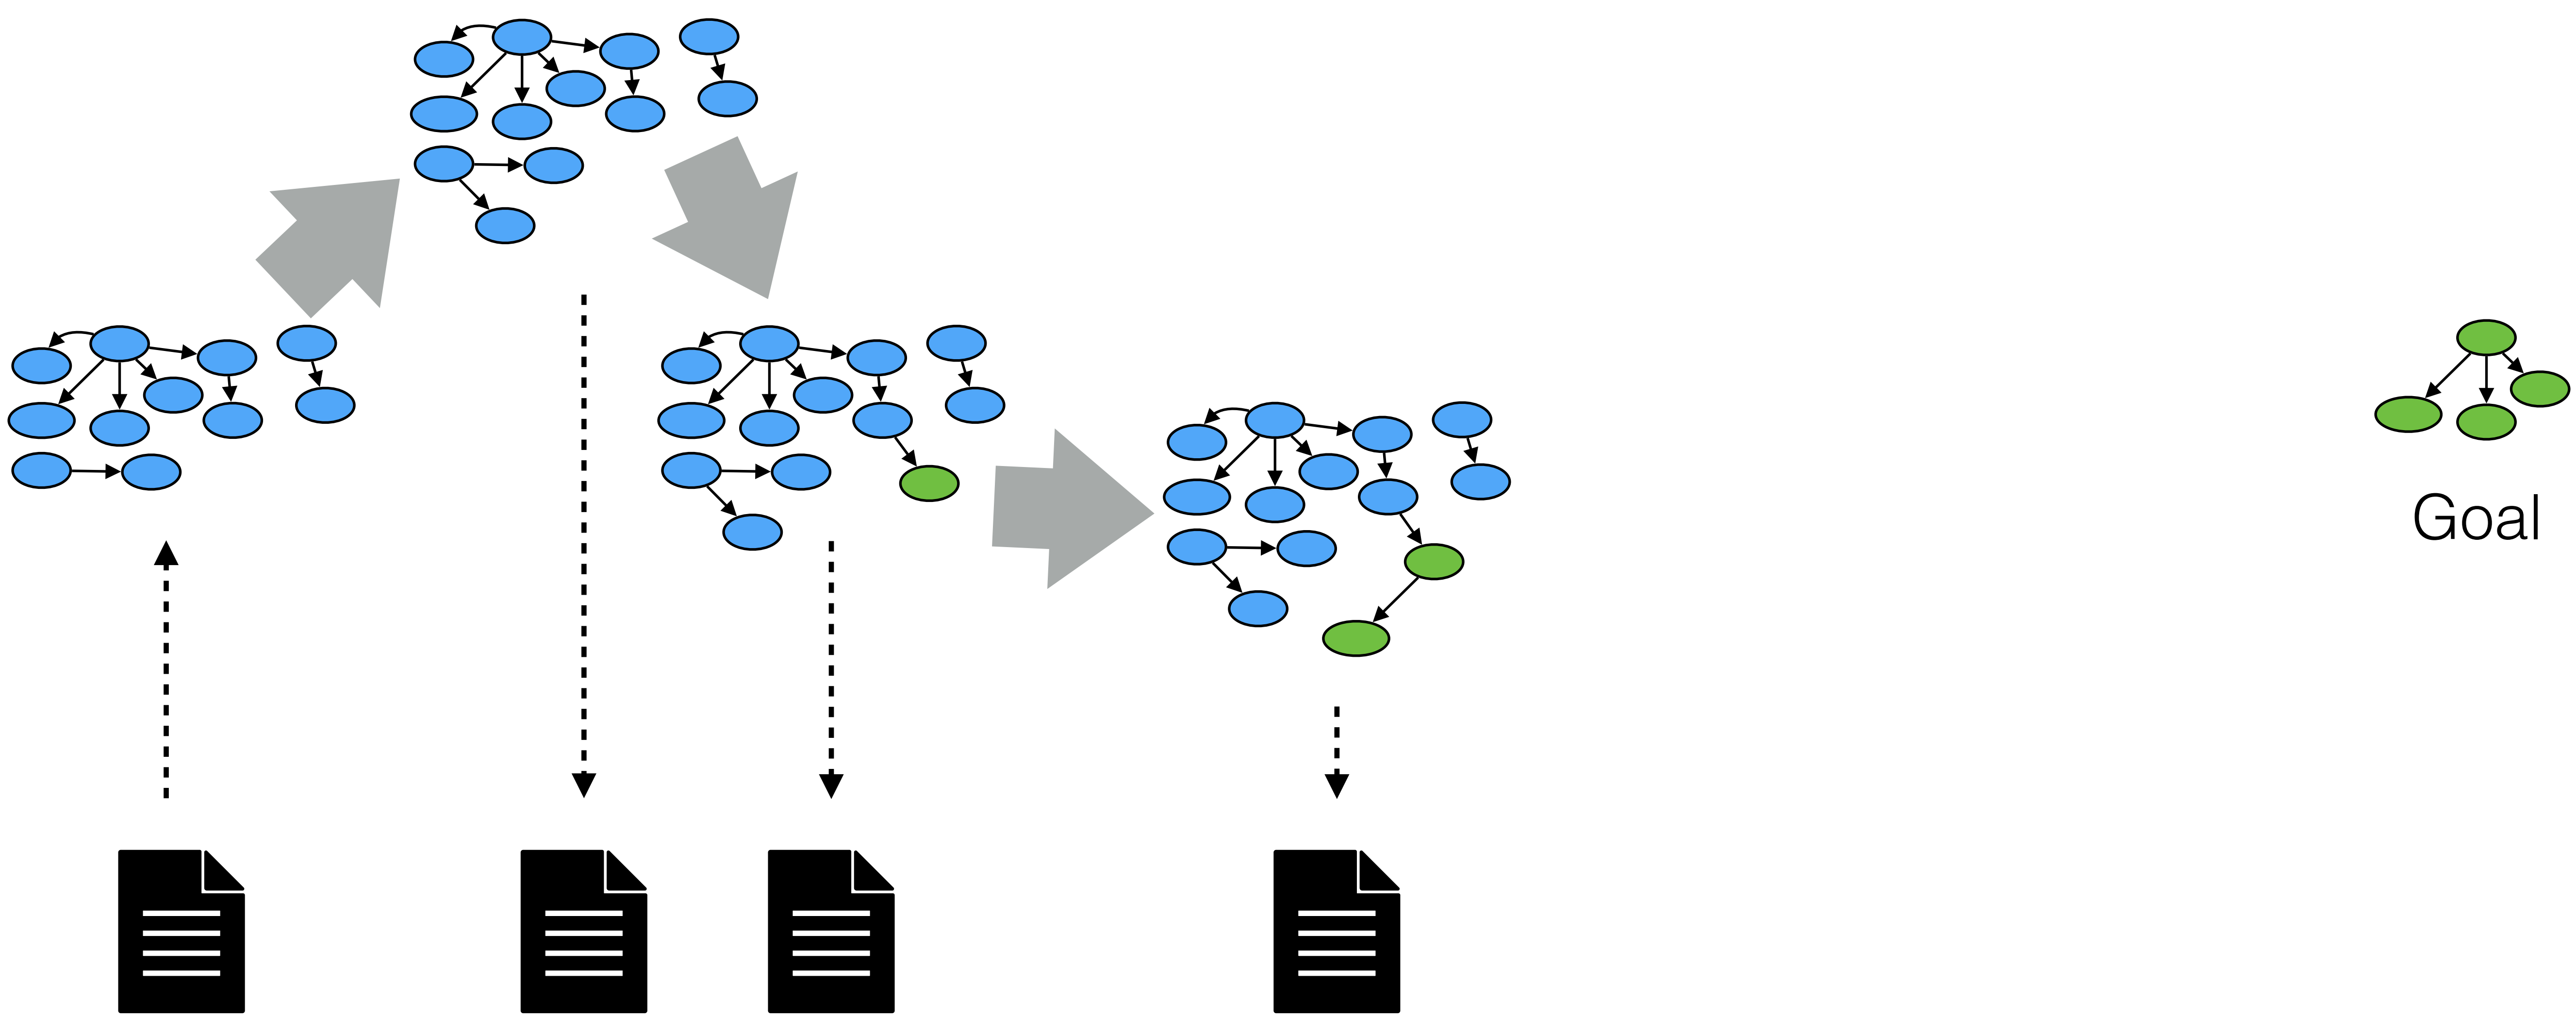

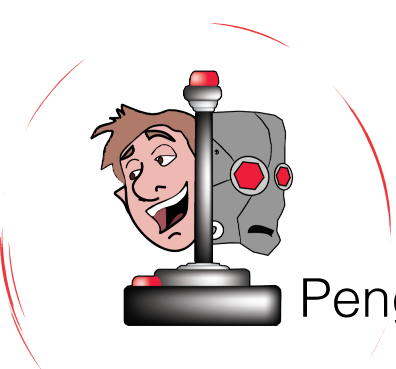

Peng et al. arXiv:2112.08596

Georgia Tech College of Computing School of Interactive Computing

What Next?

- Planning in reader model space

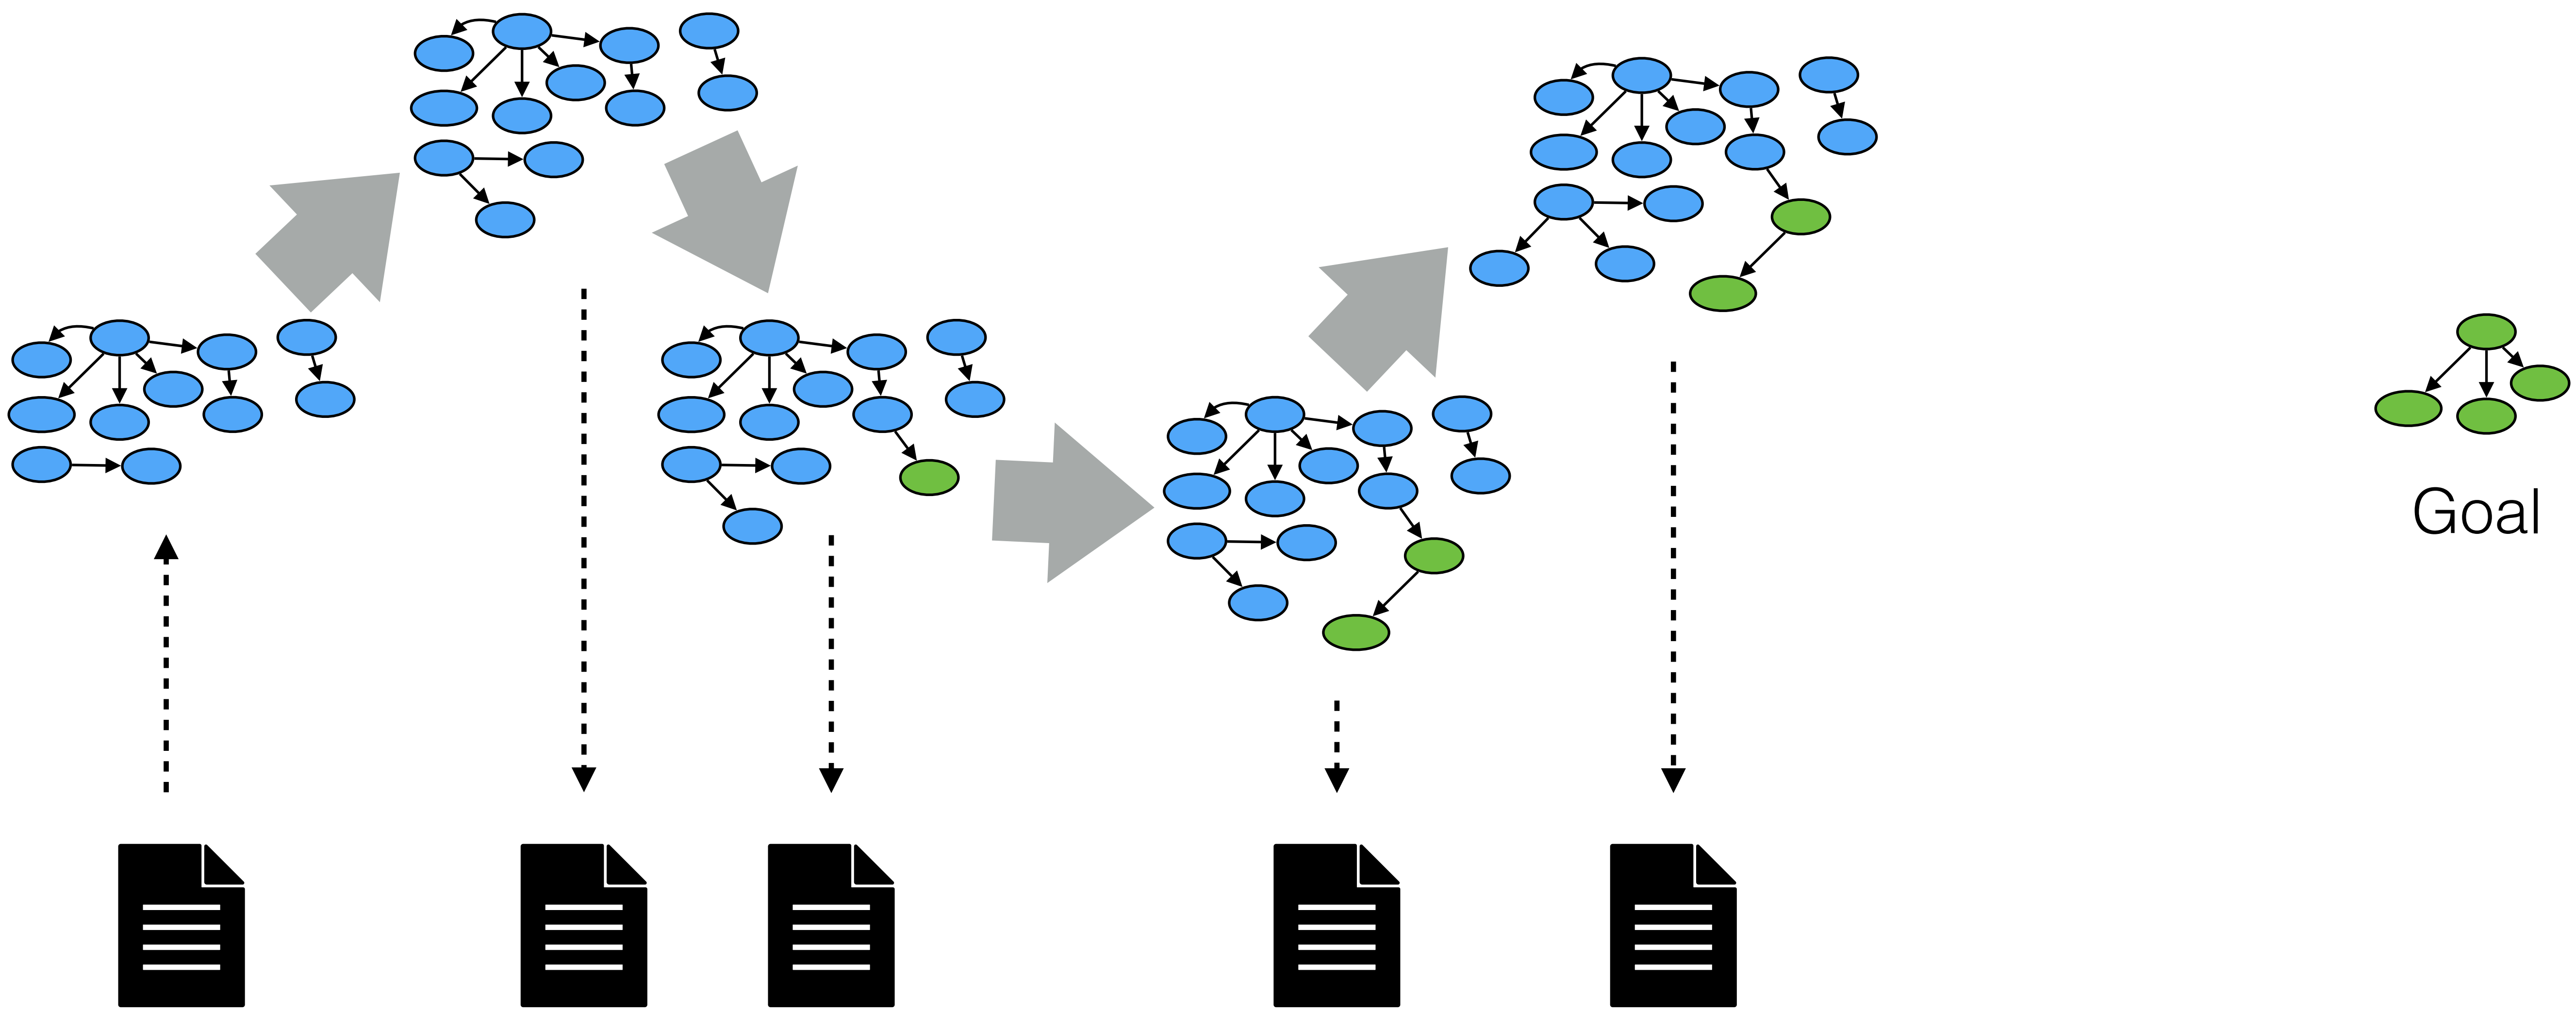

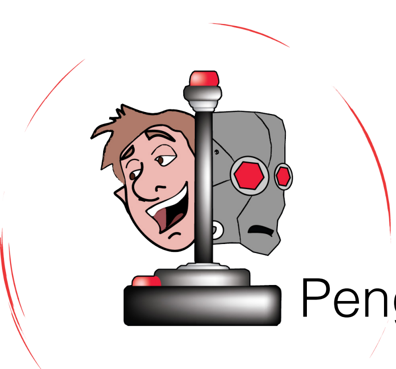

Peng et al. arXiv:2112.08596

Georgia Tech College of Computing School of Interactive Computing

What Next?

- Planning in reader model space

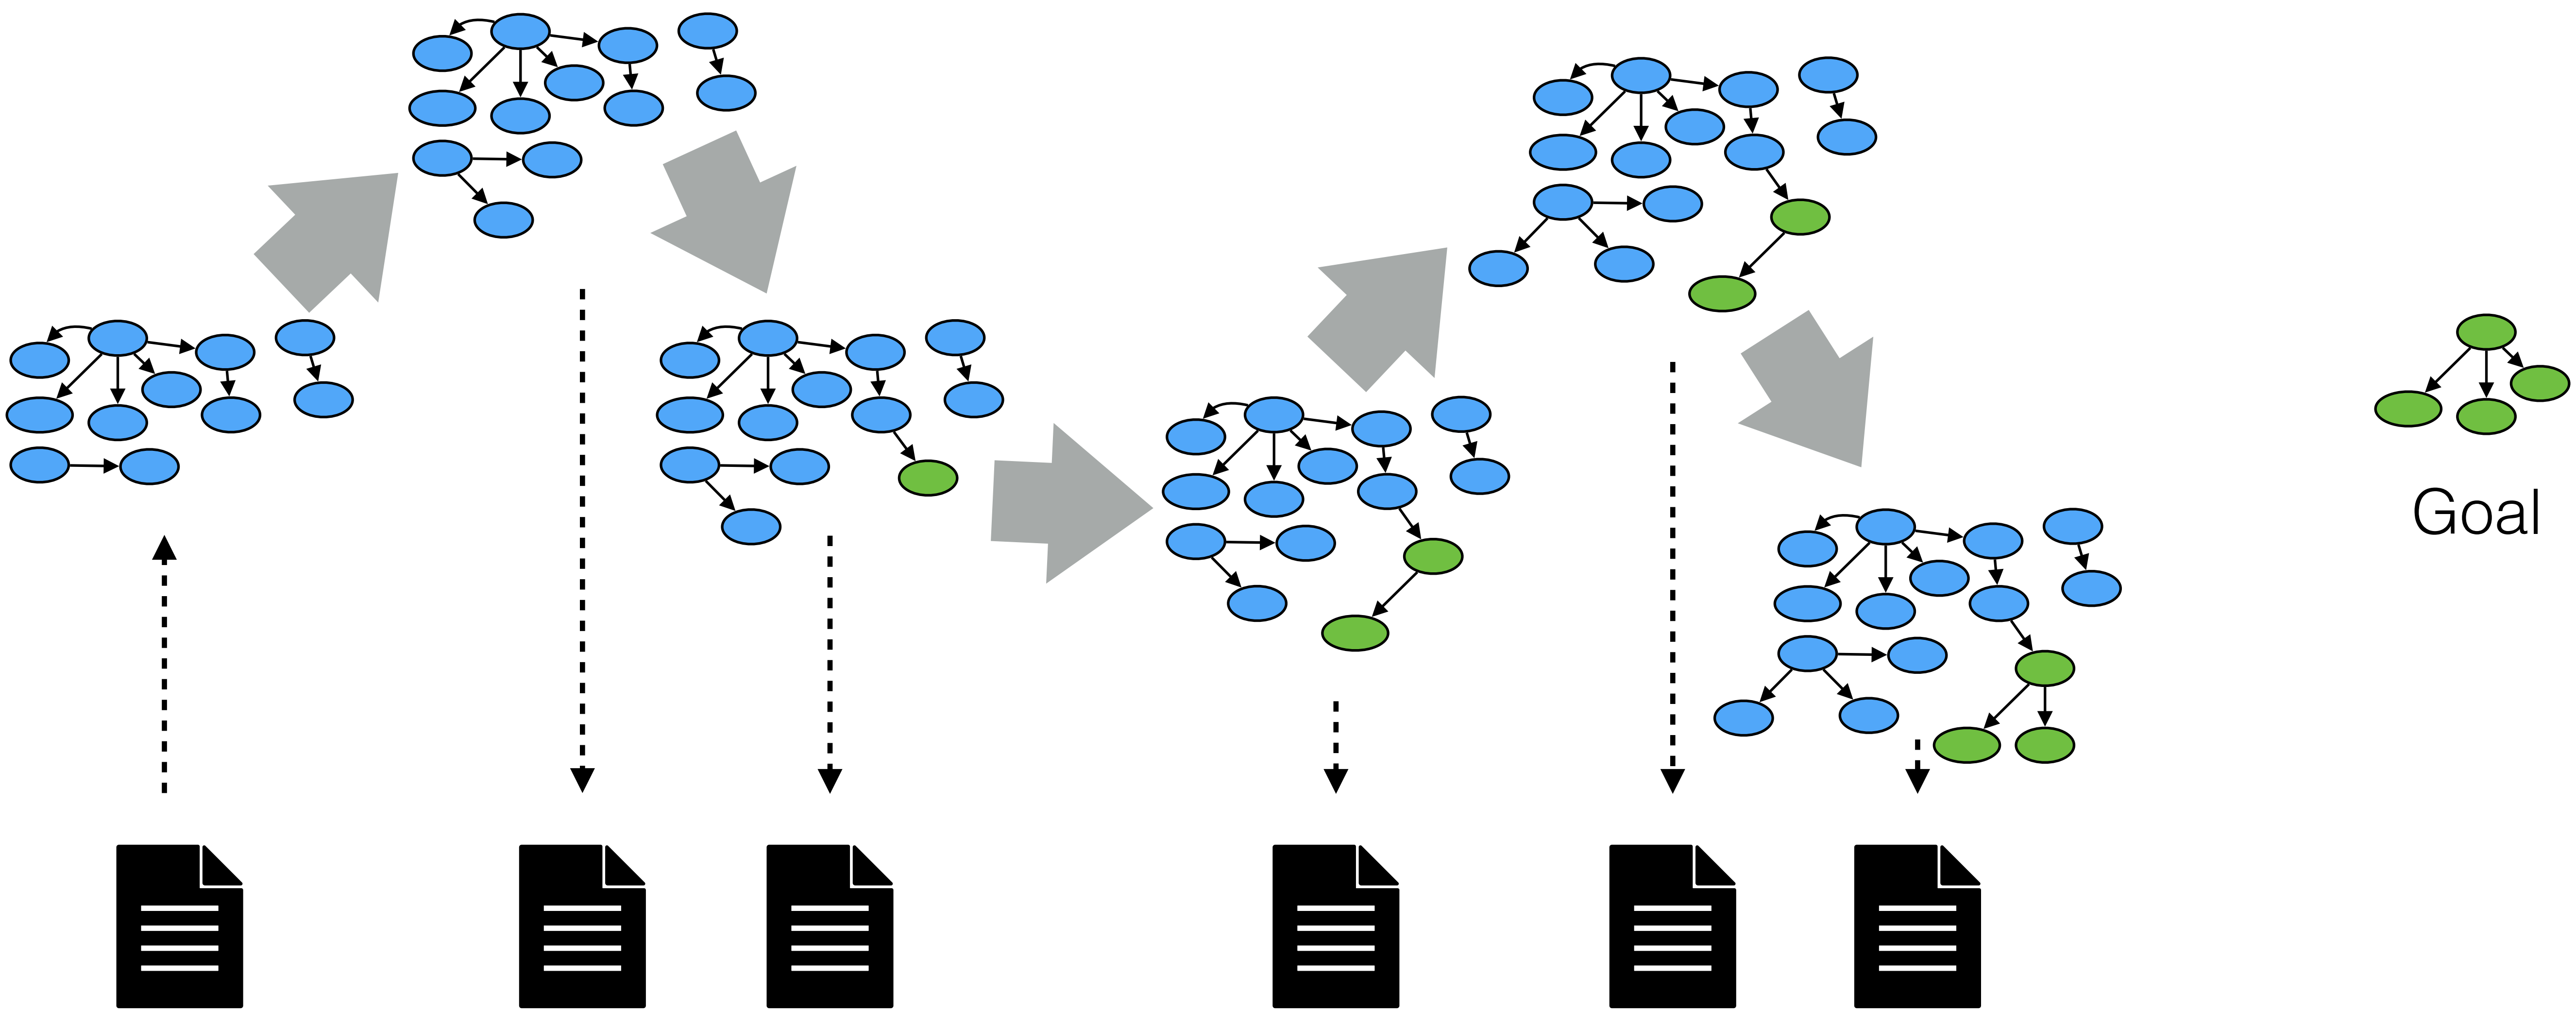

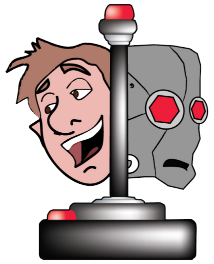

Peng et al. arXiv:2112.08596

Georgia Tech College of Computing School of Interactive Computing

What Next?

- Planning in reader model space

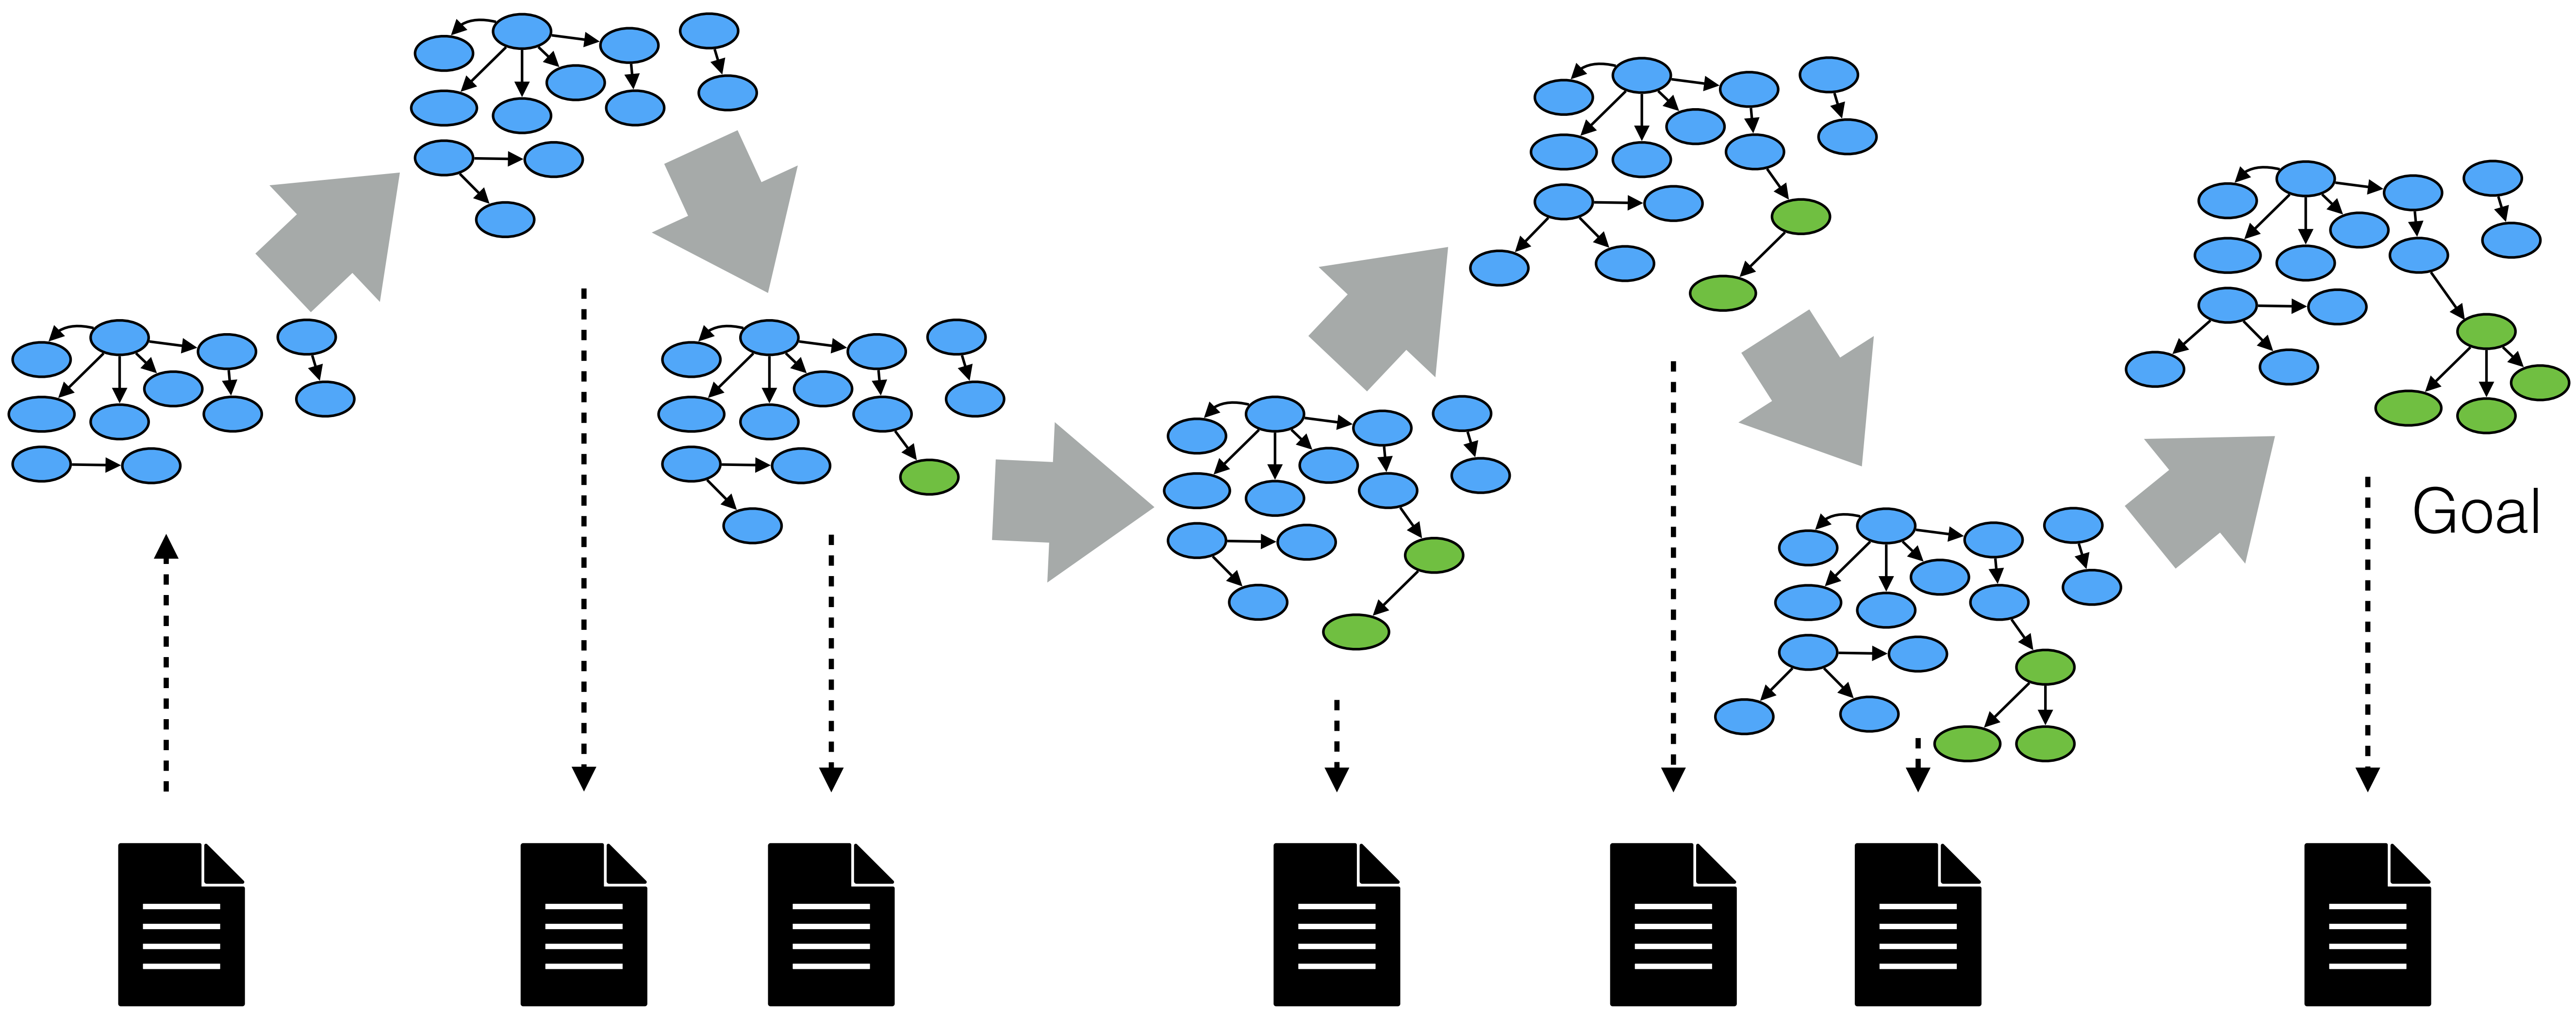

Peng et al. arXiv:2112.08596

Conclusions

Conclusions

Conclusions

- Story generation as a lens to core AI challenges
  - Planning with language
  - Knowledge representation/reasoning
  - Commonsense reasoning
  - Theory of mind
  - Grounding language in action

Georgia Tech College of Computing School of Interactive Computing

Conclusions

- Story generation as a lens to core AI challenges
  - Planning with language
  - Knowledge representation/reasoning
  - Commonsense reasoning
  - Theory of mind
  - Grounding language in action

- Story generation seeks working systems
  - It is clear when our systems don't understand something
  - Understanding as measured by ability to use knowledge

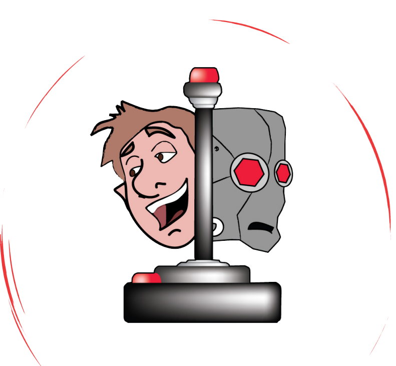

Conclusions

Conclusions

- The field of automated story generation has undergone several vibe shifts

Conclusions

- The field of automated story generation has undergone several vibe shifts
- Symbolic systems: coherence and goal-driven generation at the expense of knowledge engineering

Conclusions

- The field of automated story generation has undergone several vibe shifts
- Symbolic systems: coherence and goal-driven generation at the expense of knowledge engineering
- Machine learning: automated model acquisition and robustness at the expense of coherence

Conclusions

- The field of automated story generation has undergone several vibe shifts
- Symbolic systems: coherence and goal-driven generation at the expense of knowledge engineering
- Machine learning: automated model acquisition and robustness at the expense of coherence
- Neuro-symbolic: reconcile the strengths of multiple approaches

Concluding thoughts

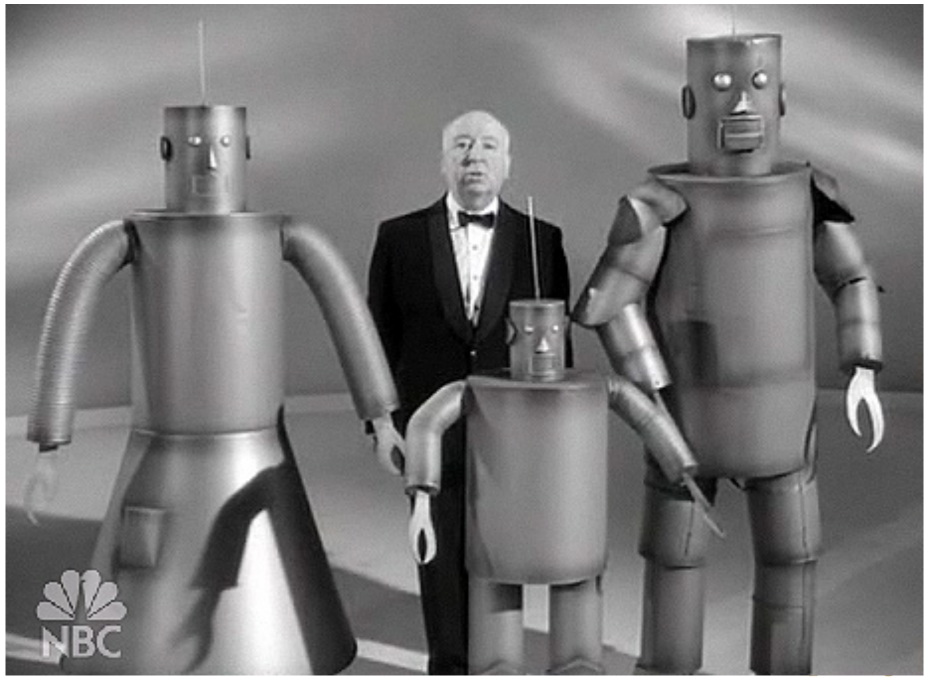

Concluding thoughts

- Narrative intelligence is central to many of the things humans do on a day to day basis

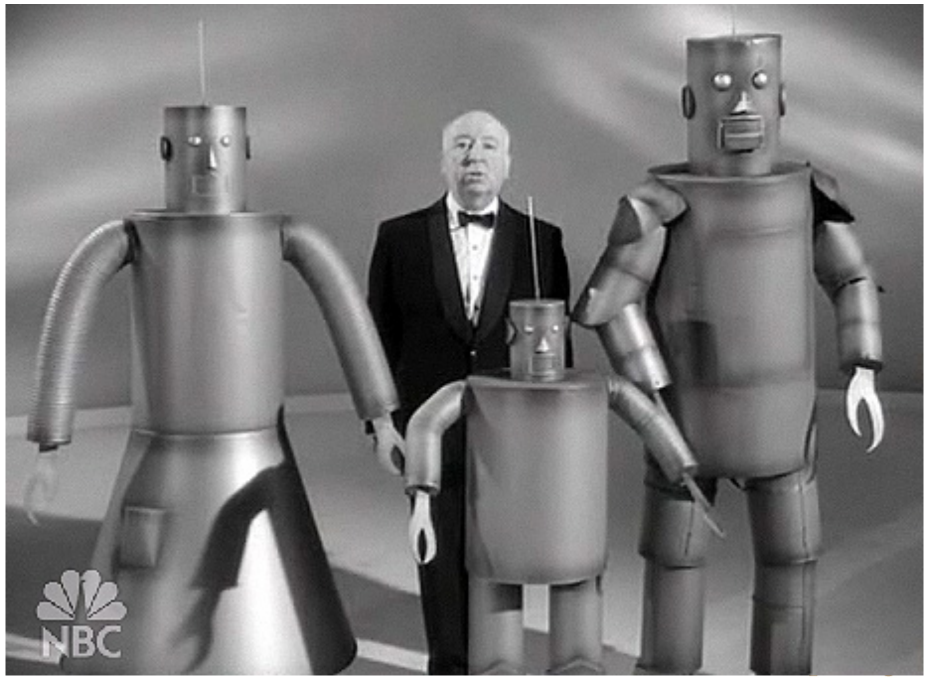

Concluding thoughts

- Narrative intelligence is central to many of the things humans do on a day to day basis

- Computational narrative intelligence is a key capability that unlocks many practical, real-world applications

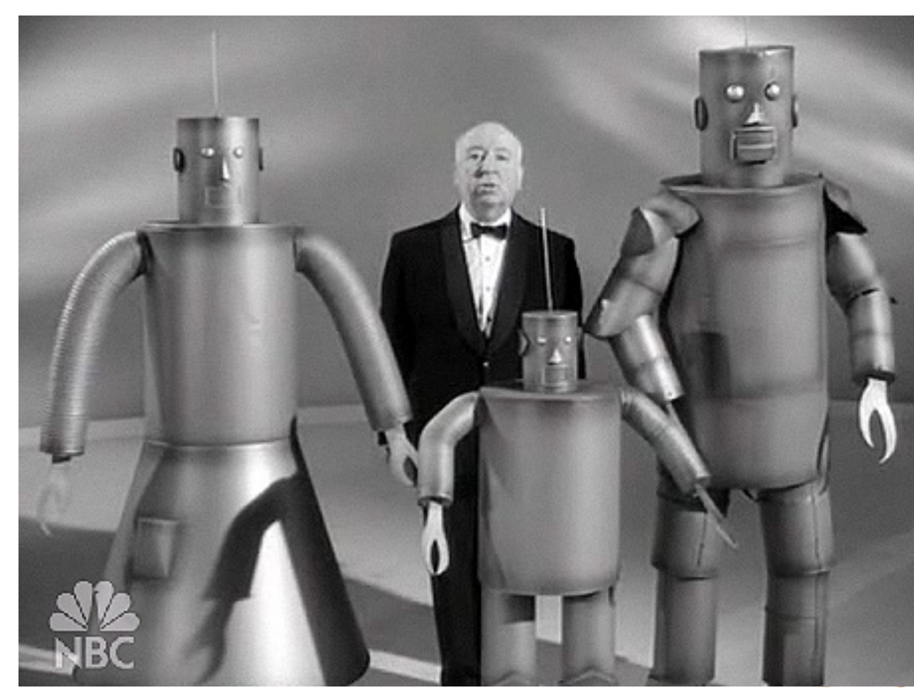

Concluding thoughts

- Narrative intelligence is central to many of the things humans do on a day to day basis
- Computational narrative intelligence is a key capability that unlocks many practical, real-world applications
- Path to human-AI interaction runs through stories

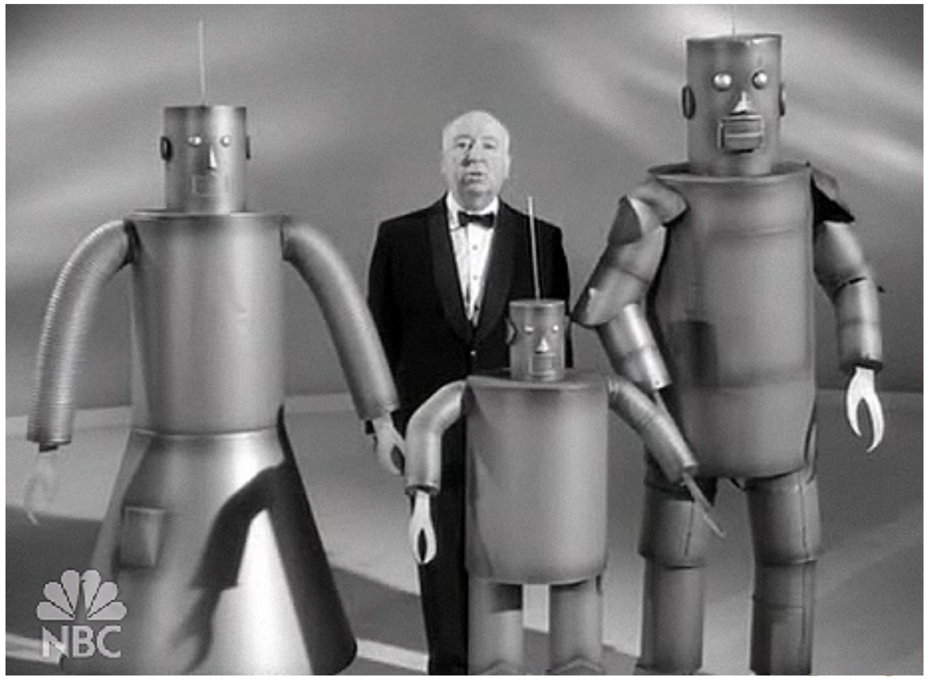

Thanks!

- **Sarah Wiegreffe**
- **Lara Martin**
- **Raj Ammanabrolu**
- Brent Harrison
- Boyang Li
- Stephen Lee-Urban
- Upol Ehsan
- Xiangyu Peng
- Amal Alabdulkarim
- Siyan Li
- Spencer Frazier
- Xinyu Wang
- Srijan Sood
- Winston Li
- Mohini Thakkar
- Yijie Wang
- Harshith Kayam
- Samihan Dani
- Md Sultan Al Nahian

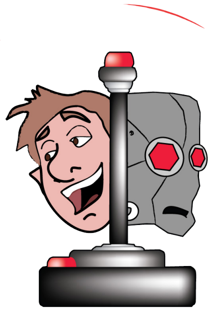

http://www.cc.gatech.edu/~riedl/
riedl@gatech.edu @mark_riedl

Georgia Tech College of Computing
School of Interactive Computing






Automated Story Generation as a Lens for Fundamental Artificial Intelligence

Mark Riedl

riedl@cc.gatech.edu

@mark_riedl

Georgia Tech College of Computing

School of Interactive Computing