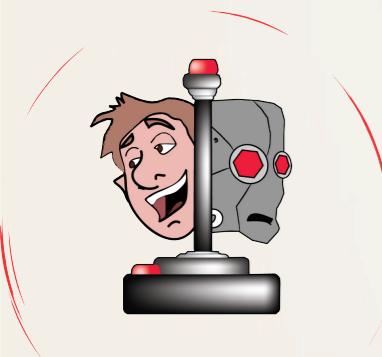
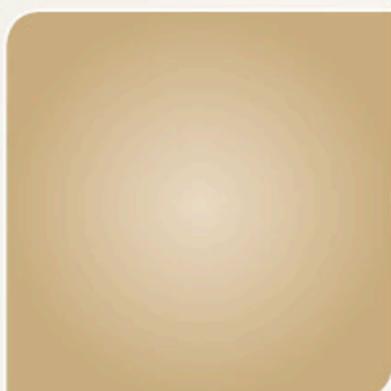


# Intelligent Narrative Generation: From Cognition to Crowdsourcing

Mark Riedl

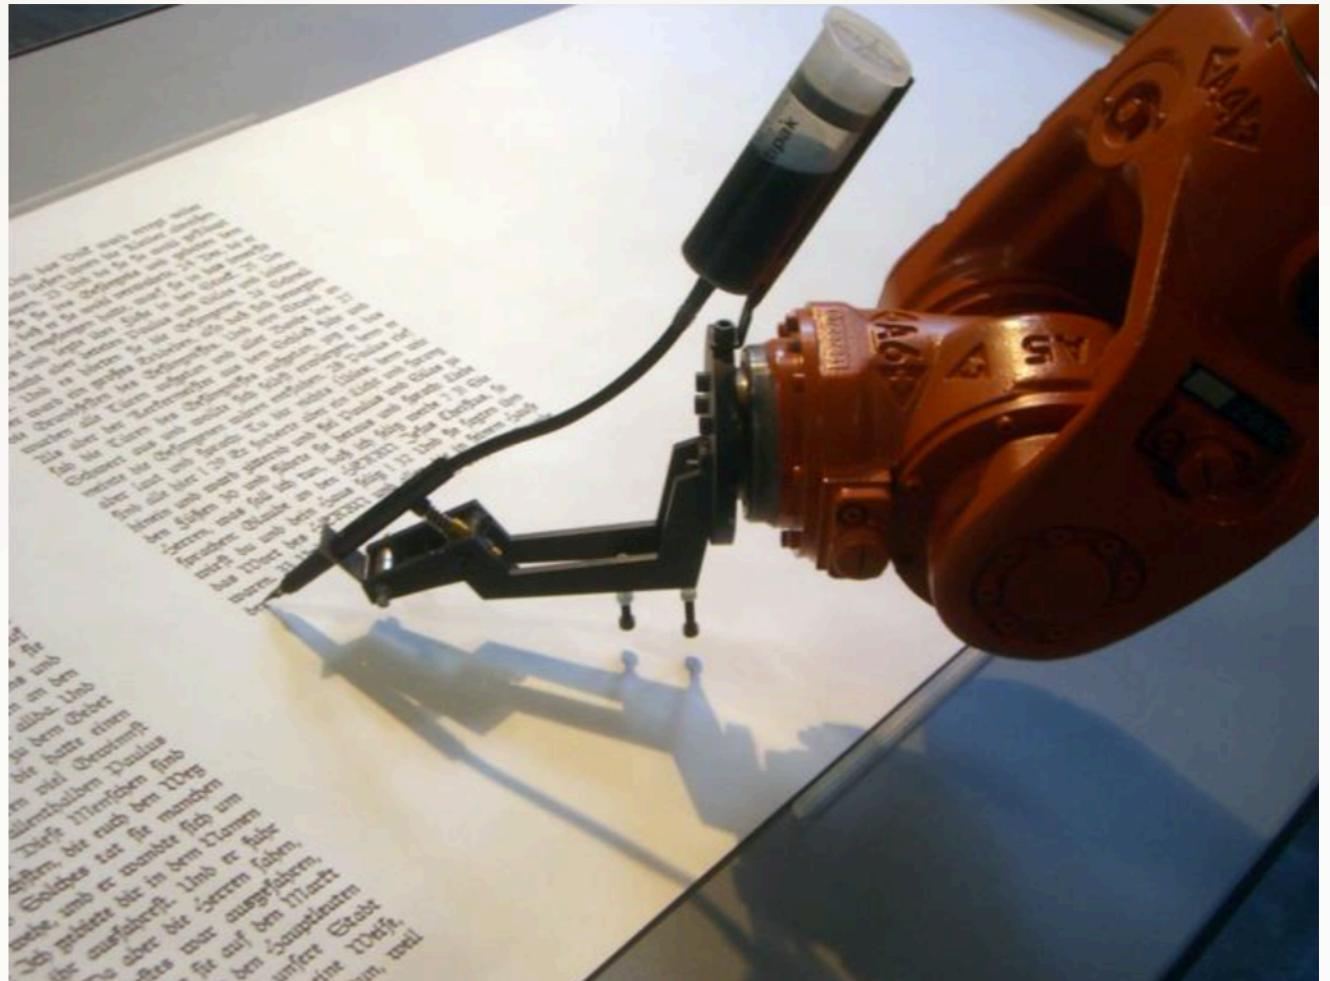
[riedl@cc.gatech.edu](mailto:riedl@cc.gatech.edu)  
[@mark\\_riedl](https://twitter.com/mark_riedl)





# Intelligent systems versus stories

- Understanding stories
  - Question answering
  - Distant reading
- Information retrieval
- Story generation





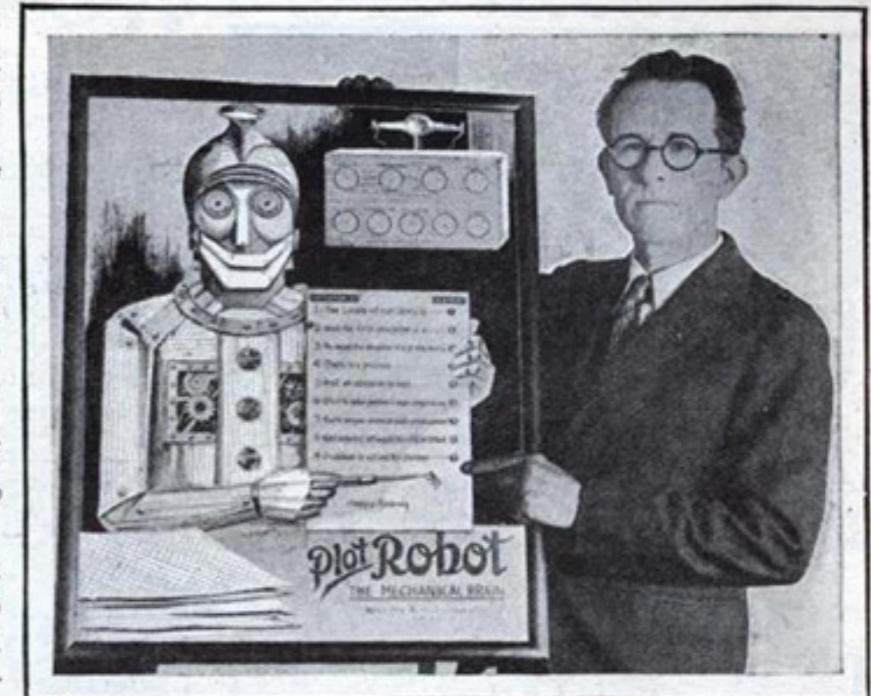
# Why study story generation?

- Stories are everywhere
- Humans make up stories all the time, but computers do not
- Cognitive science
- Computational creativity

## Robot With Mechanical Brain Thinks Up Story Plots

FORMERLY robots were merely mechanical devices that could perform a variety of stunts under the guidance of a human being, but now a robot has made its appearance that thinks, has a soul of a kind, creative imagination, and other qualities necessary for writing a modern stereotyped short story. This robot, the invention of Wycliffe Hill, a Los Angeles scenario writer, is declared to be able to build up millions of plots, no two alike, for magazine stories or movie plays.

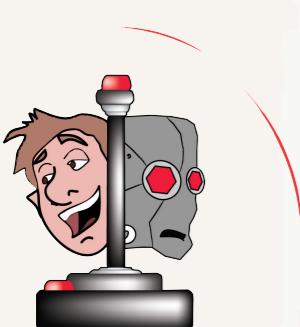
Mr. Hill has equipped his robot with an index chart, divided into eight sections, one devoted to each of the eight elements of a story—background, character, obstacle, problem, predicament, complication, crisis and climax—and with an assortment of variations. The robot selects the material as required from this inexhaustible source and builds plots that could never be imagined by the author



Mr. Wycliffe Hill demonstrating his new story writing robot, which can think up any kind of plot with its mechanical brains.

without the aid of the mechanical brain. Now if you want to become a successful author simply obtain a robot and put it to work.

© Popular Mechanics, 1931



A photograph of a paved road with a yellow double line running down the center. A metal fork is standing upright in the middle of the road, pointing downwards. In the background, there are trees and a small, blurry green object on the right side.

**Stories exist on the page and  
in the mind**

**NLP and story generation  
took different paths**

**Mea culpa**



There is a woman named Jasmine. There is a king named Jafar. This is a story about how King 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.

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

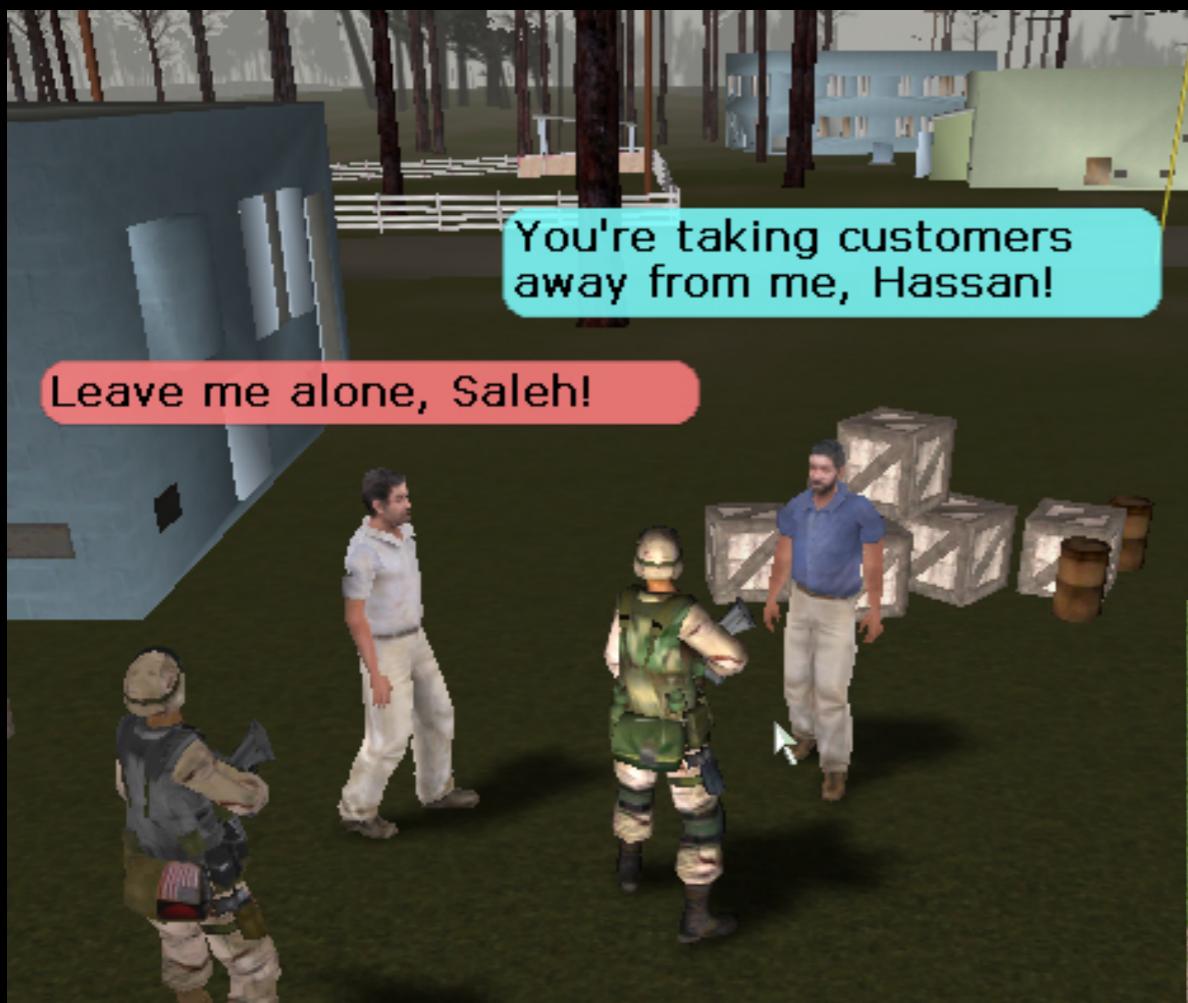
The genie is dead. King Jafar and Jasmine are married. The end.





You'll never believe what  
happened to me when I went to  
the bank yesterday. A man who...











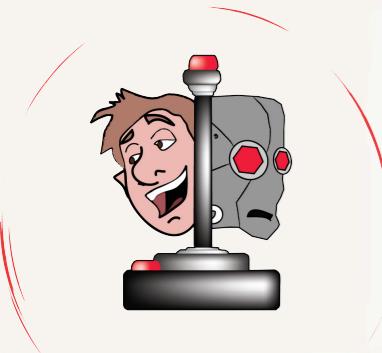




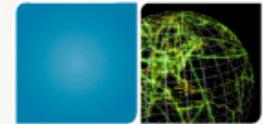


# Outline

- A brief history of story generation
  - Cognitive influences and narratology
  - Computer game influences
- Some work on story generation systems
- Next steps

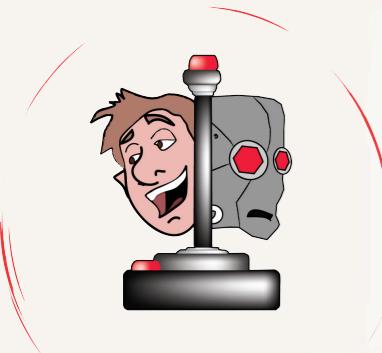


Once upon a time...

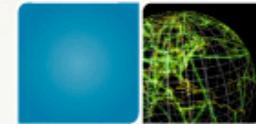


# Down the rabbit hole

- Want computers to interact and communicate with humans
- “I was on my way back from Lisbon. When I got to my gate, I discovered that my airplane was surrounded by firetrucks. Needless to say, I missed my connection in Amsterdam...”
- Story understanding: how does one make sense of stories?



9



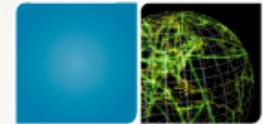
# Cognitive influence

- Studies show that readers remember the general gist of stories but do not remember the specific words (or images)
- Focus on concepts instead of words
- Narrative intelligence is knowledge-intensive
- Conceptual dependency theory?



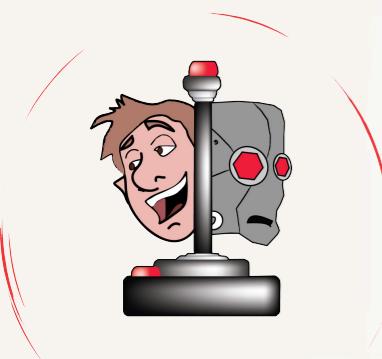
## *The Ordering Scene of the Restaurant Script*

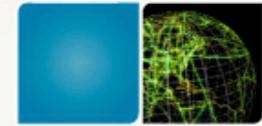
Sequence of events	Action
S MTRANS signal to W	
W PTRANS W to table	S asks for the menu
S MTRANS "need menu" to W	
W PTRANS W to menu	
W ATRANS menu to W	
W PTRANS W to table	W brings the menu
W ATRANS menu to S	
S MTRANS food list to mind	
S MBUILD choice of Food F	S decides what to order
S MTRANS signal to W	
W PTRANS W to table	
S MTRANS "I want F" to W	S tells W
W PTRANS W to C	
W MTRANS (ATRANS F) to C	W tells C
C DO (prepare F script)	



# Story generation is born

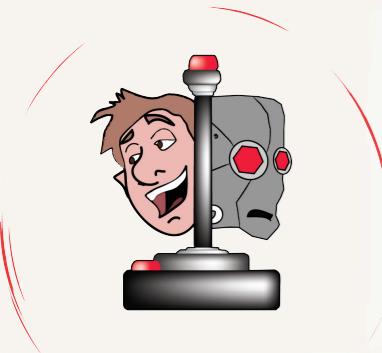
- Tale-Spin (Meehan, 1976)
  - Physical and social simulation
- Author (Dehn, 1981)
  - Considers story generation from author's perspective
- Universe (Lebowitz, 1985)
  - Early version of hierarchical task network planning
- Minstrel (Turner, 1992)
  - Early version of case-based reasoning



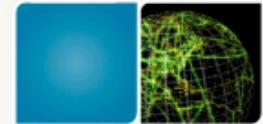


# Modern influences on story generation

- Narratology
- Cognitive studies of story comprehension



12



# Structural narratology

- **Fabula**

- The events that occur in the fictional world (the ground truth)
- Includes inferred events not presented to the audience
- Temporally ordered and can have co-occurring events

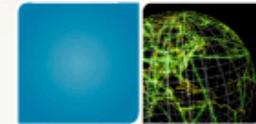
- **Sjuzet/discourse**

- A selection of the fabula to be presented to the audience
- Linear recounting
- Can re-order the events for dramatic effect

- **Text/media**

- Words, images, etc. that the audience directly observes

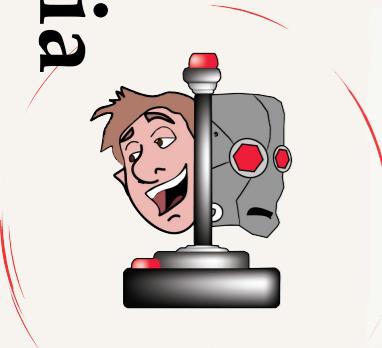


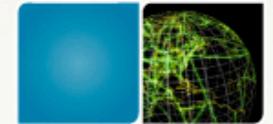


Fabula

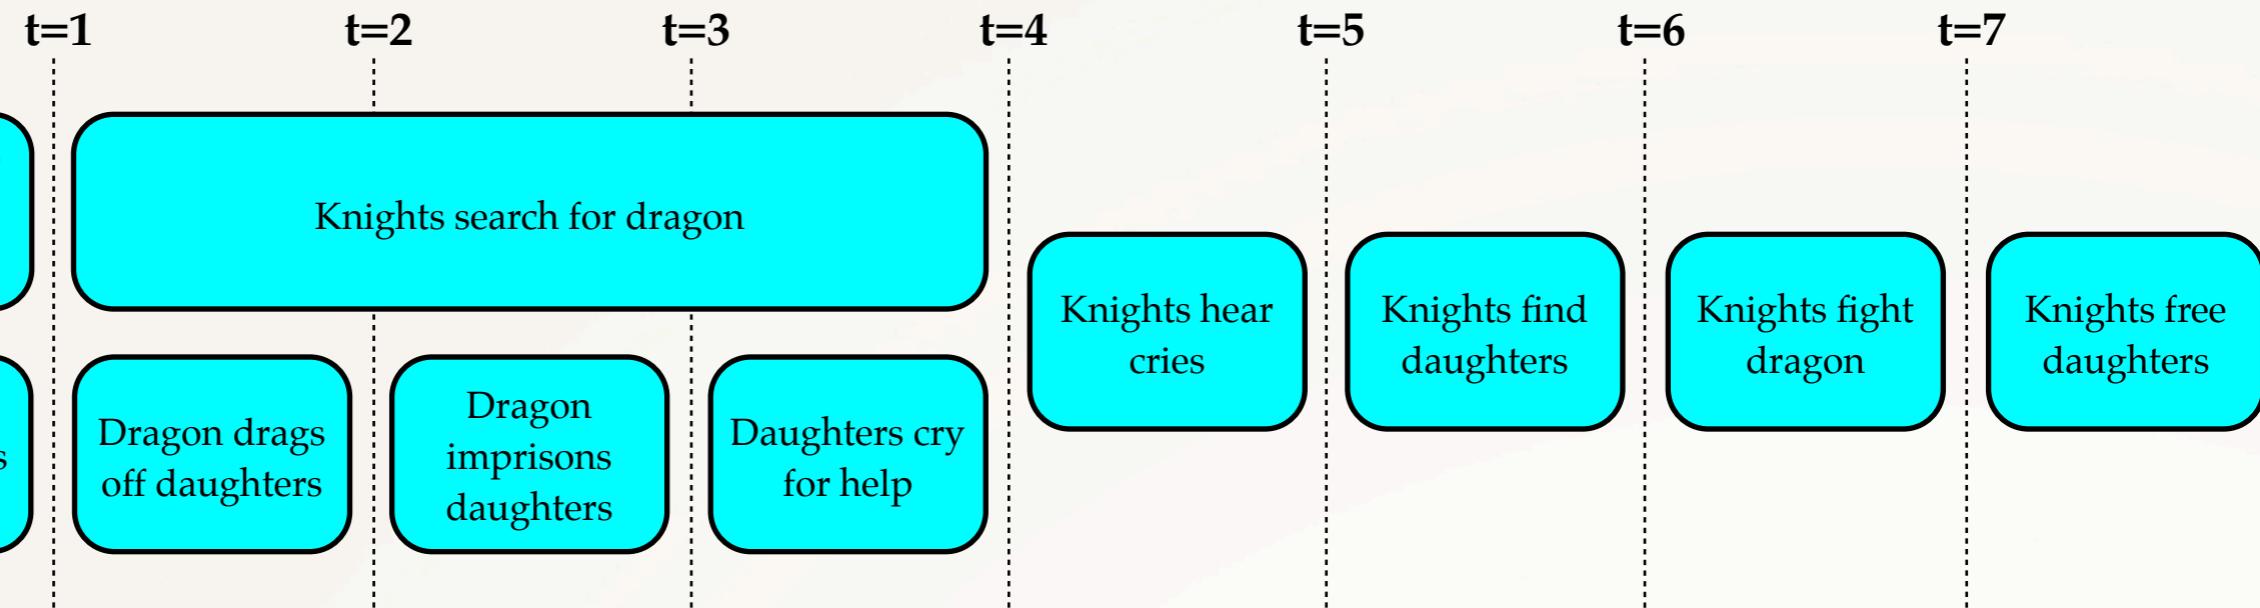
Sjuzet

Text/Media





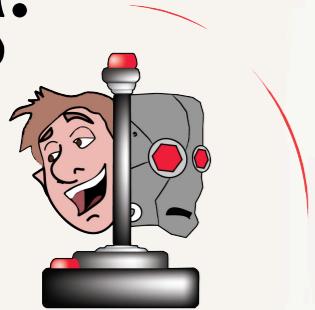
## Fabula

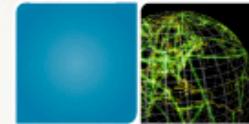


## Sjuzet

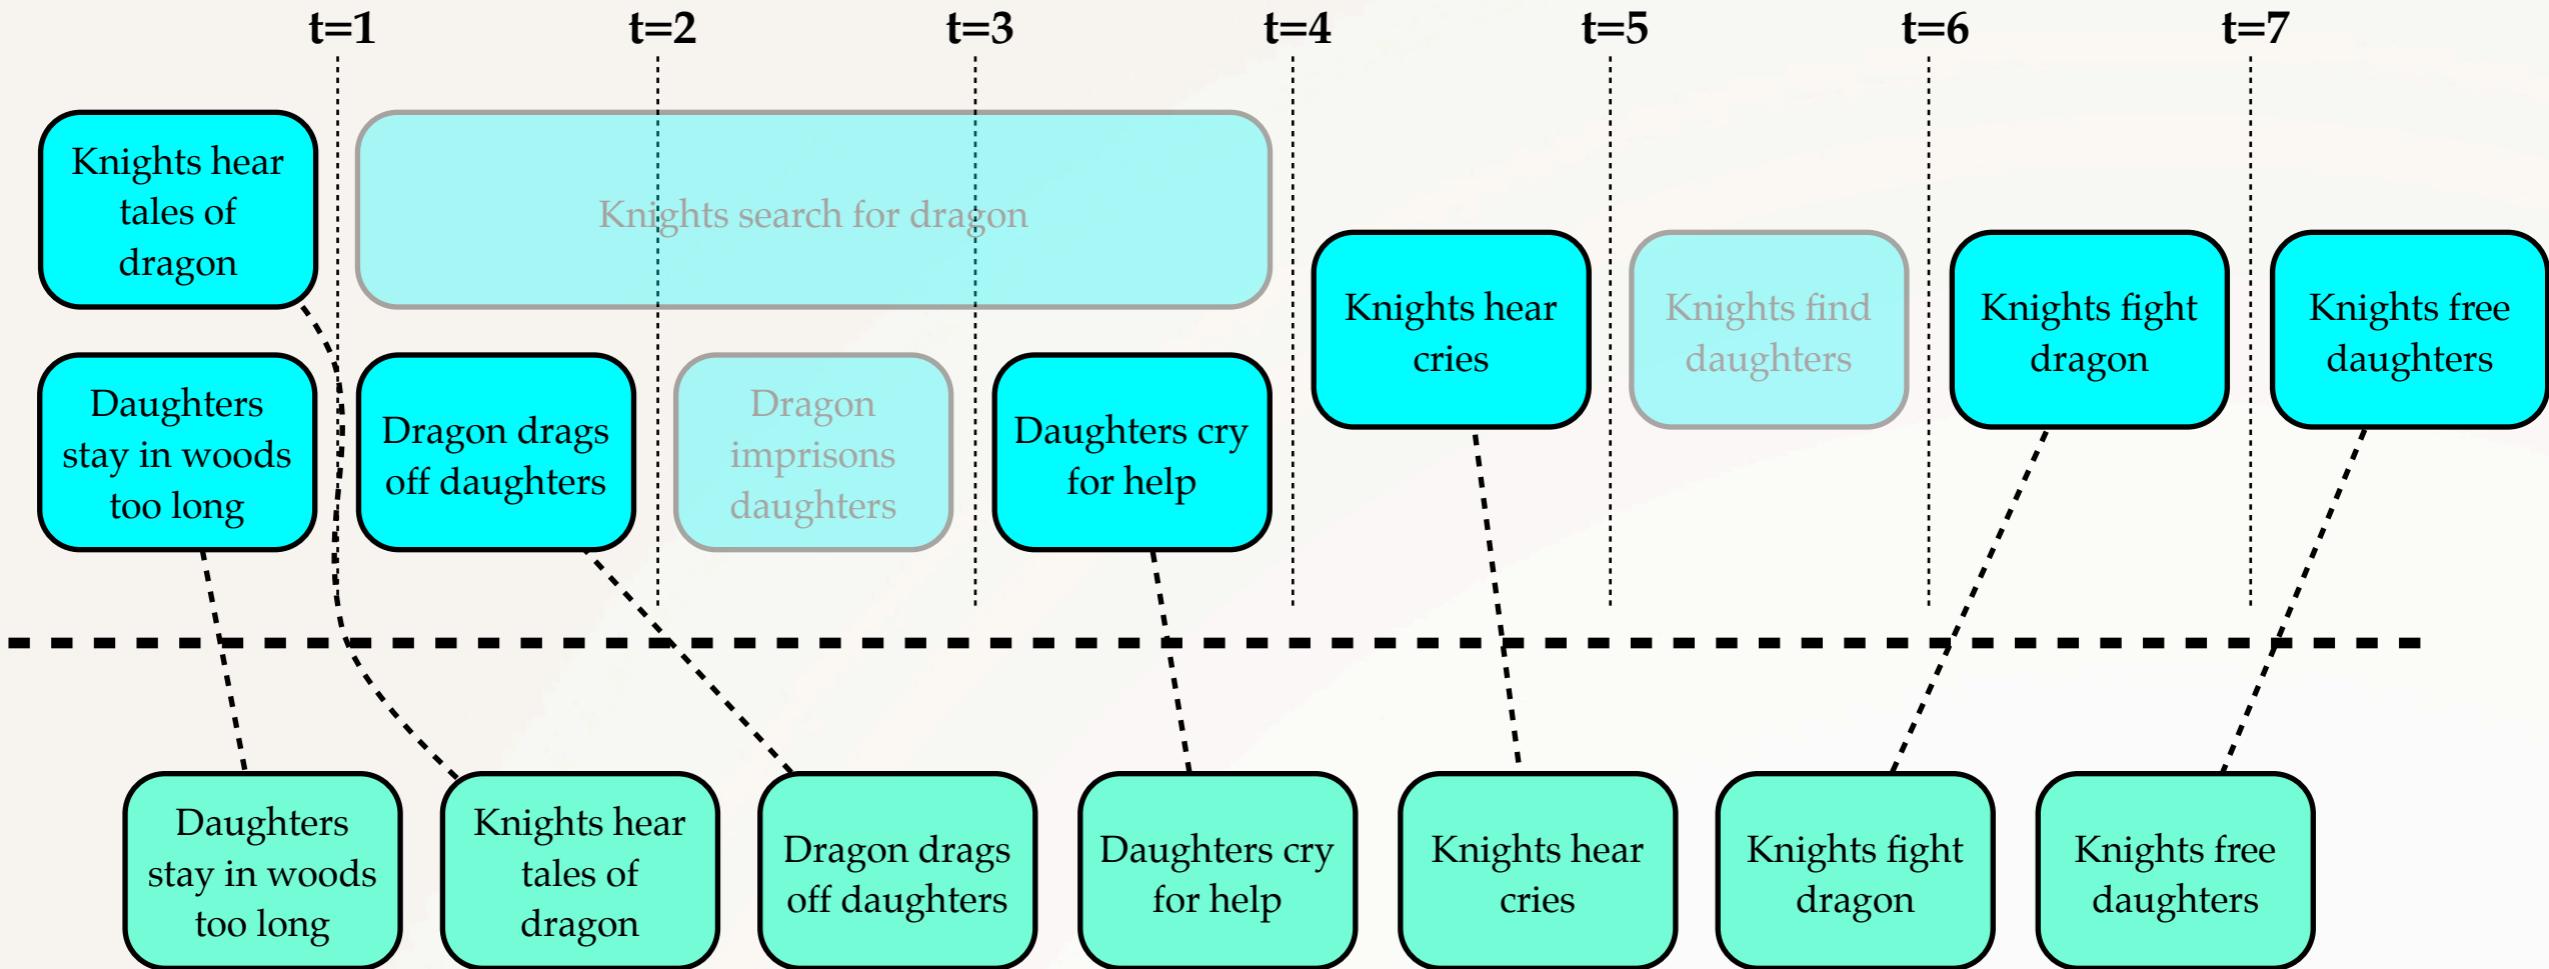
---

## Text/Media





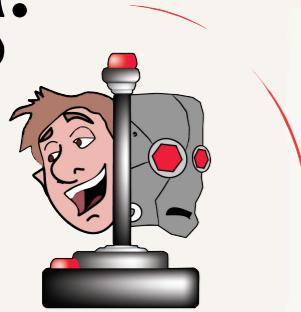
# Fabula

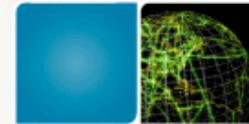


# Sjuzet

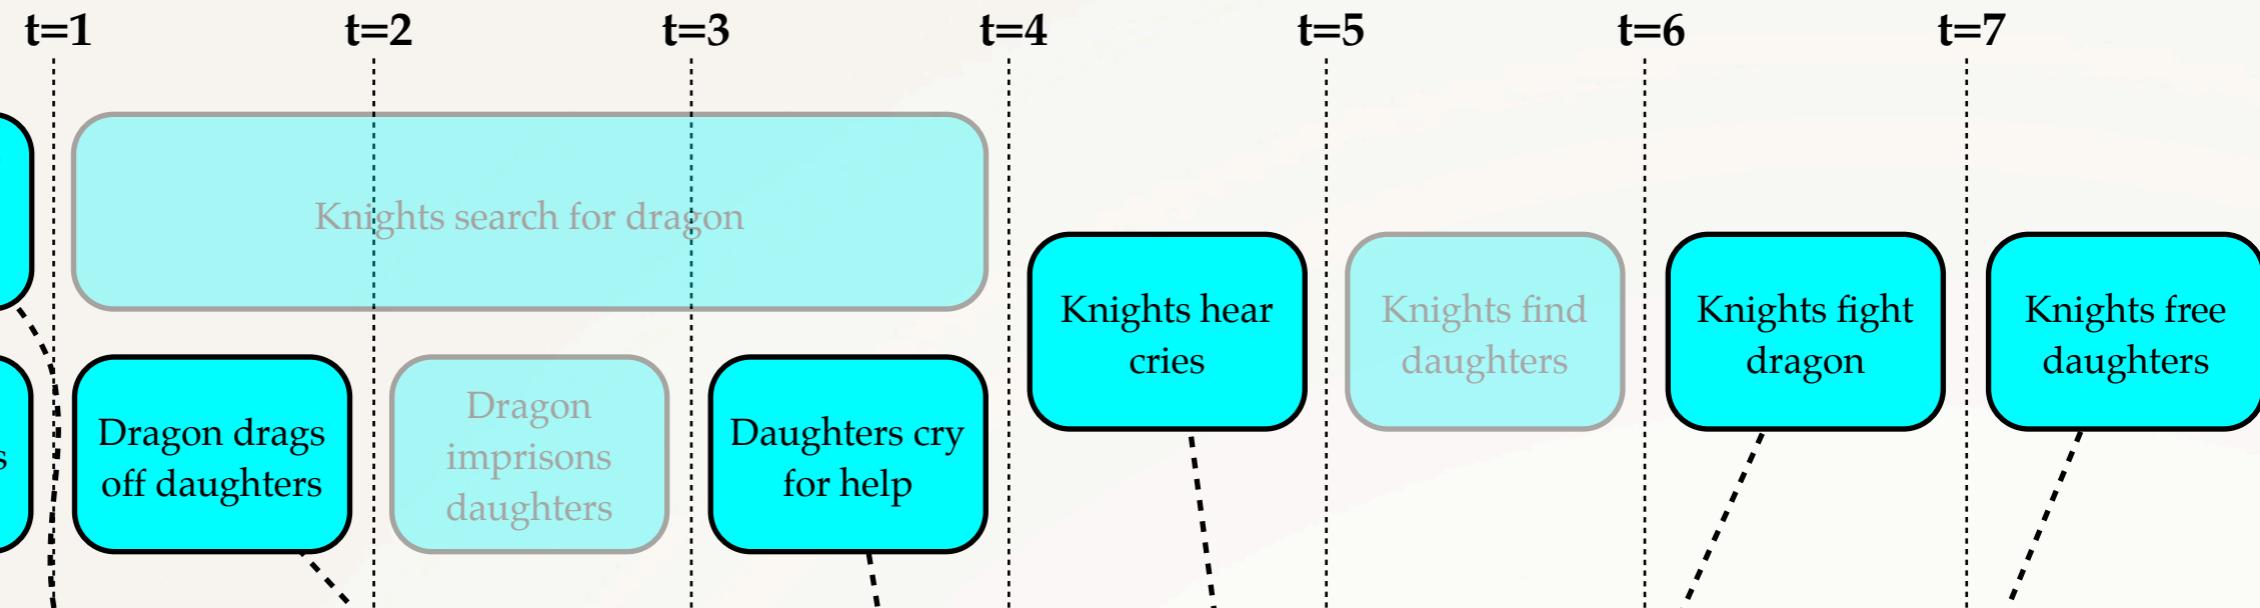
-----

# Text/Media





# Fabula



# Sjuzet

Daughters stay in woods too long

Knights hear tales of dragon

Dragon drags off daughters

Daughters cry for help

Knights hear cries

Knights fight dragon

Knights free daughters

The daughters were enjoying themselves in the woods so much that they lost track of time.

Meanwhile, some knights heard rumors of a dragon and set off to hunt it.

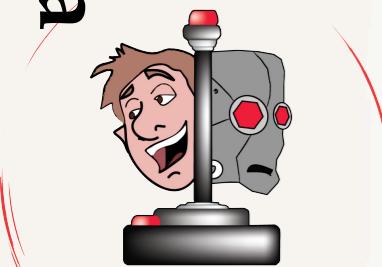
The dragon dragged off the daughters...

Having been imprisoned, the daughters cry out for help.

The knights, unsuccessful so far, heard the cries.

The knights fight the dragon...

...and free the daughters.

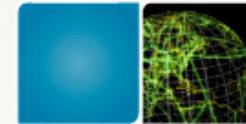
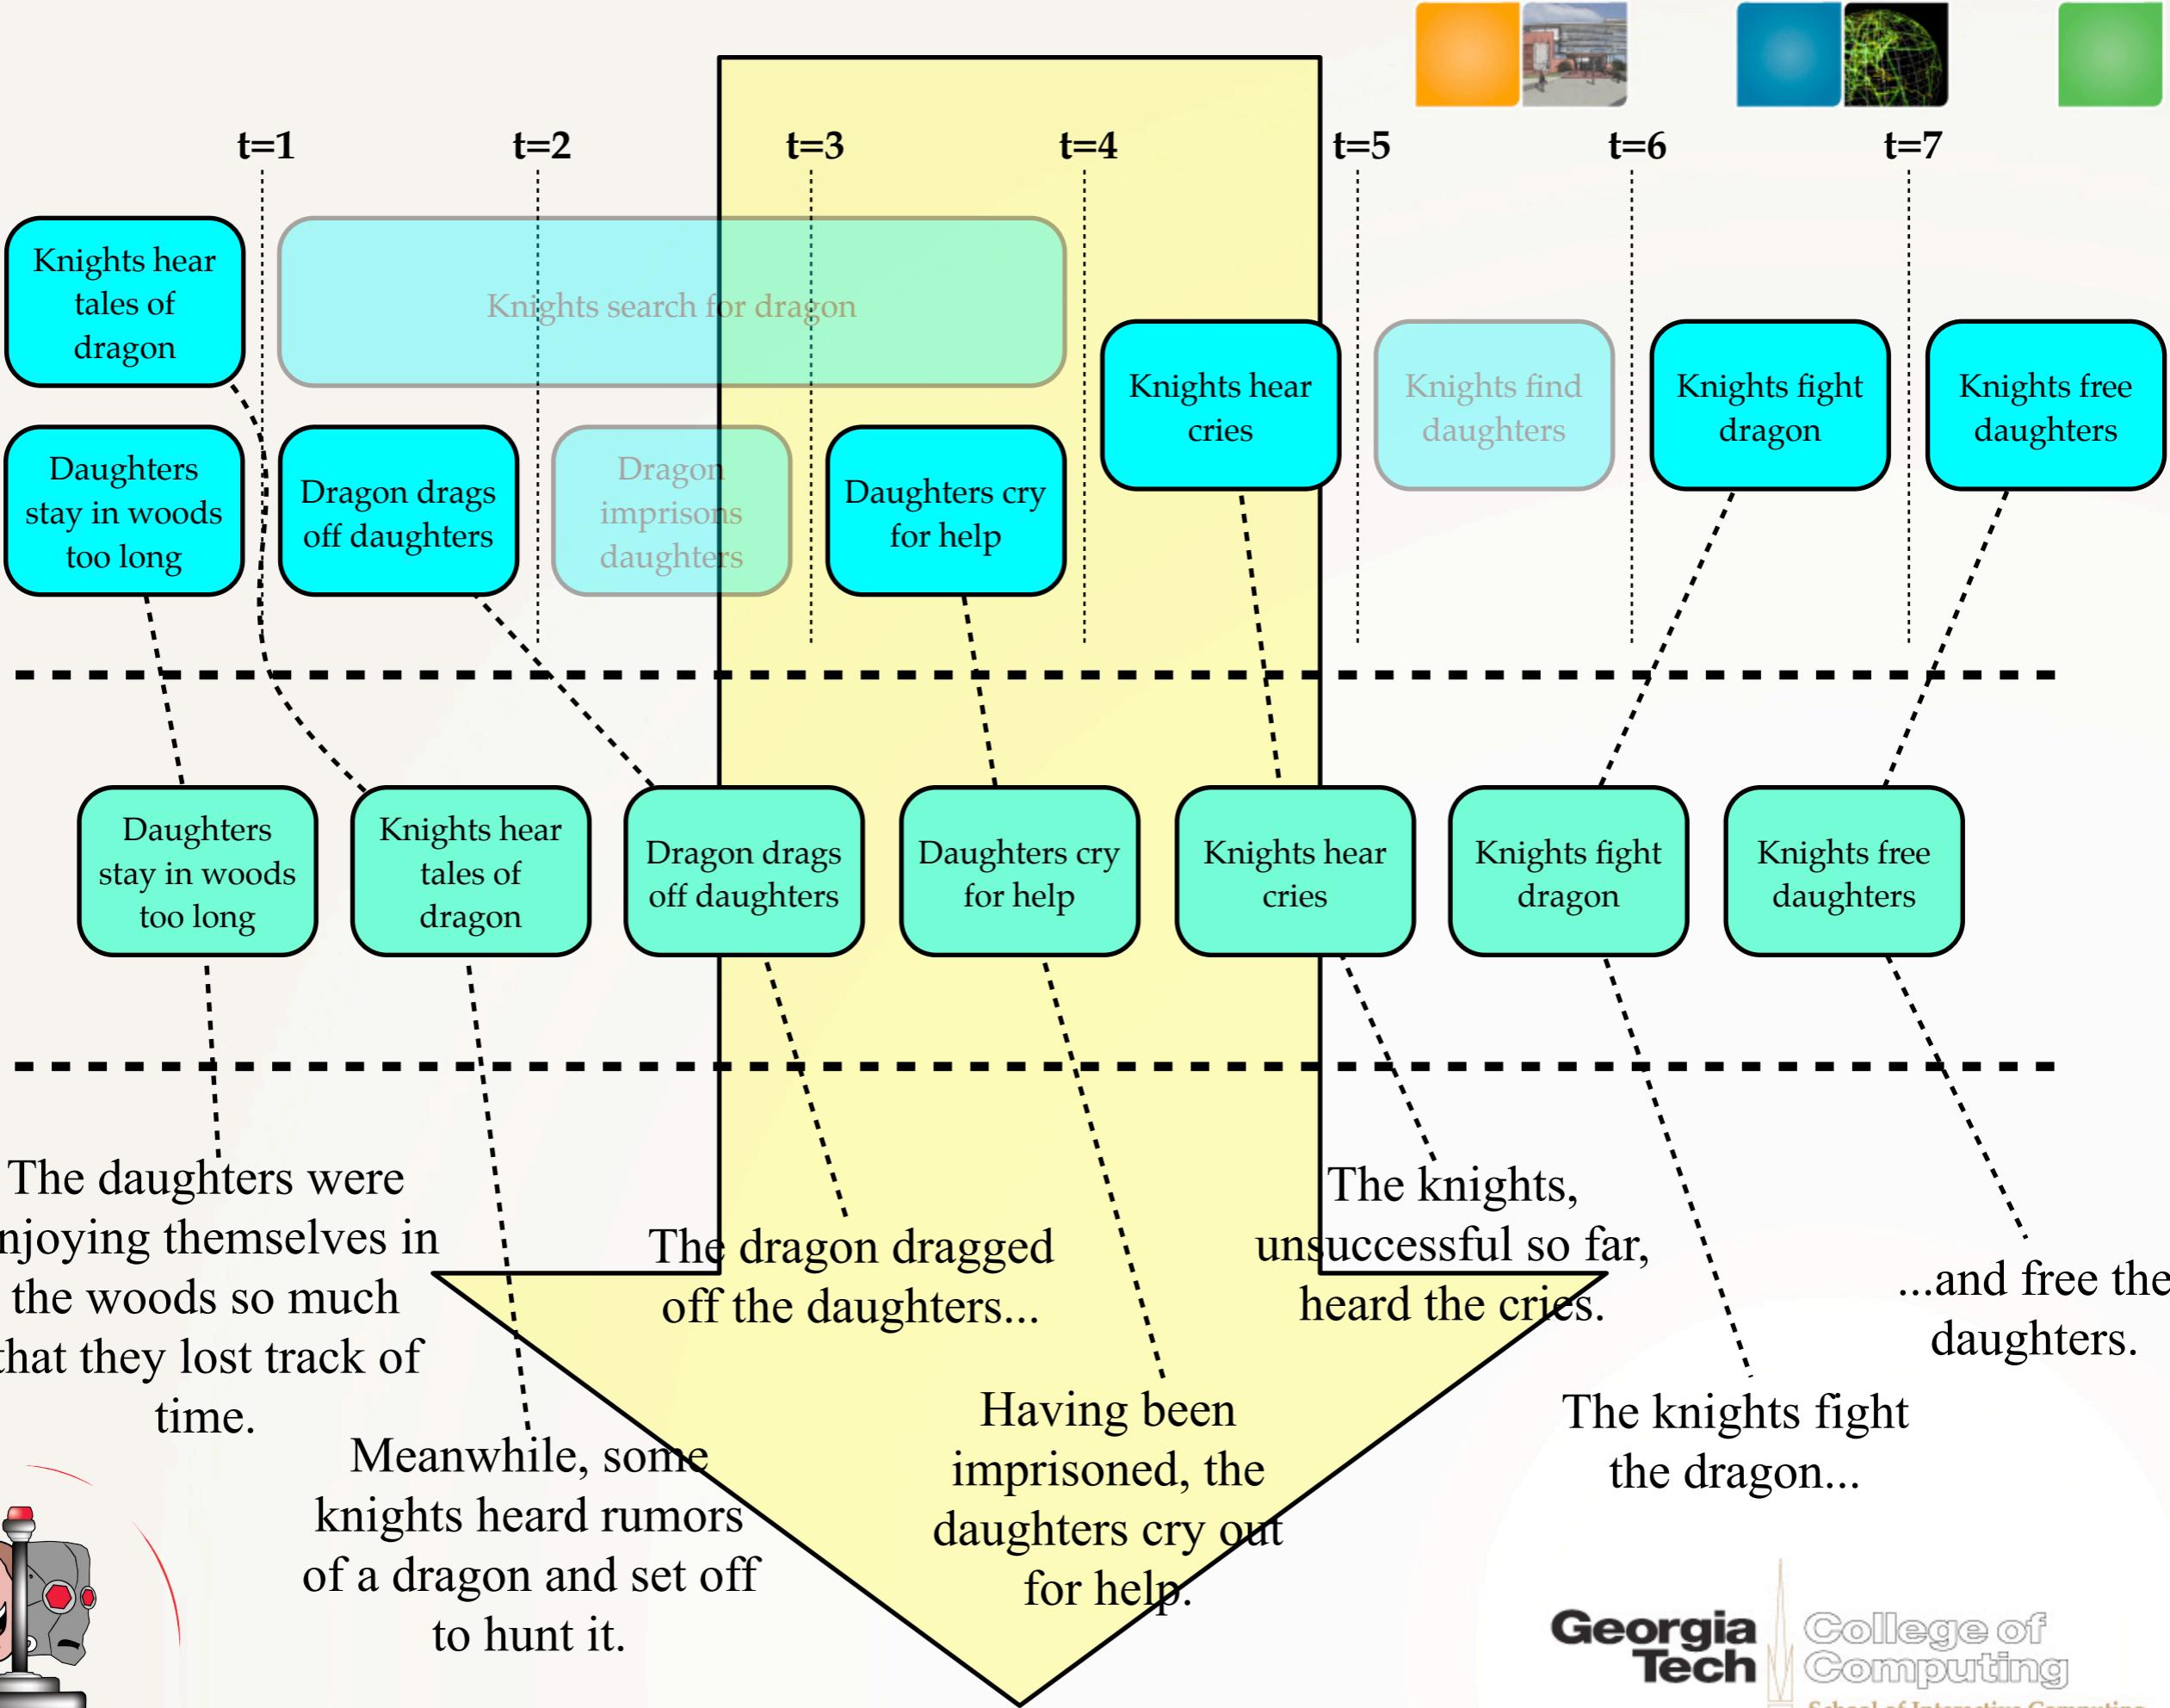


# Text/Media

# Fabula

# Sjuzet

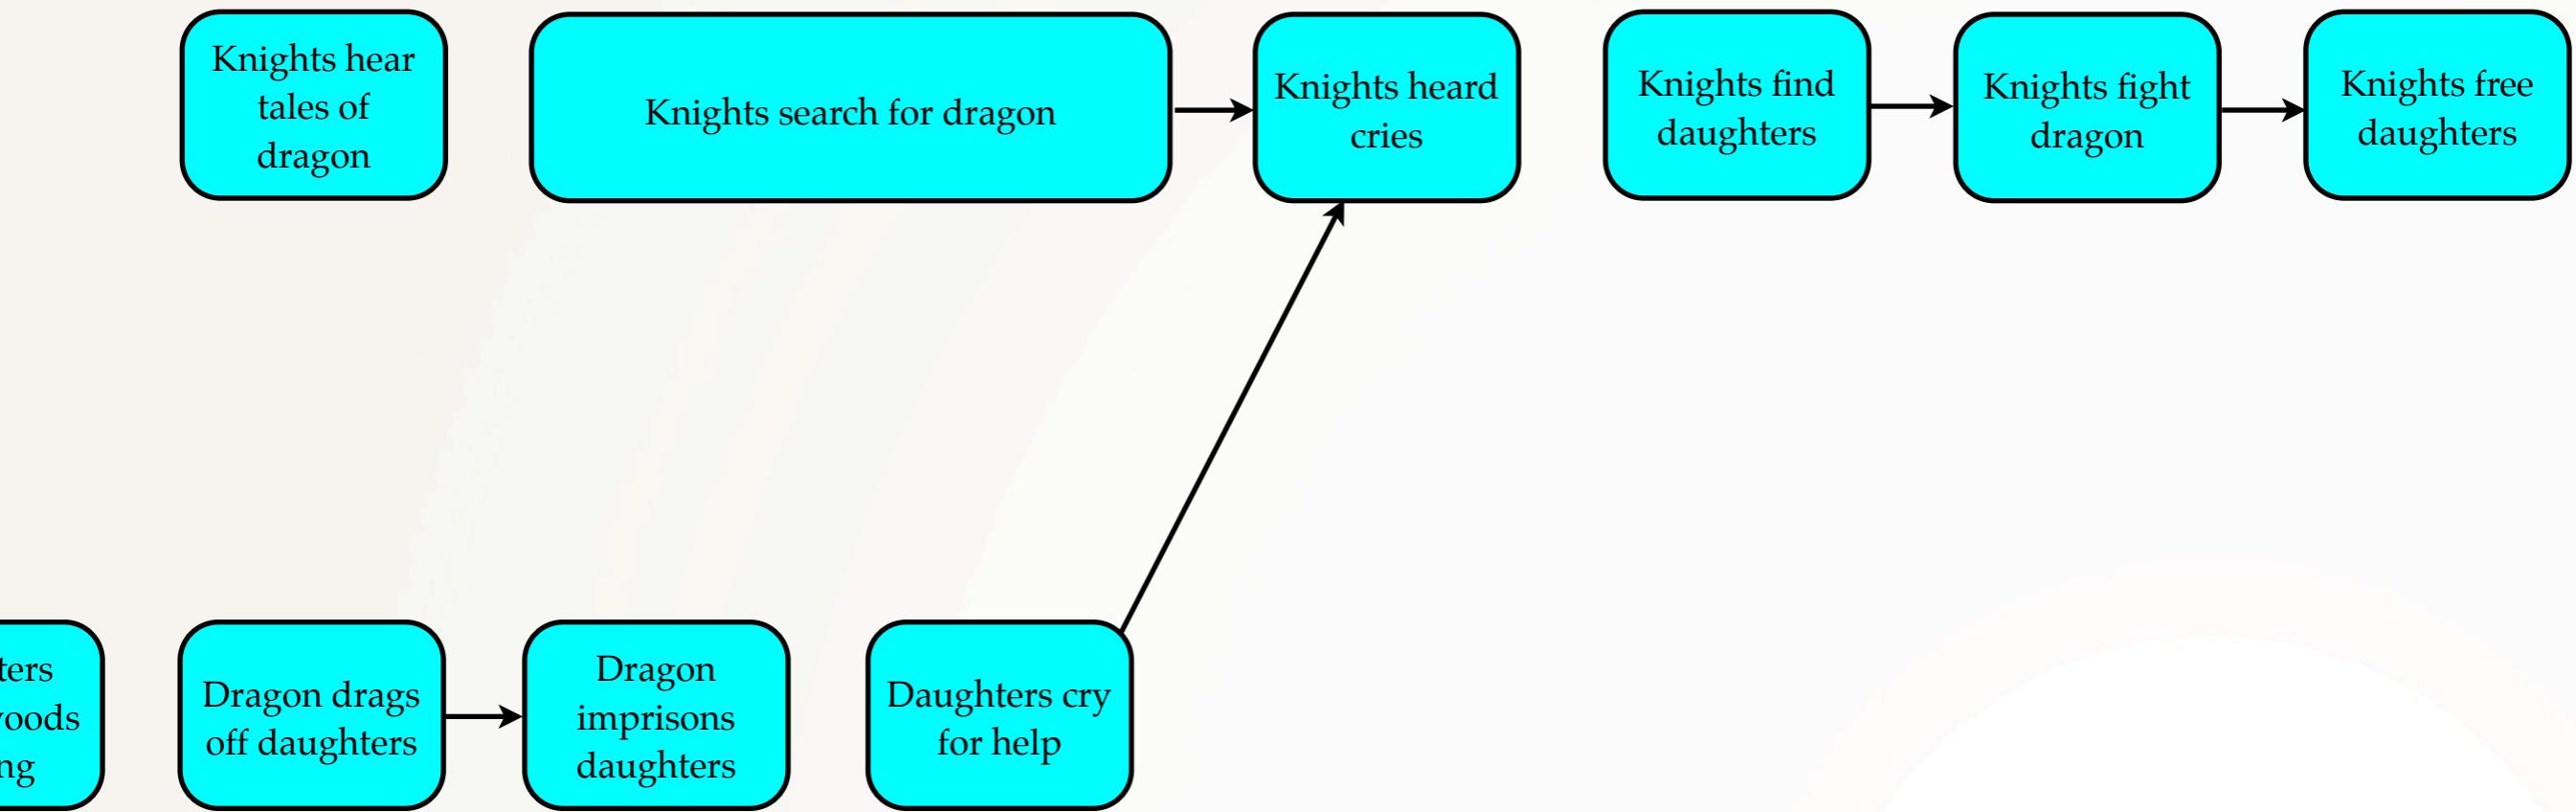
# Text/Media



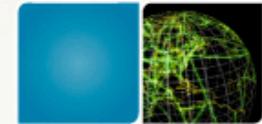


# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

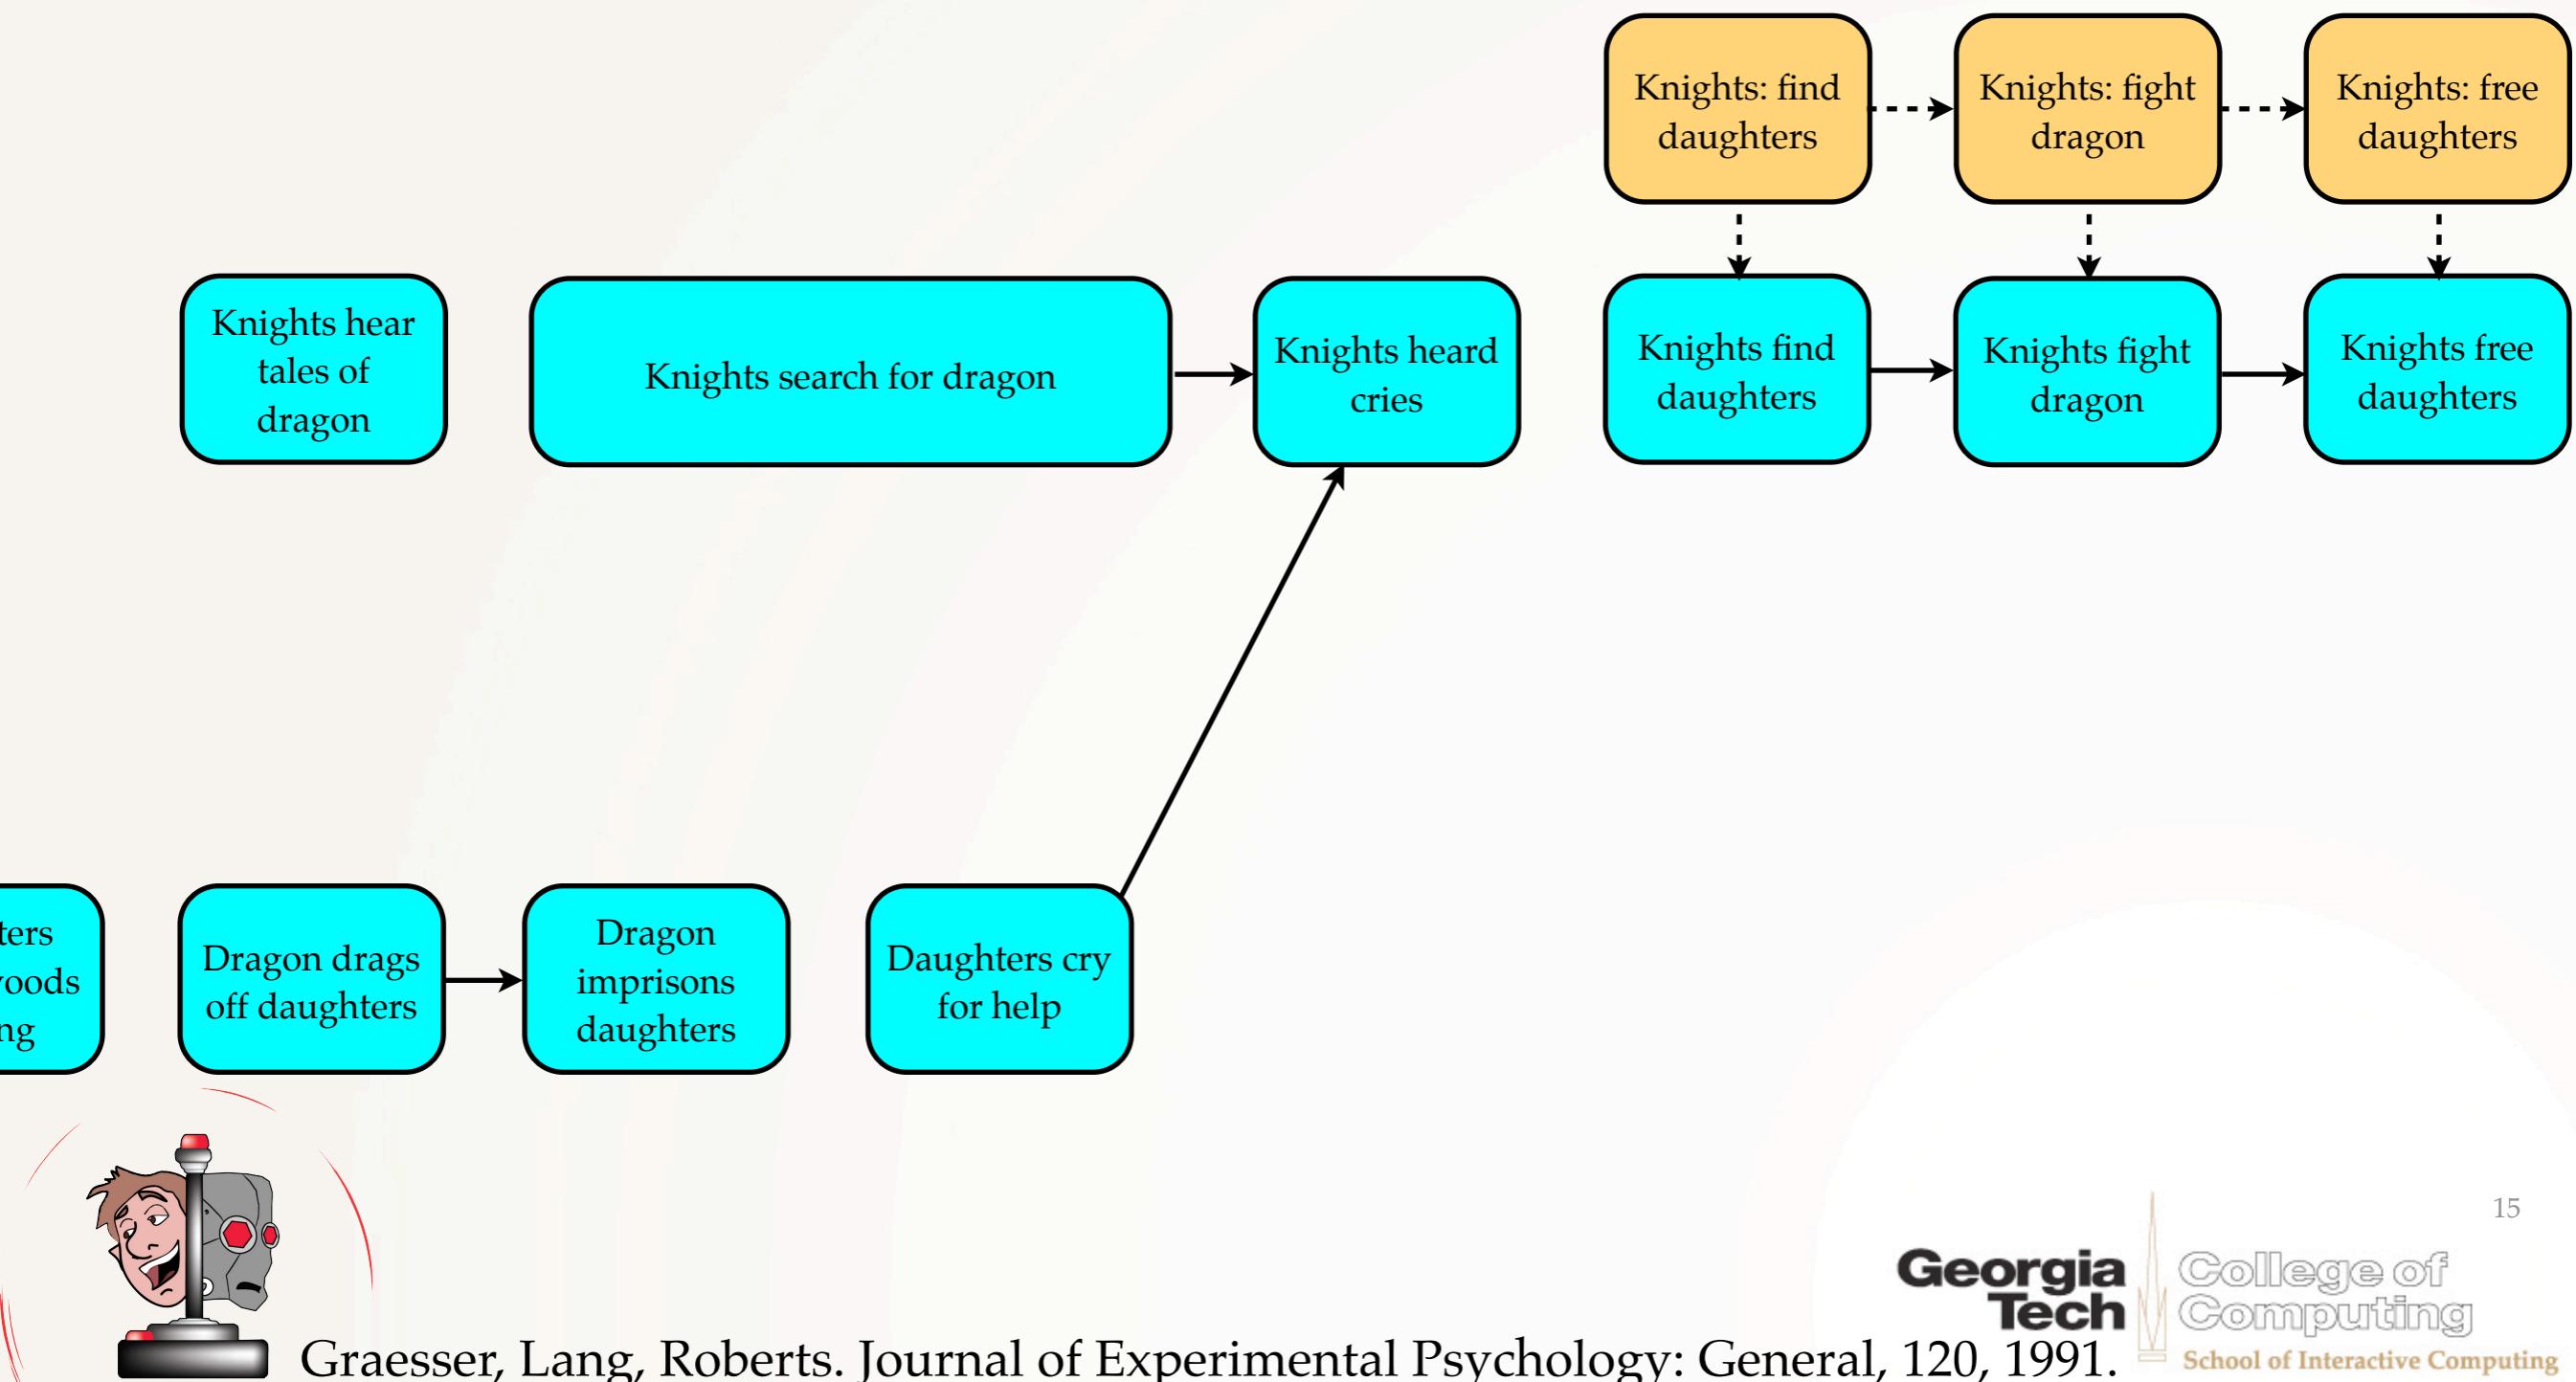


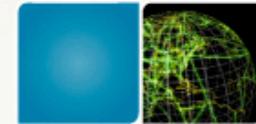
Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.



# Cognitive models of story comprehension

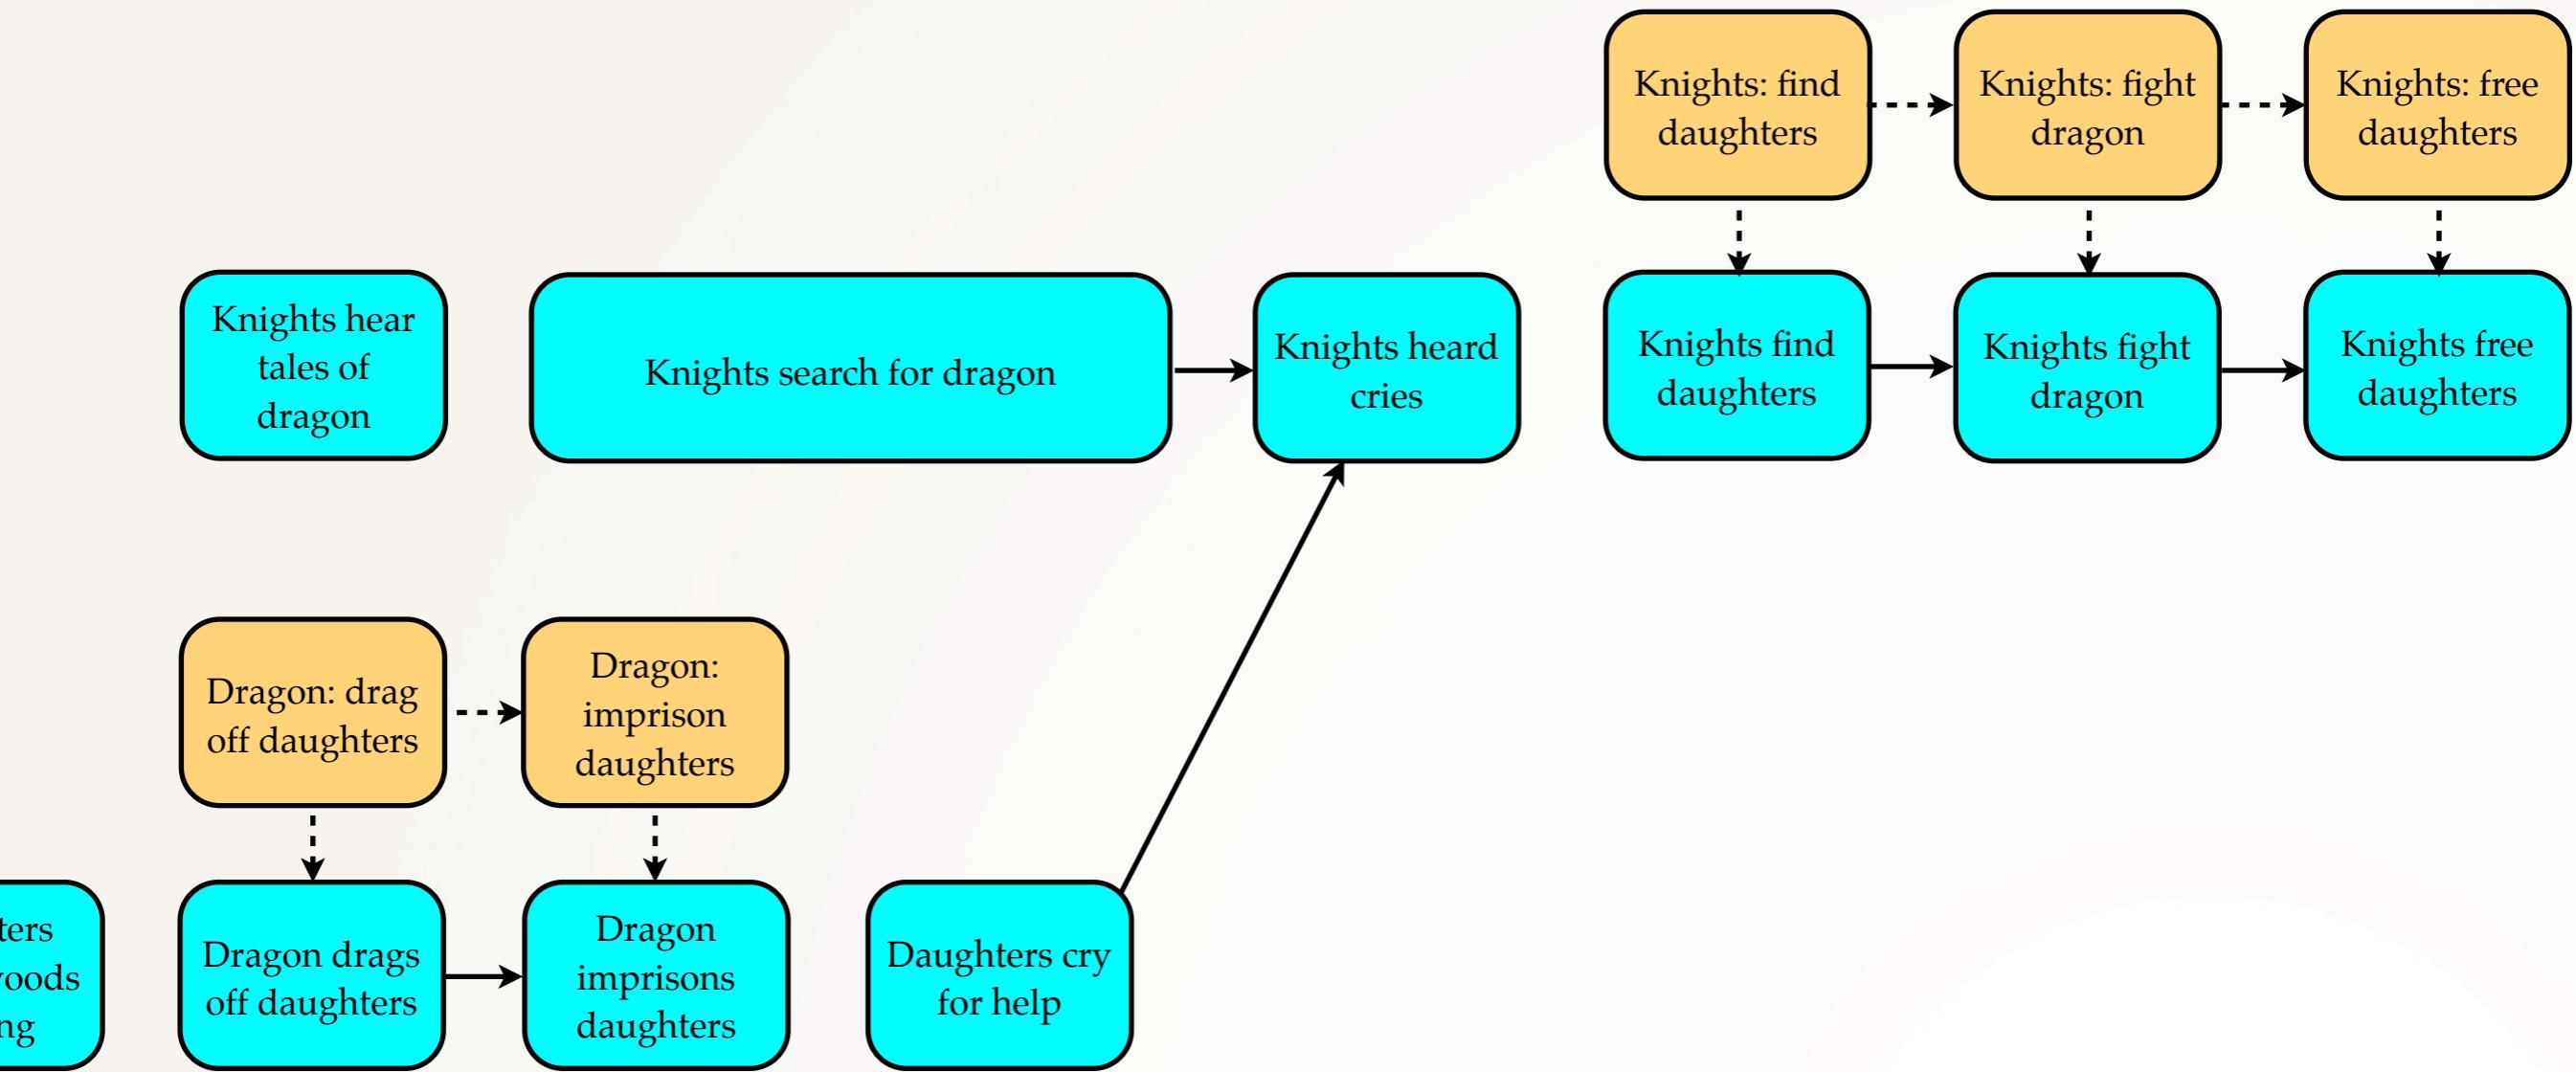
- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols





# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

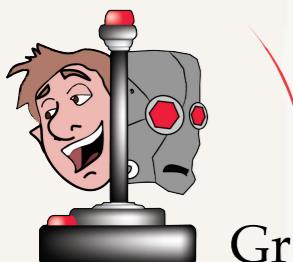
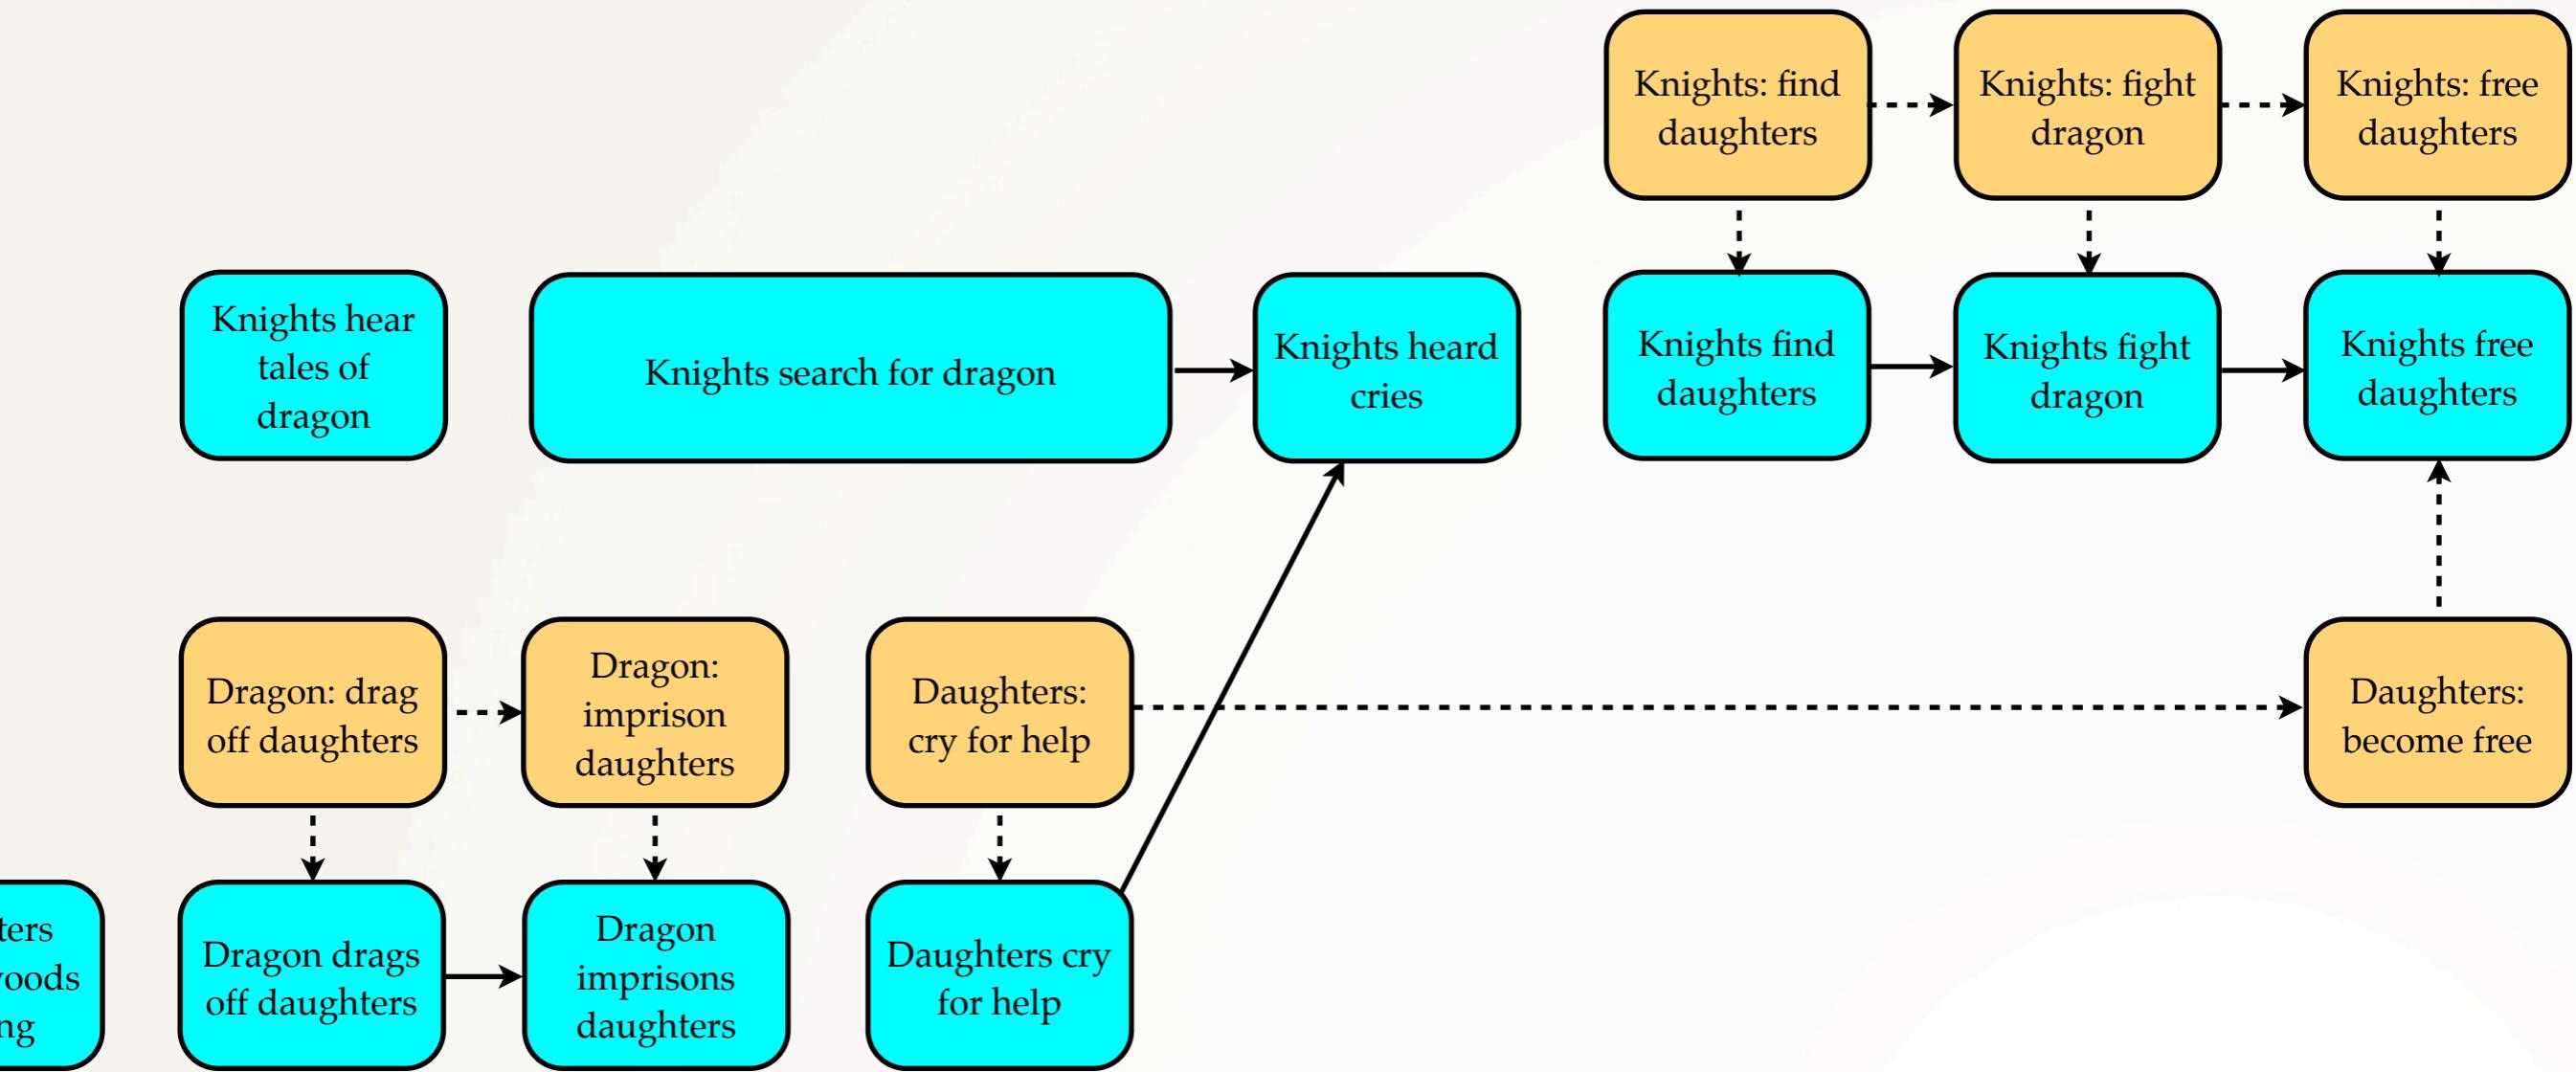


Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.



# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

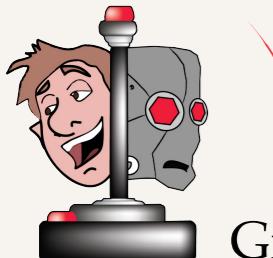
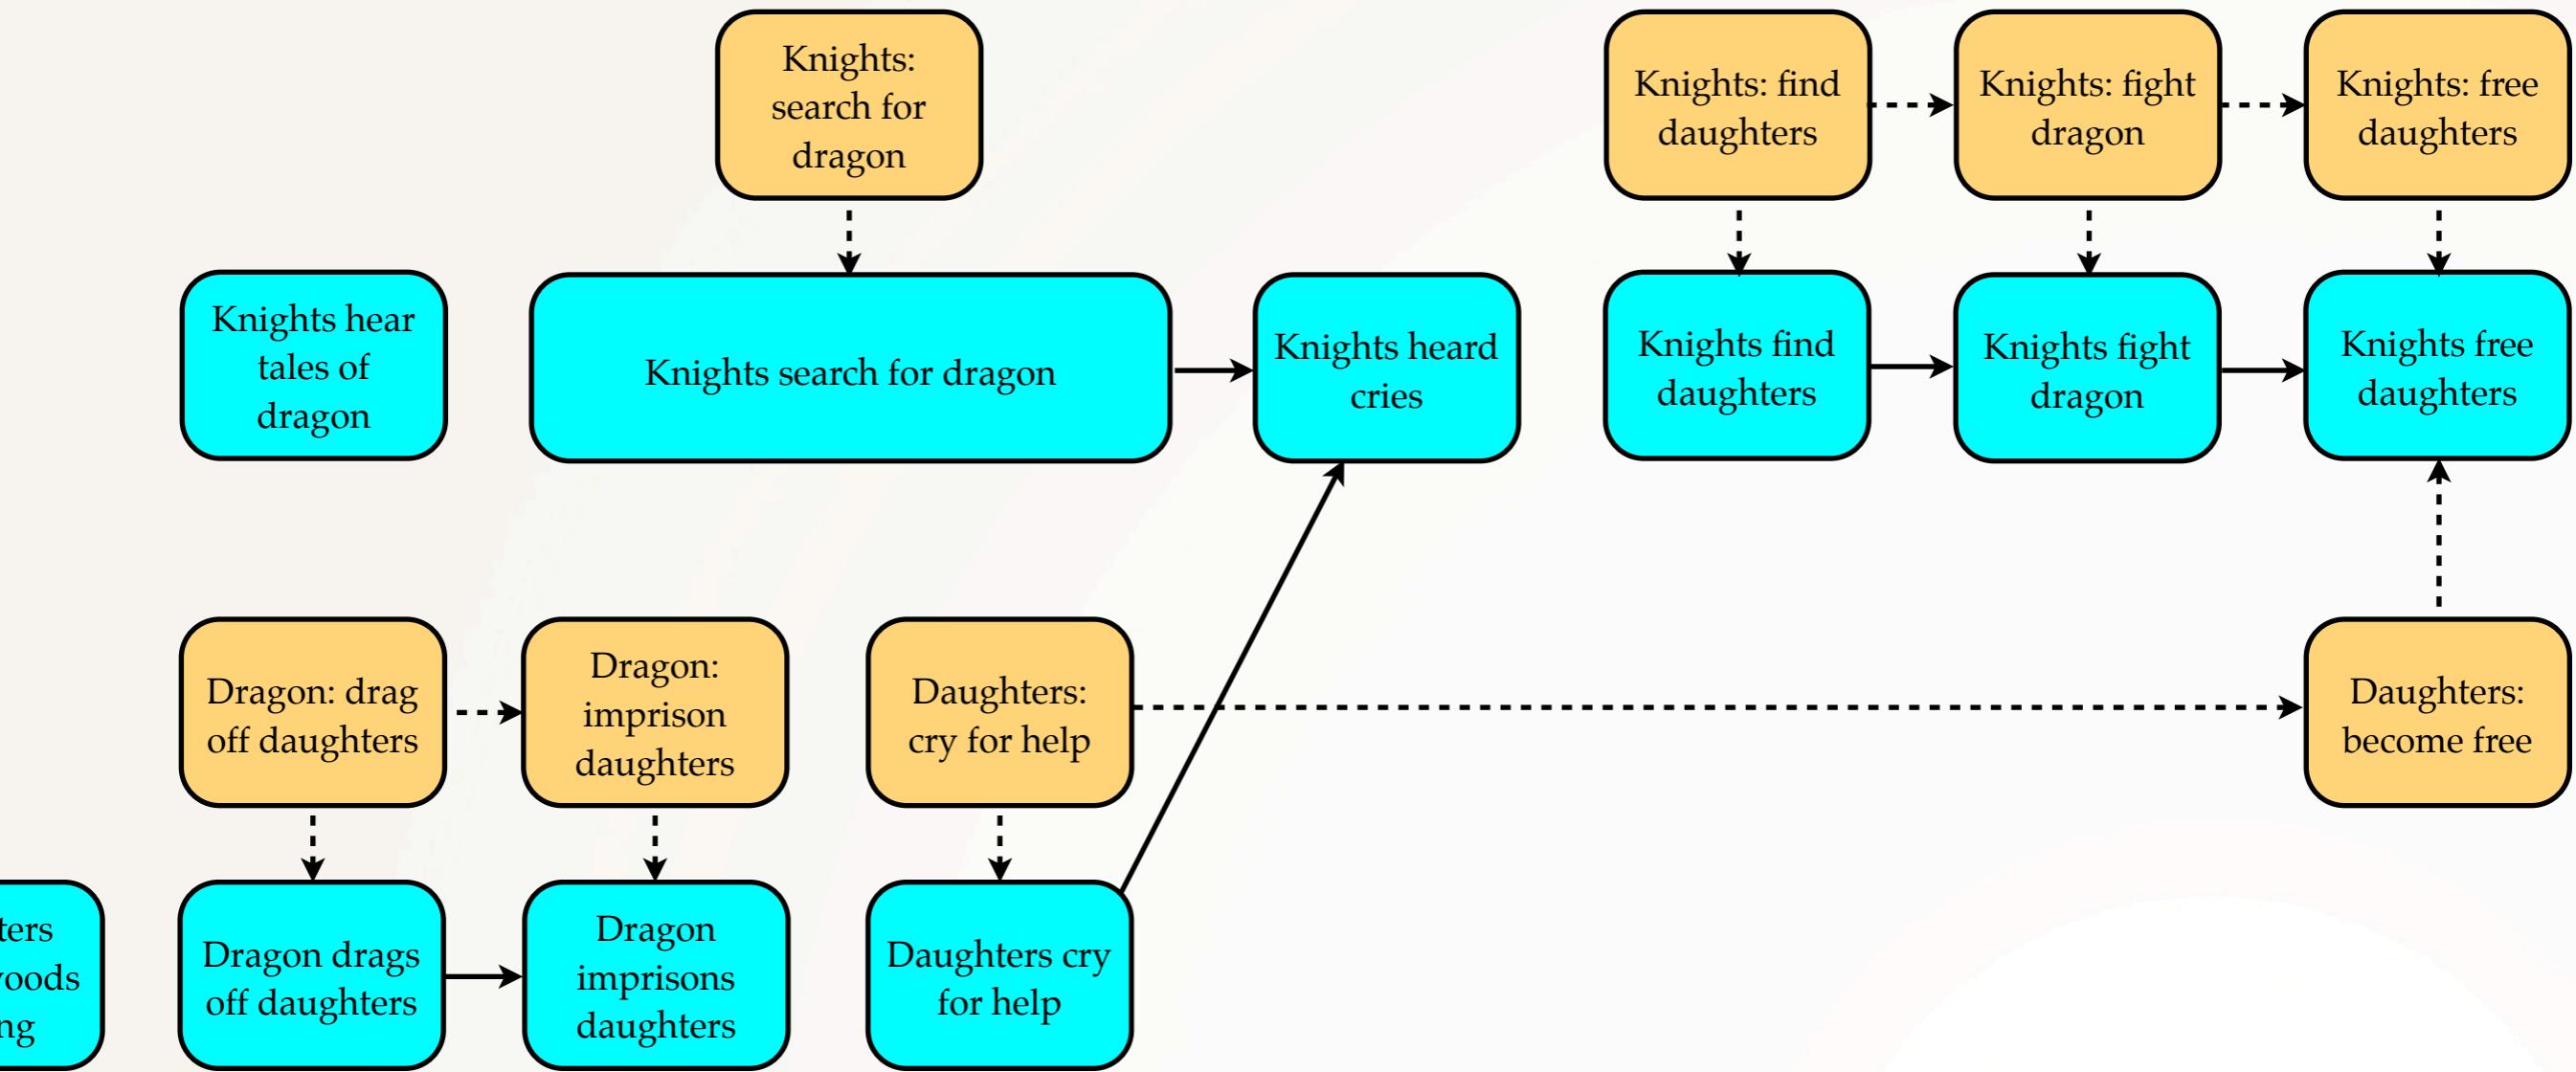


Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.

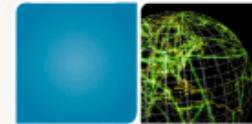


# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

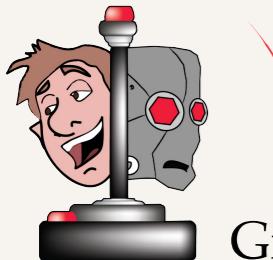
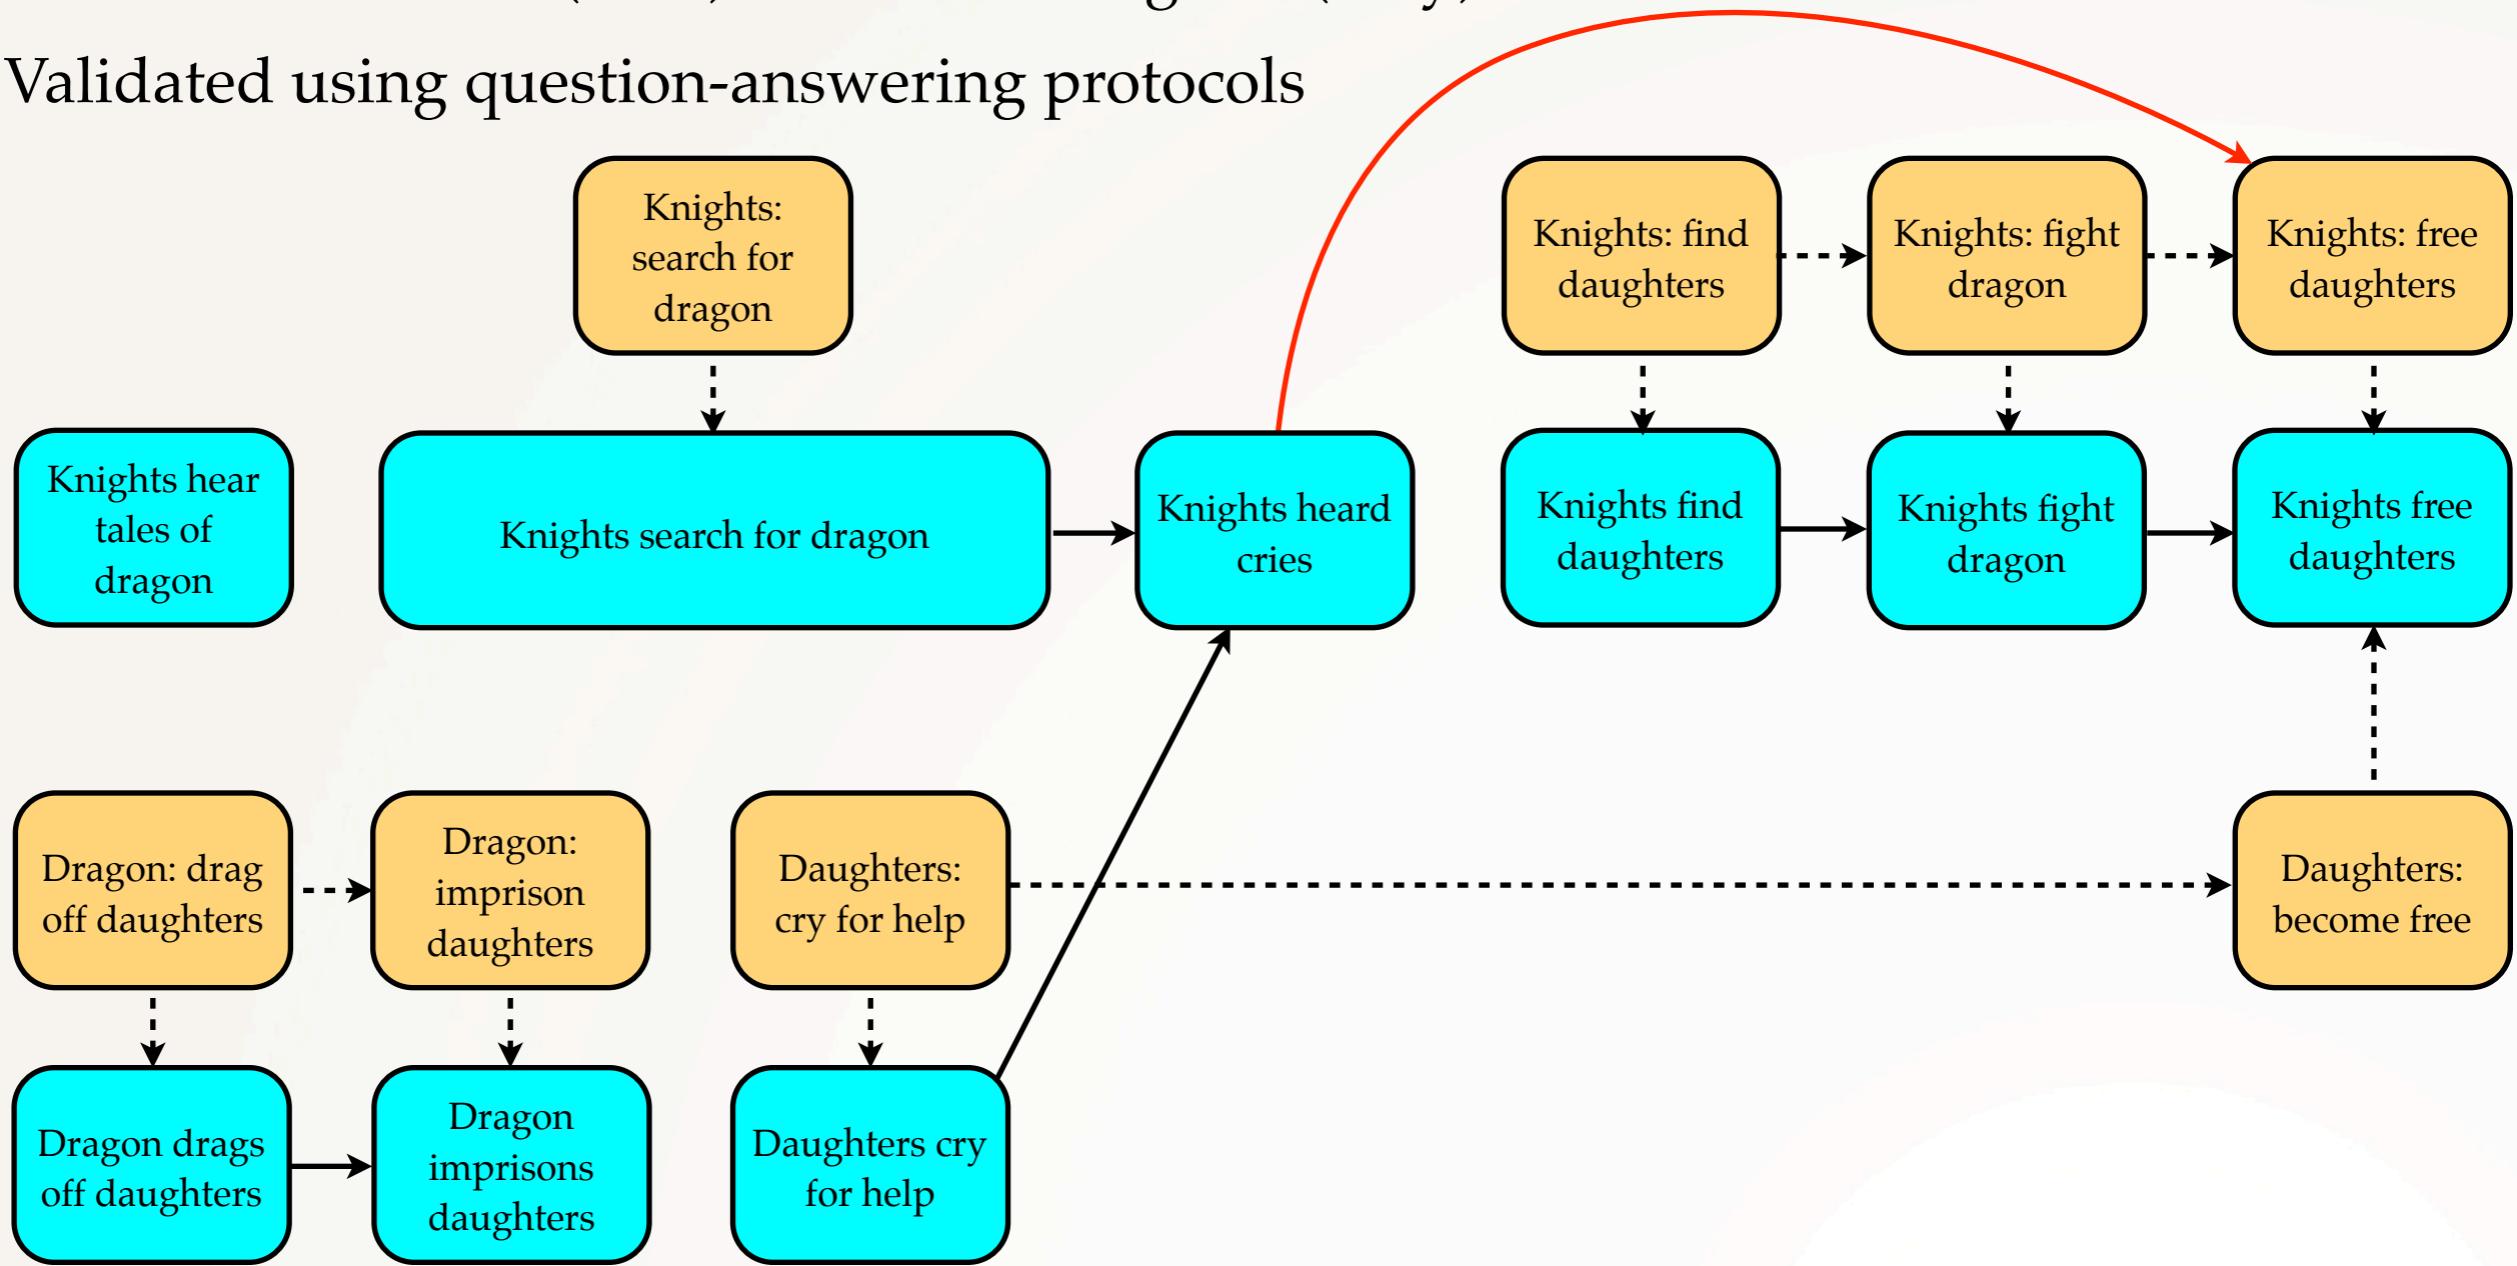


Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.

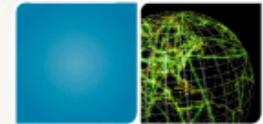


# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

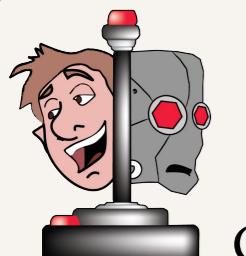
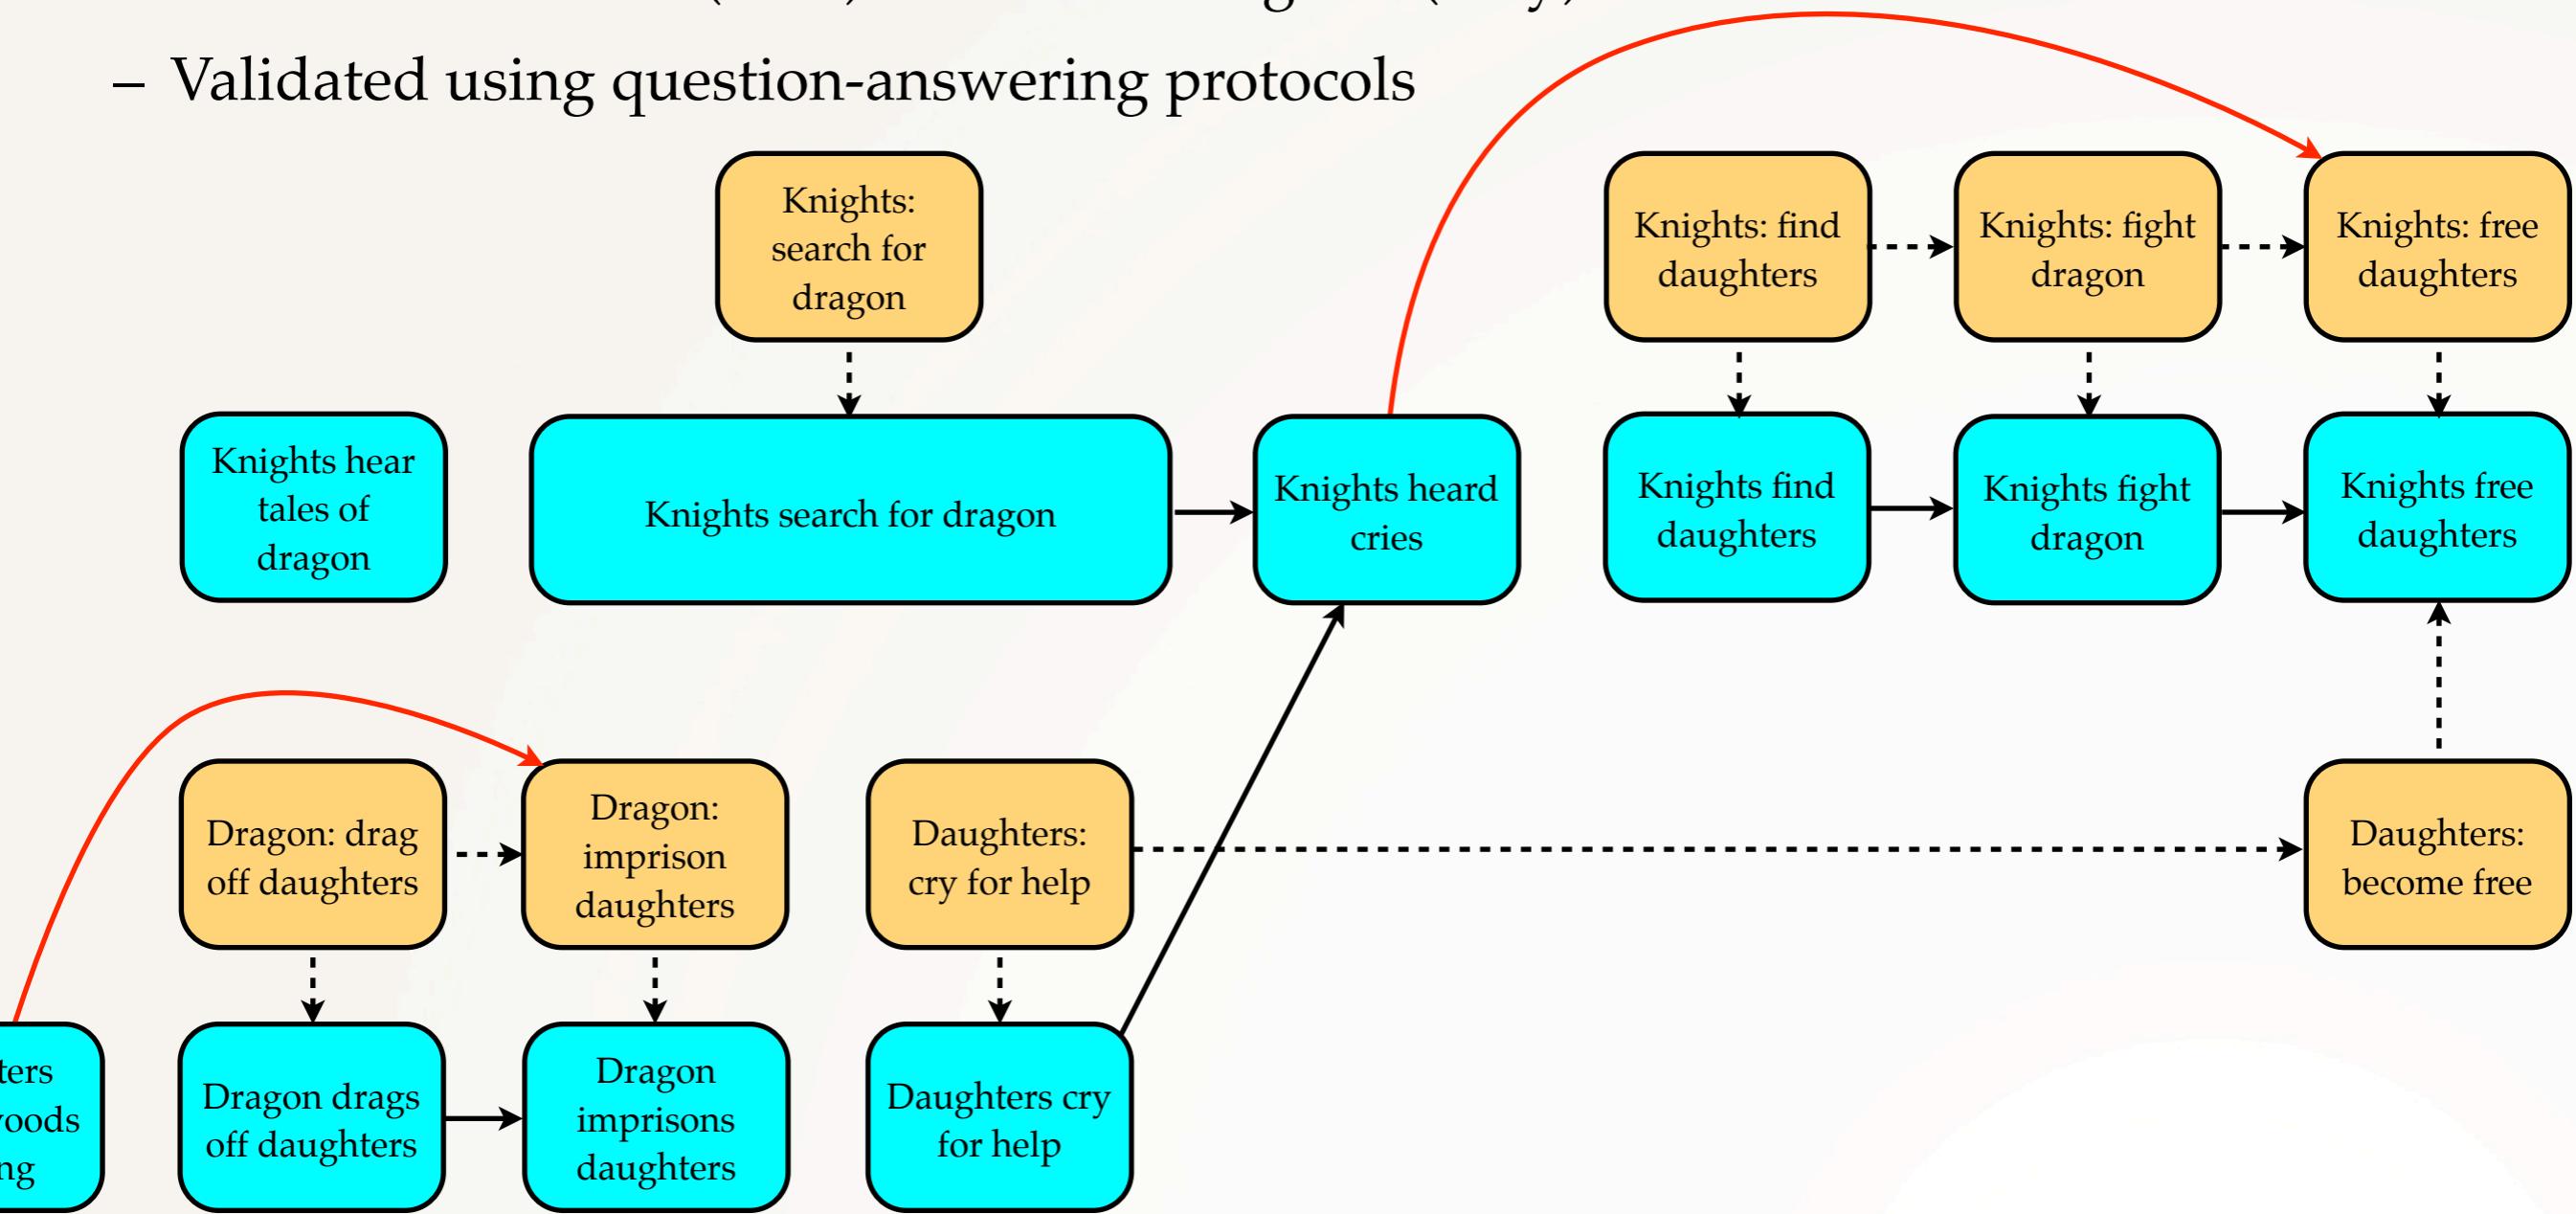


Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.

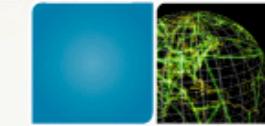


# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

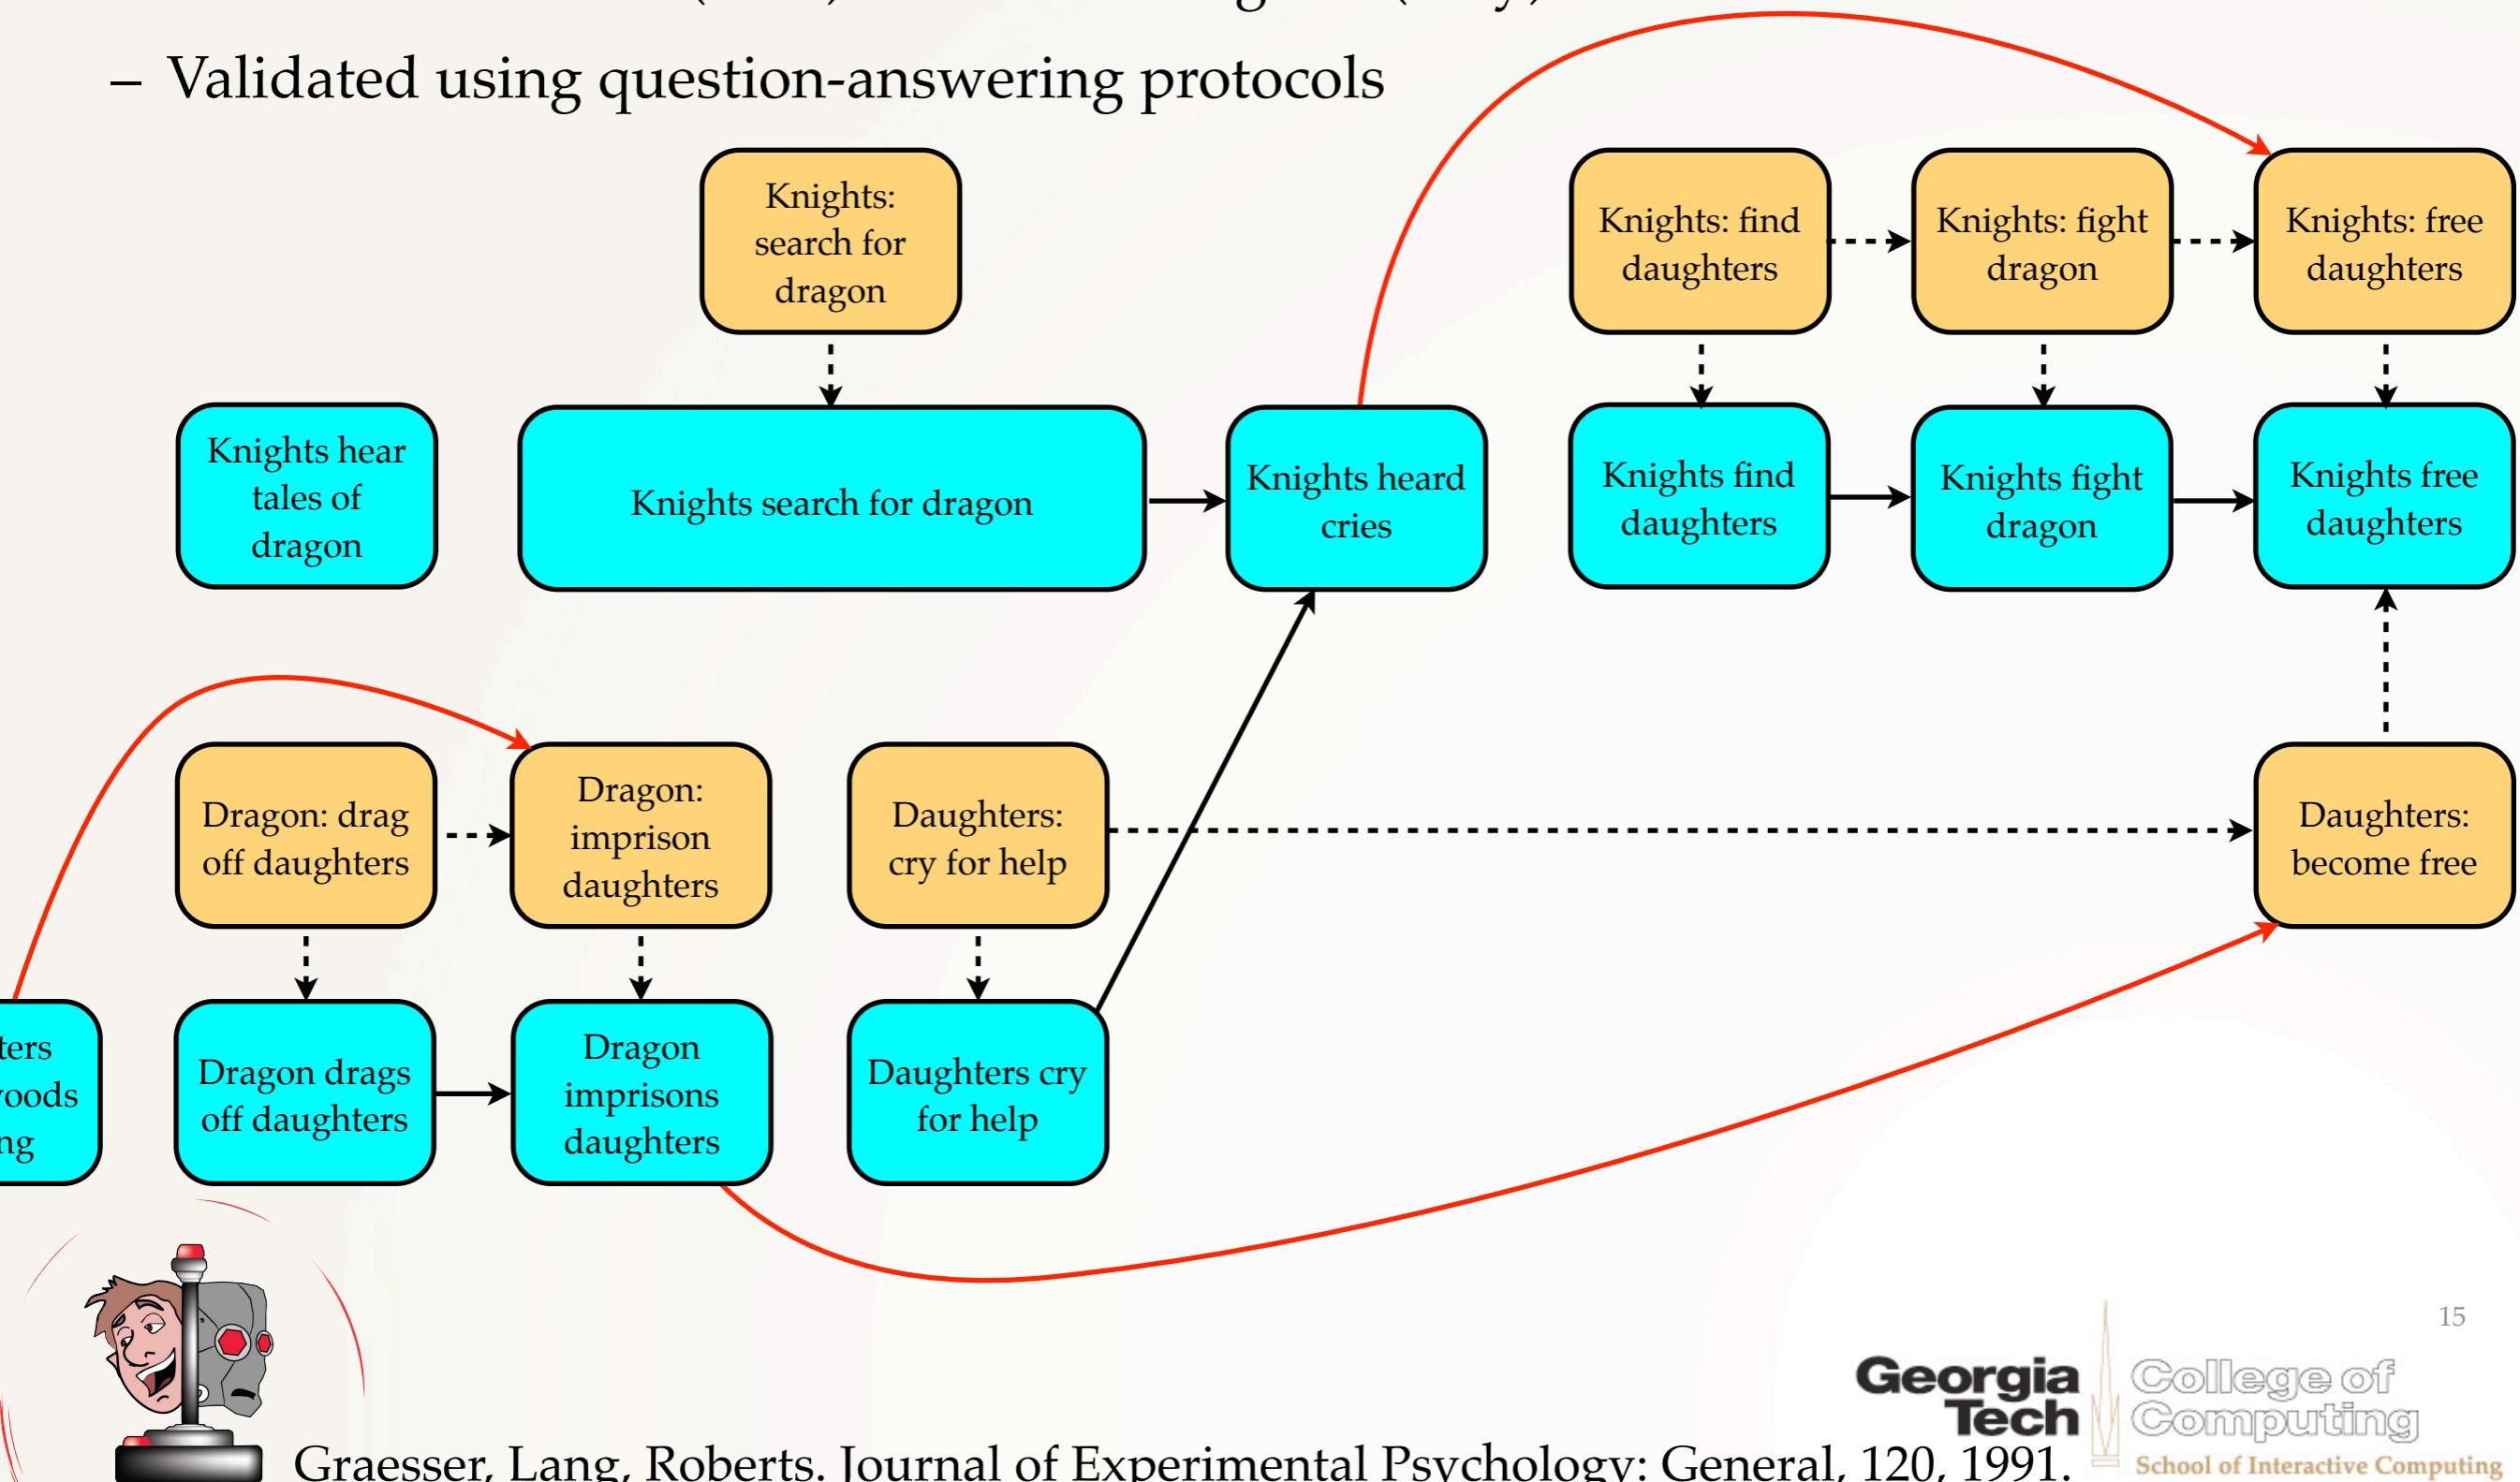


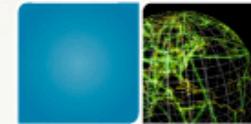
Graesser, Lang, Roberts. Journal of Experimental Psychology: General, 120, 1991.



# Cognitive models of story comprehension

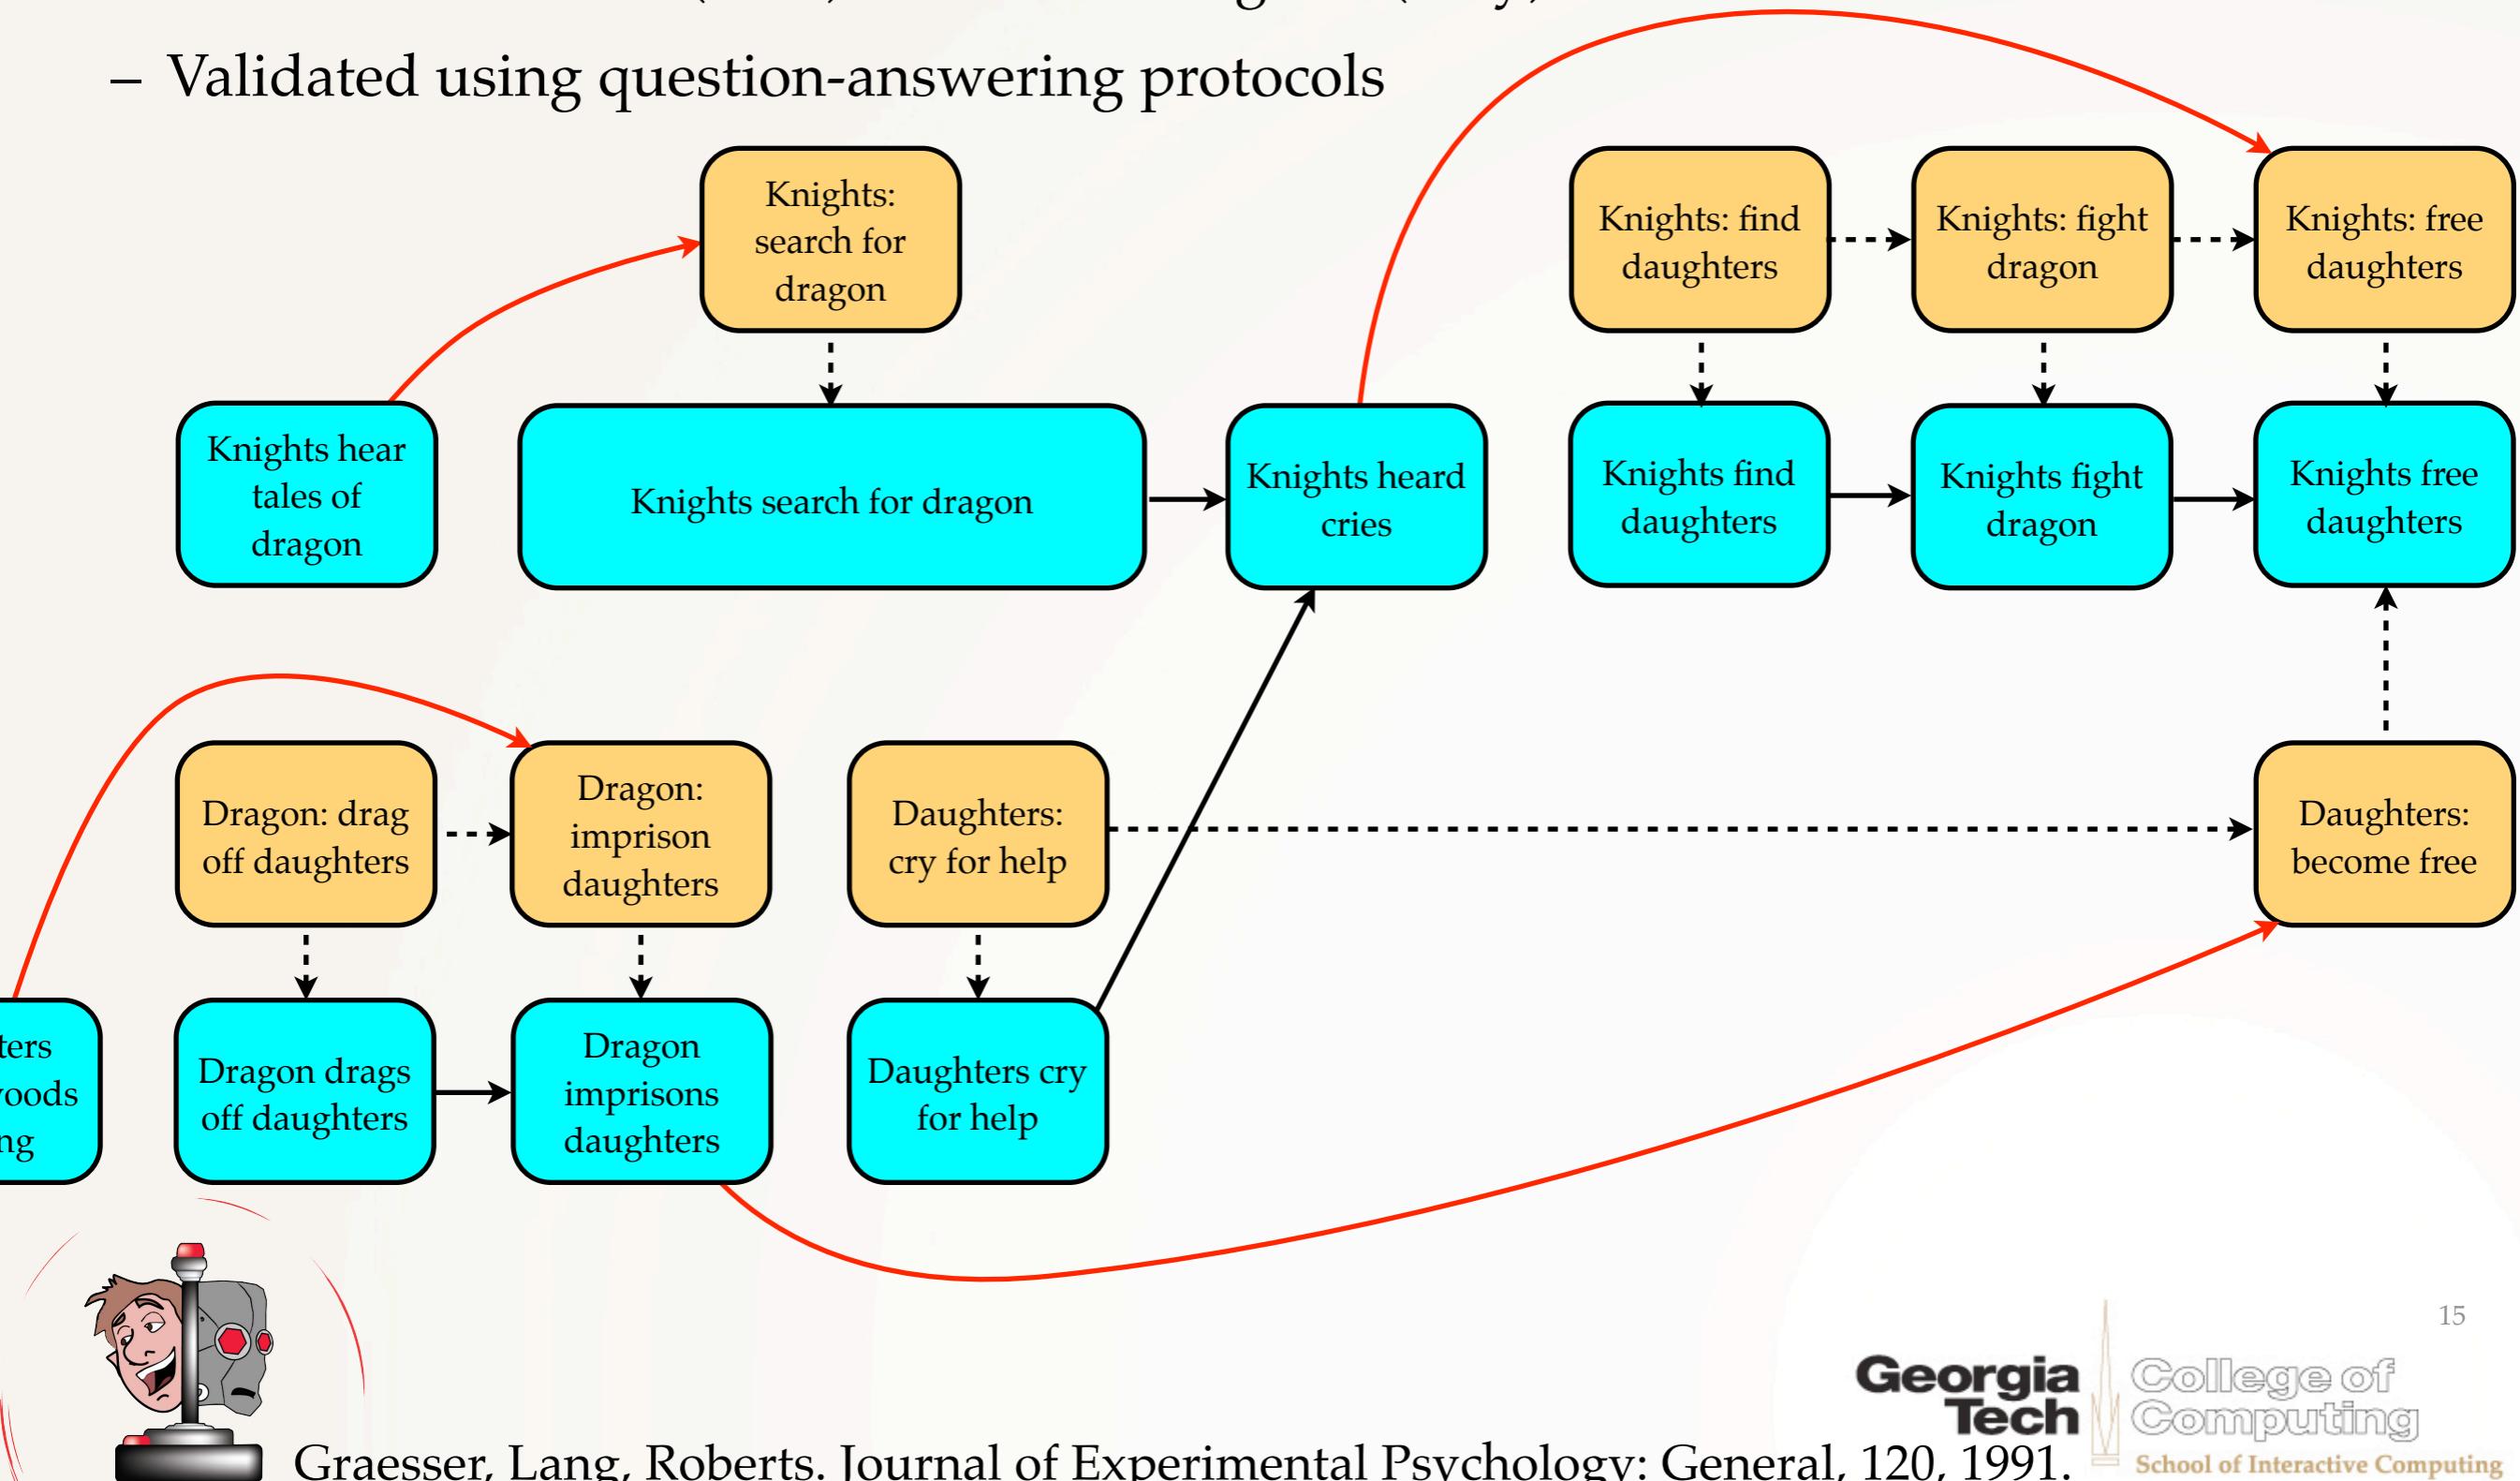
- Causal enablement (how) and character goals (why)
  - Validated using question-answering protocols





# Cognitive models of story comprehension

- Causal enablement (how) and character goals (why)
- Validated using question-answering protocols

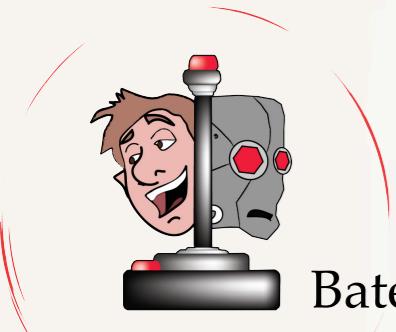


**Story generation reloaded**

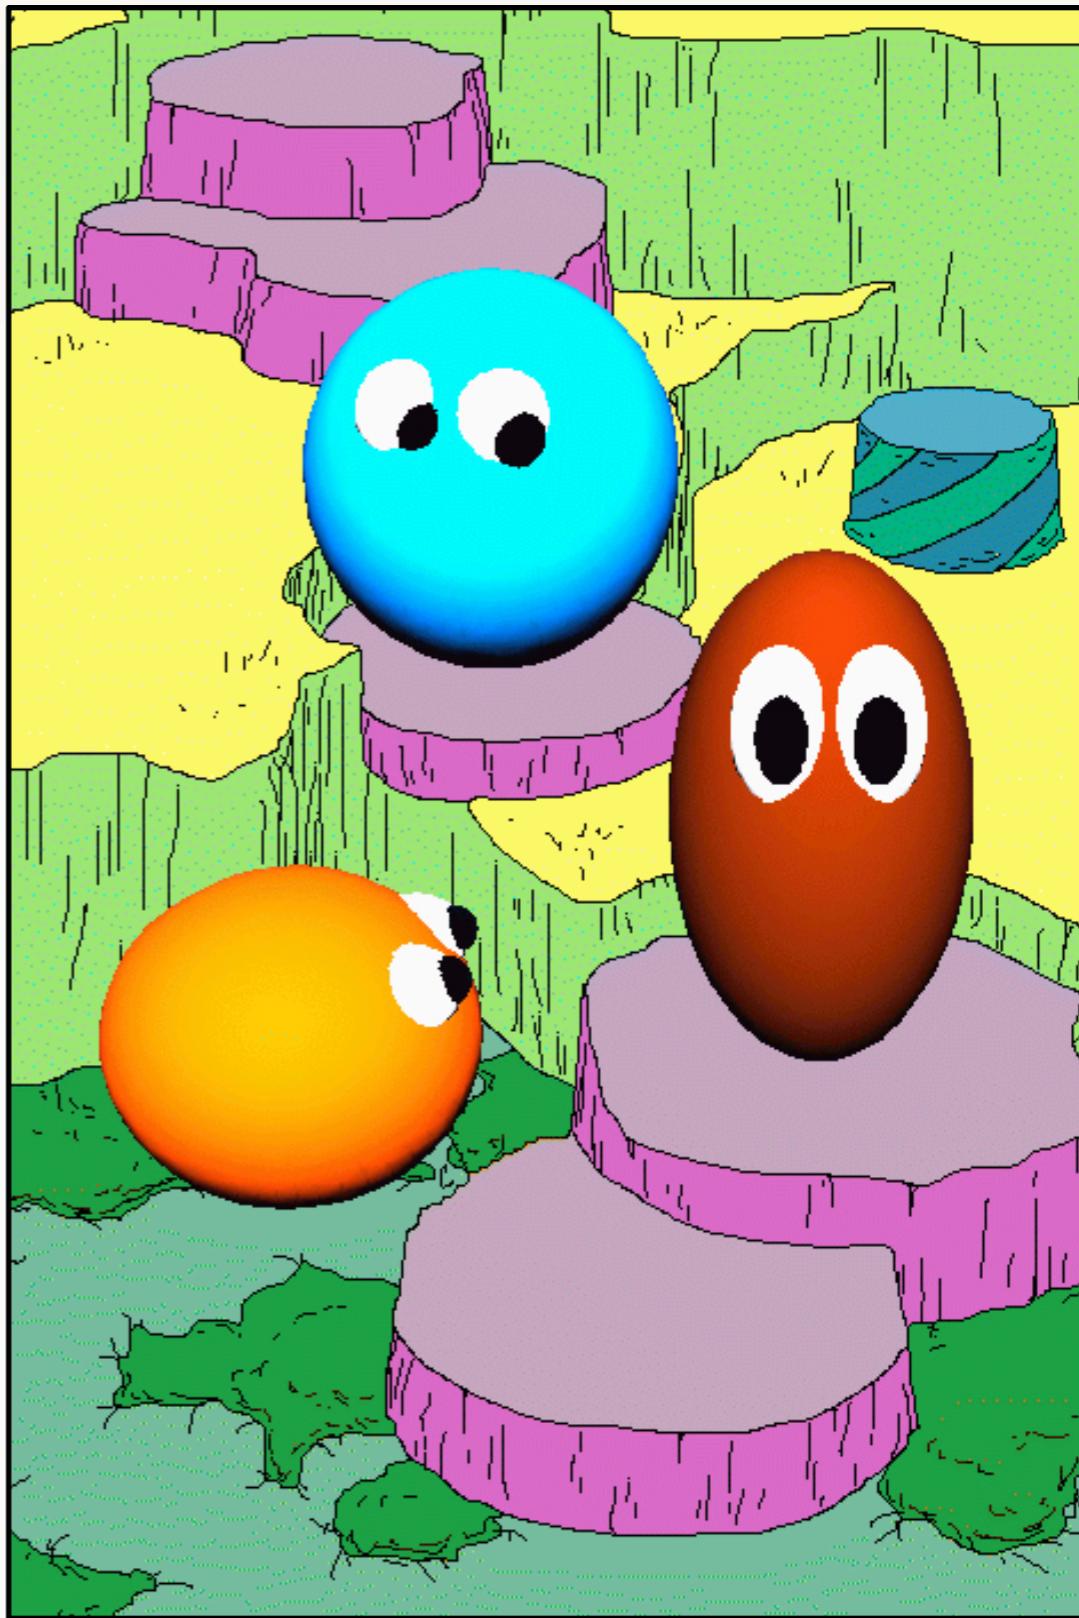


# Computer games

- Sense of immersion
- Interaction with virtual characters
- Stories in games are mostly linear  
(no significant player agency)
- Could humans be participants in  
an interactive drama?
- Can the player join the bad guys?



Bates et al. CMU Oz Project, c.1992



# Interactive narrative



A form of digital entertainment in which the player influences a dramatic storyline through actions





Riedl, Automated Story Director



Mateas & Stern, Facade



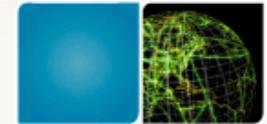
Magerko, Haunt II



Thue, PaSSAGE

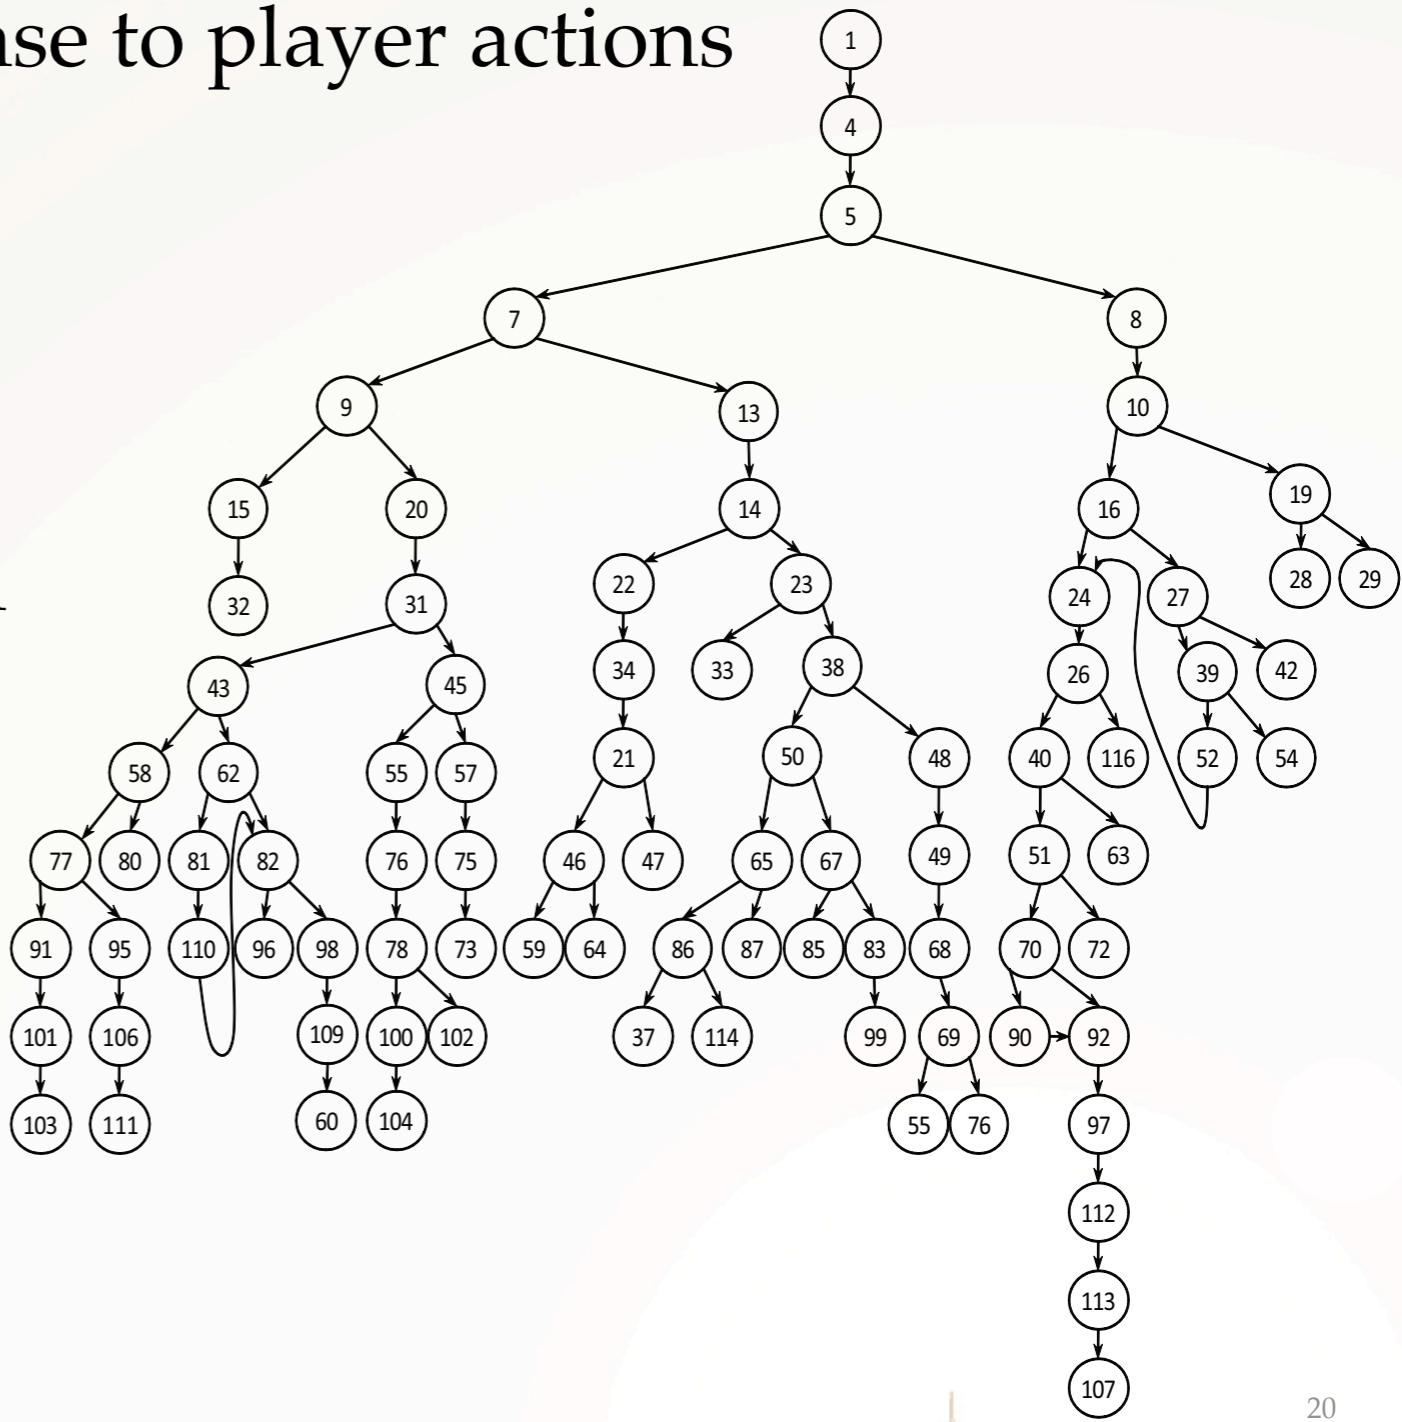


Cavazza, Merchant of Venice



# Behind the screen

- Narrative branches in response to player actions
- Combinatorics of authoring
- Story generator: elaborate on human designer's authorial intent



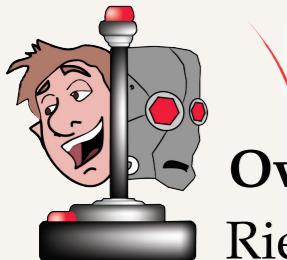
Overview article:

Riedl, Bulitko. AI Magazine, 34(1), 2013.



# Experience (drama) manager

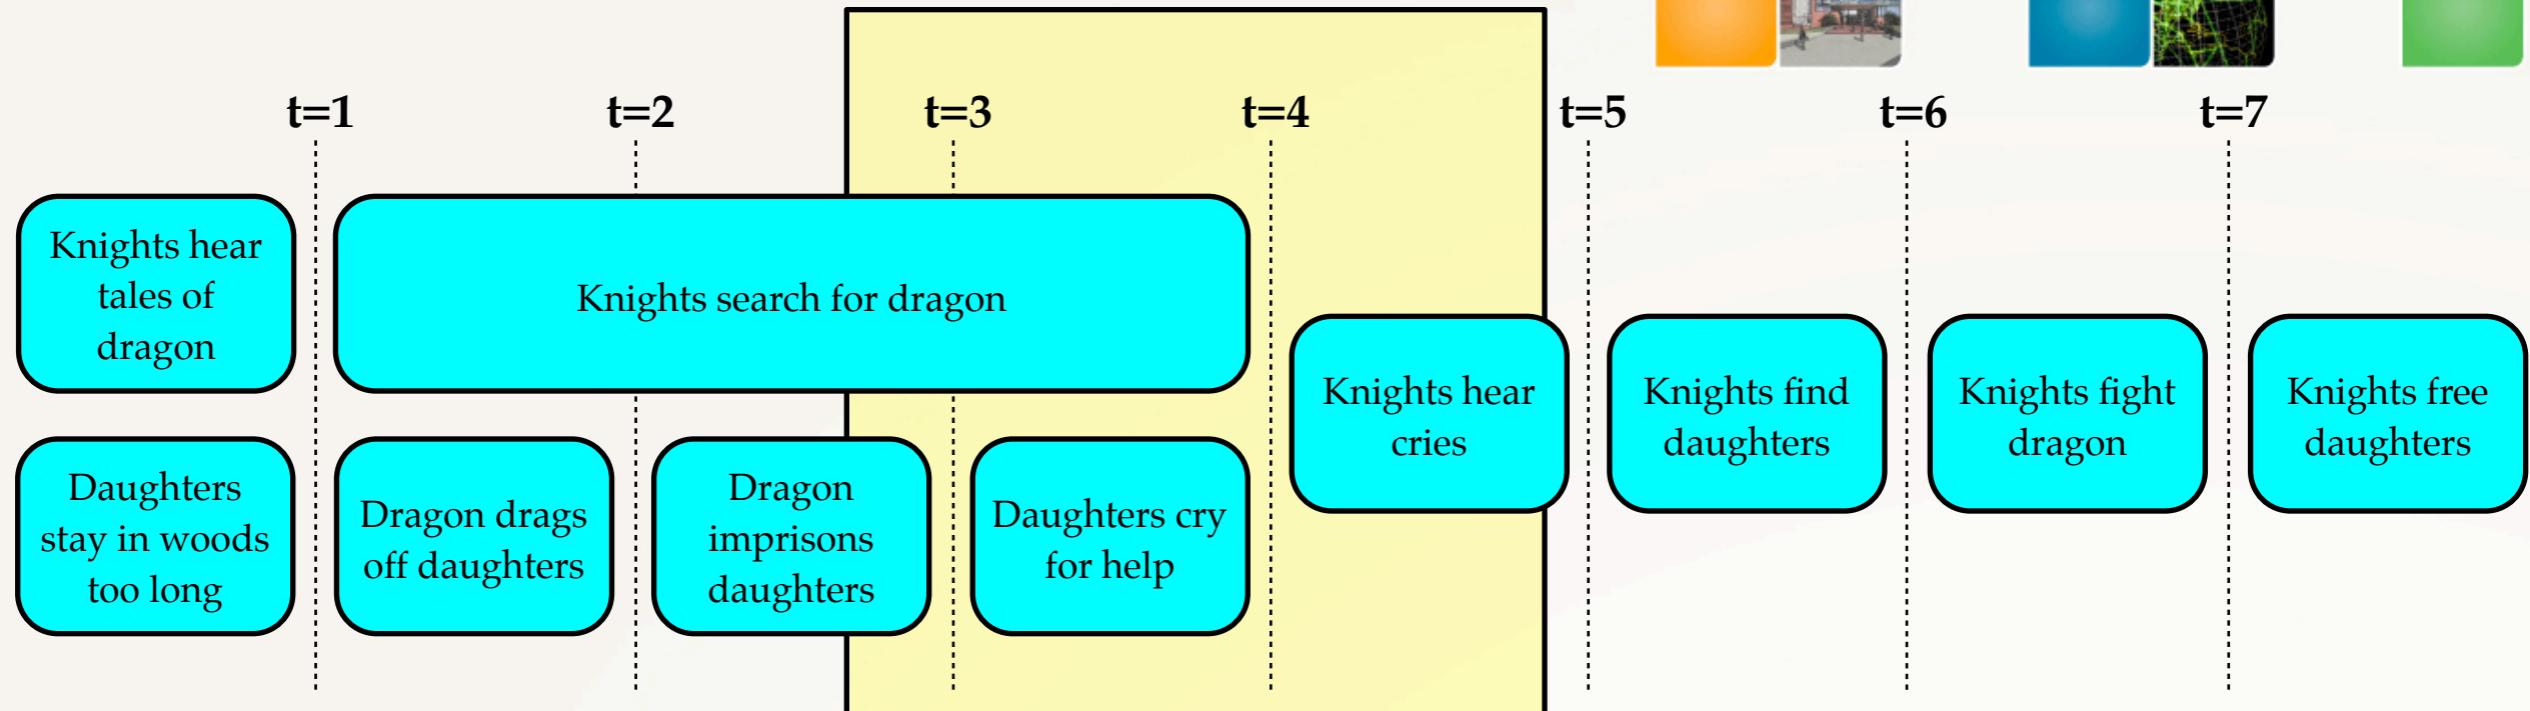
- An omniscient, disembodied agent
- Observes a virtual world
- Intervenes to drive the player's narrative experience forward...
- ...according to some model of quality of experience
- Experience manager is a surrogate for the human storyteller



Overview article:

Riedl, Bulitko. AI Magazine, 34(1), 2013.

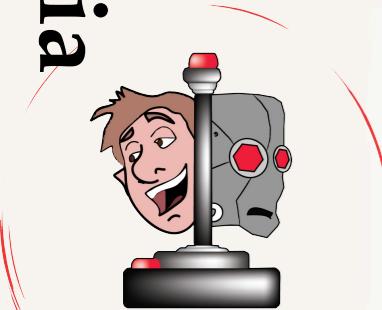
# Fabula



# Sjuzet

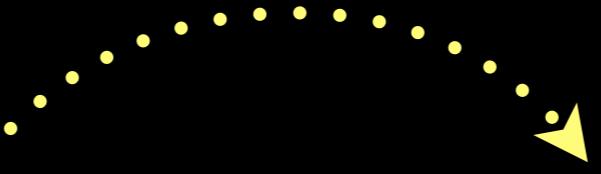
Finish this with pieces of code controlling the virtual avatars, bypassing sjuzet.

# Text/Media



# Story Generation

---

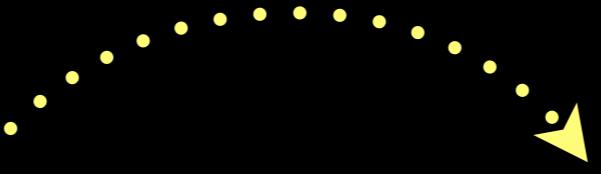


# Interactive Narrative

1. Cognitively inspired planning
2. Crowdsourcing

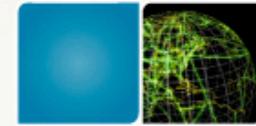
# Story Generation

---



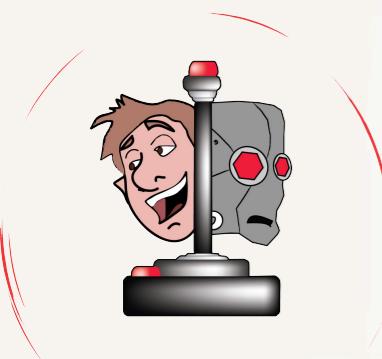
# Interactive Narrative

1. Cognitively inspired planning
2. Crowdsourcing



# Cognitively inspired story generation

- Construct a sequence of events that can be told as a story---that meets some set of given criteria for success
  - Need computational model of “good”
  - Dramatic arc? Emotional impact? Makes a point? Sense of engagement?
- Two nearly-universal properties of story:
  - Causal progression
  - Character intentionality

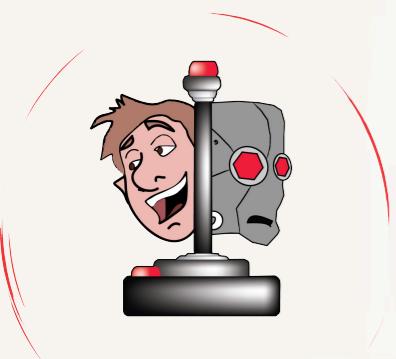
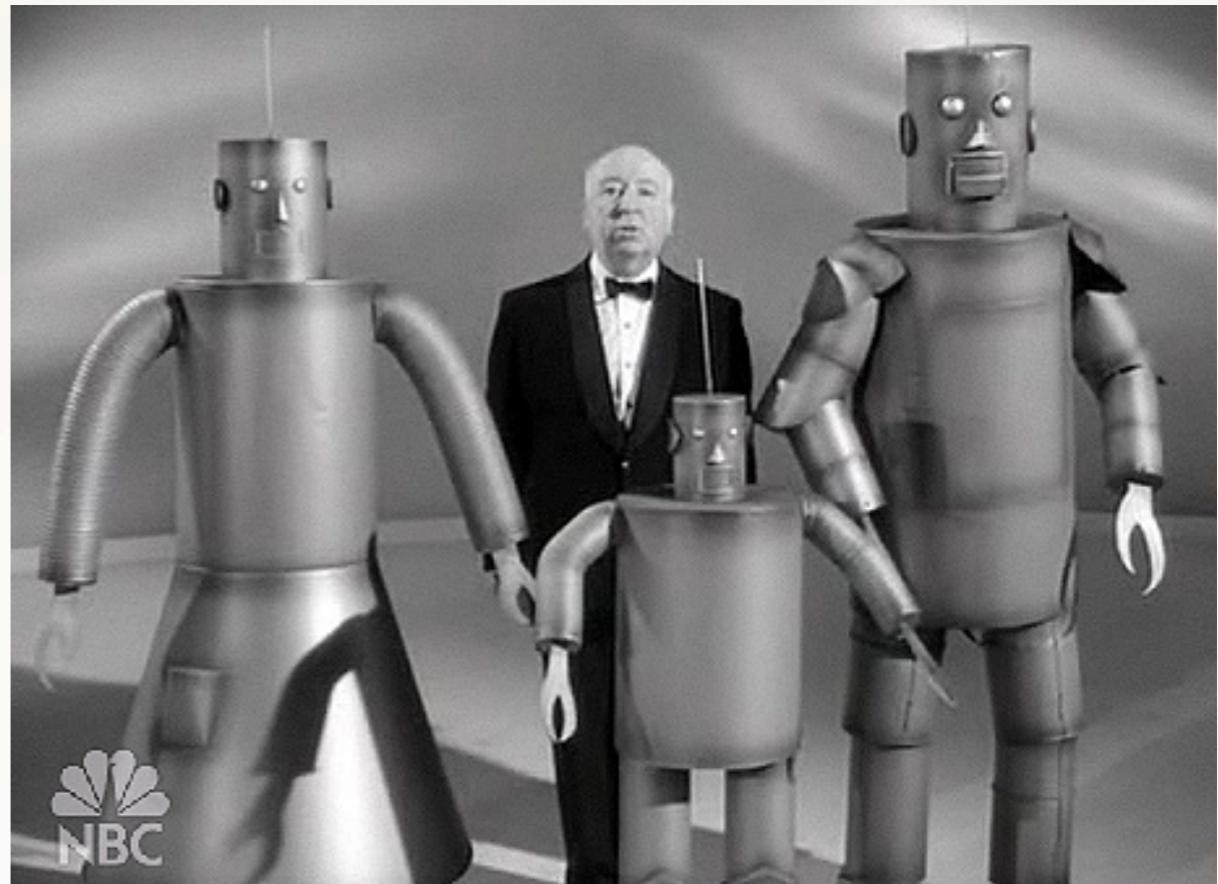


25

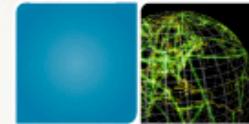


# Computer as author

- Creative writing is a problem-solving activity
- Author goals vs. character goals
- Model: plan out the events that should occur in the narrative

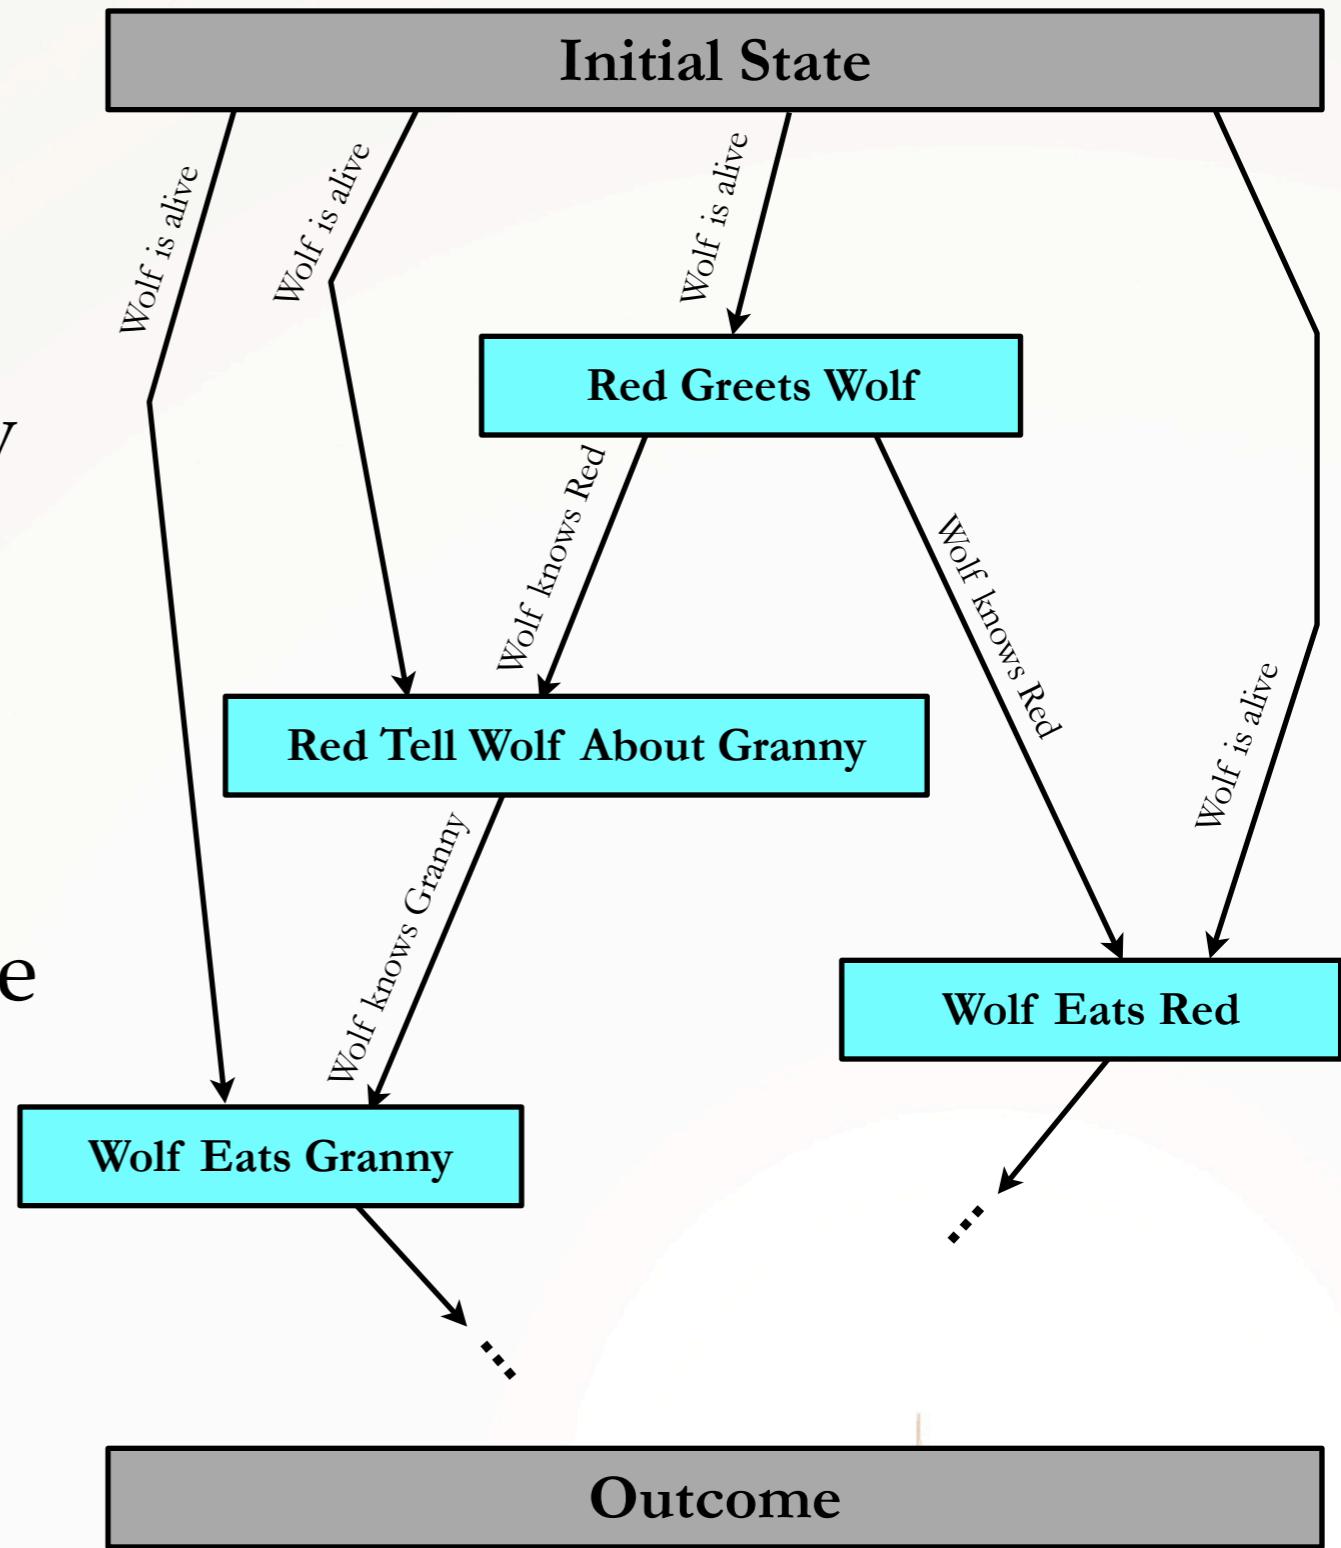


26



# Narratives as plans

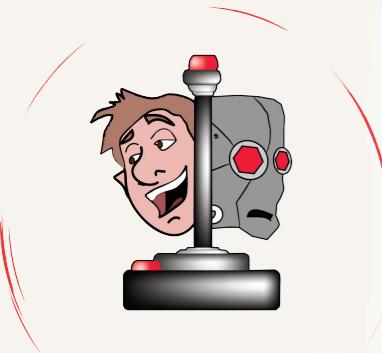
- Partial-order plan is a good representation of plot
  - Action, temporality, causality
  - Functionally equivalent to story comprehension models
- Planning: find a sound and coherent sequence of actions that transforms the initial state into one in which the goal situation holds





# Planning stories

- But, is planning a good model of story creation?



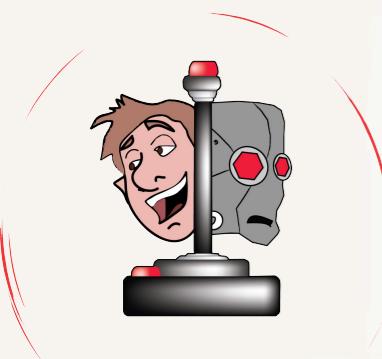
28



# Planning stories

- But, is planning a good model of story creation?

- Narrative Planning:
  - Multiple characters
  - Goal state describes outcome of the story
  - Outcome is not necessarily intended by any characters



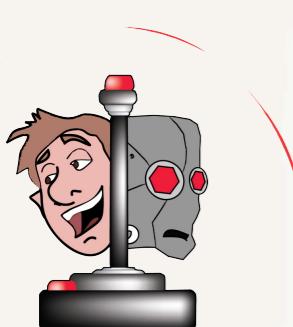


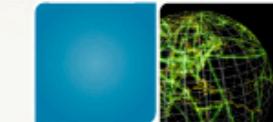
# Planning stories

- But, is planning a good model of story creation?

- Narrative Planning:

- Multiple characters
- Goal state describes outcome of the story
- Outcome is not necessarily intended by any characters





# Planning stories

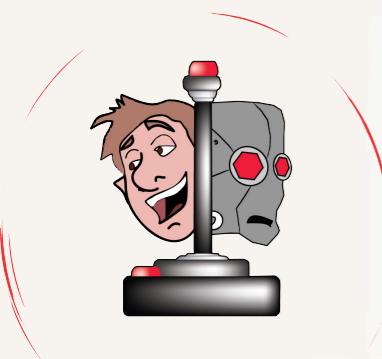
- But, is planning a good model of story creation?

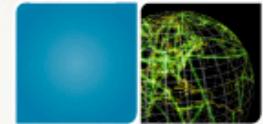
- Narrative Planning:

- Multiple characters
- Goal state describes outcome of the story
- Outcome is not necessarily intended by any characters



- Augment planning algorithm to reason about author goals and character goals



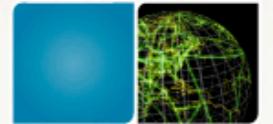


# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals



Riedl & Young. Journal of Artificial Intelligence Research, 39, 2010.



# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression

has (Villain, \$\$\$)

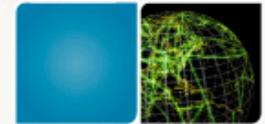
Outcome

- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals



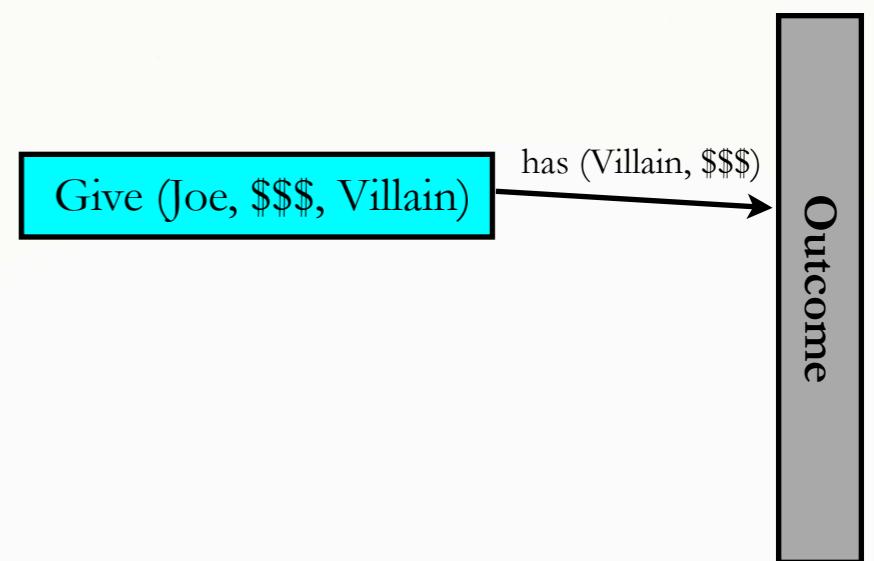
Riedl & Young. Journal of Artificial Intelligence Research, 39, 2010.

29



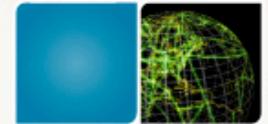
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



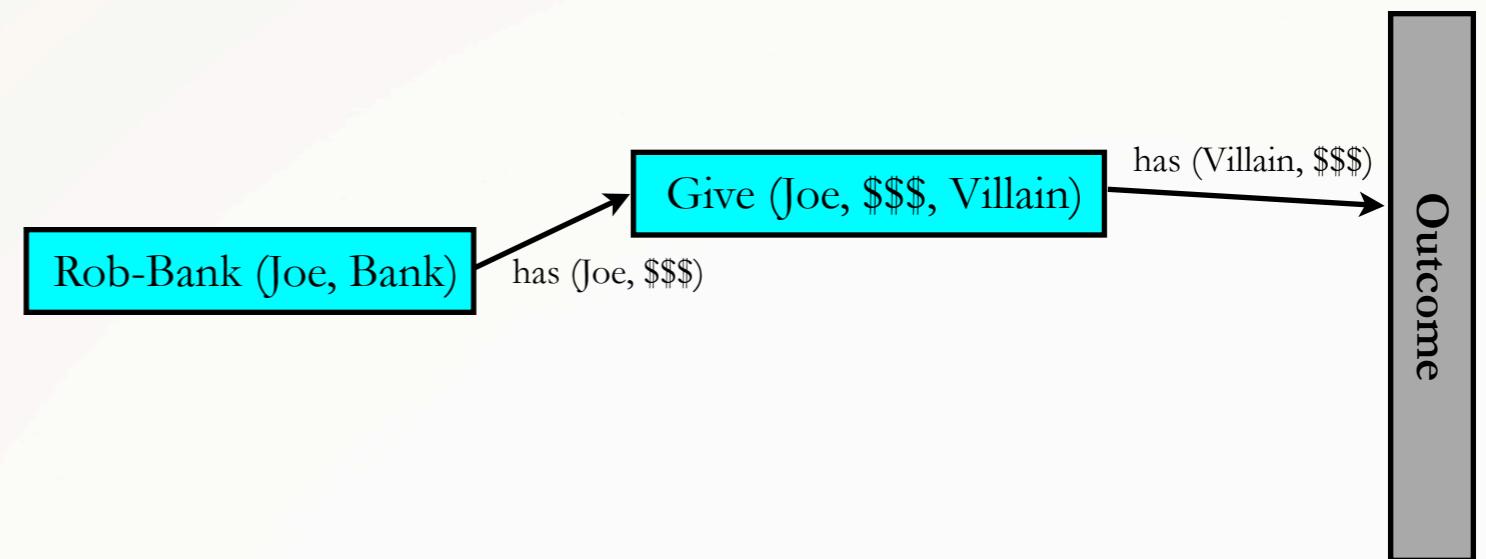
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals





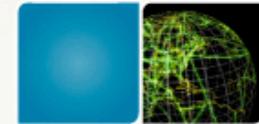
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



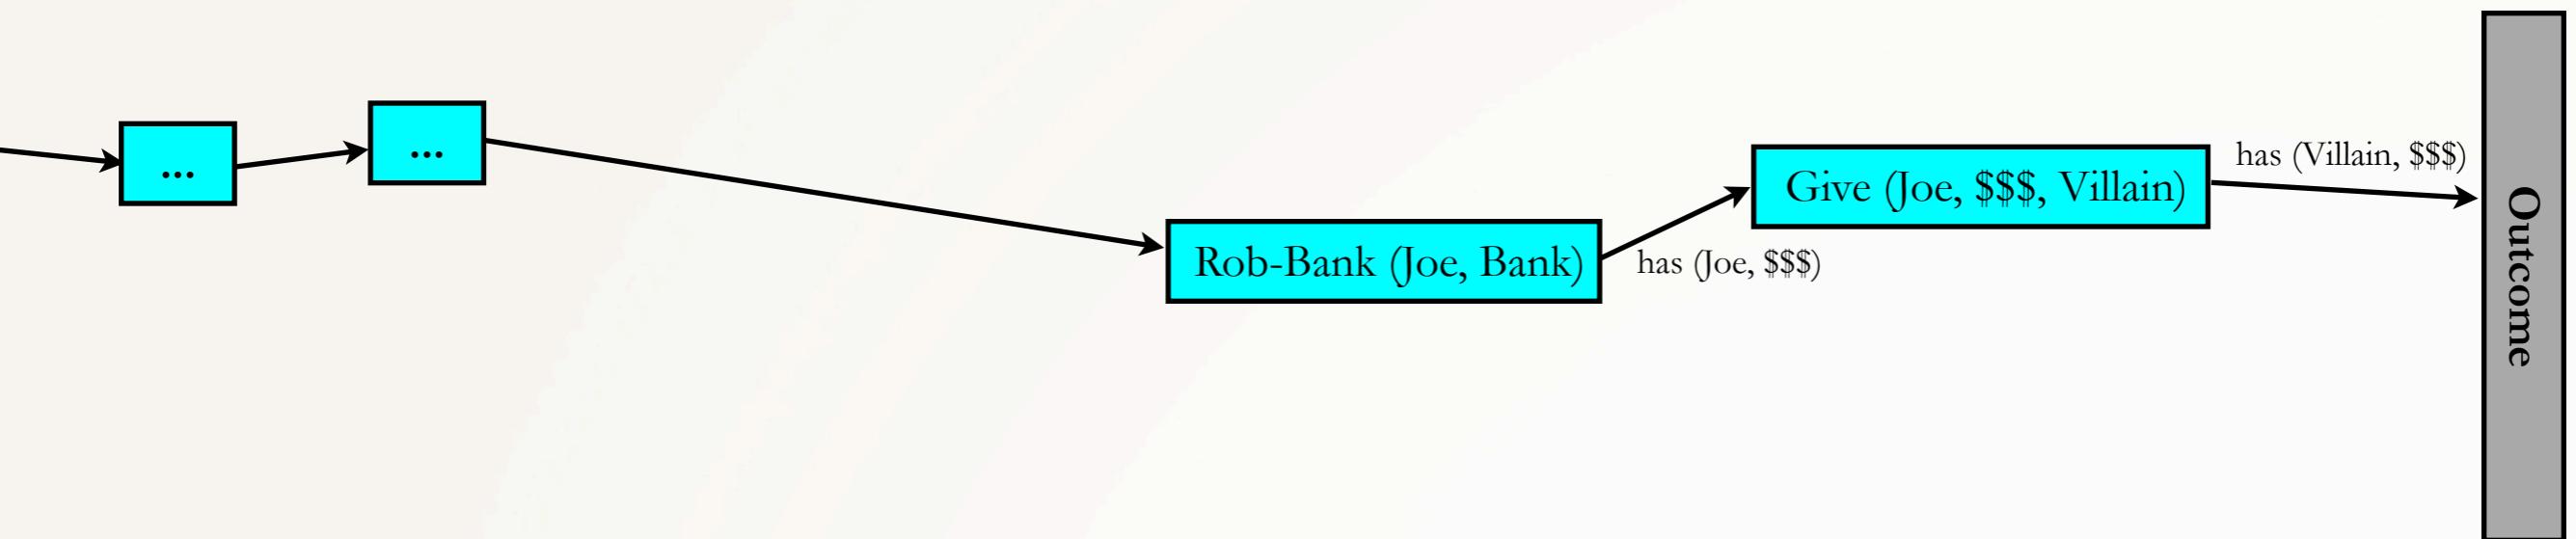
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals



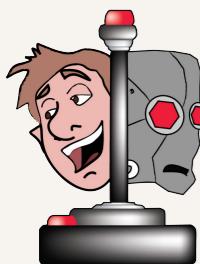


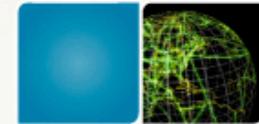
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



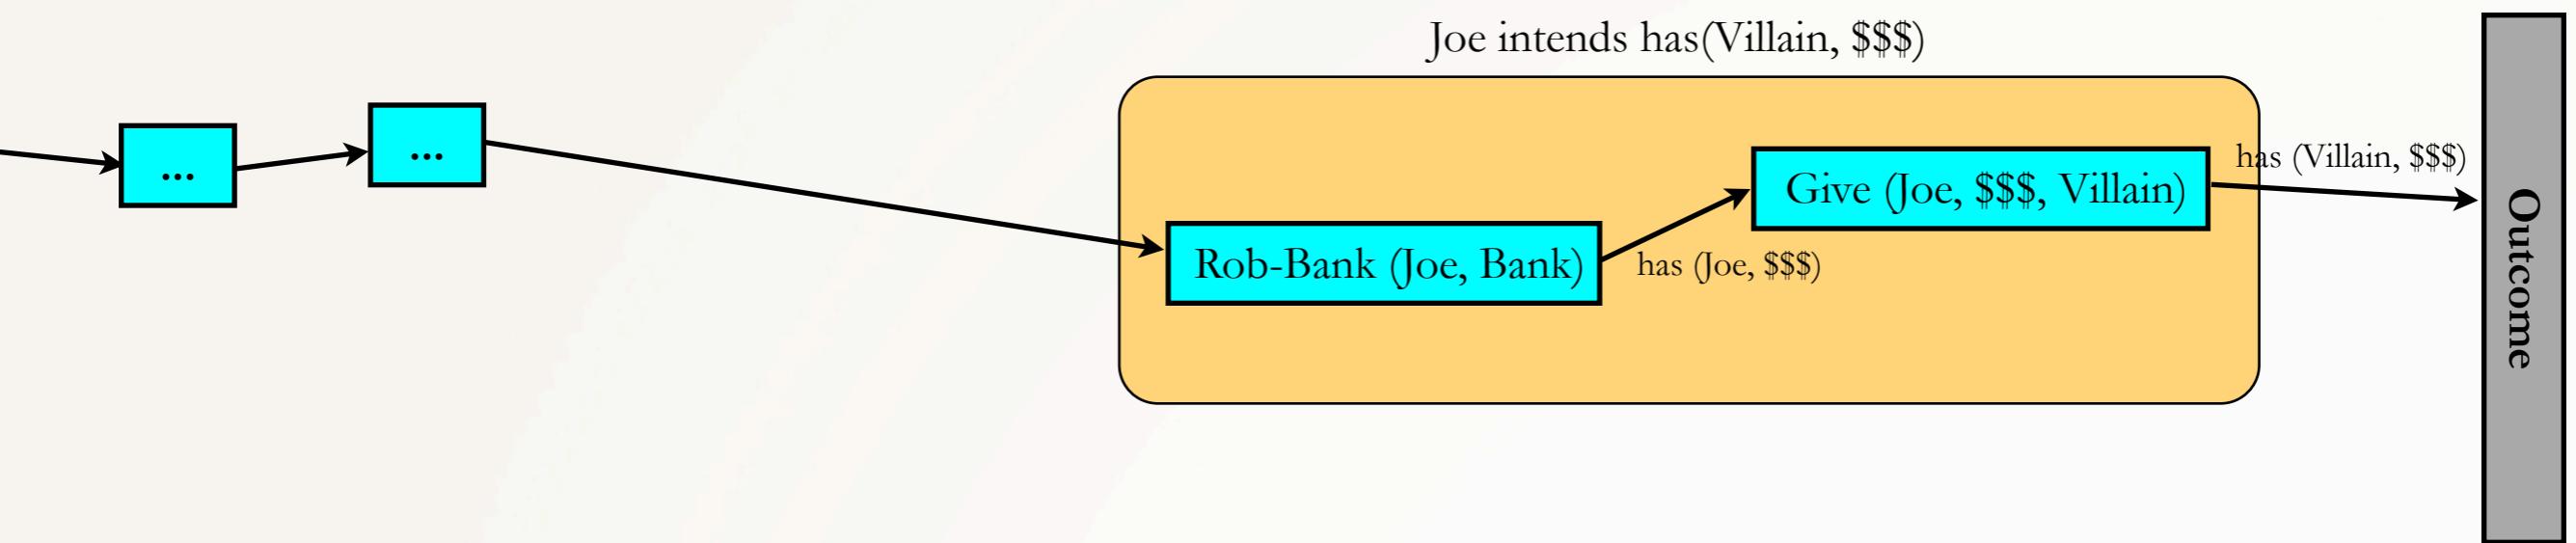
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals





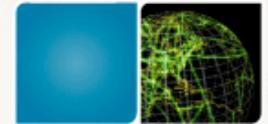
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



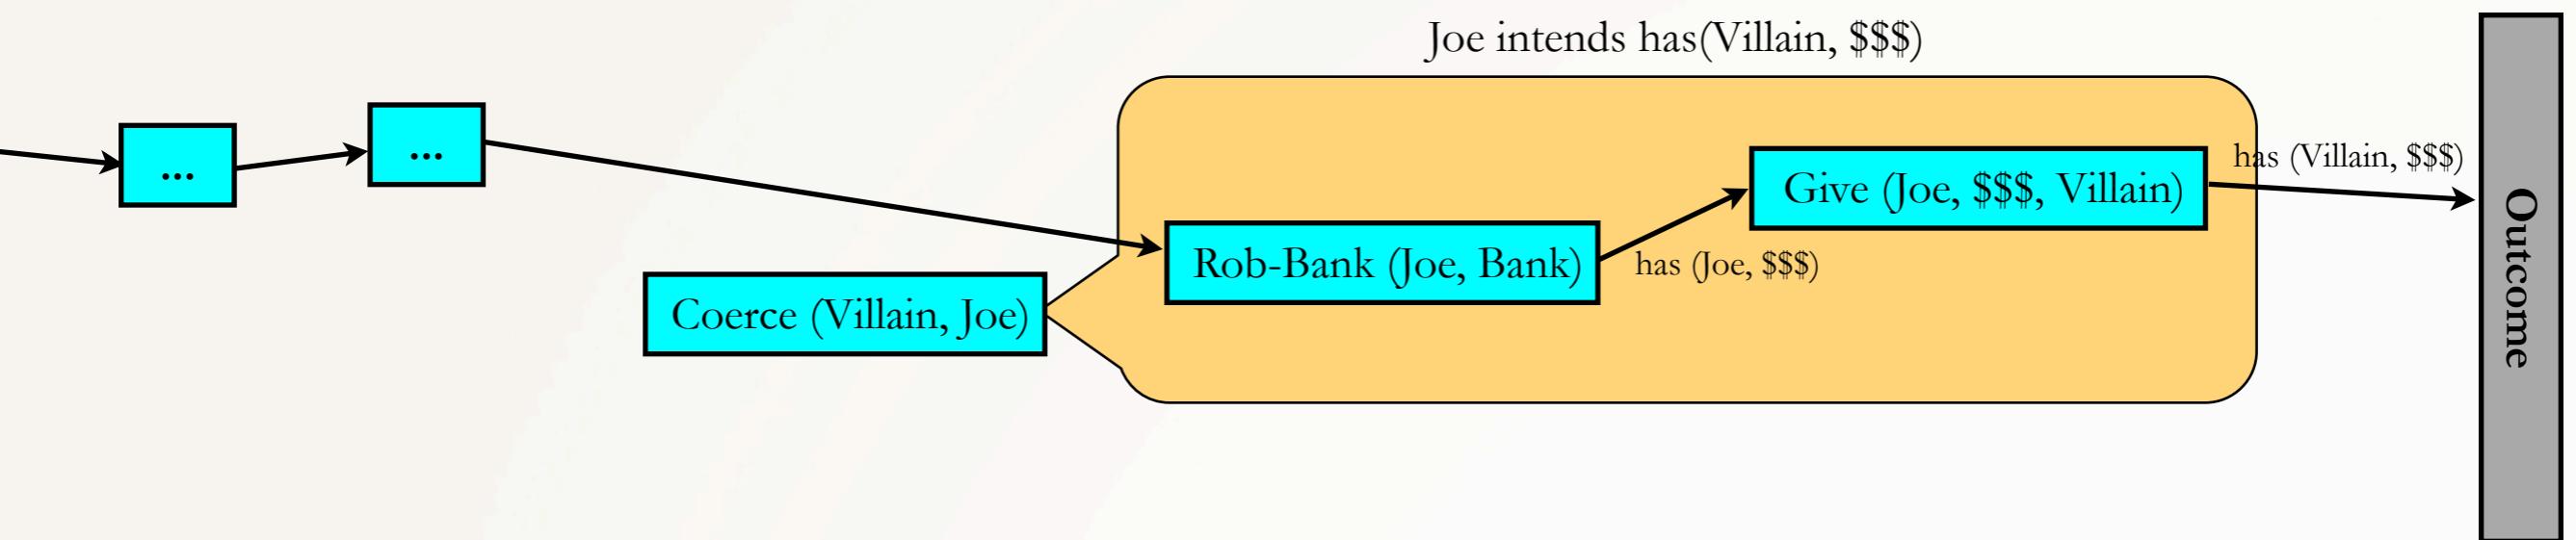
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals



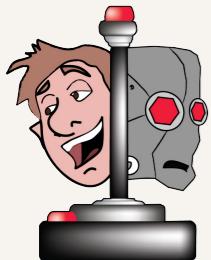


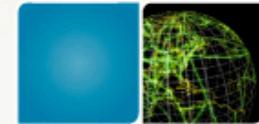
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



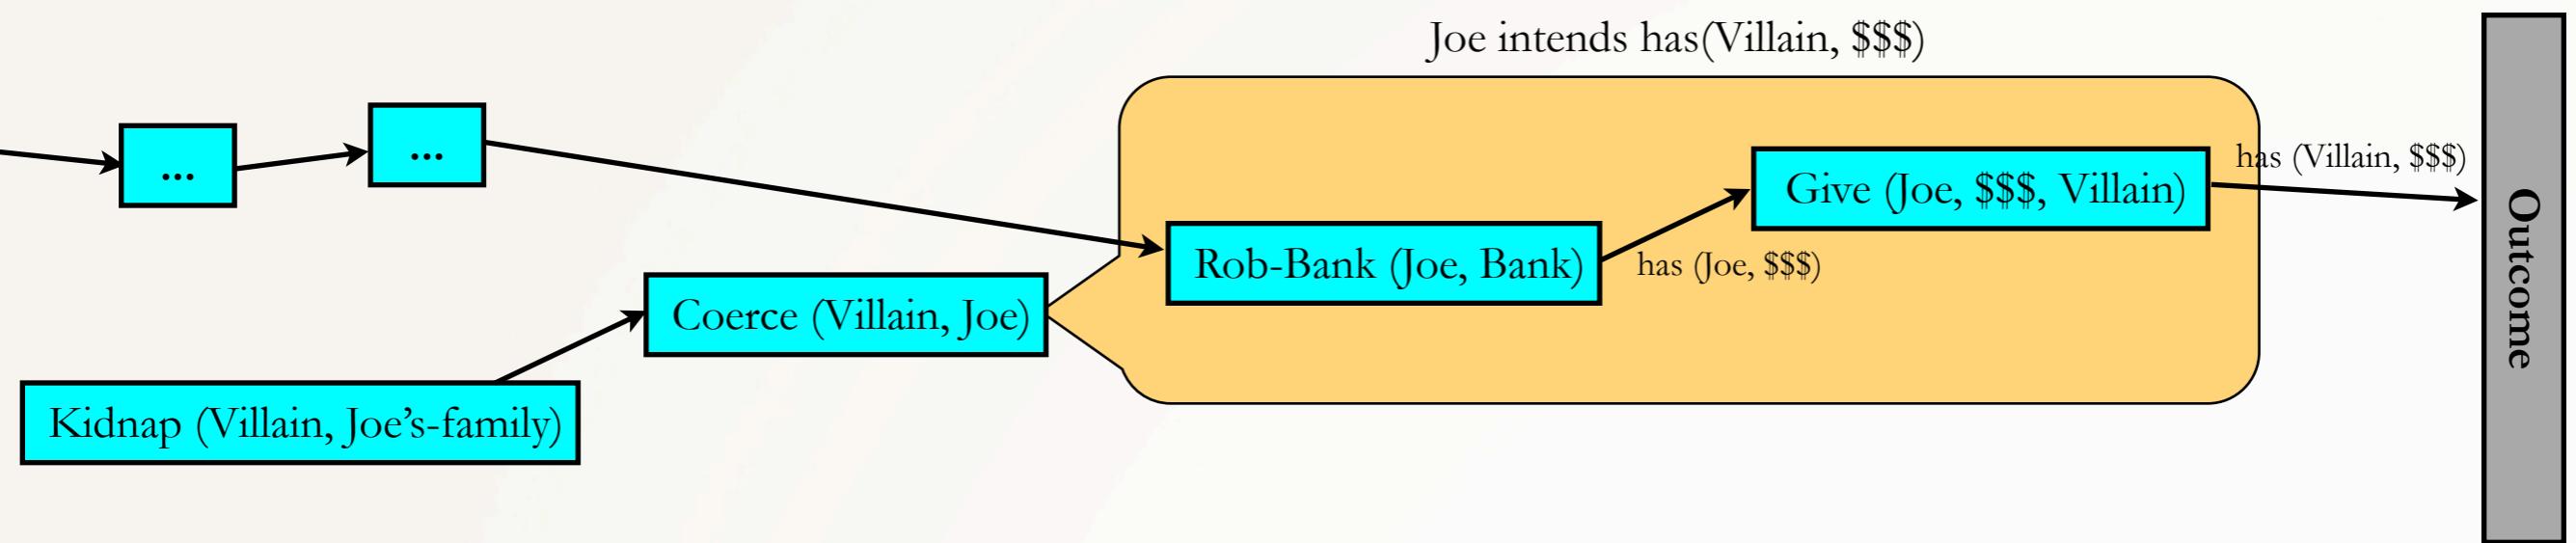
- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals





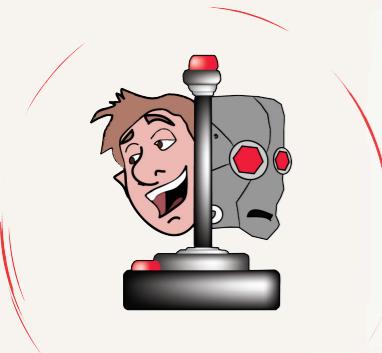
# Fabulist

- Conventional causal dependency planning
  - Provides logical causal progression



- Reasoning about character intentions
  - Use a cognitive model to determine whether characters appear intentional and revise the plan otherwise
  - Insert actions that explain character goals

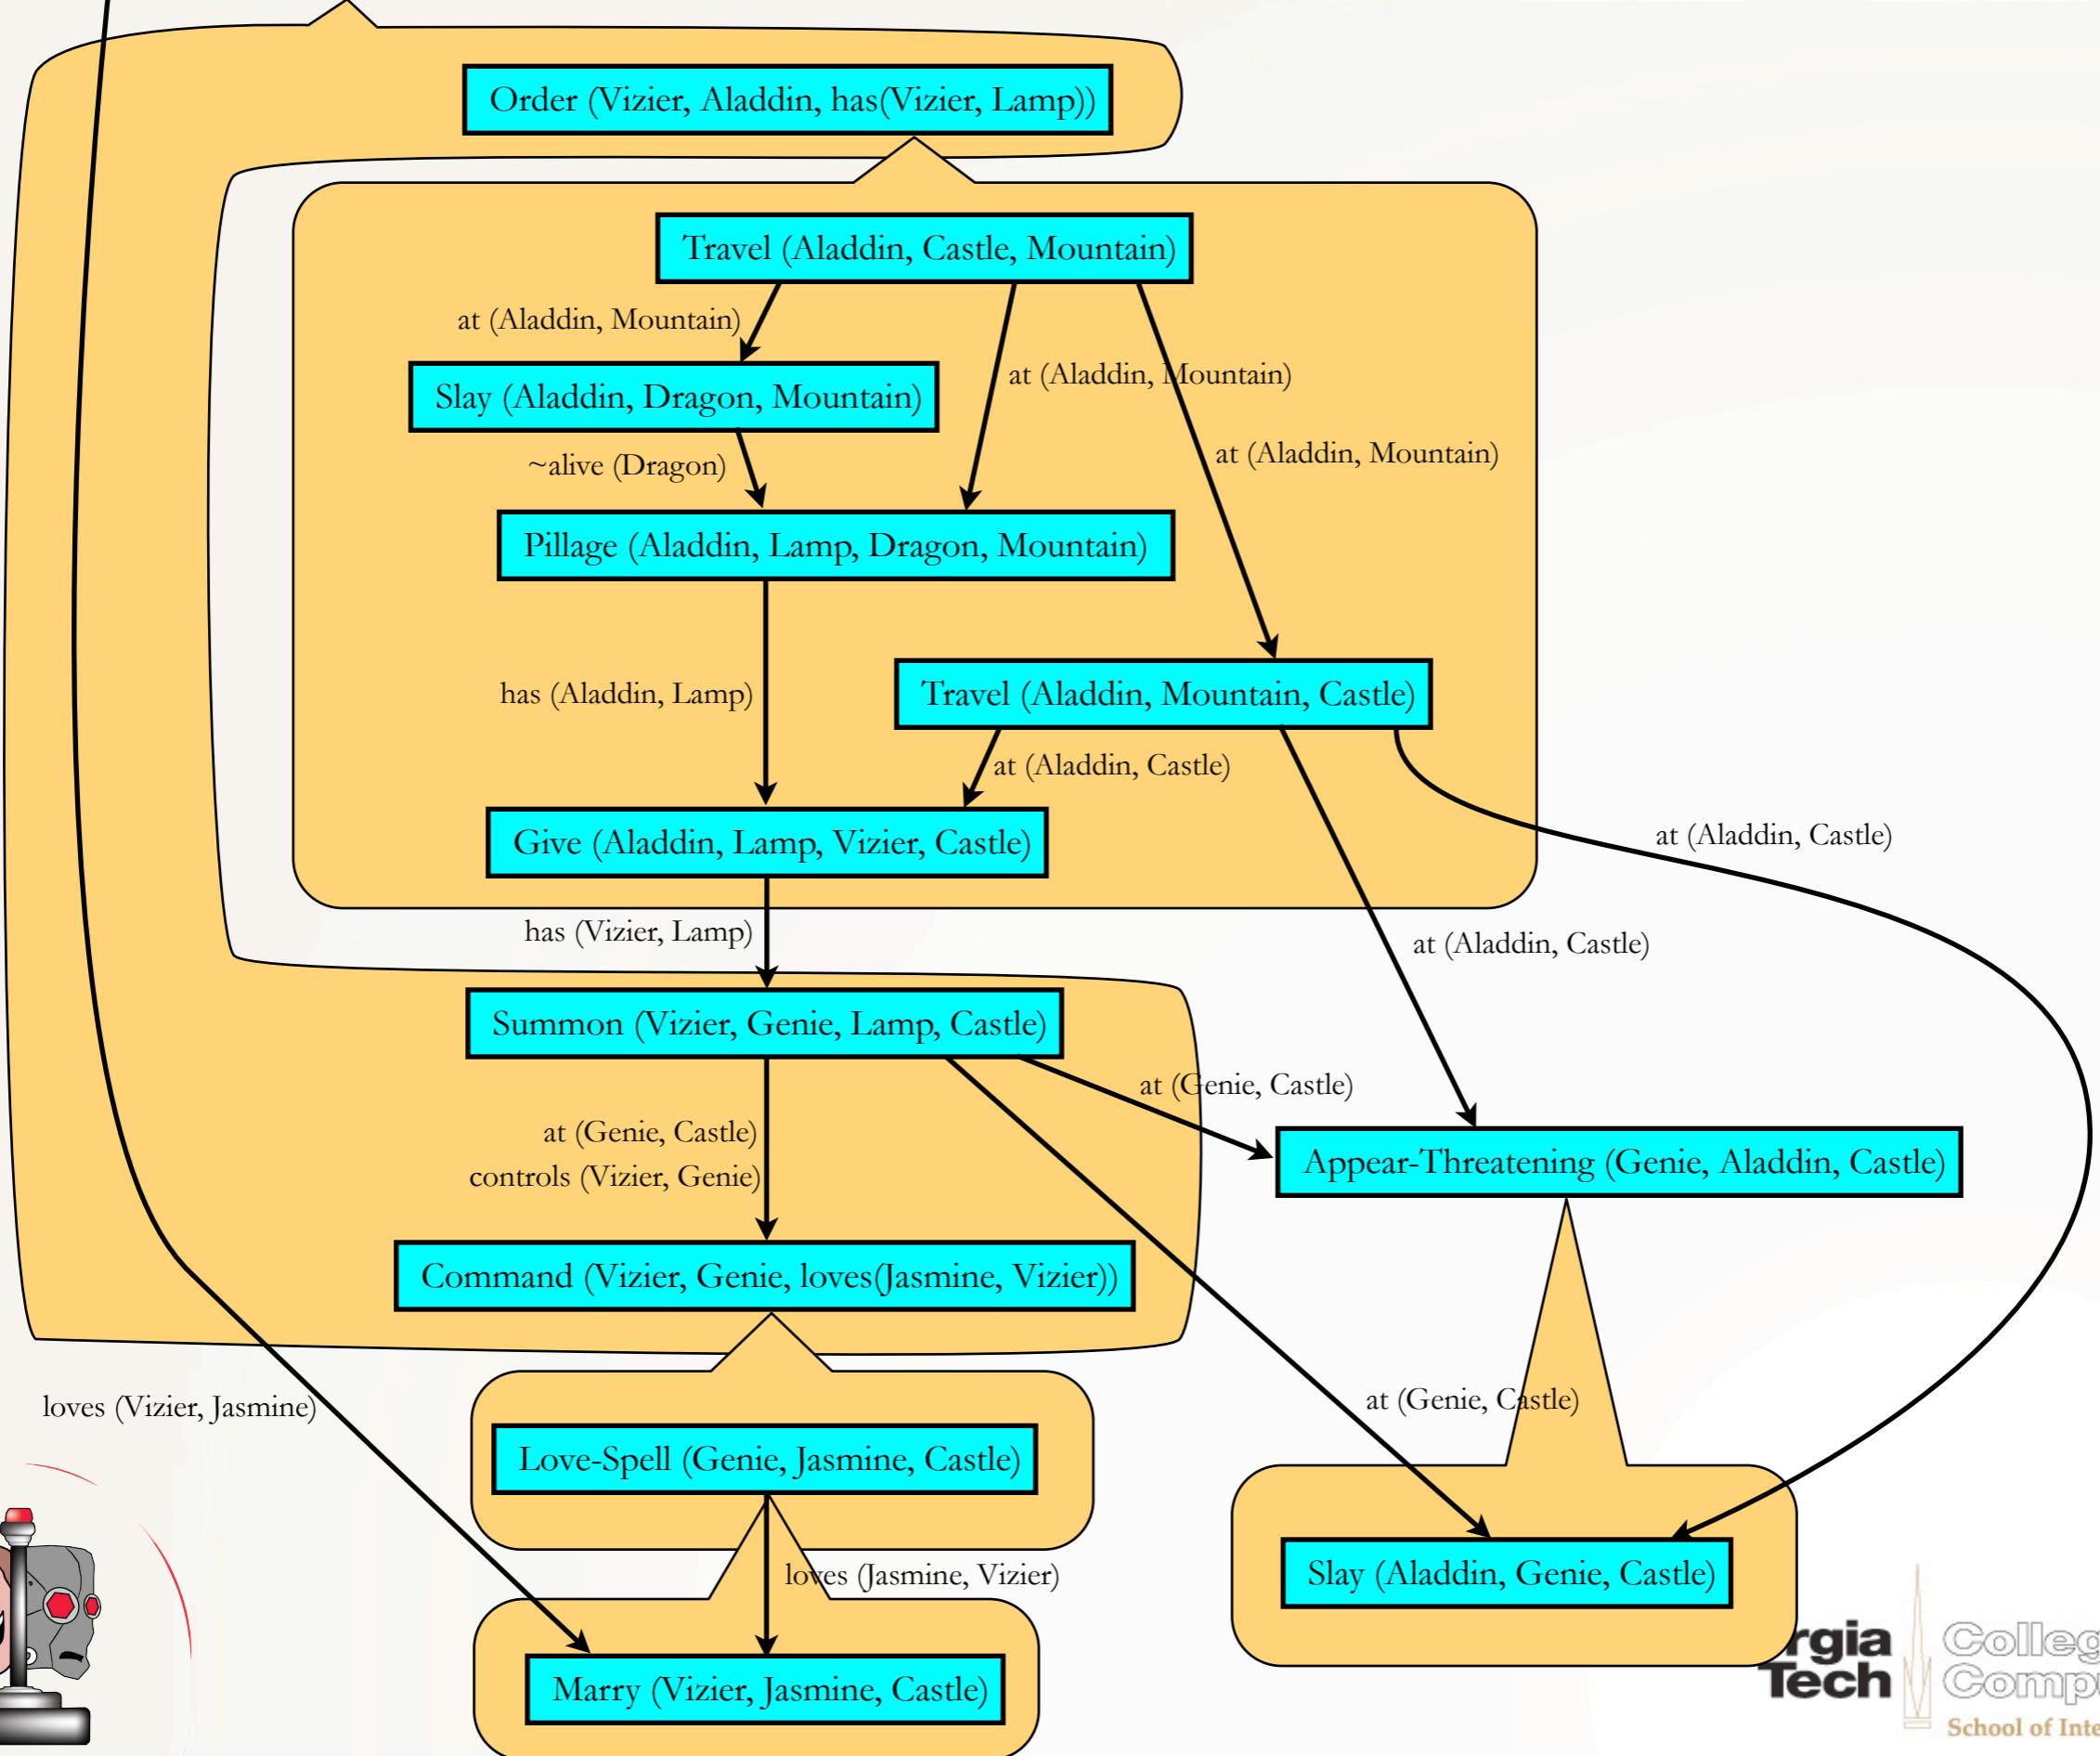




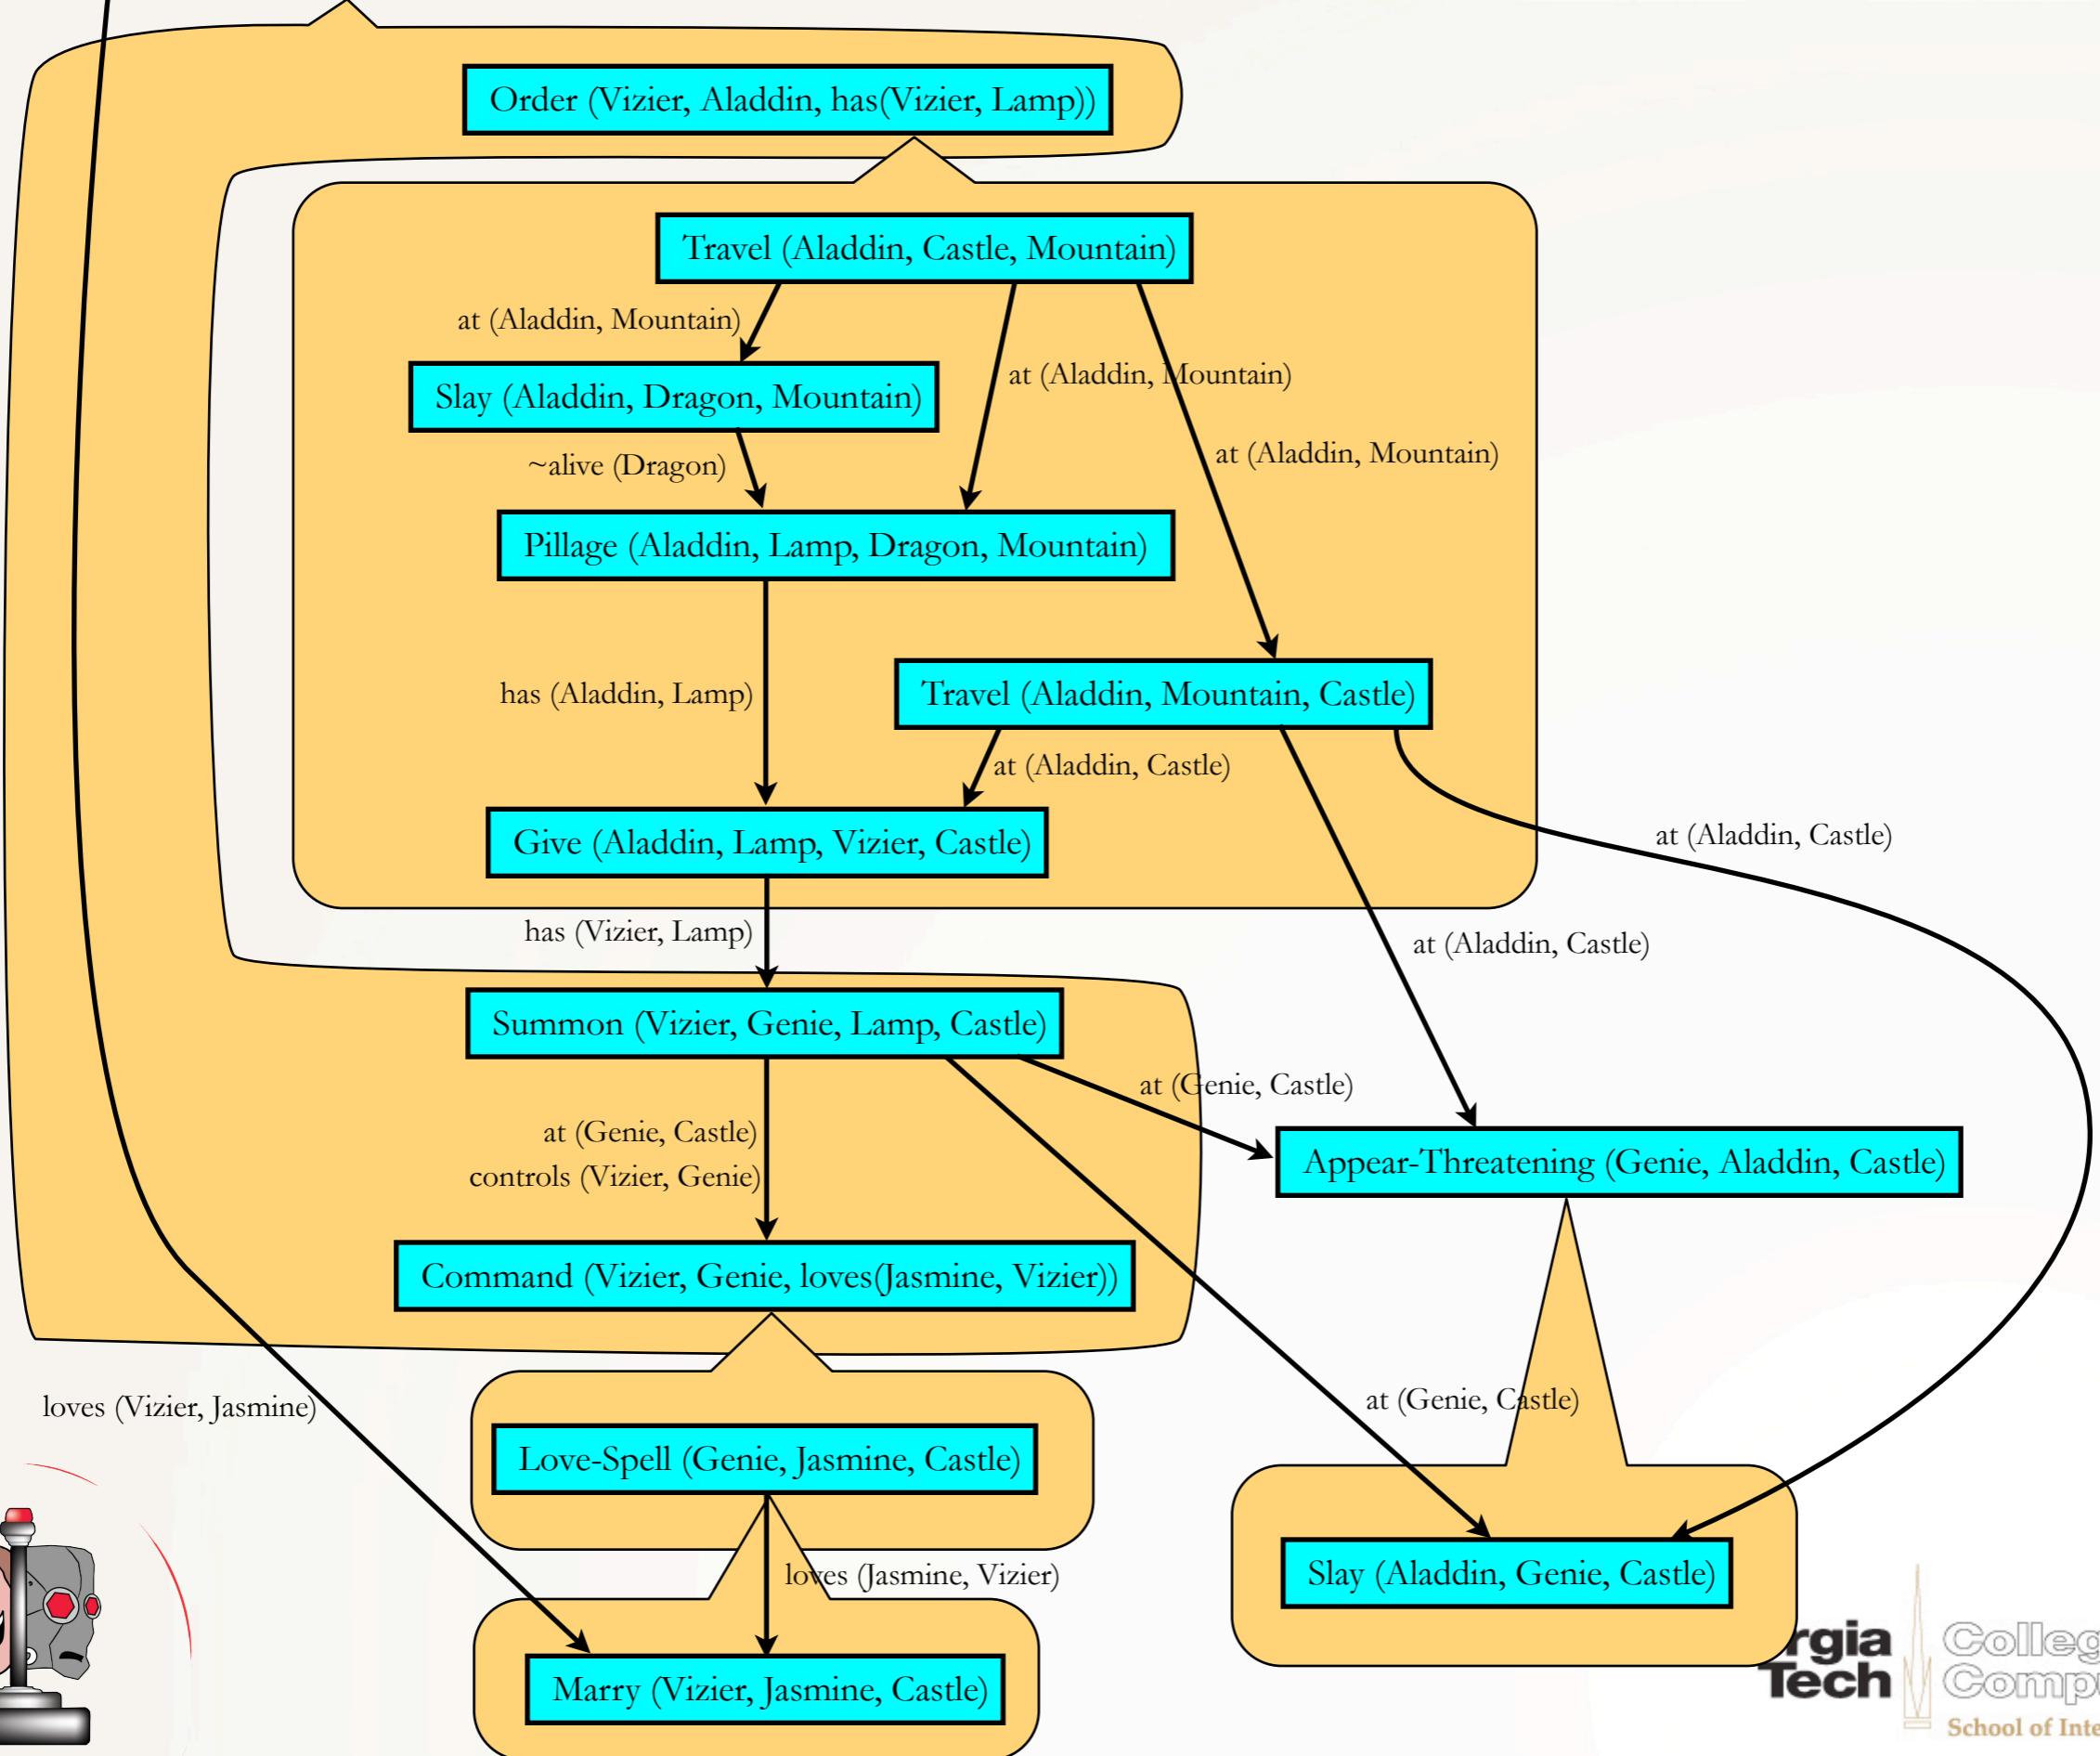
Marry (Vizier, Jasmine, Castle)

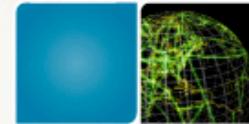
Slay (Aladdin, Genie, Castle)

Falls-in-Love (Vizier, Jasmine, Castle)



Falls-in-Love (Vizier, Jasmine, Castle)



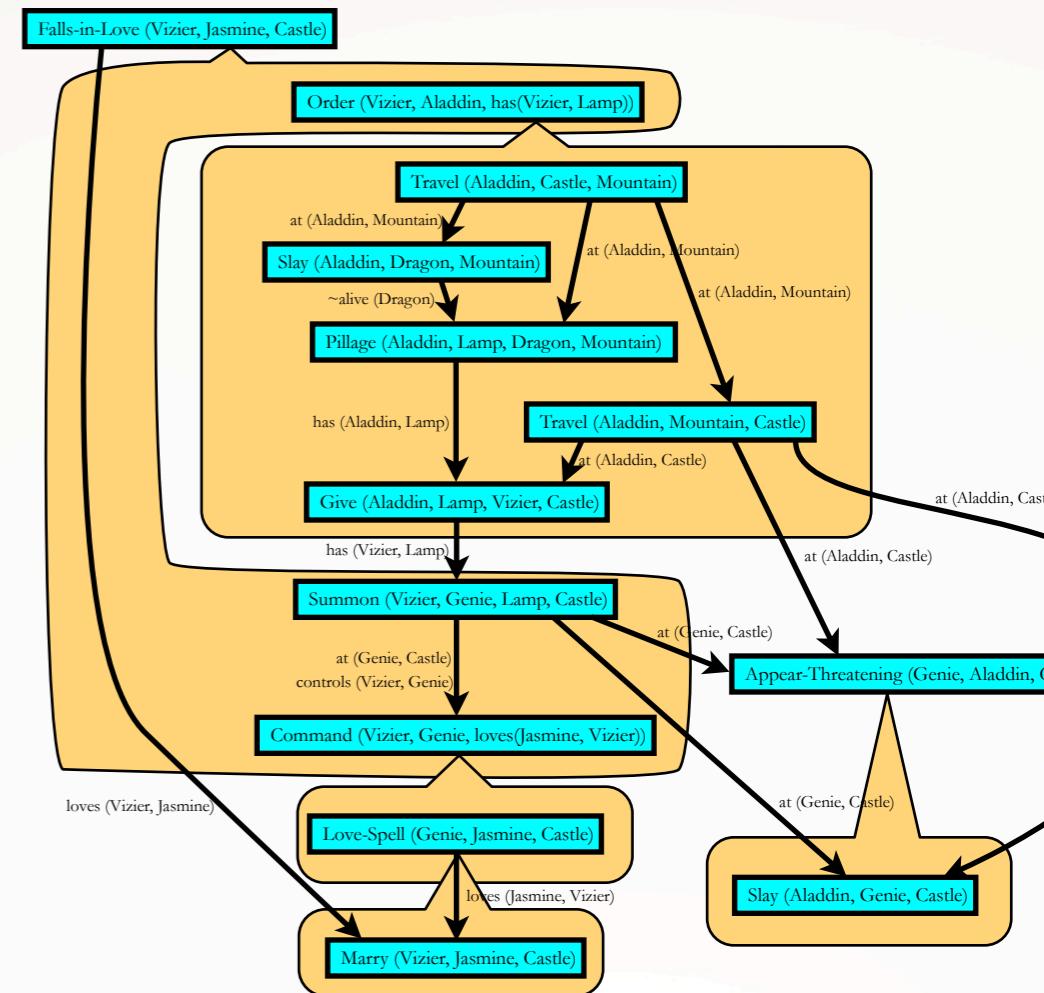


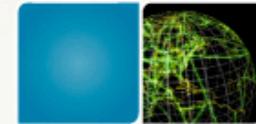
There is a woman named Jasmine. There is a vizier named Jafar. This is a story about how Jafar becomes married to Jasmine. There is a magic genie. This is also a story about how the genie dies.

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.





# Effect on reader comprehension?

- Fabulist vs. conventional planner
- Question-answering protocol to elicit readers' mental models
  - Enablement (how questions)
  - Intention (why questions)
- Compared readers' answers to the computer's answers

**Q:** Why did Aladdin travel from the mountains to the castle?

**A:** Because the genie wanted to cast a spell on Jasmine.

4: Very  
Good

3: Good

2: Bad

1: Very  
Bad

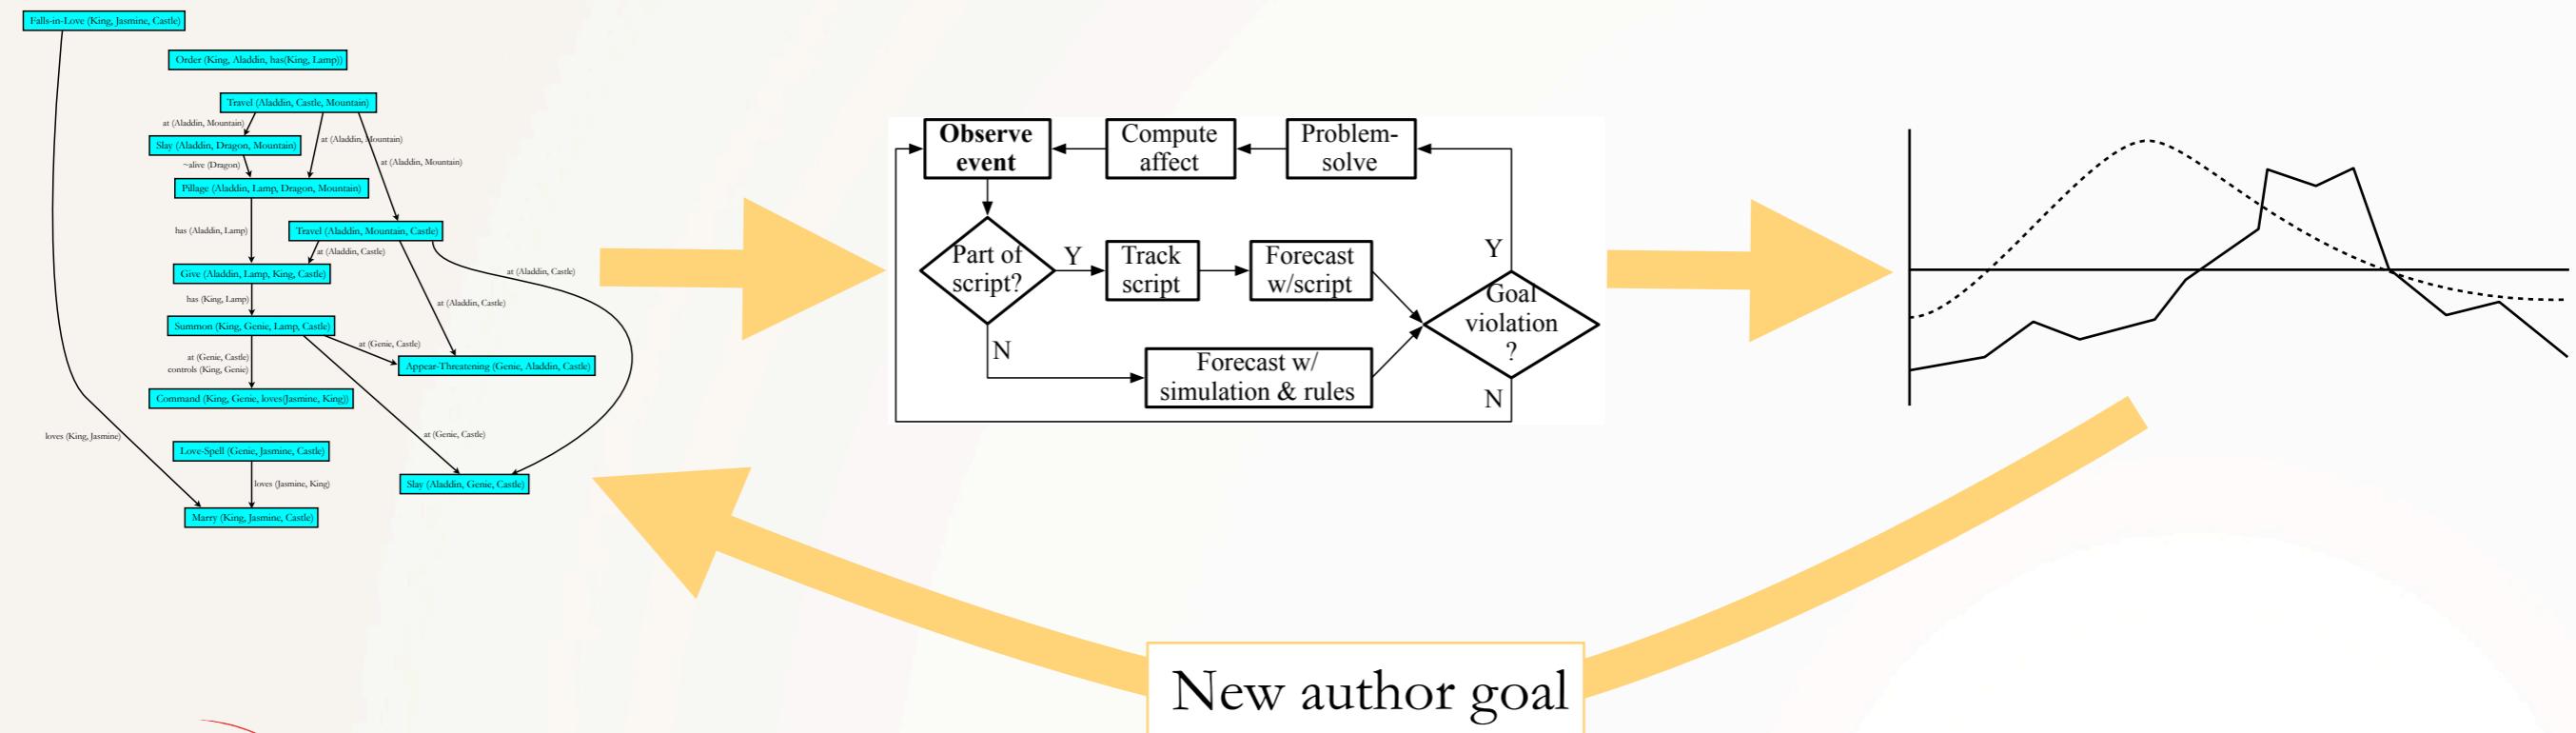
	Questions computer thinks are "good" (closer to 4)	Questions computer thinks are "bad" (closer to 1)
<b>Fabulist</b>	3.20	1.19
<b>Control</b>	2.99	1.30





# The “goodness” question

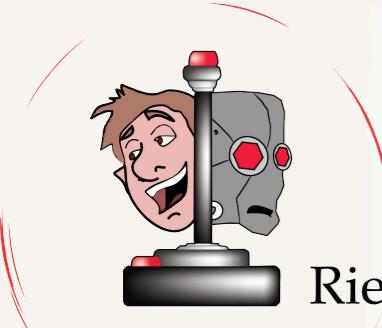
- Does the story generator know that it is creating something good?
- How can a computational system generate a suspenseful story?





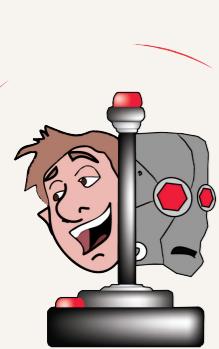
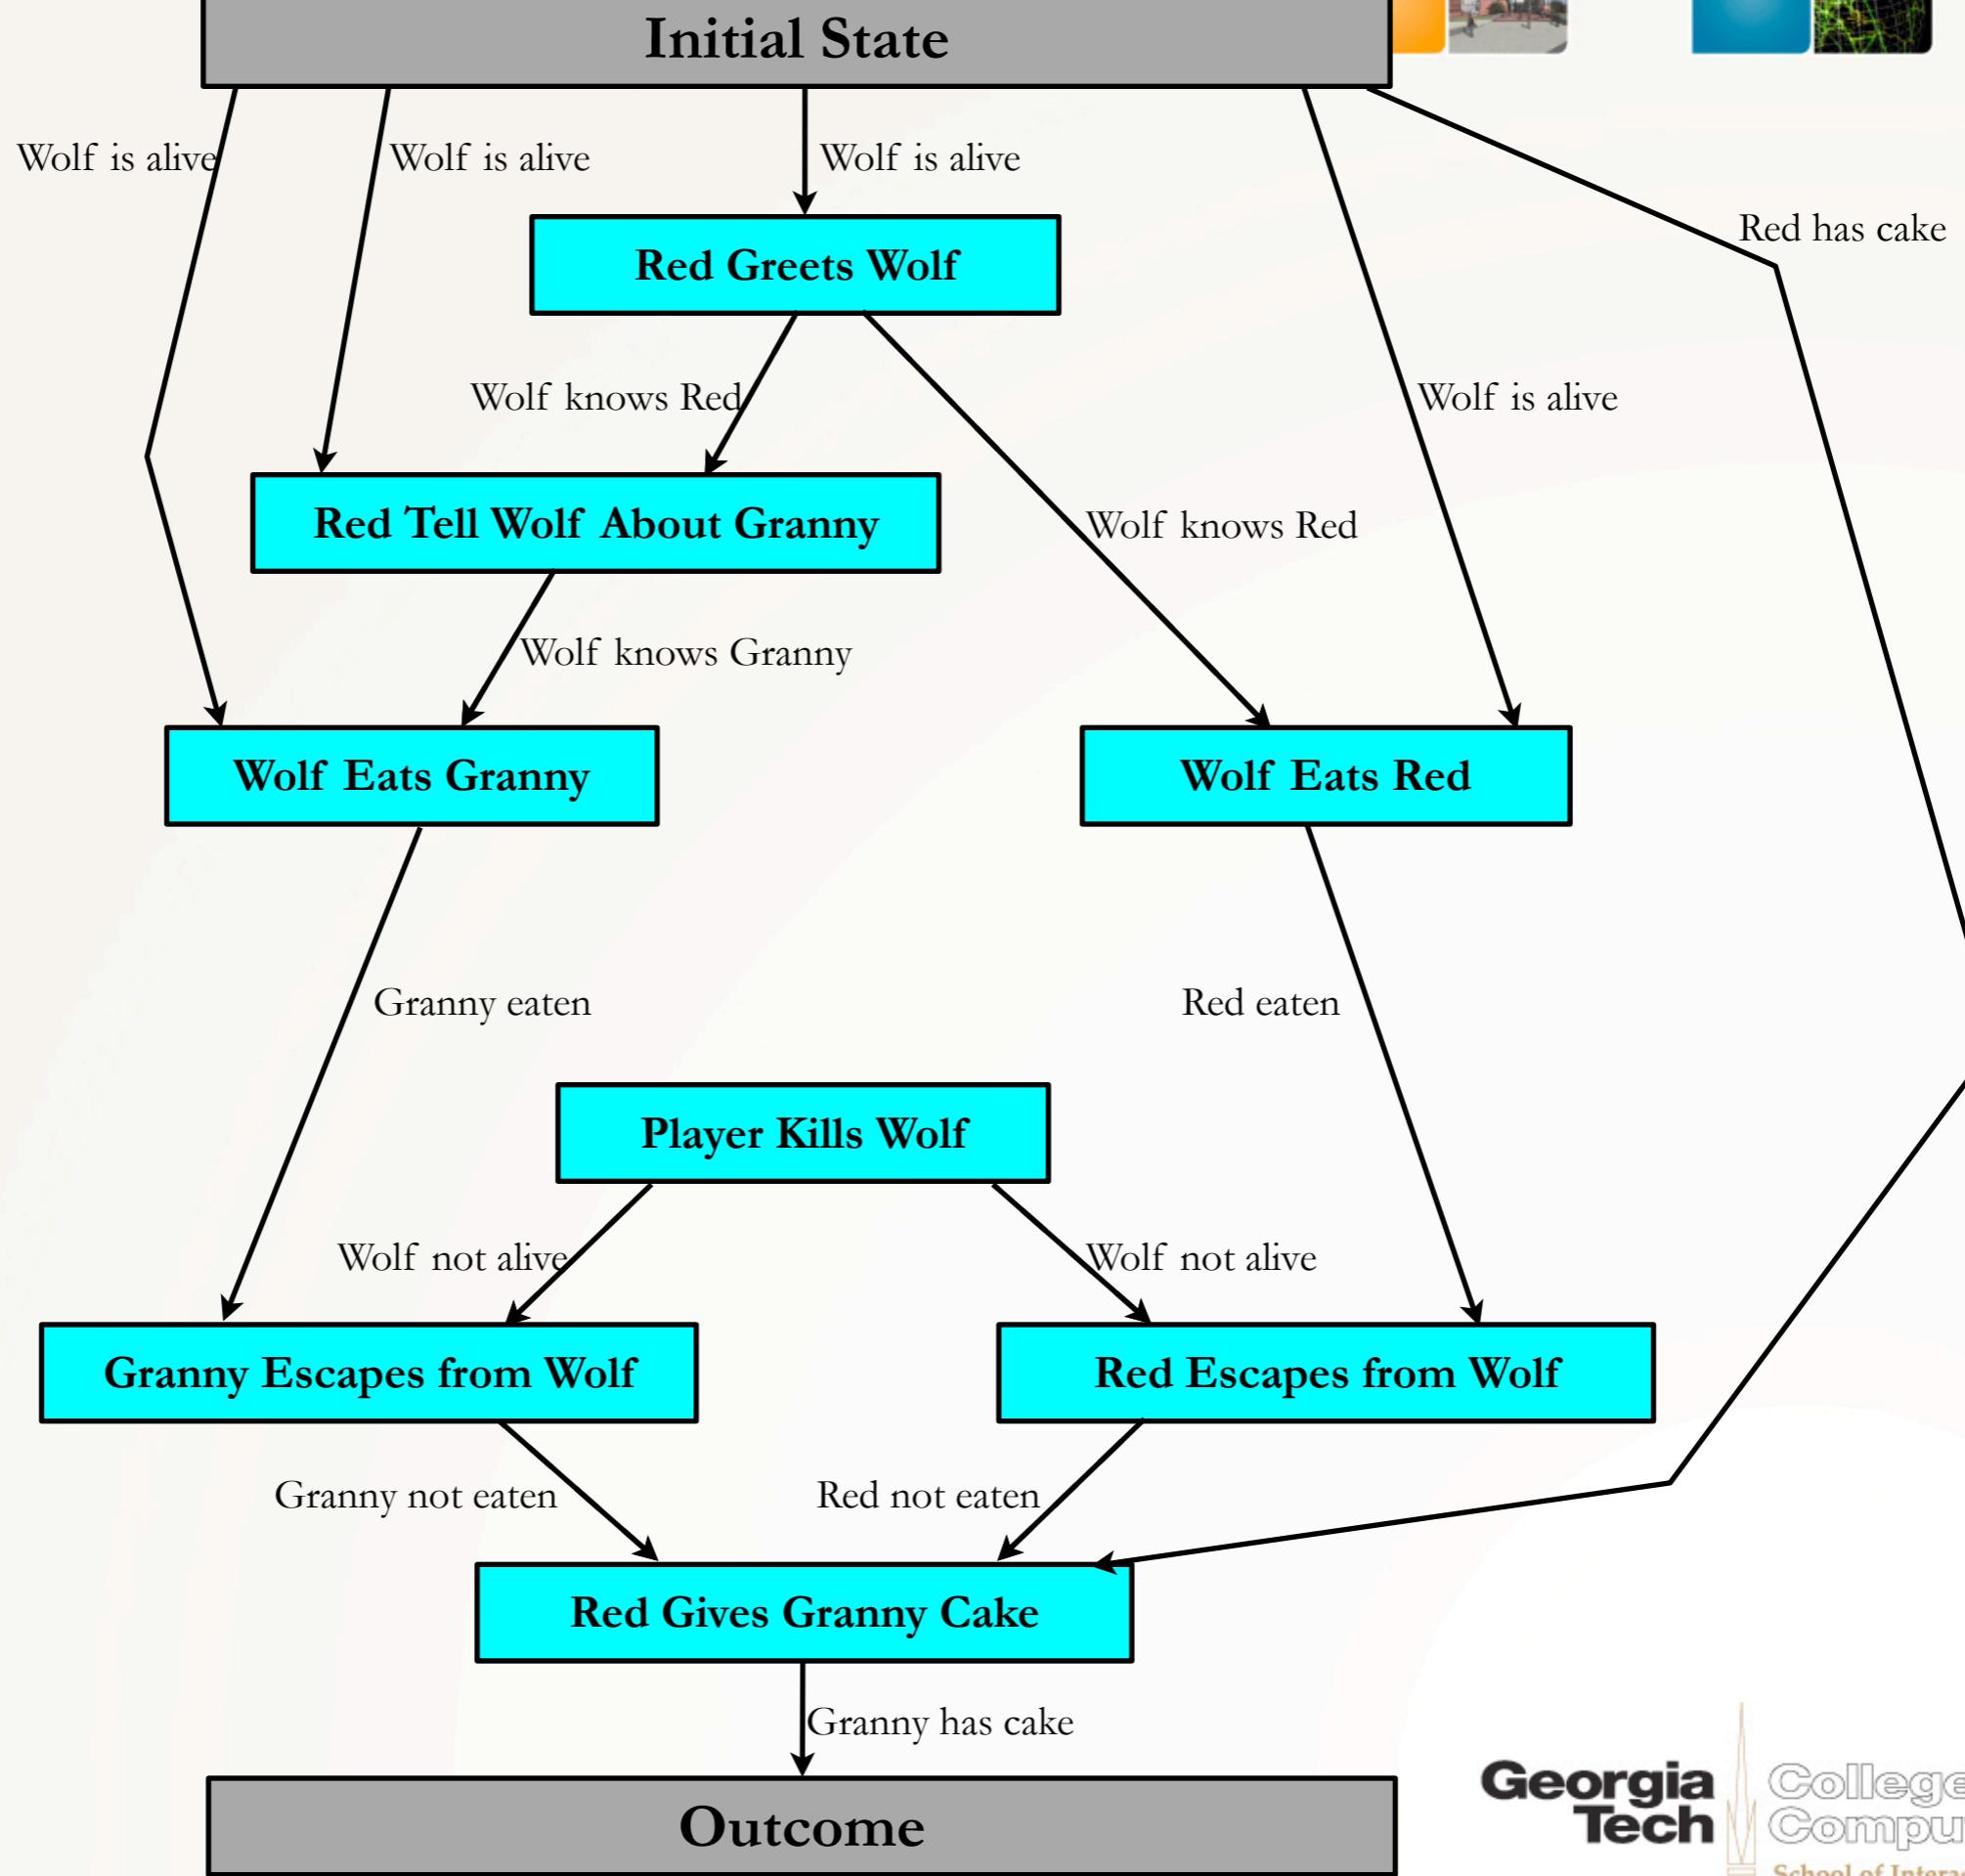
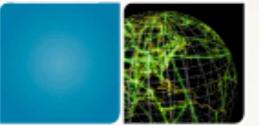
# Automated story director

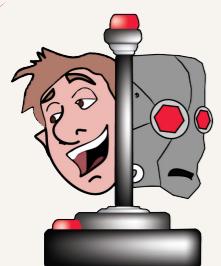
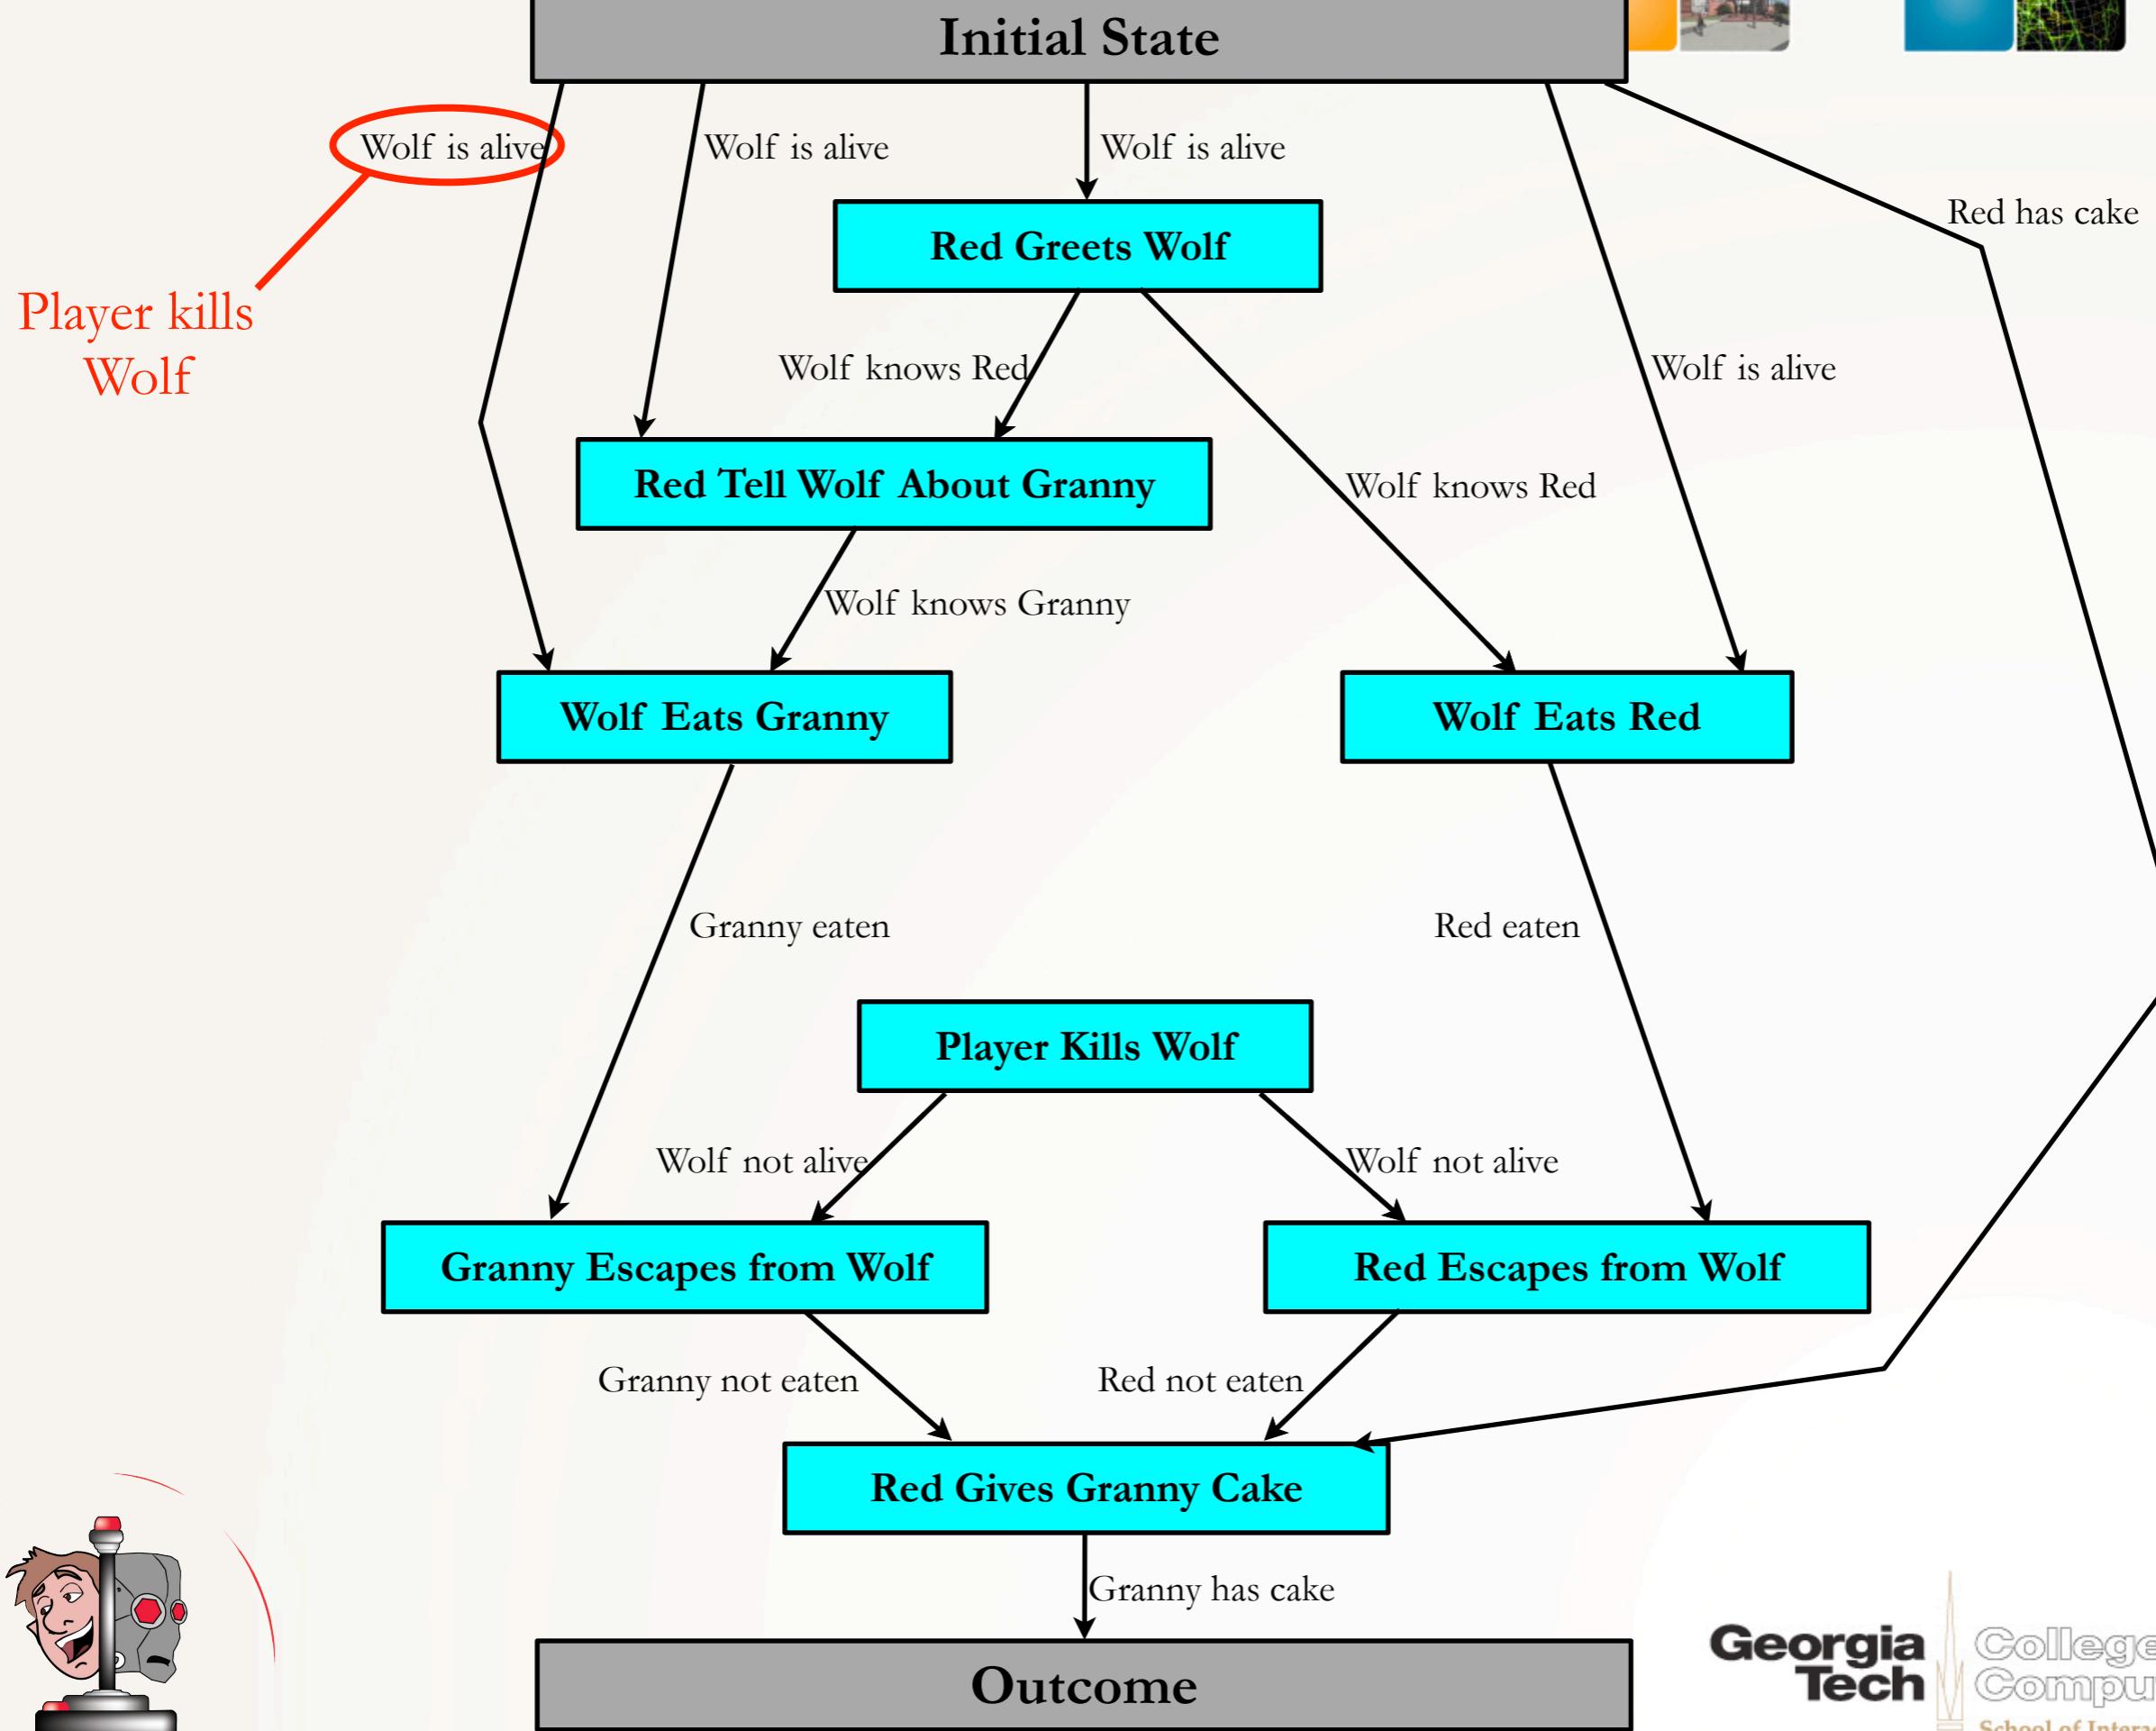
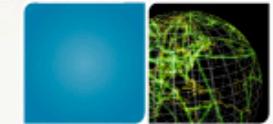
- Use a story generator to construct alternative branches
- Start with an exemplar narrative (generated or hand-authored)
- Exceptions: player actions that derail the narrative plan
- How to detect exceptions?
- How to respond to exceptions?

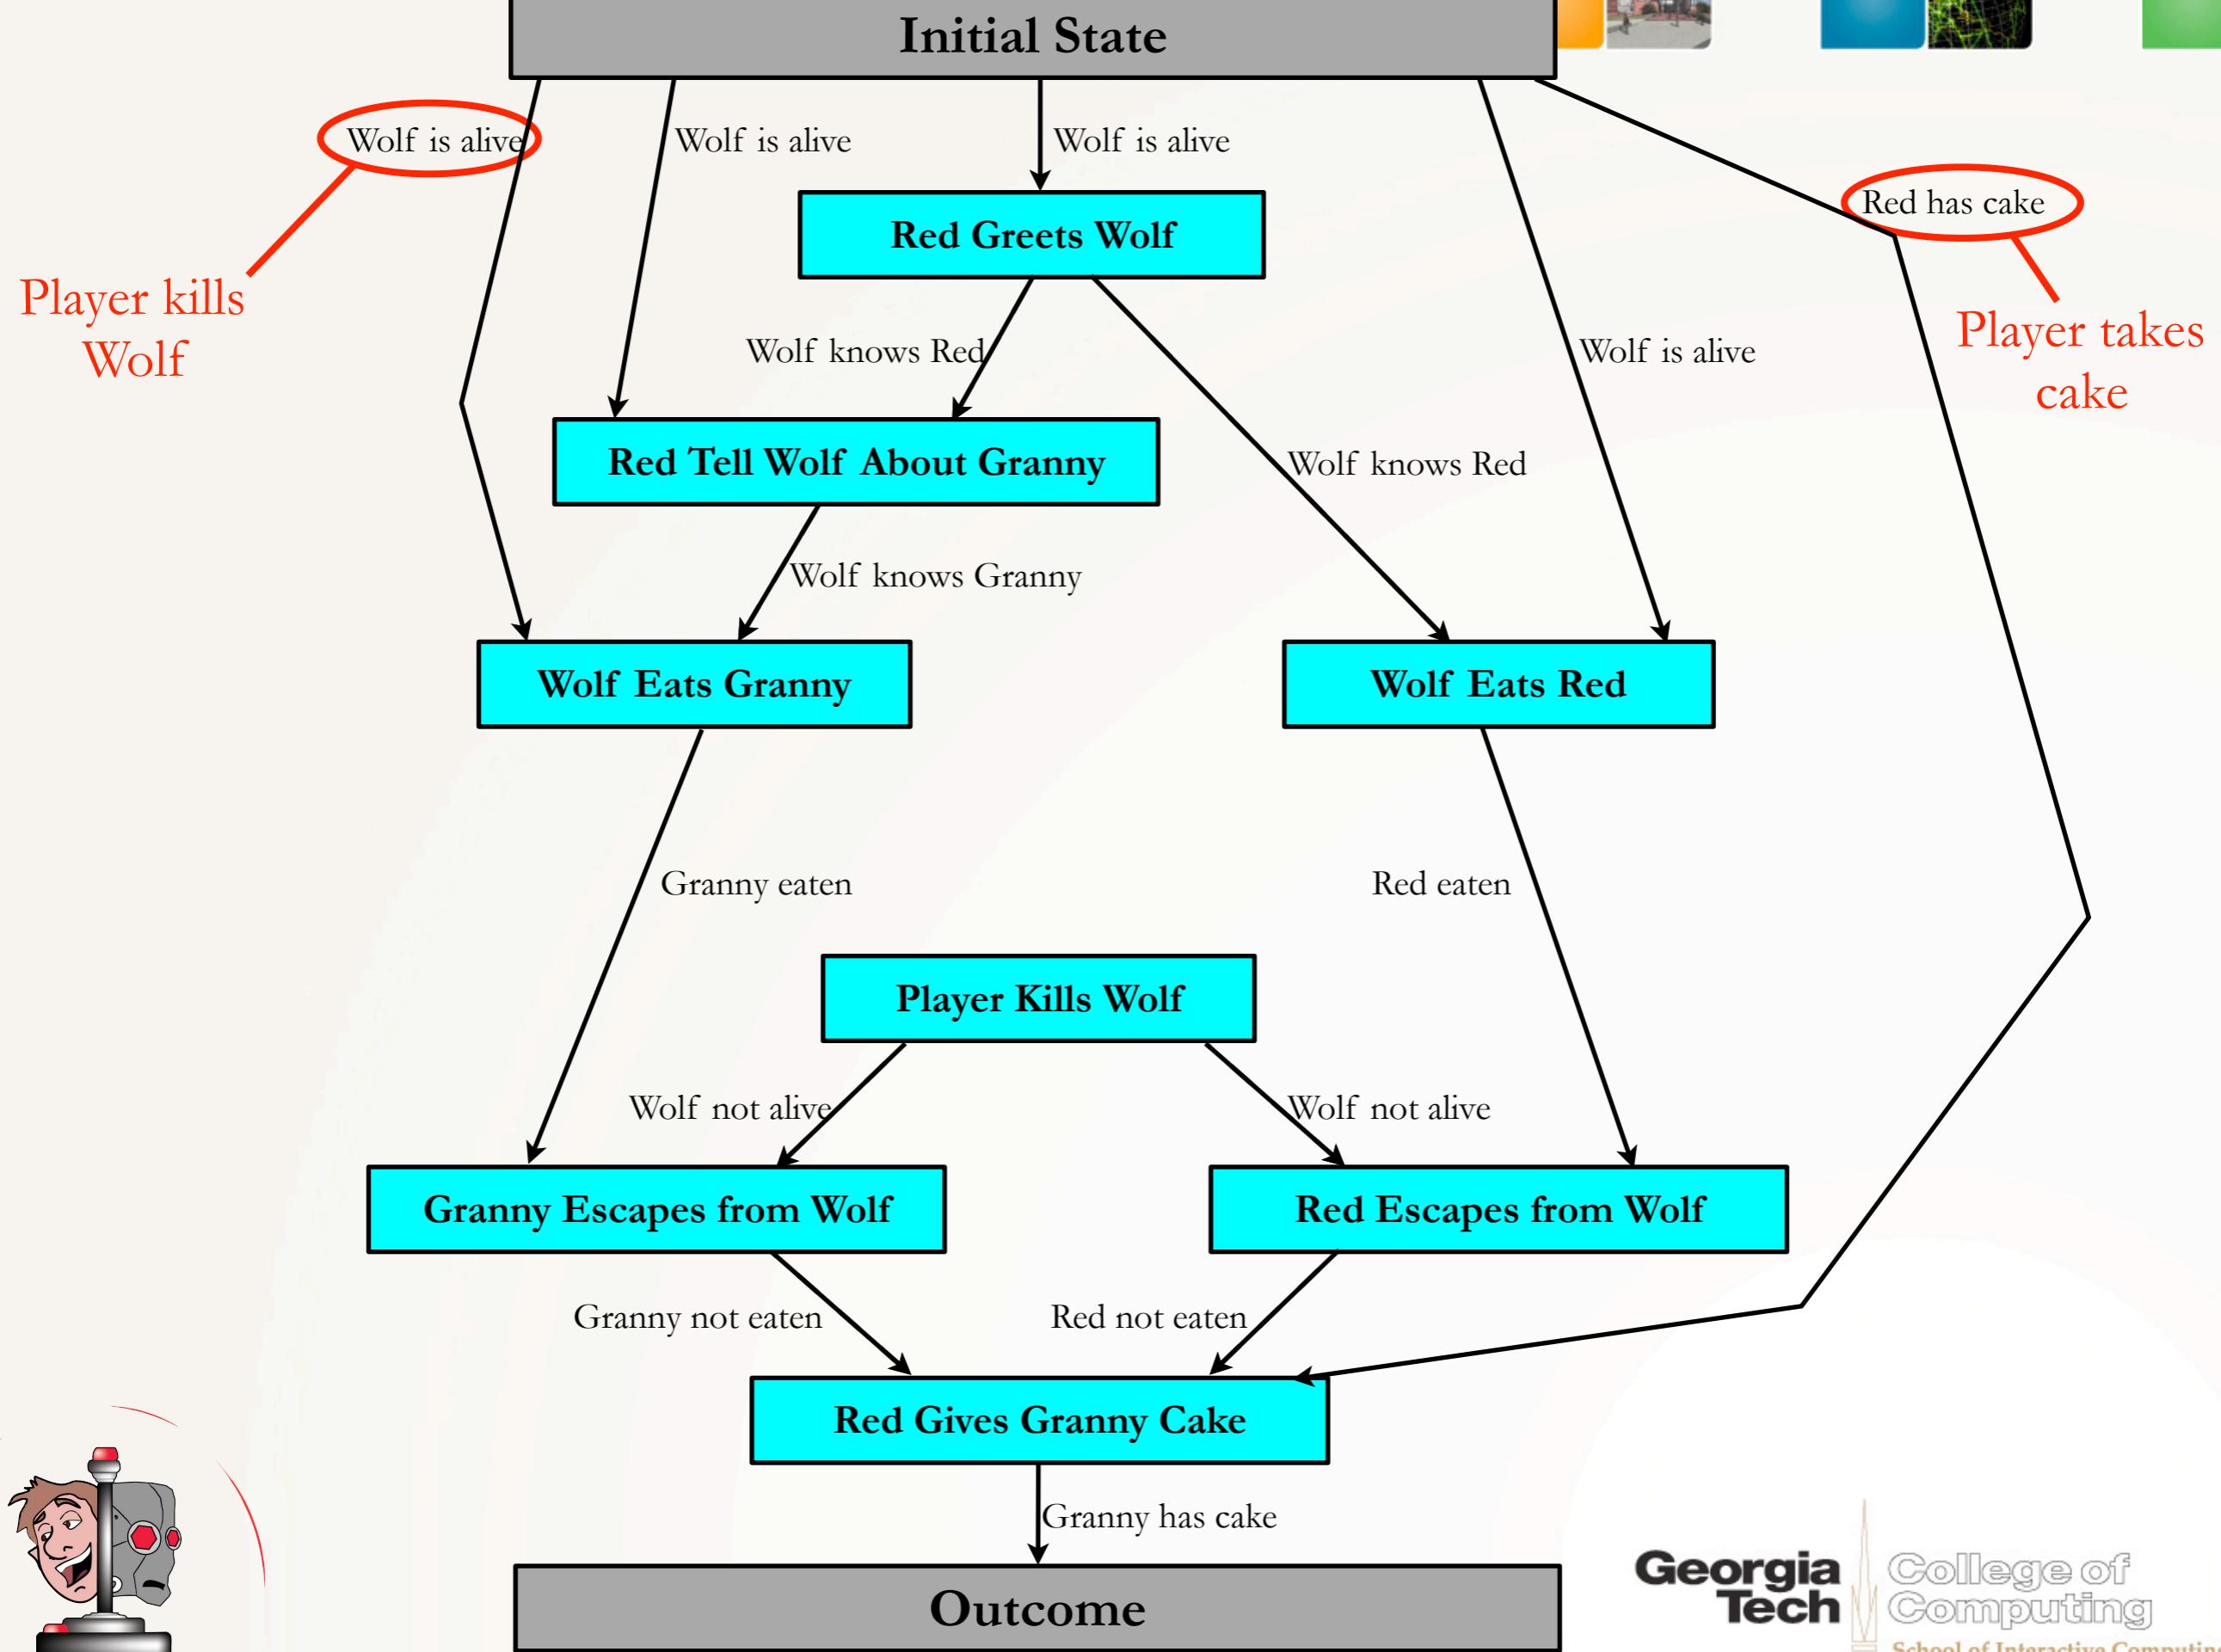
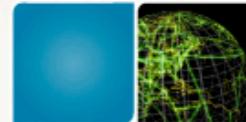


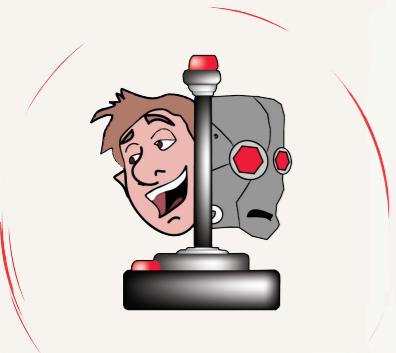
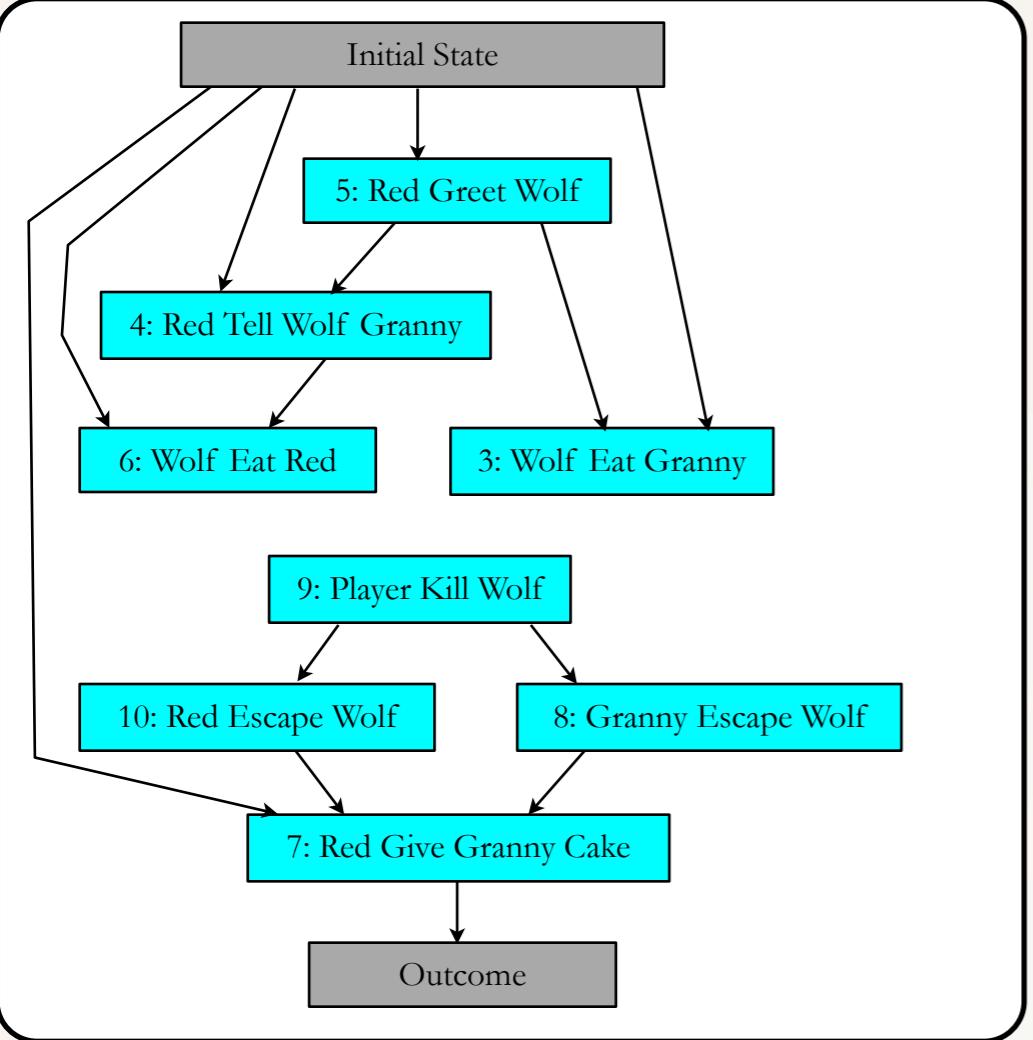
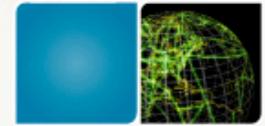
Riedl, Stern, Dini & Alderman. ITSSA, 3, 2008.

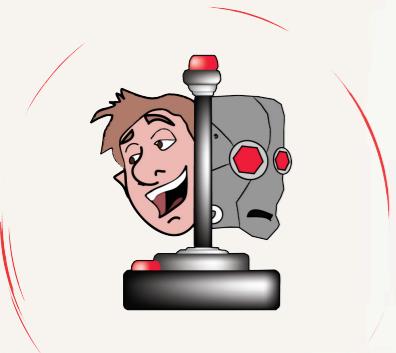
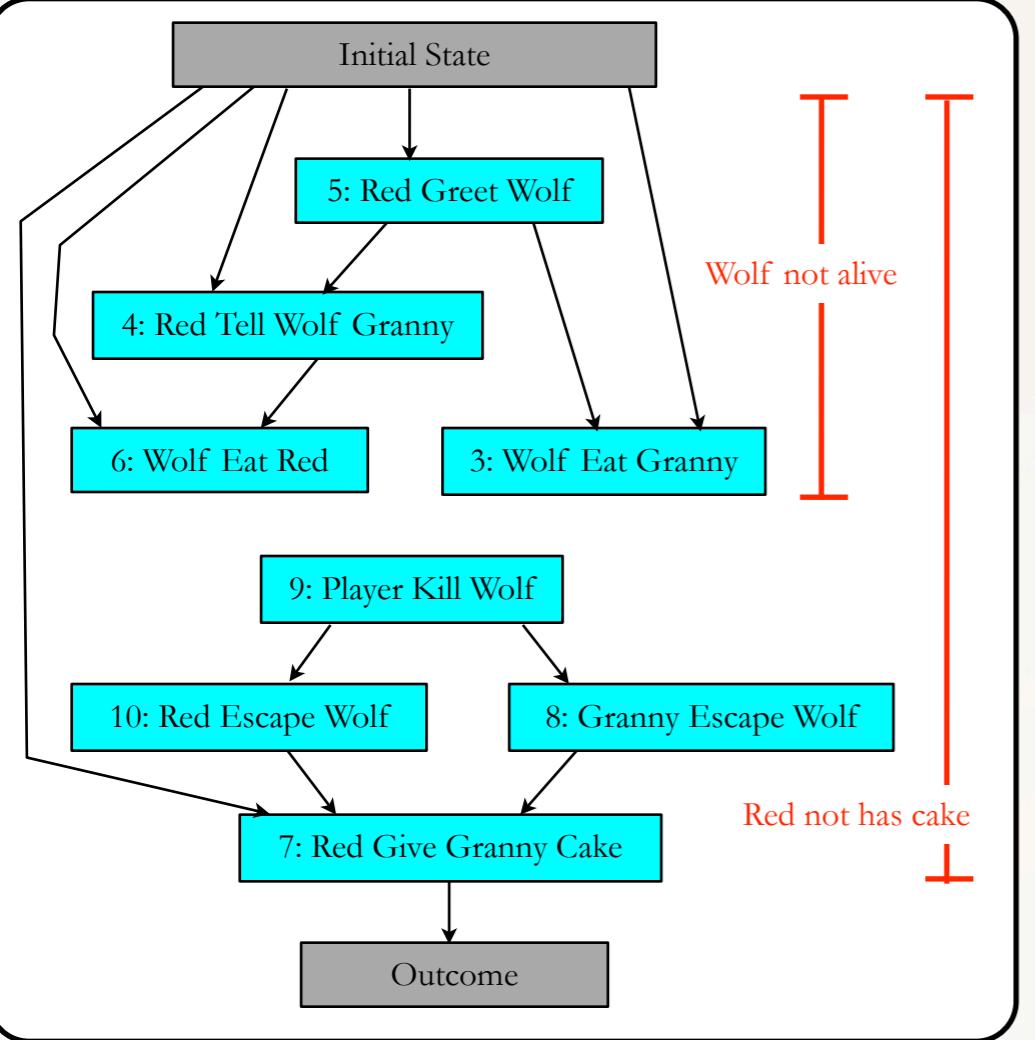


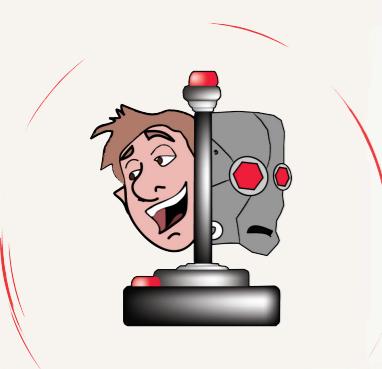
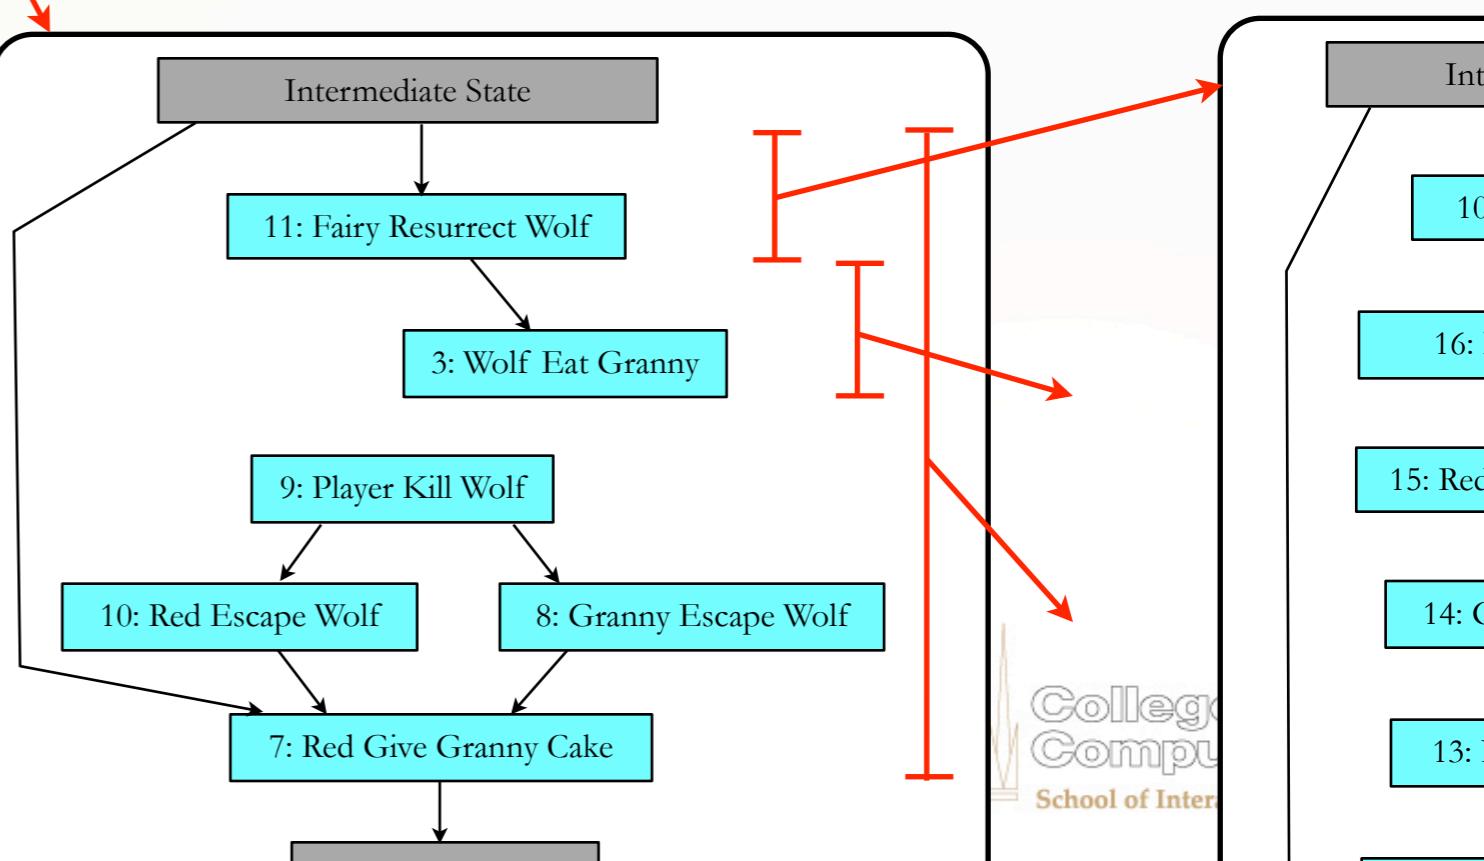
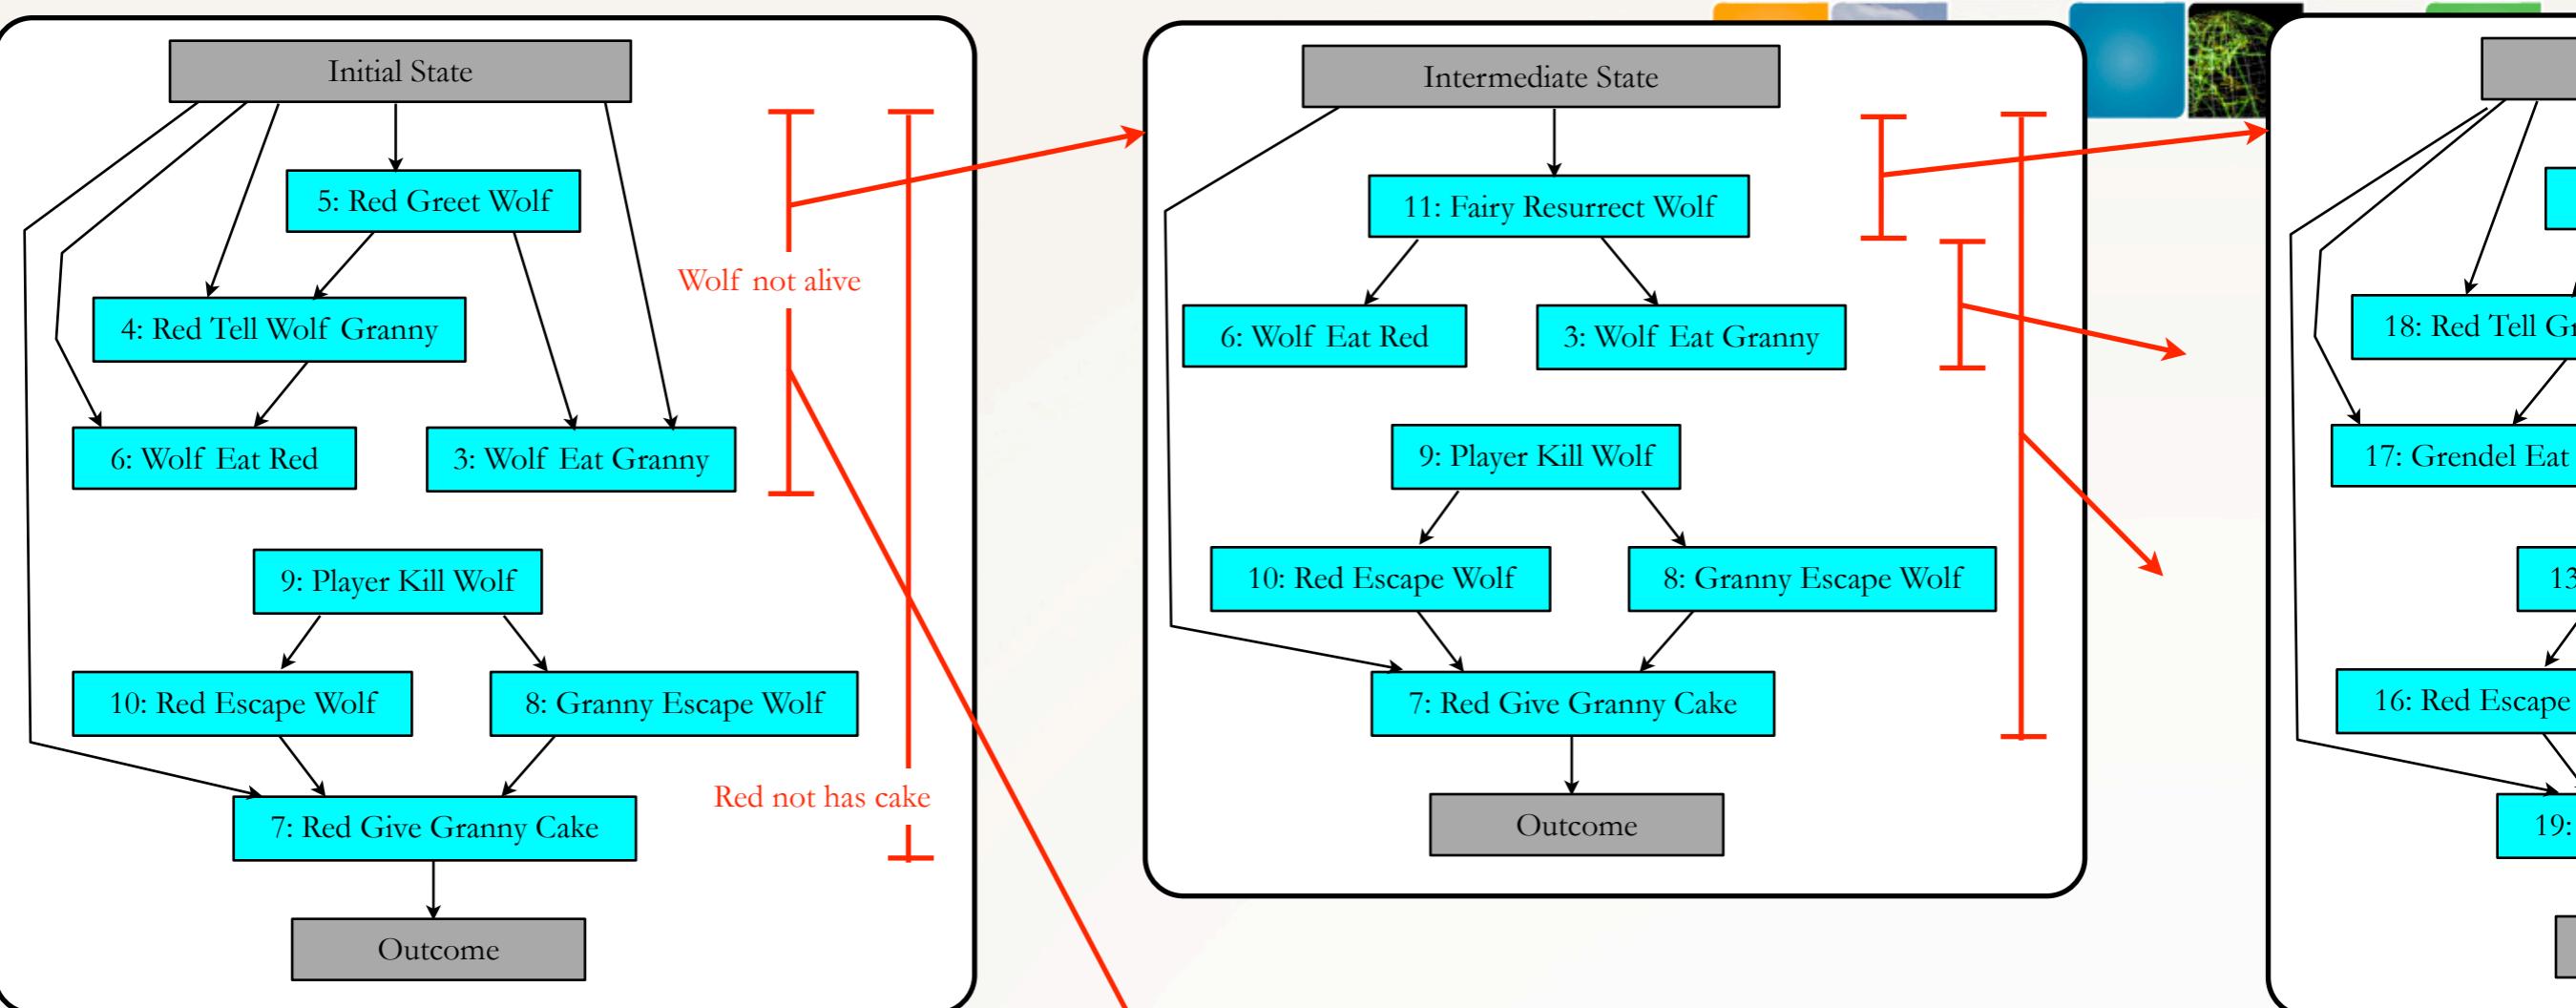










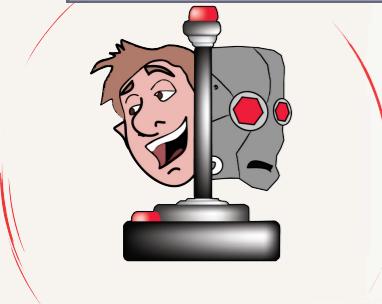




# Automated story director

- Pre-compute tree of contingency narratives
- Semi-autonomous virtual characters

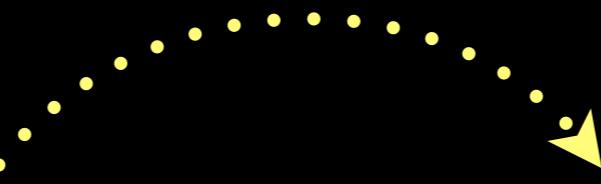
```
Wolf sits and waits.  
Wolf walks in a circle and sits back down.  
Wolf walks in a circle and sits back down.  
Red has arrived.  
Wolf eyes Little Red hungrily.  
Red looks closely at Wolf.  
Red says, "My what big teeth you have!"  
Wolf says, "The better to eat you with, my dear!"  
Red screams.  
Wolf swallows Little Red in one big gulp.  
Red disappears suddenly for parts unknown.  
Wolf is feeling very drowsy.  
Wolf falls asleep.  
:kills wolf  
Fred kills wolf  
Granny arrives.  
Granny claws her way out of the wolf.
```



# Story Generation

---

2. Crowdsourcing



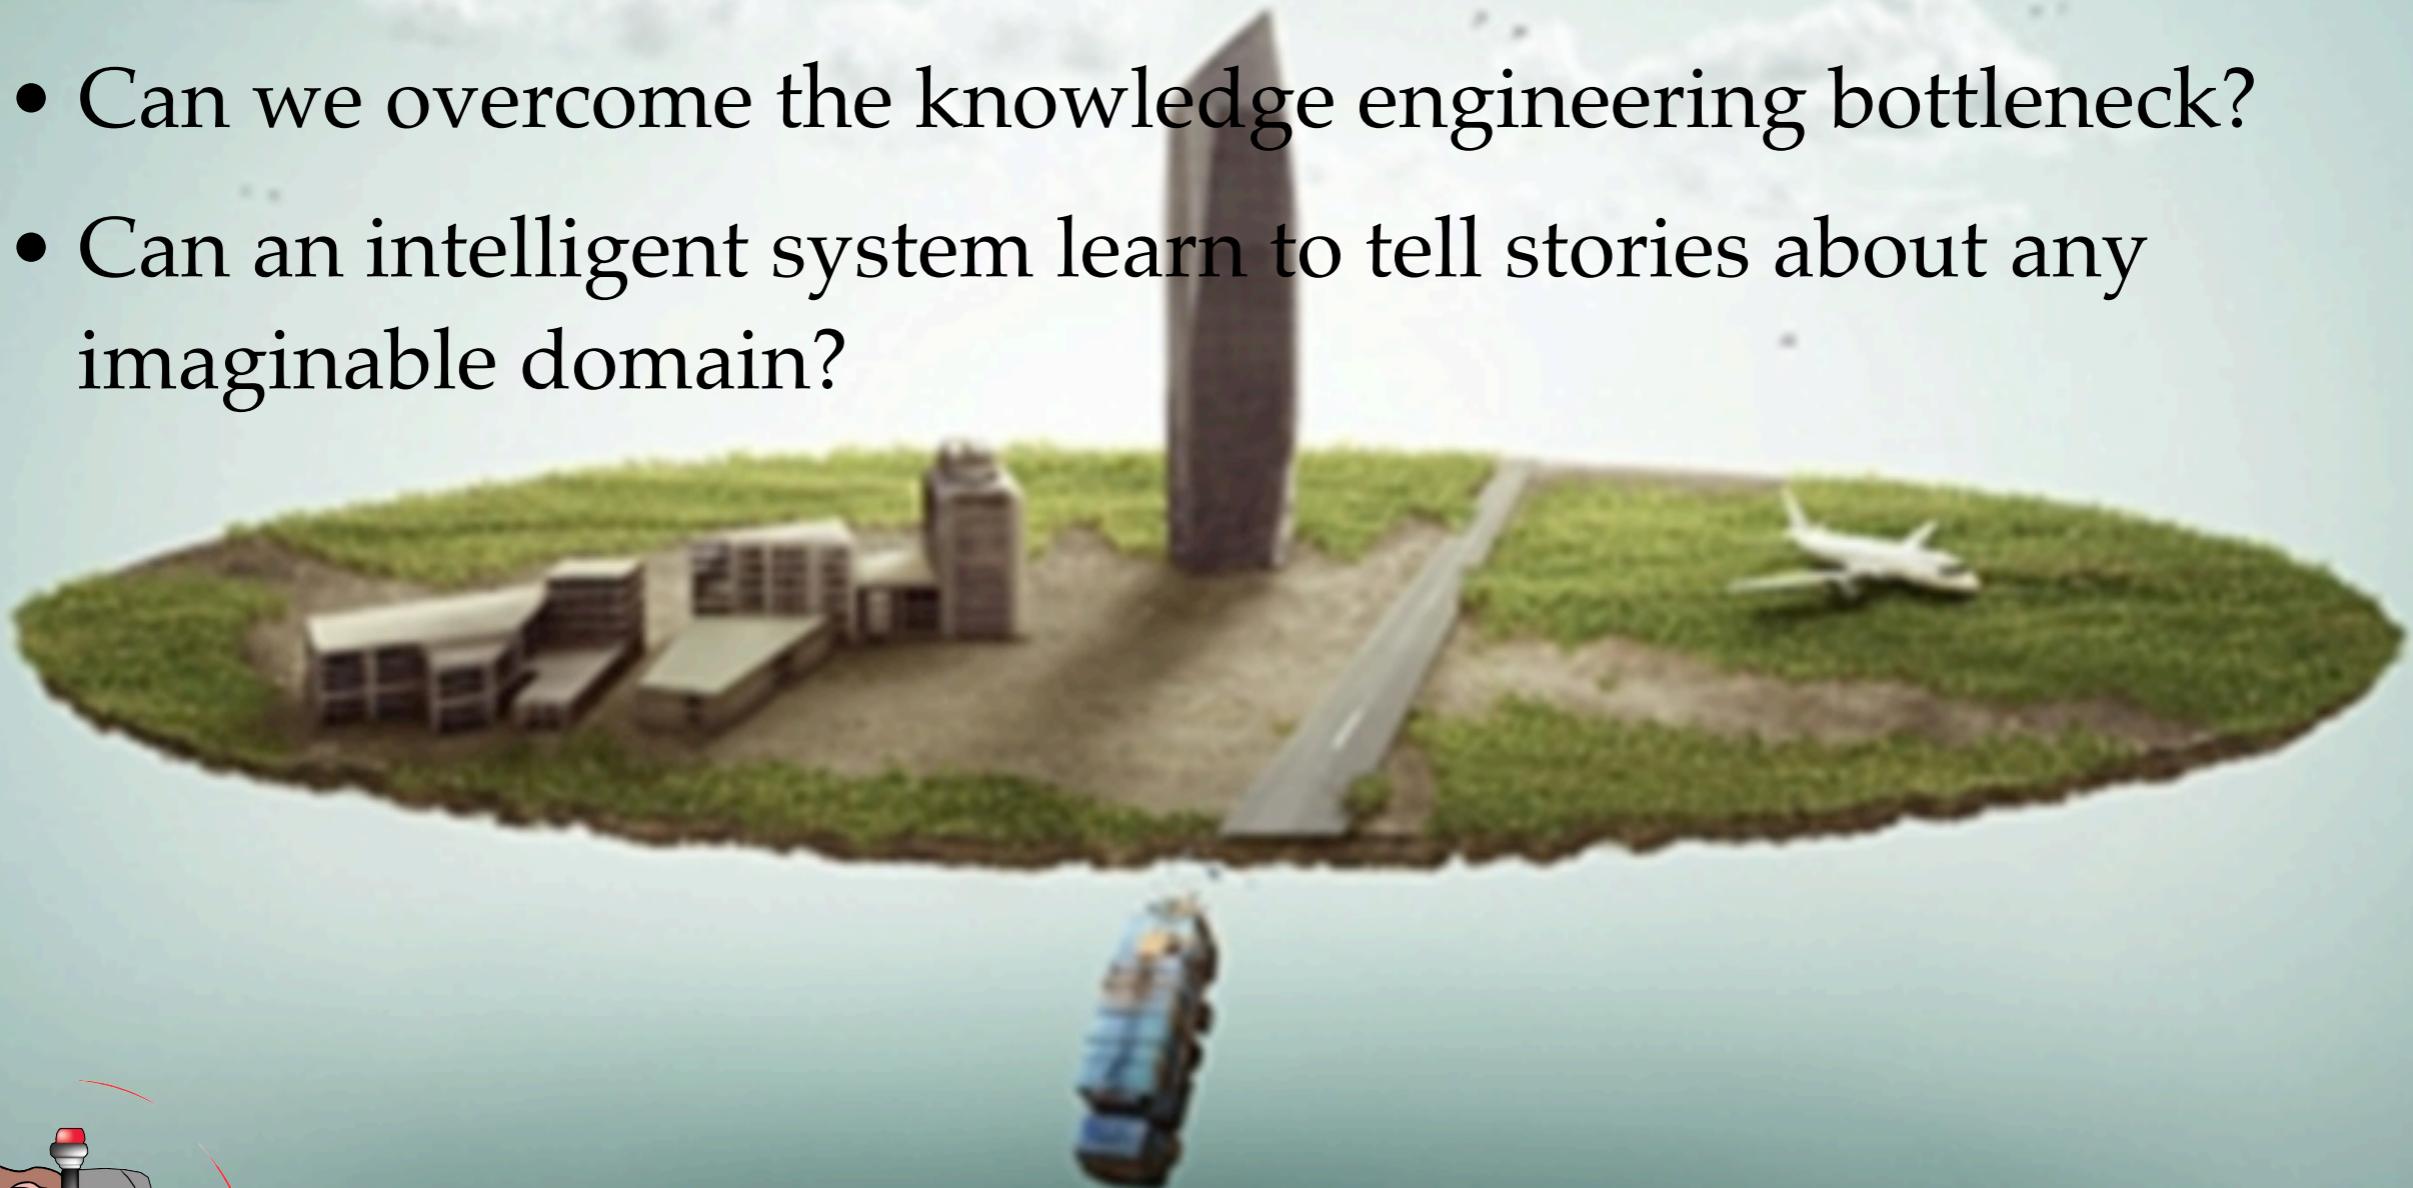
# Interactive Narrative

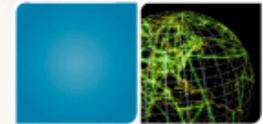
1. Cognitively inspired planning



# Open story generation

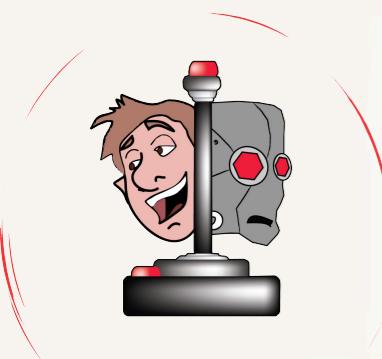
- Narrative intelligence is knowledge-intensive, resulting in micro-worlds
- Can we overcome the knowledge engineering bottleneck?
- Can an intelligent system learn to tell stories about any imaginable domain?



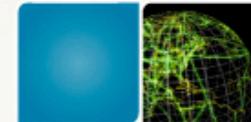


# Sociocultural storytelling

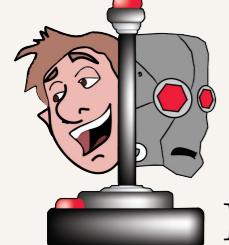
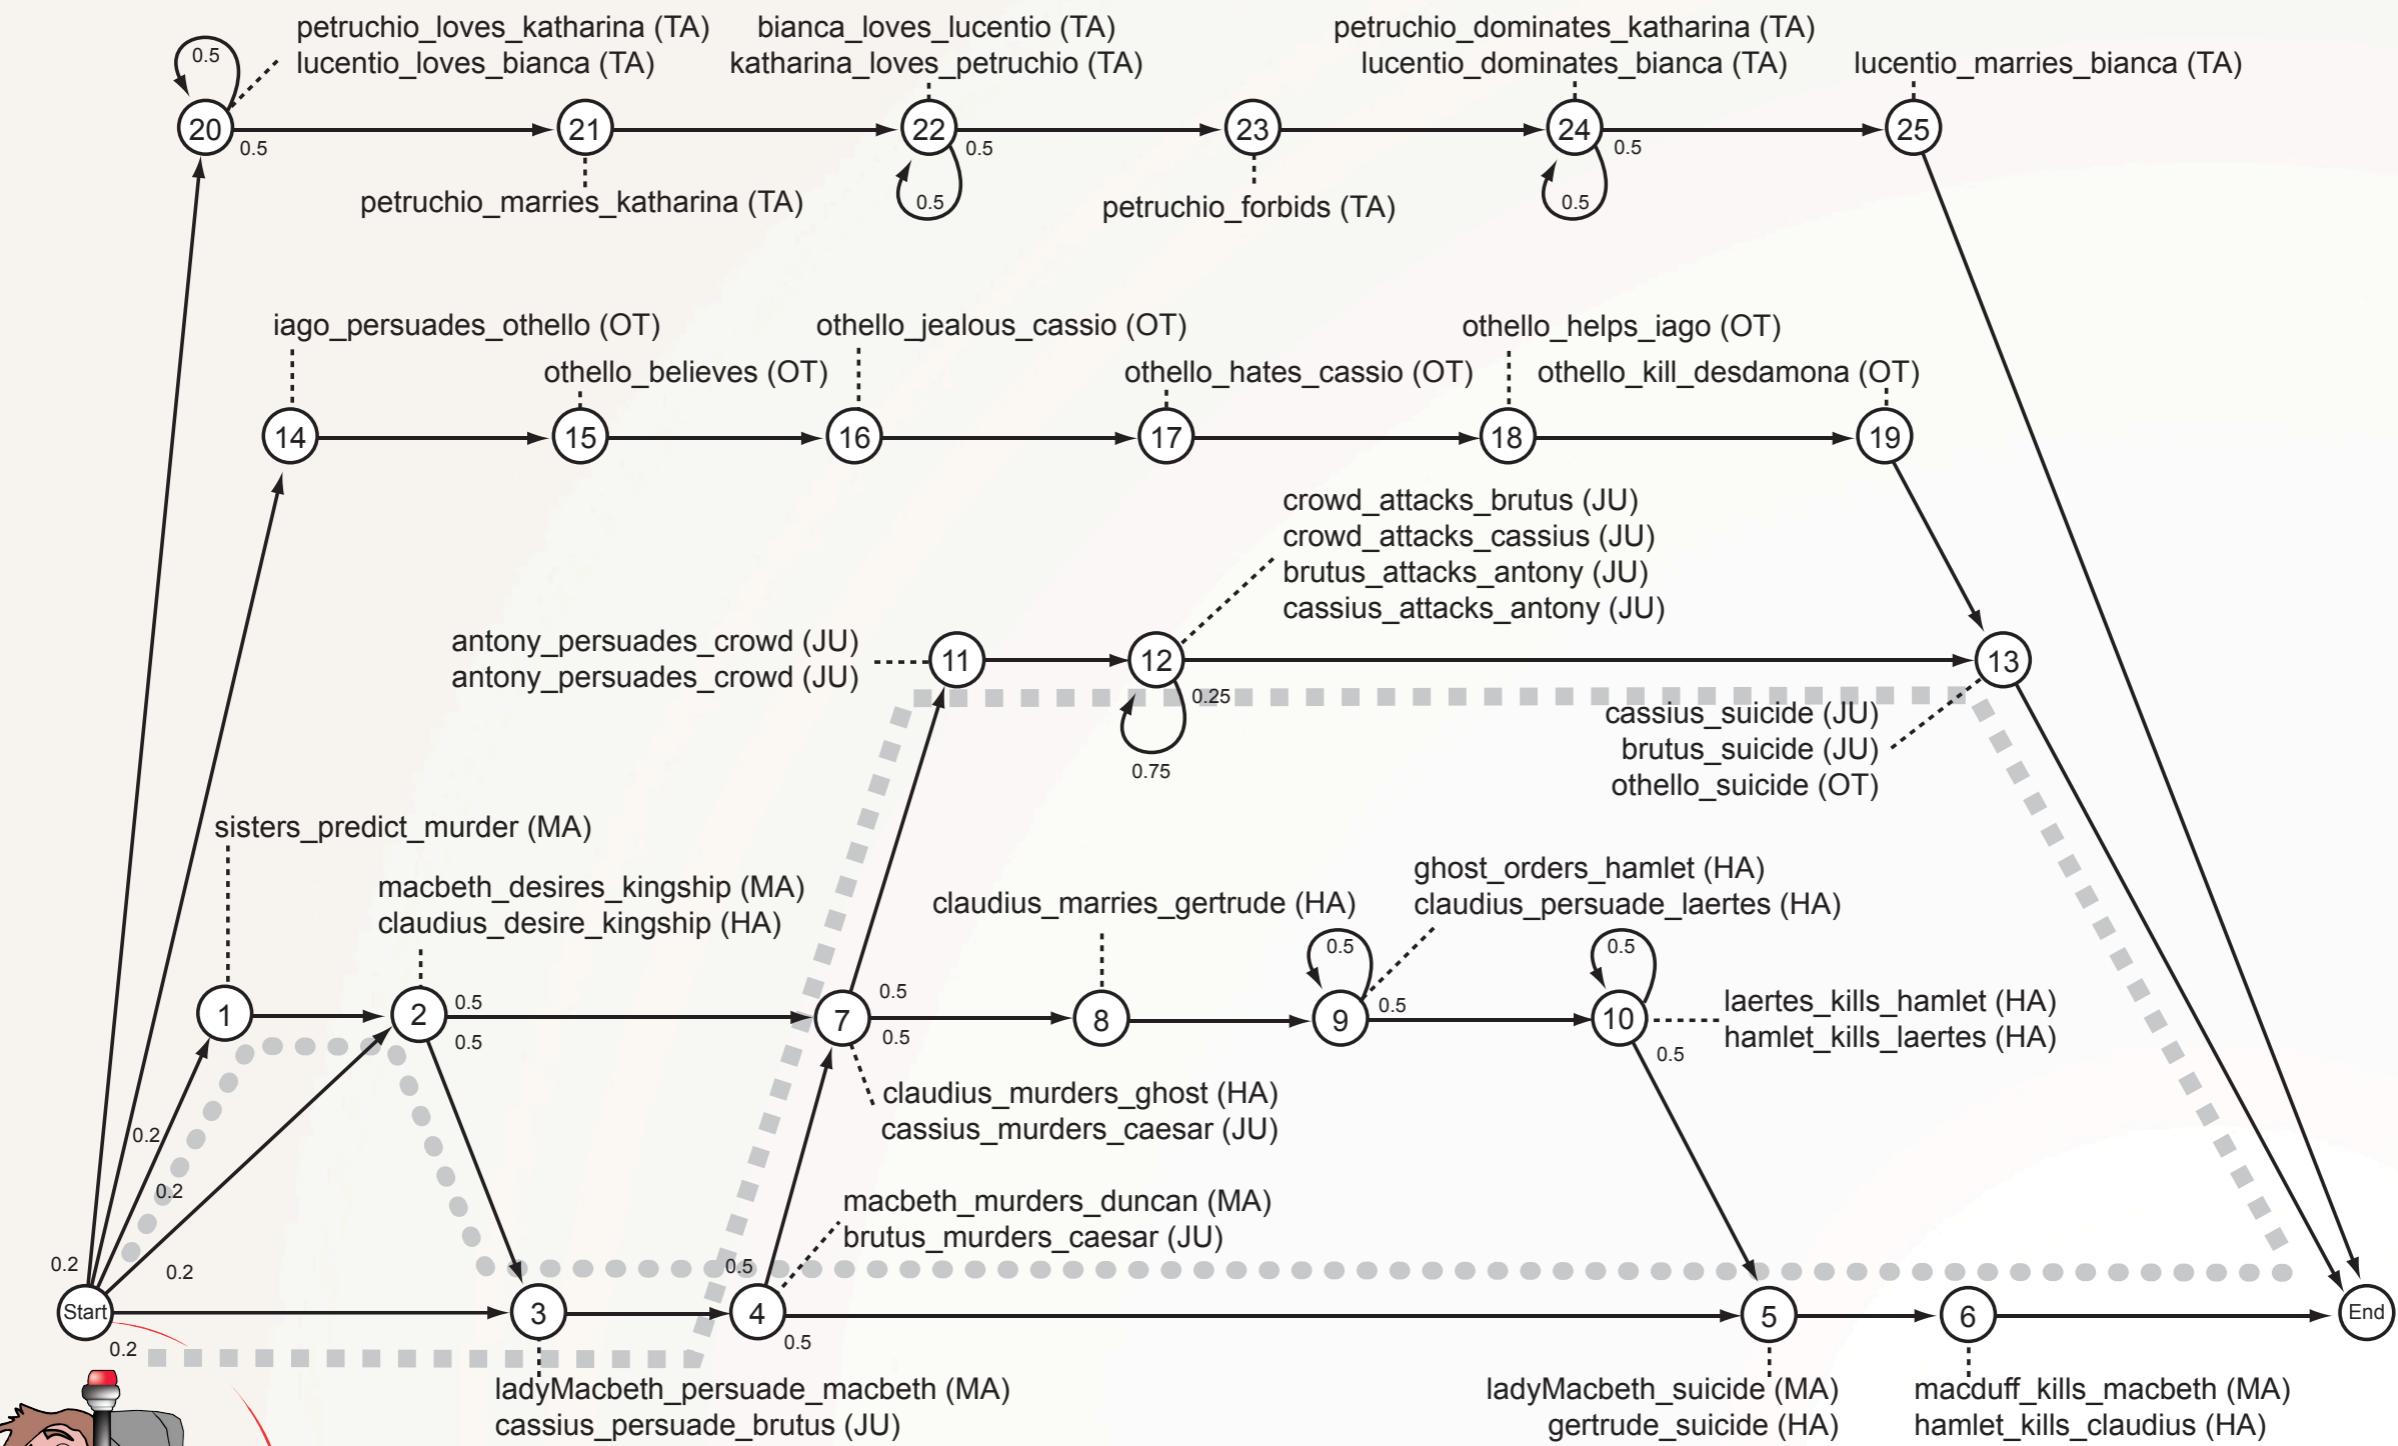
- Humans rely on a lifetime of experiences from which to explain stories, tell stories, or act in the real-world
- **Just-in-time domain model learning**
- Approaches:
  - Read natural language corpora & websites
  - Mine commonsense knowledge bases
  - Learn from humans



41

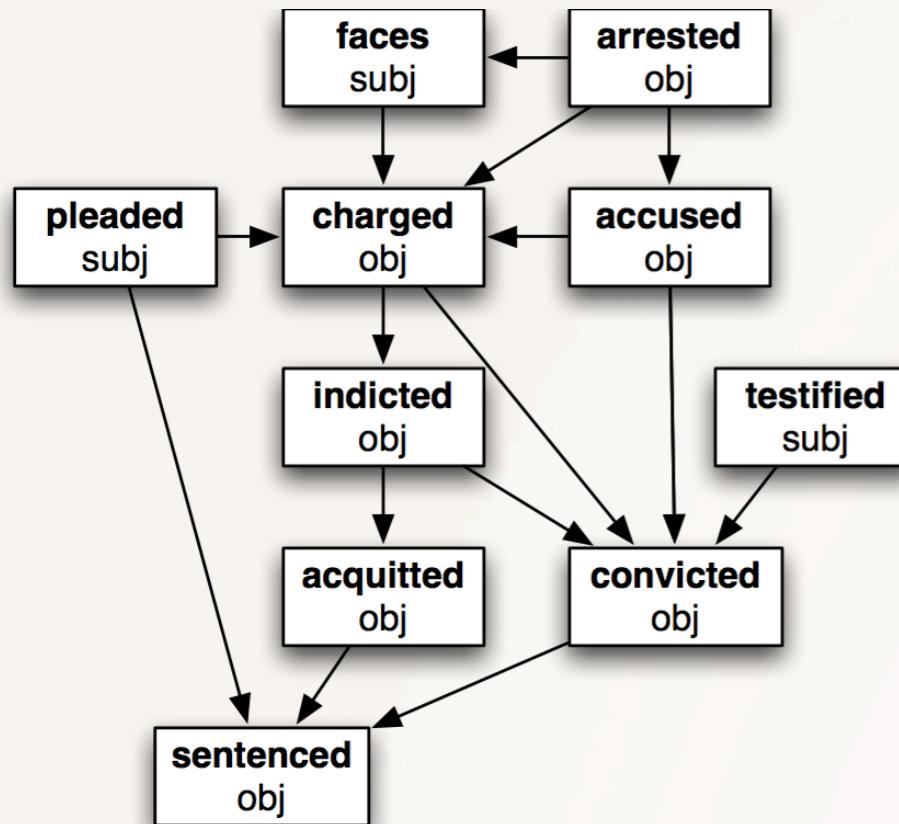


# Story grammar learning



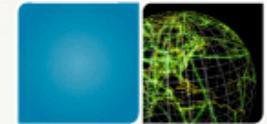


# Unsupervised script learning



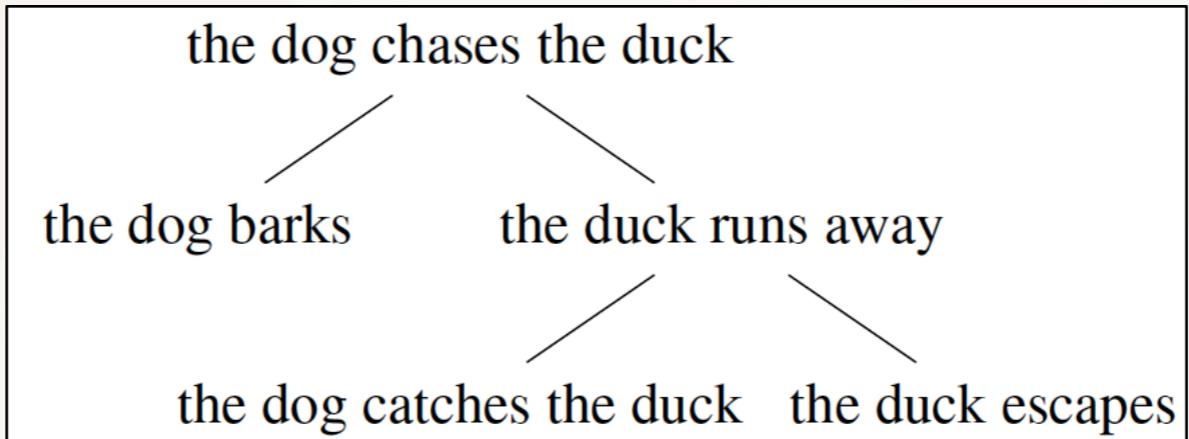
row	s <sub>1</sub>	s <sub>2</sub>	s <sub>3</sub>	s <sub>4</sub>
1	∅	walk into restaurant	∅	enter restaurant
2	∅	∅	walk to the counter	go to counter
3	∅	find the end of the line	∅	∅
4	∅	stand in line	∅	∅
5	look at menu	look at menu board	∅	∅
6	decide what you want	decide on food and drink	∅	make selection
7	order at counter	tell cashier your order	place an order	place order
8	∅	listen to cashier repeat order	∅	∅
9	pay at counter	∅	pay the bill	pay for food
10	∅	listen for total price	∅	∅
11	∅	swipe credit card in scanner	∅	∅
12	∅	put up credit card	∅	∅
13	∅	take receipt	∅	∅
14	∅	look at order number	∅	∅
15	∅	take your cup	∅	∅
16	∅	stand off to the side	∅	∅
17	∅	wait for number to be called	wait for the ordered food	∅
18	receive food at counter	get your drink	get the food	pick up order
19	∅	∅	∅	pick up condiments
20	take food to table	∅	move to a table	go to table
21	eat food	∅	eat food	consume food
22	∅	∅	∅	clear tray
22	∅	∅	exit the place	∅





# Story generation from corpora

- Fairy tale corpus



- Blogs

What happens next?

Your story so far

I saw it behind the door. The guy near the door put his gun at his side and reached for the knife. We would have to cut it open instead. I was overly excited and came very close to slicing open my finger, too. Luckily the blade narrowly missed my finger and stuck in the wall next to me.

Click to select the sentence that should go next.

"A grown man unwilling to sleep when he clearly needs it?"  
I regretted unbuttoning my shirt because his fingers were now gently caressing my collar bone where the chain traveled.  
A girl in a "little black party dress" and a domino mask came out and was like, "Hey, who are you?" and I thought that question was, like, totally inappropriate.  
Jack ripped the gun from her hands.  
**"I panicked."**  
Or potentially even getting involved in it.  
Now there are lots of spiders around our front door, so they're not necessarily black widows, but I'm still scared.  
She looks up at me and gives me a tired smile.  
"Aw... you missed!" I whined. "That was supposed to end my affected inner anguish!"  
This time I thought I was smarter and finangled it to lay down the stairs, so it wouldn't get stuck in the ceiling again.

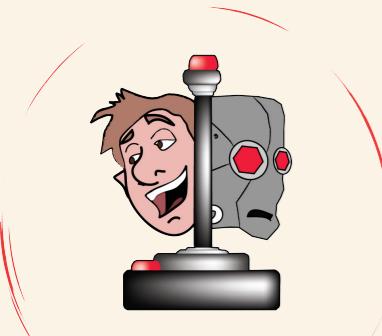
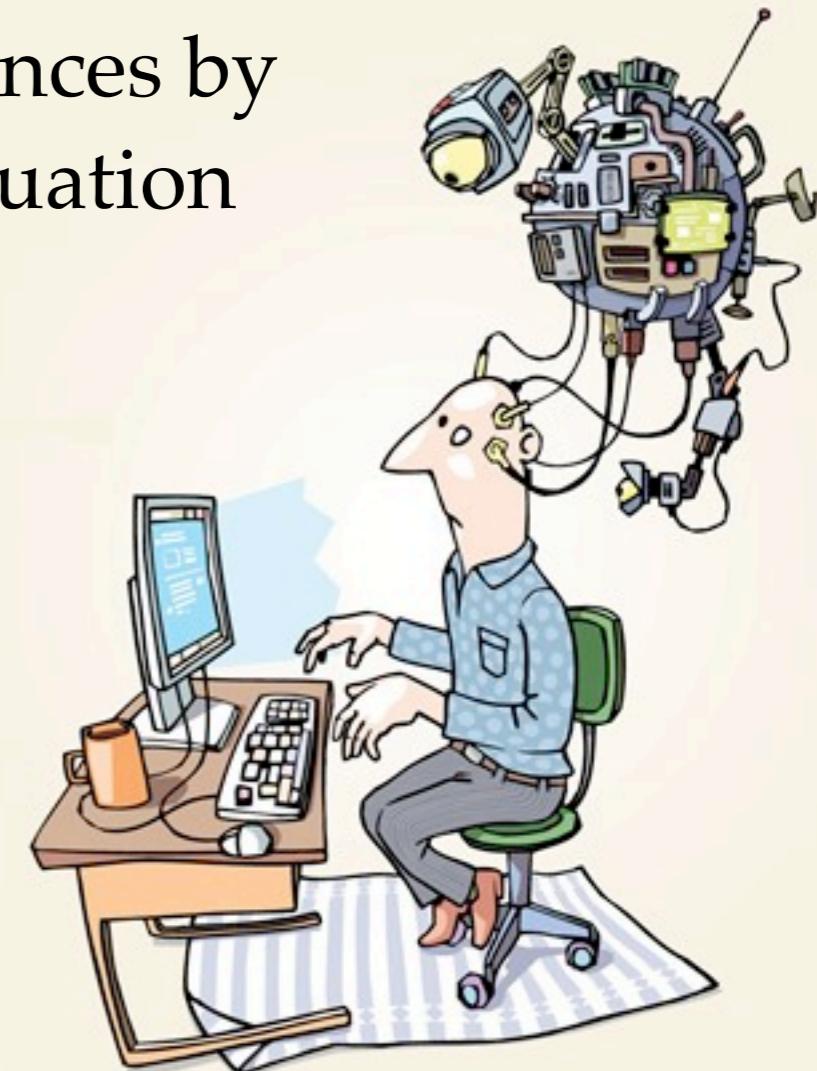
Accept



McIntyre & Lapata. Proc. ACL 2009 Conference.  
Swanson & Gordon. ACM TIIS, 2(3), 2012.

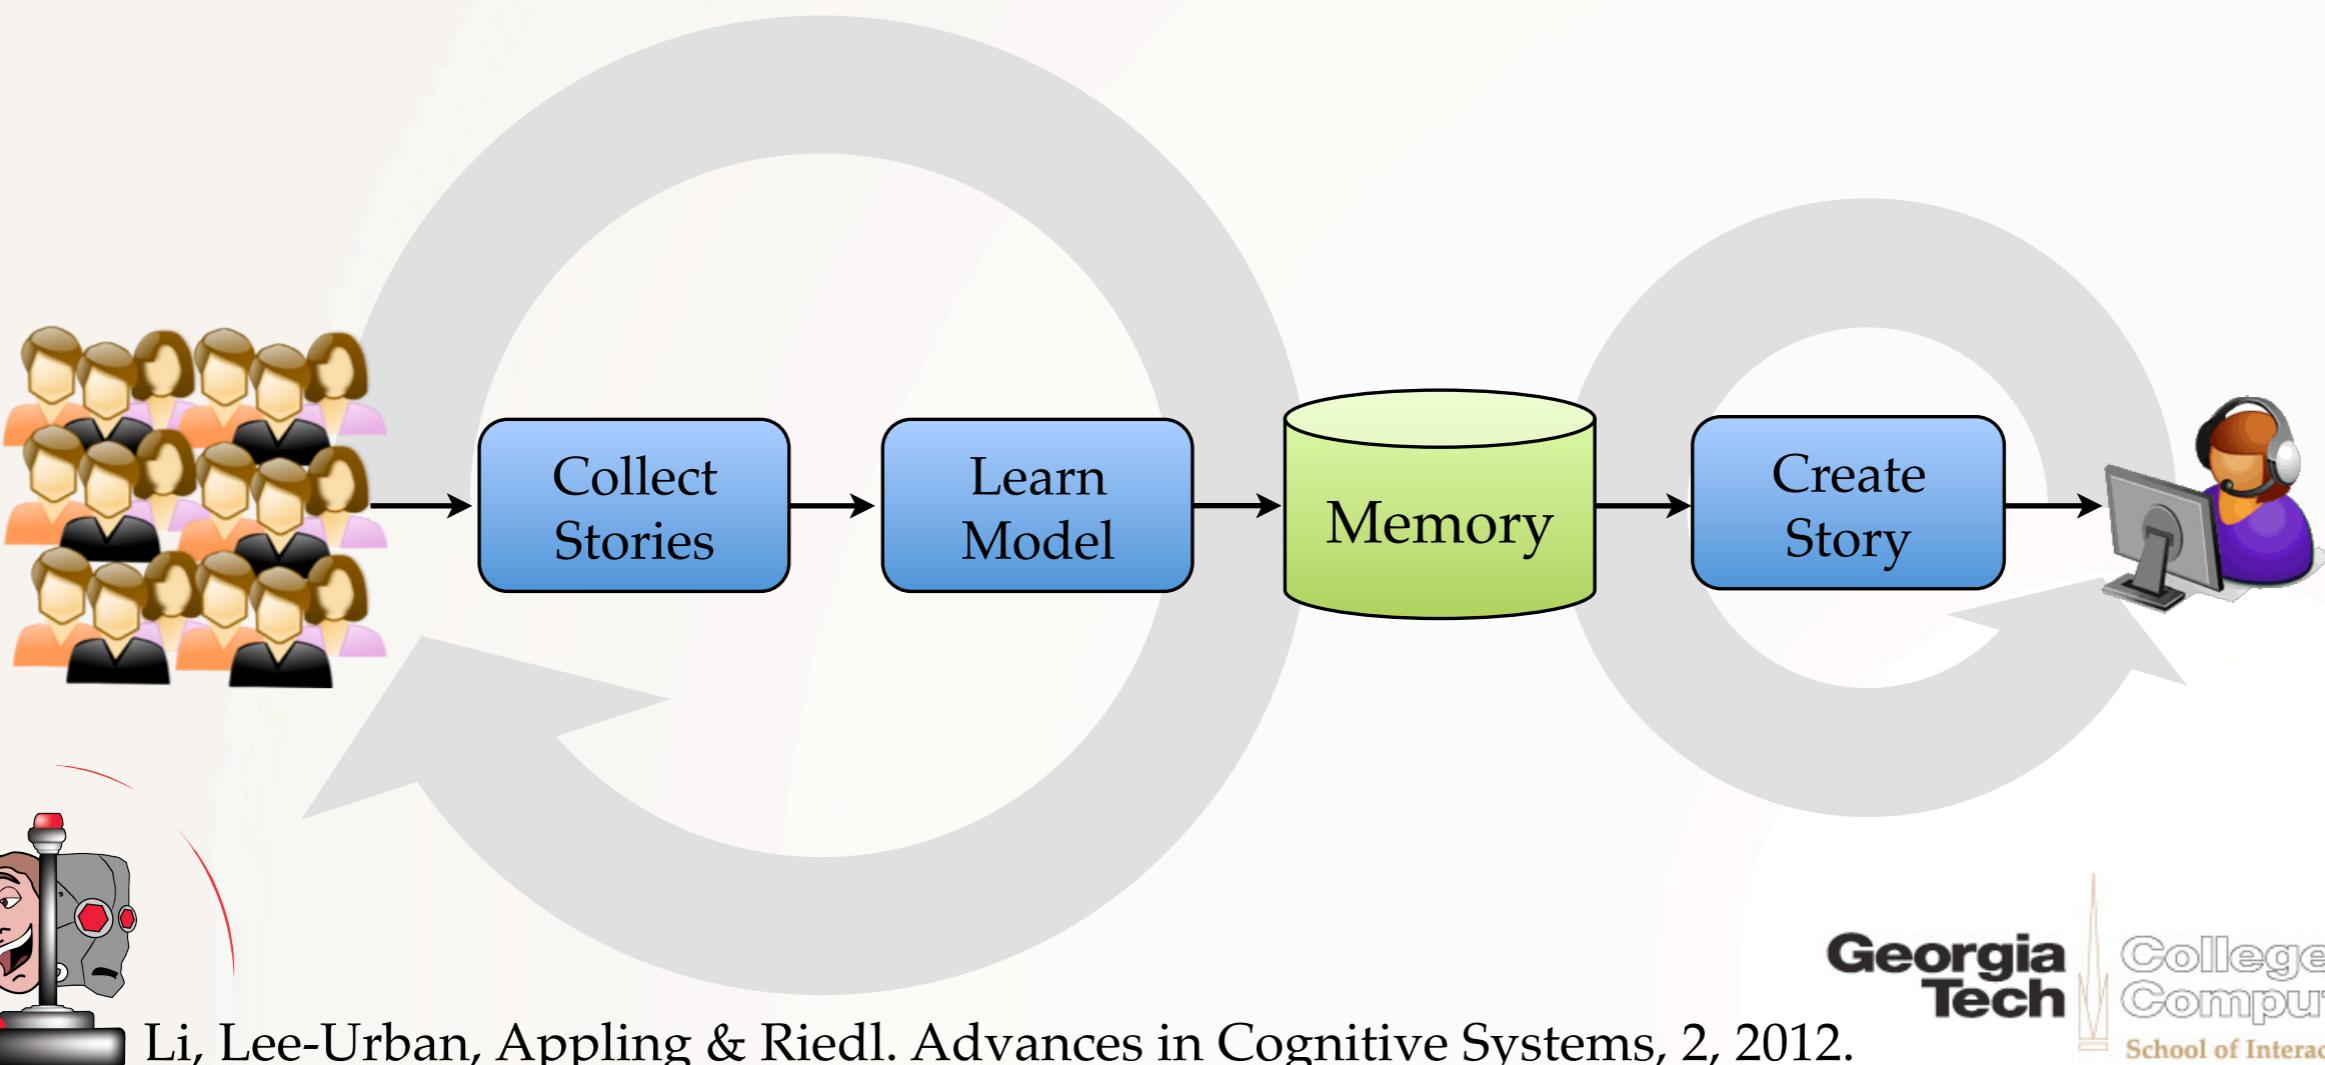
# Crowdsourcing narrative intelligence

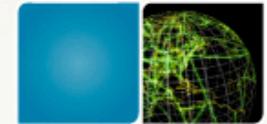
- A crowd of humans on the web → a supercomputer
- Crowd as a massive distributed knowledge base
- **Insight:** humans learn from stories
- Use a crowd to simulate a lifetime of experiences by asking people to tell stories about a given situation
- Crowdsource a highly specialized corpus of narrative examples and learn a generalized model of sociocultural situations





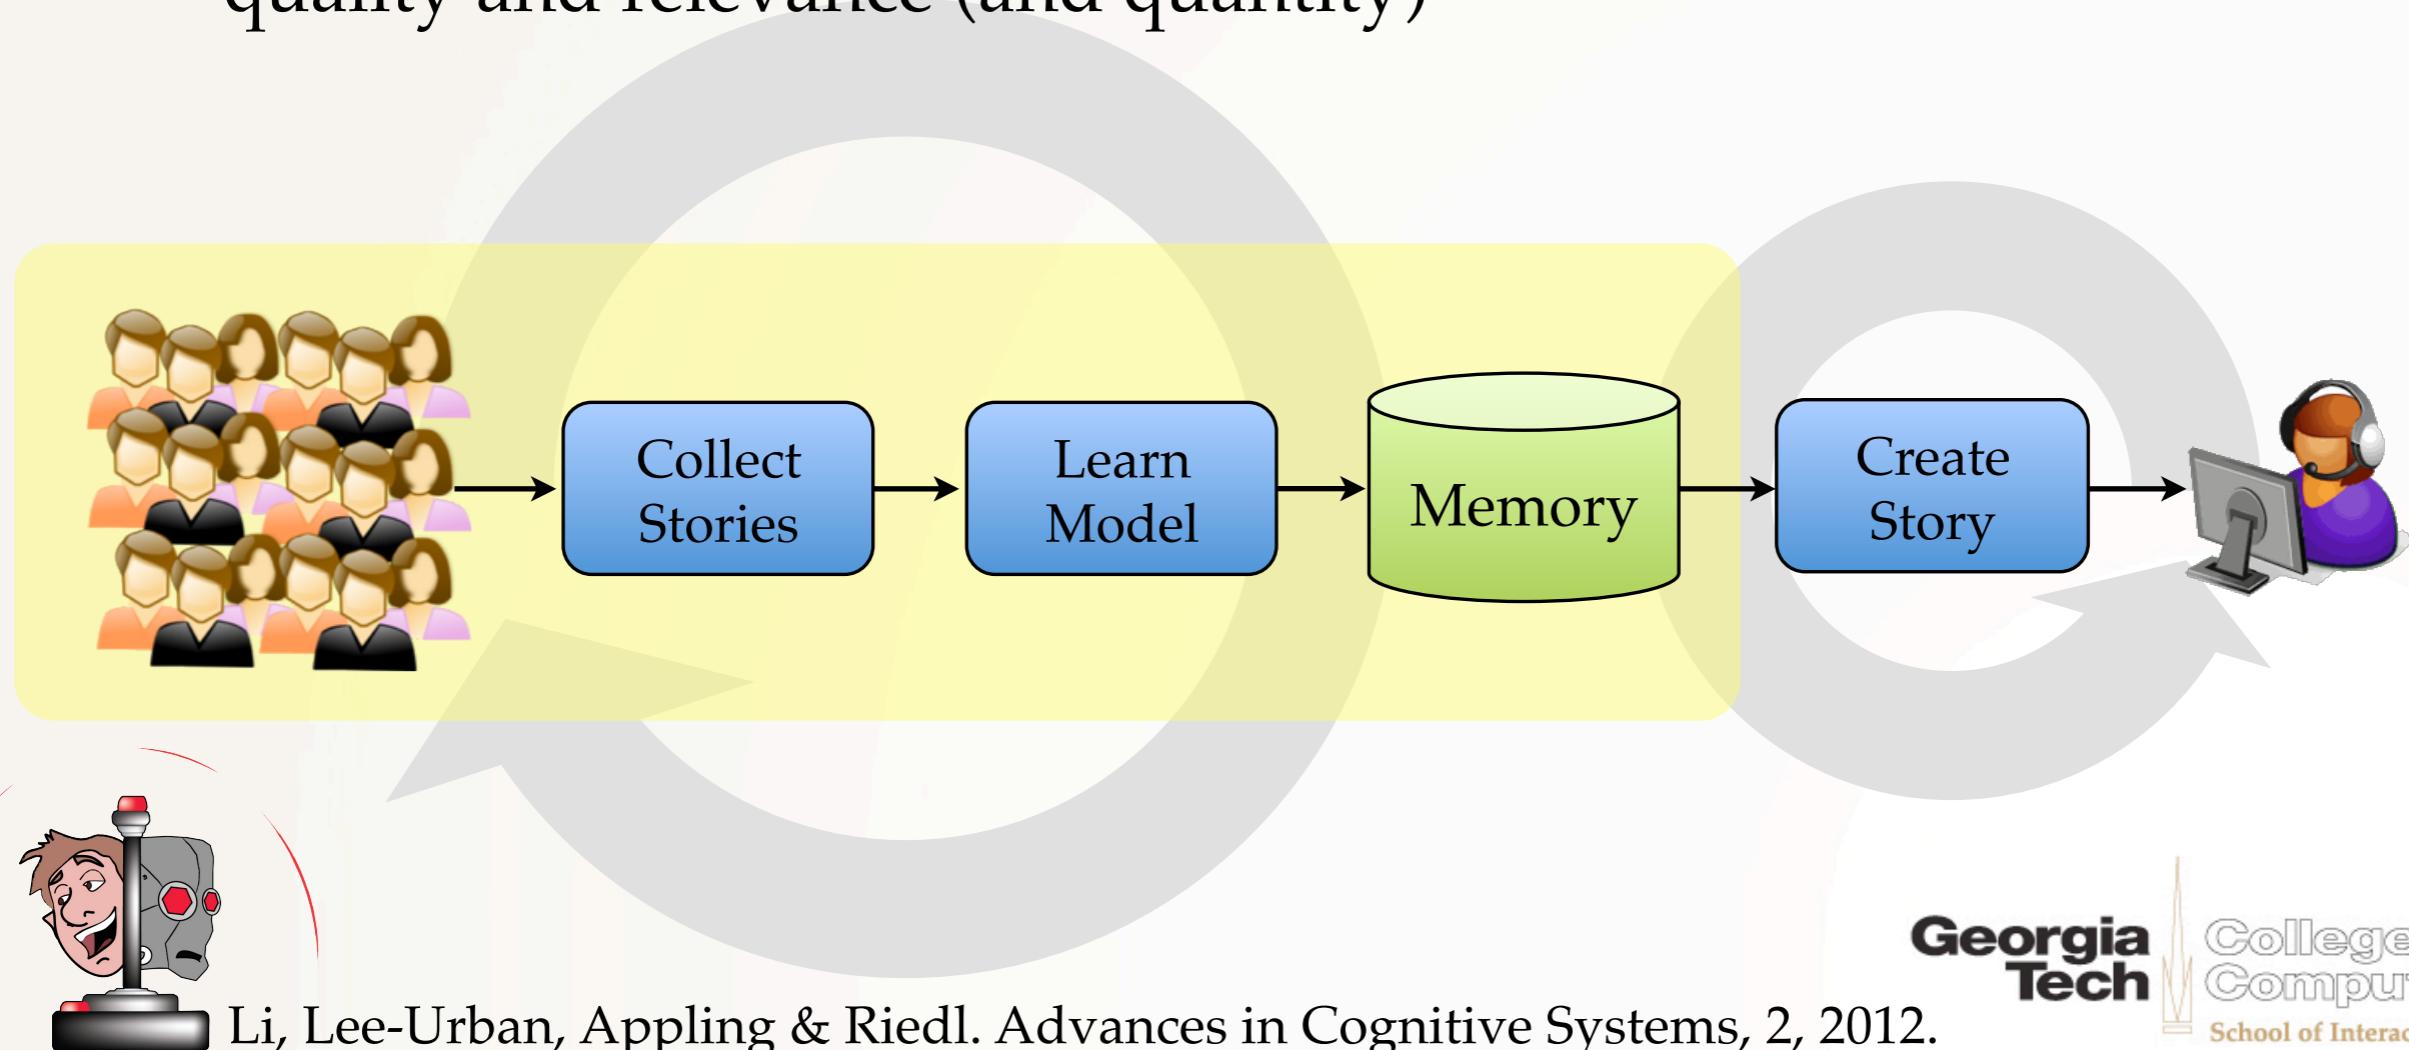
# Story generation pipeline

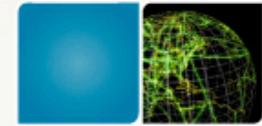




# Story generation pipeline

- Just-in-time model learning
  - Crowd workers write typical stories about a given situation
  - Learn domain model from what was written
  - Human narratives implicitly encode causality and obey the maxims of quality and relevance (and quantity)





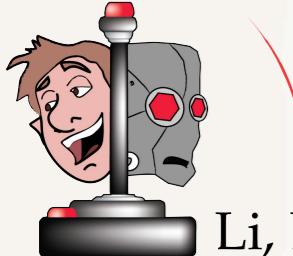
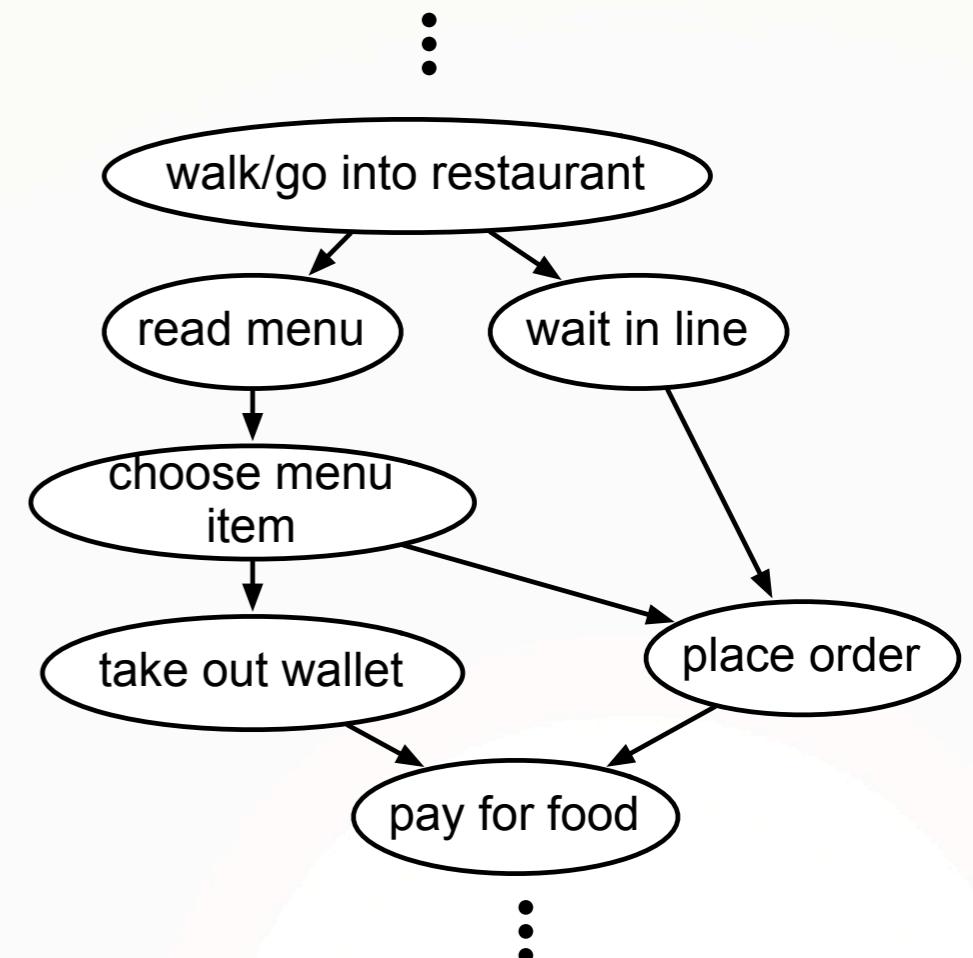
# Sociocultural knowledge representation

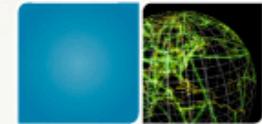
- Model a situation as a **script**
  - Representation of procedural knowledge
  - Tells one what to do and when to do it
  - Correlated with expertise

- Set of temporally ordered events

- What are the primitive events?

- How are events ordered?



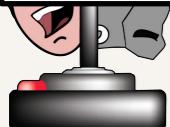


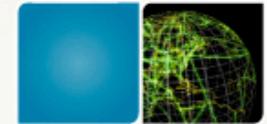
# Learning to tell stories

1. Query the crowd
2. Identify the salient events
3. Determine event ordering
4. Mutually exclusive events

- Crowd control:
  - Segment narrative
  - Use one verb per sentence
  - Avoid conditionals and compound structures
  - Avoid using pronouns

<b>Story A</b>	<b>Story B</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.	...
...	

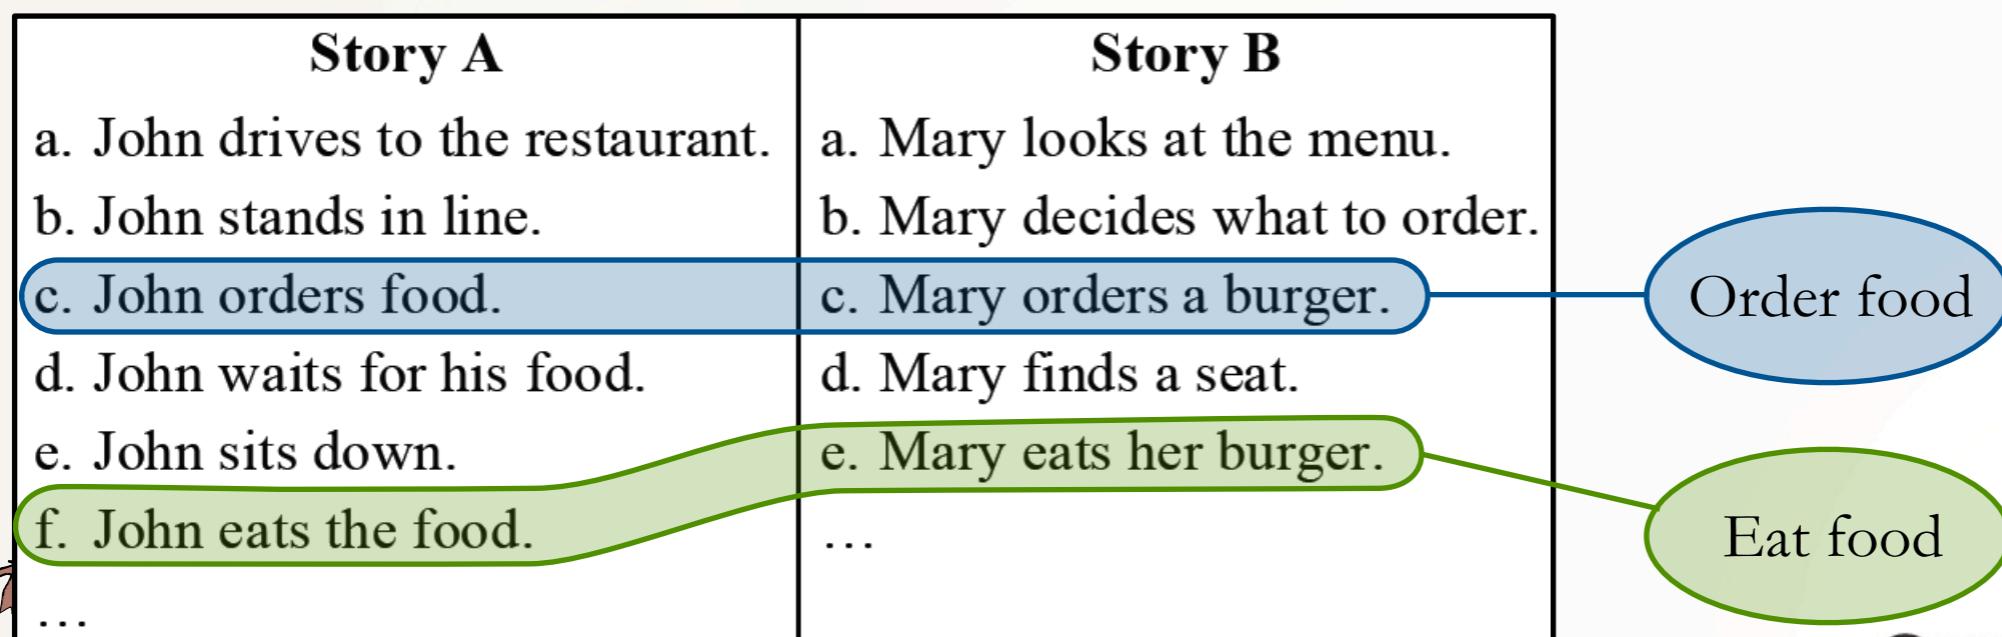




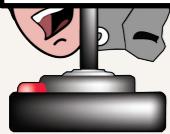
# Learning to tell stories

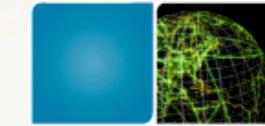
1. Query the crowd
2. Identify the salient events
3. Determine event ordering
4. Mutually exclusive events

- Crowd control simplifies NLP
- Compute semantic similarity between sentences
- Cluster sentences into events



\* Precision  $\approx 85\%$

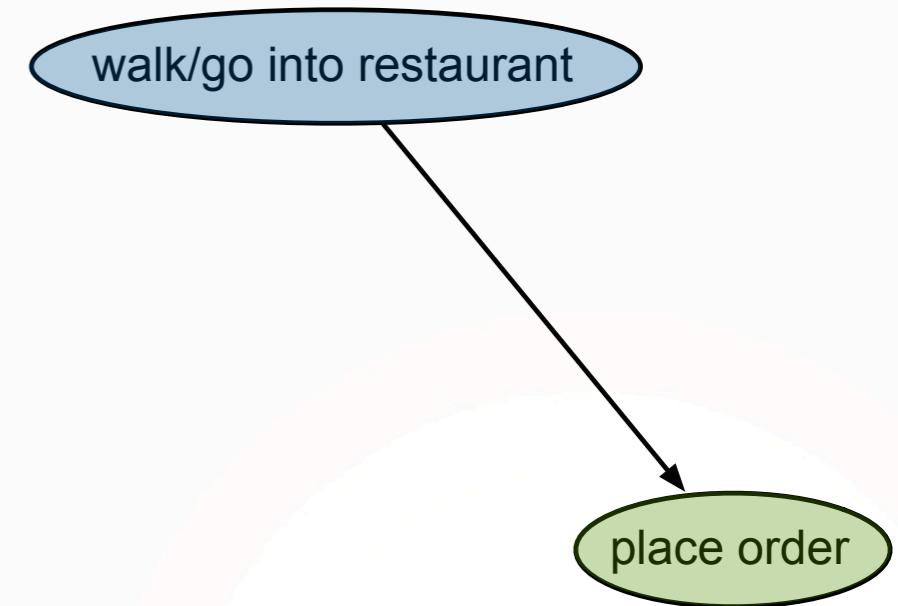
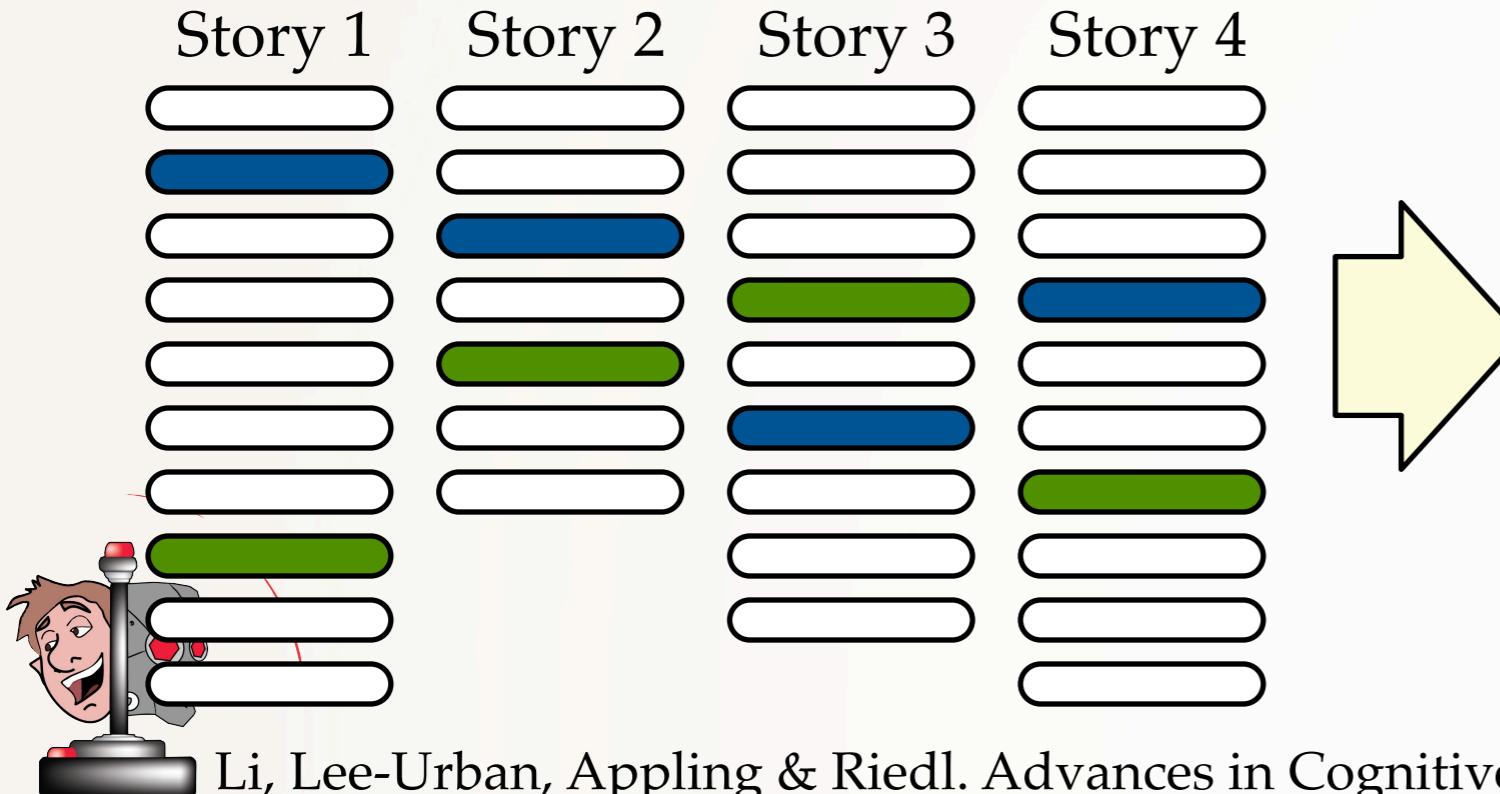


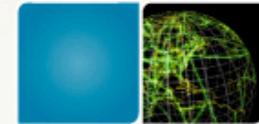


# Learning to tell stories

1. Query the crowd
  2. Identify the salient events
  - 3. Determine event ordering**
  4. Mutually exclusive events

- Seek evidence for temporal relations
  - Binomial confidence testing
  - Search for the most compact graph that explains the stories

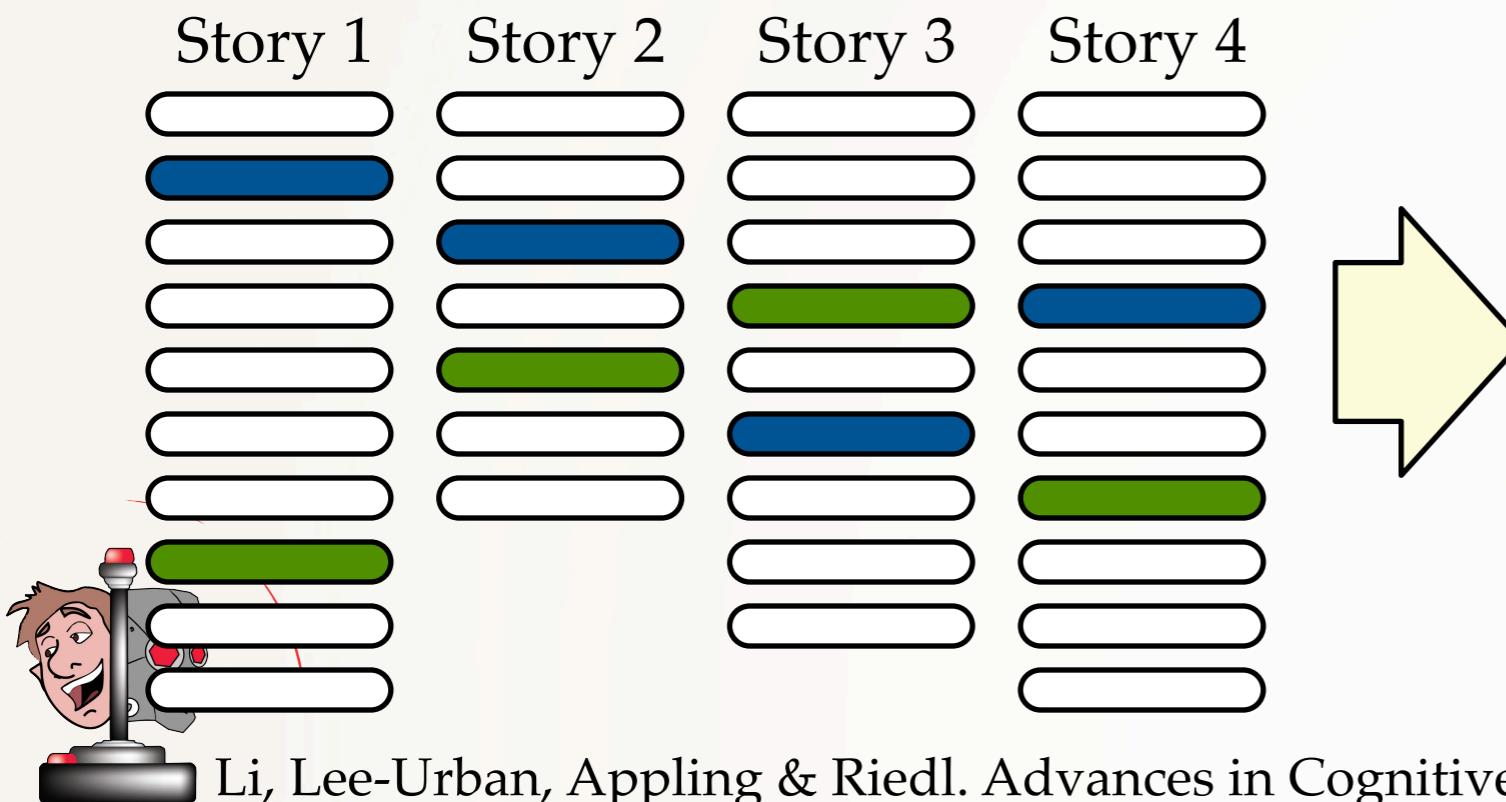


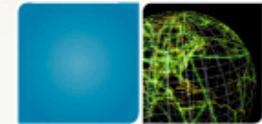


# Learning to tell stories

1. Query the crowd
2. Identify the salient events
3. Determine event ordering
4. Mutually exclusive events

- Seek evidence for temporal relations
- Binomial confidence testing
- Search for the most compact graph that explains the stories

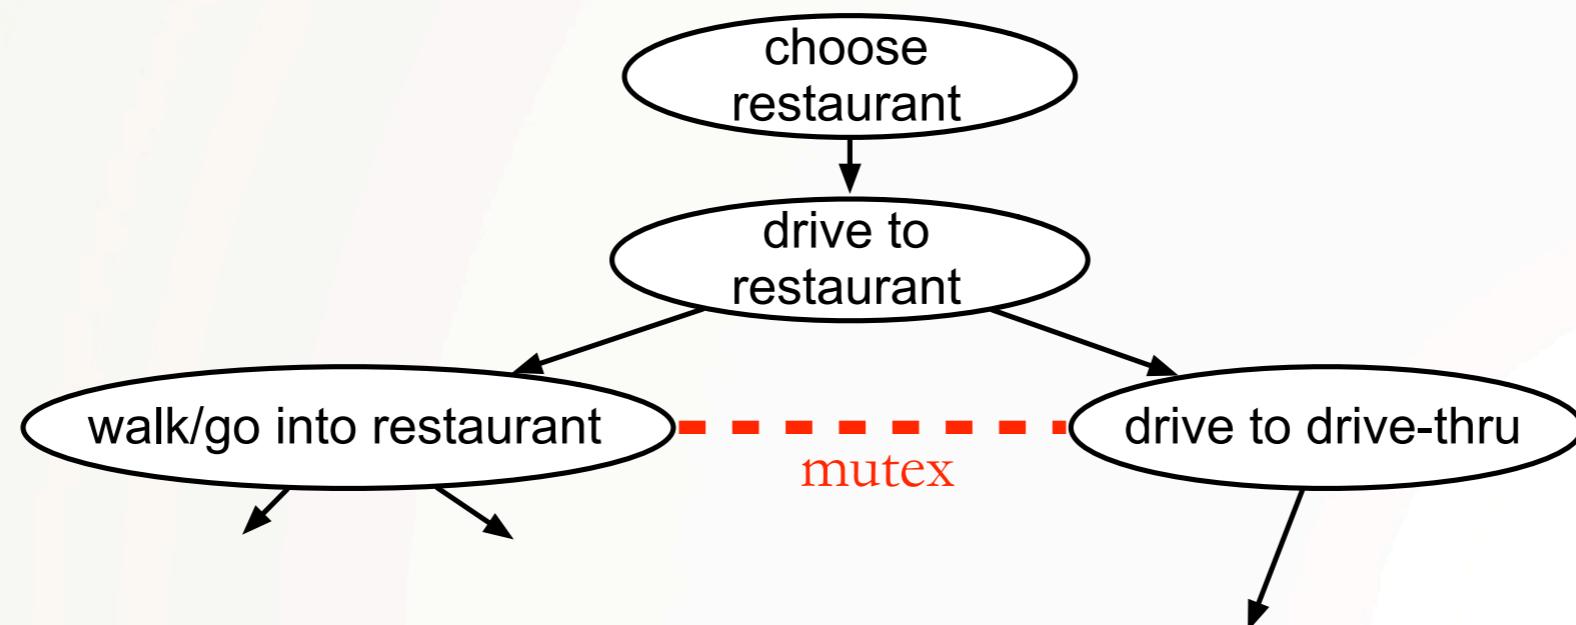




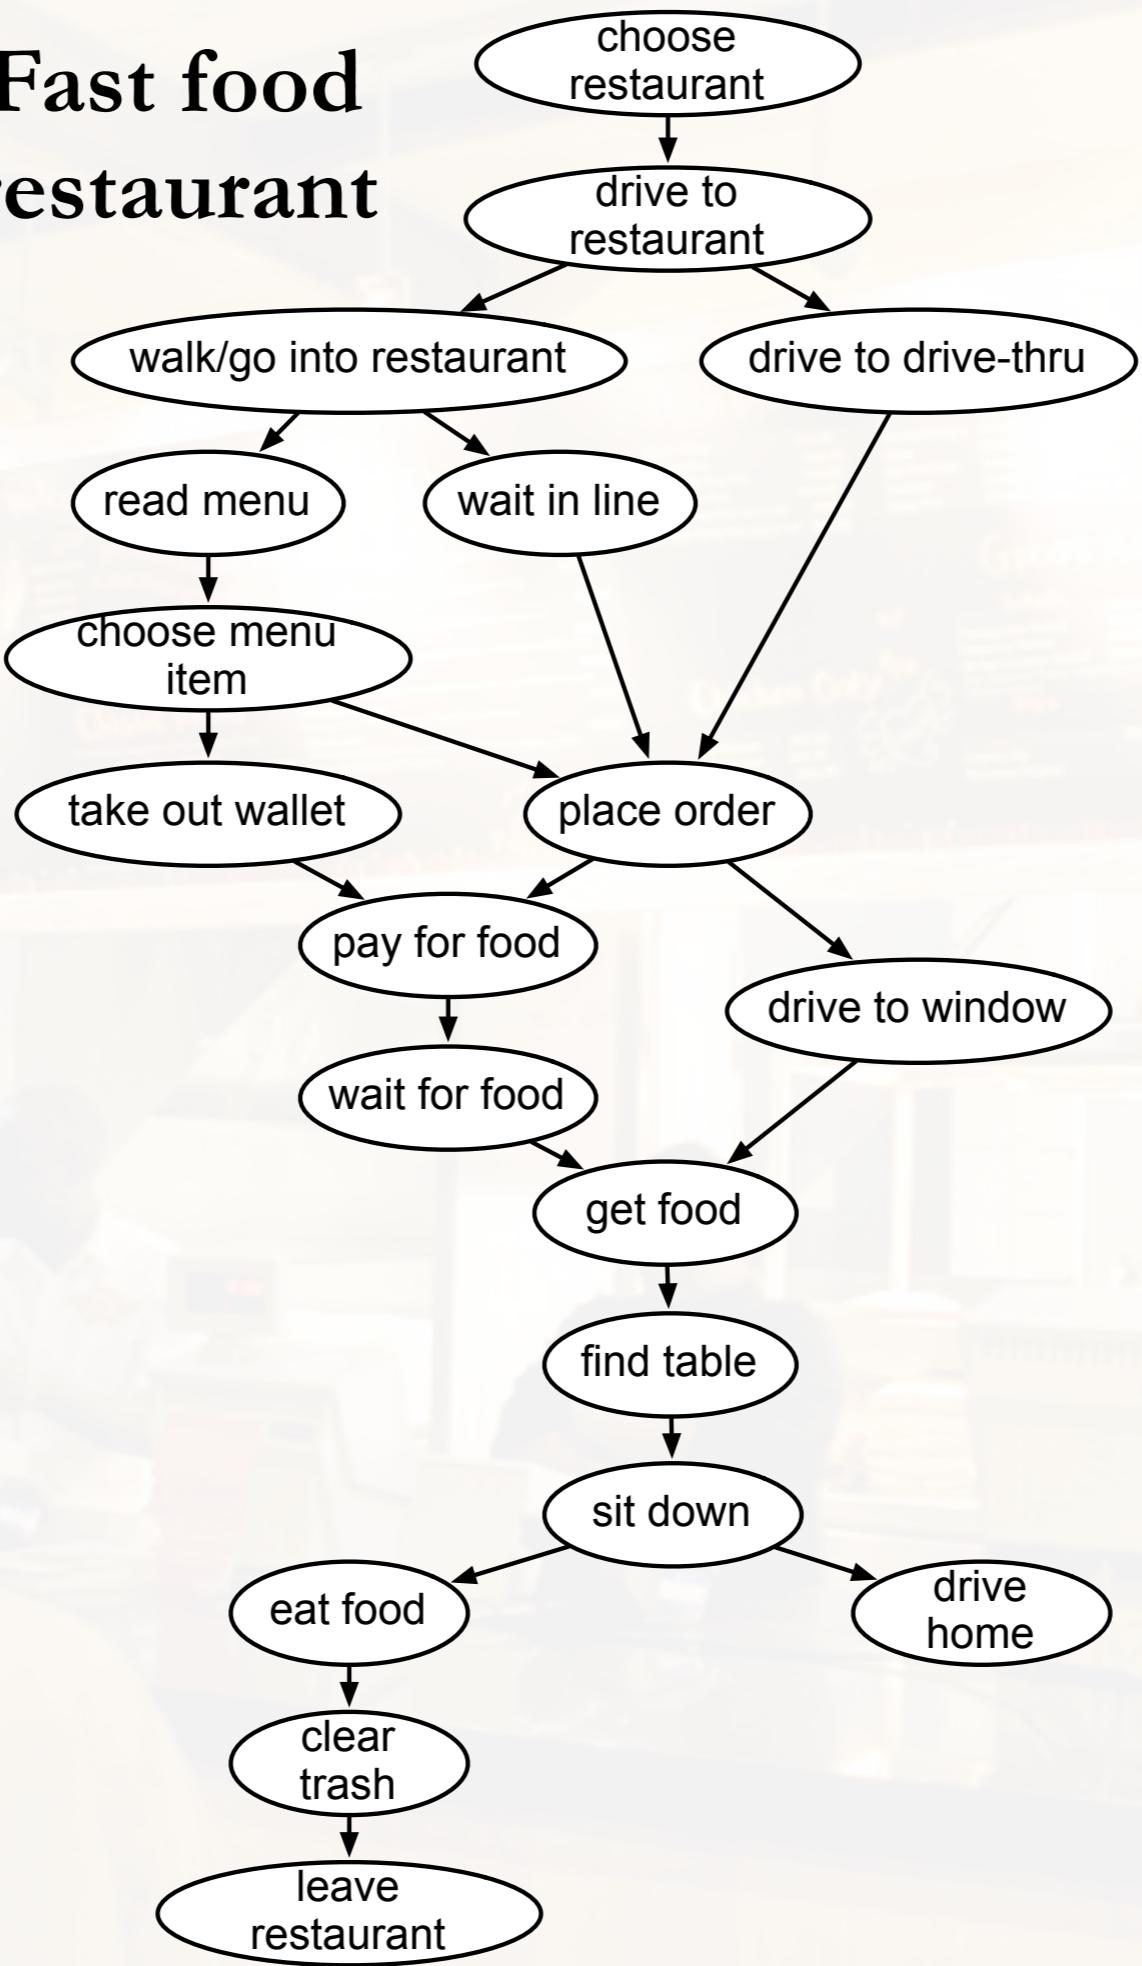
# Learning to tell stories

1. Query the crowd
2. Identify the salient events
3. Determine event ordering
4. Mutually exclusive events

- Measure mutual information between events
- Mutual information is high and co-occurrence is low
- Generalization of “or” relations



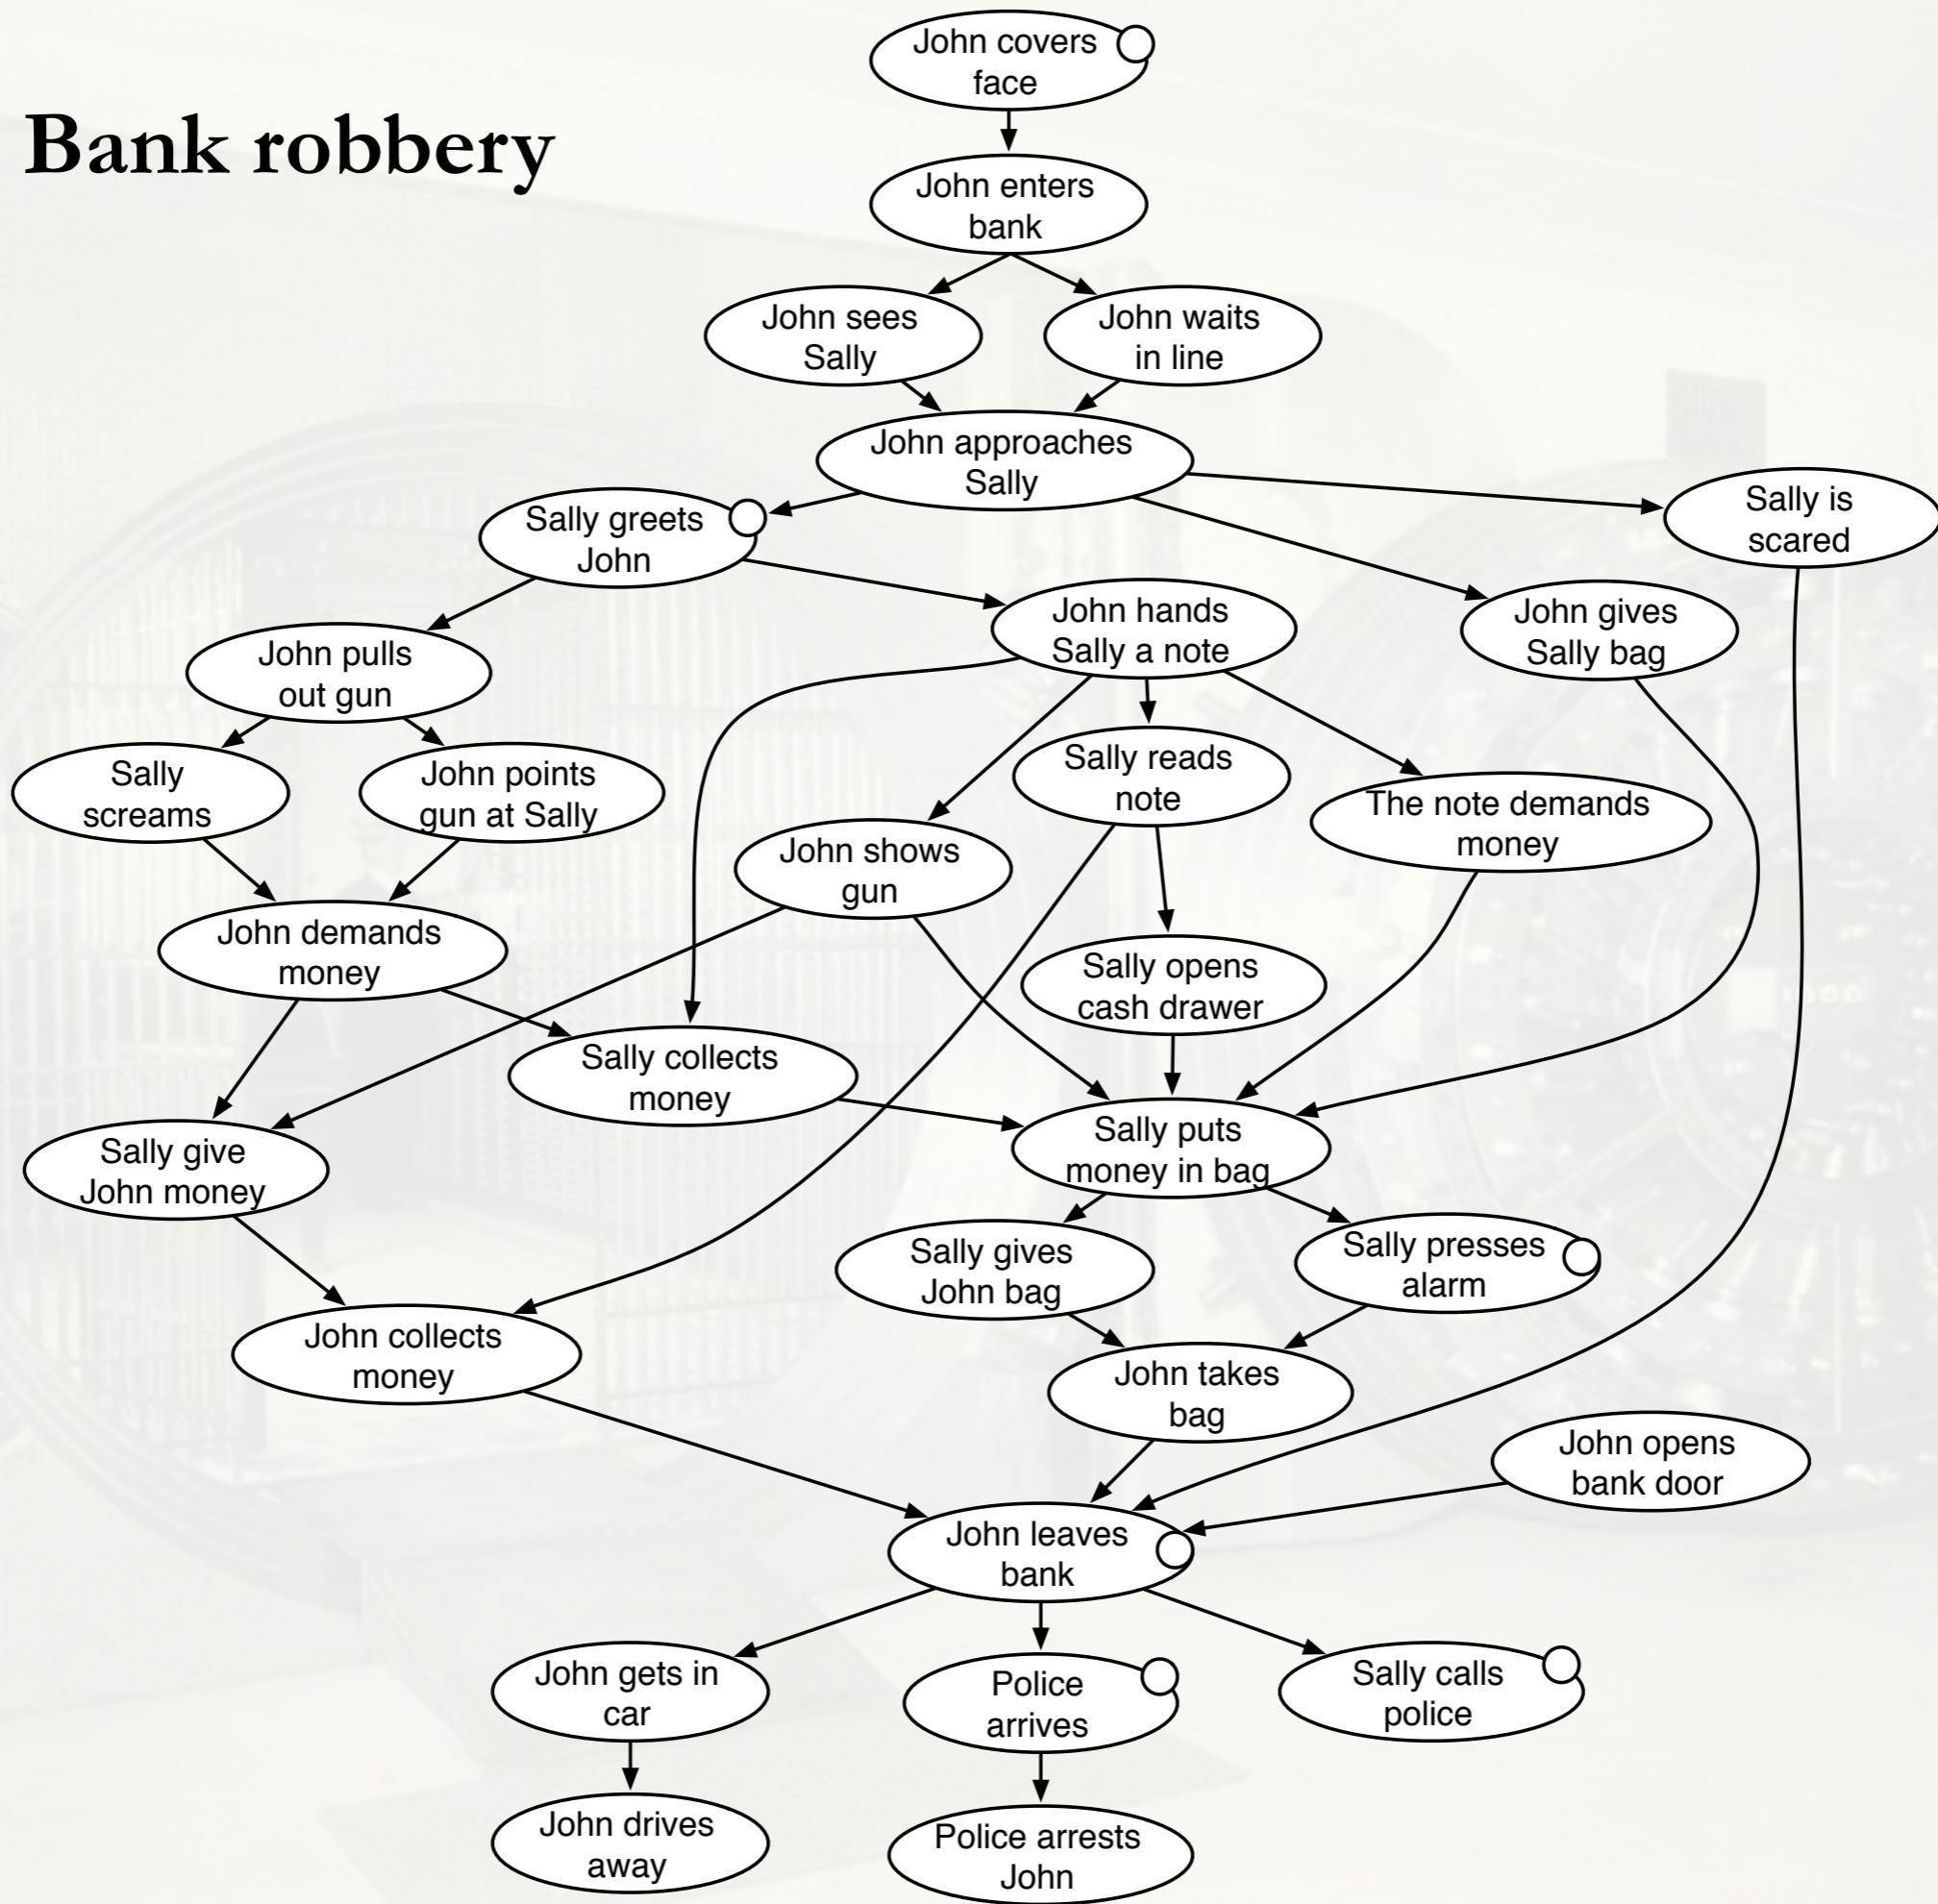
# Fast food restaurant

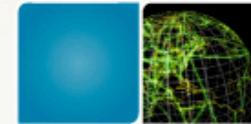




## Going on a date to the movies

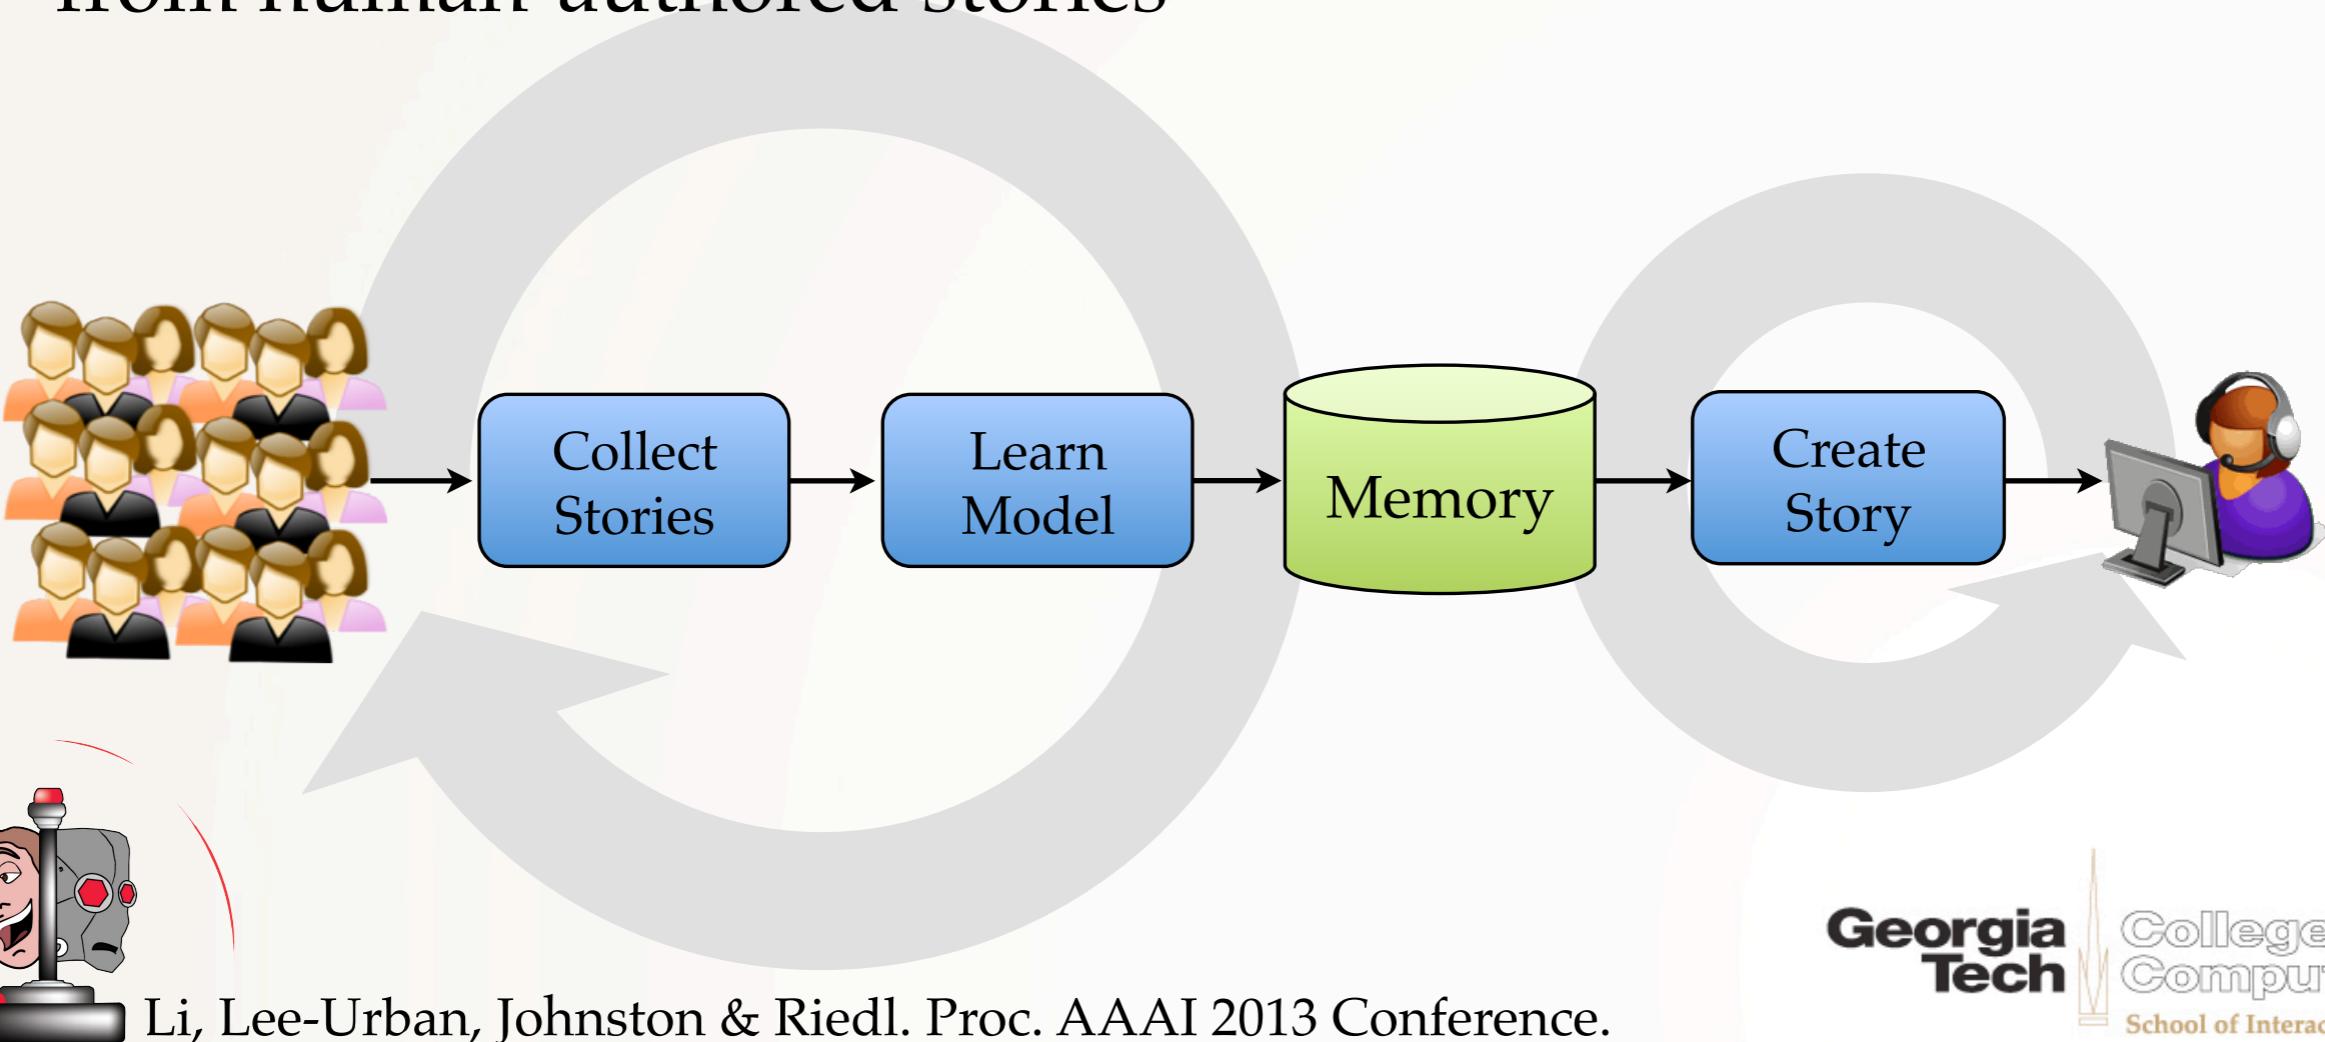
# Bank robbery

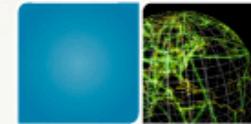




# Narrative generation

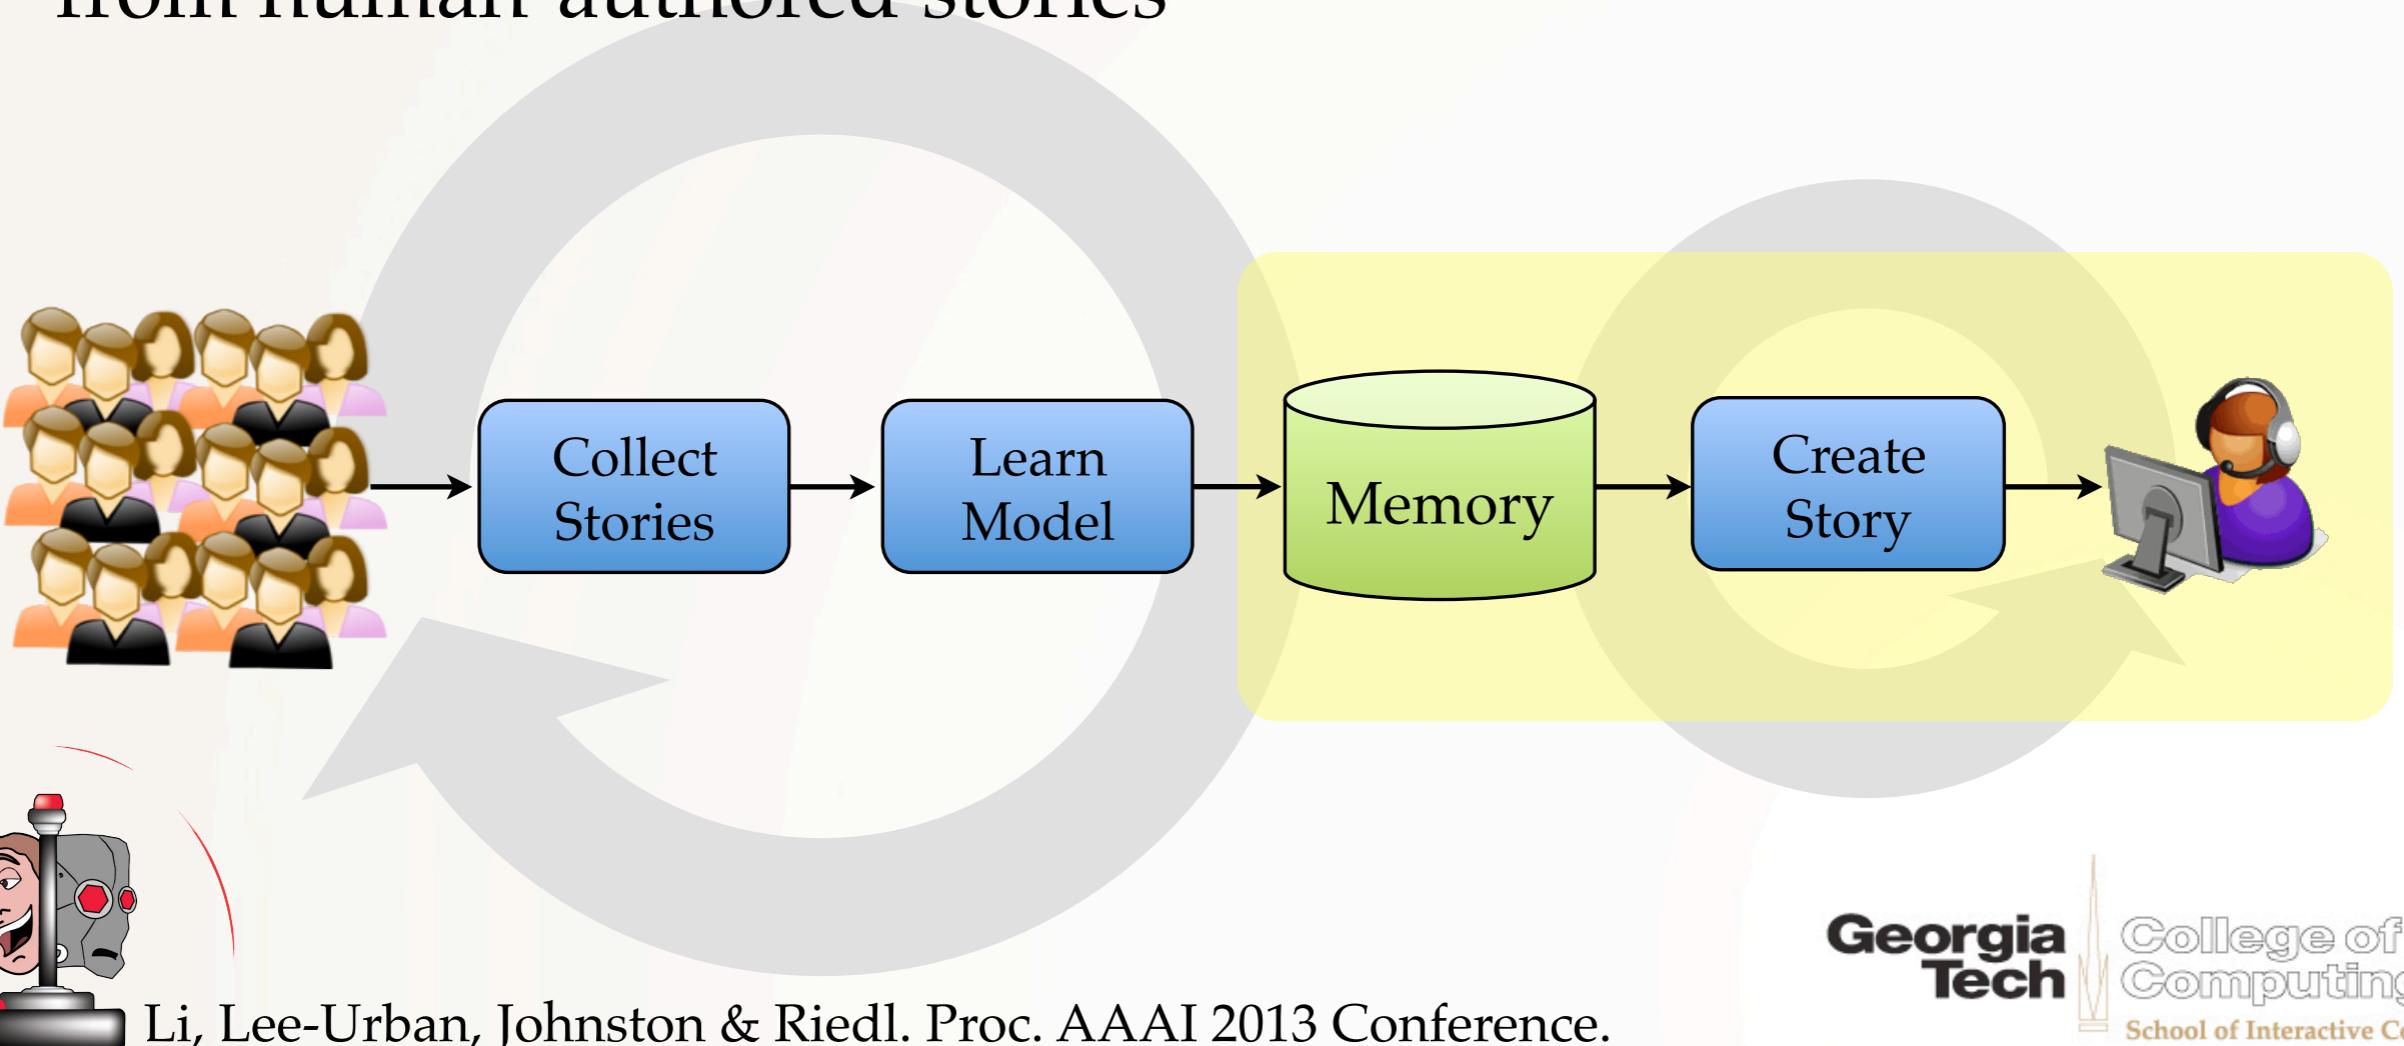
- Script defines a space of linear sequences
- Random walk, preserving temporal order and mutual exclusion
- Evaluation shows generated stories statistically indistinguishable from human-authored stories



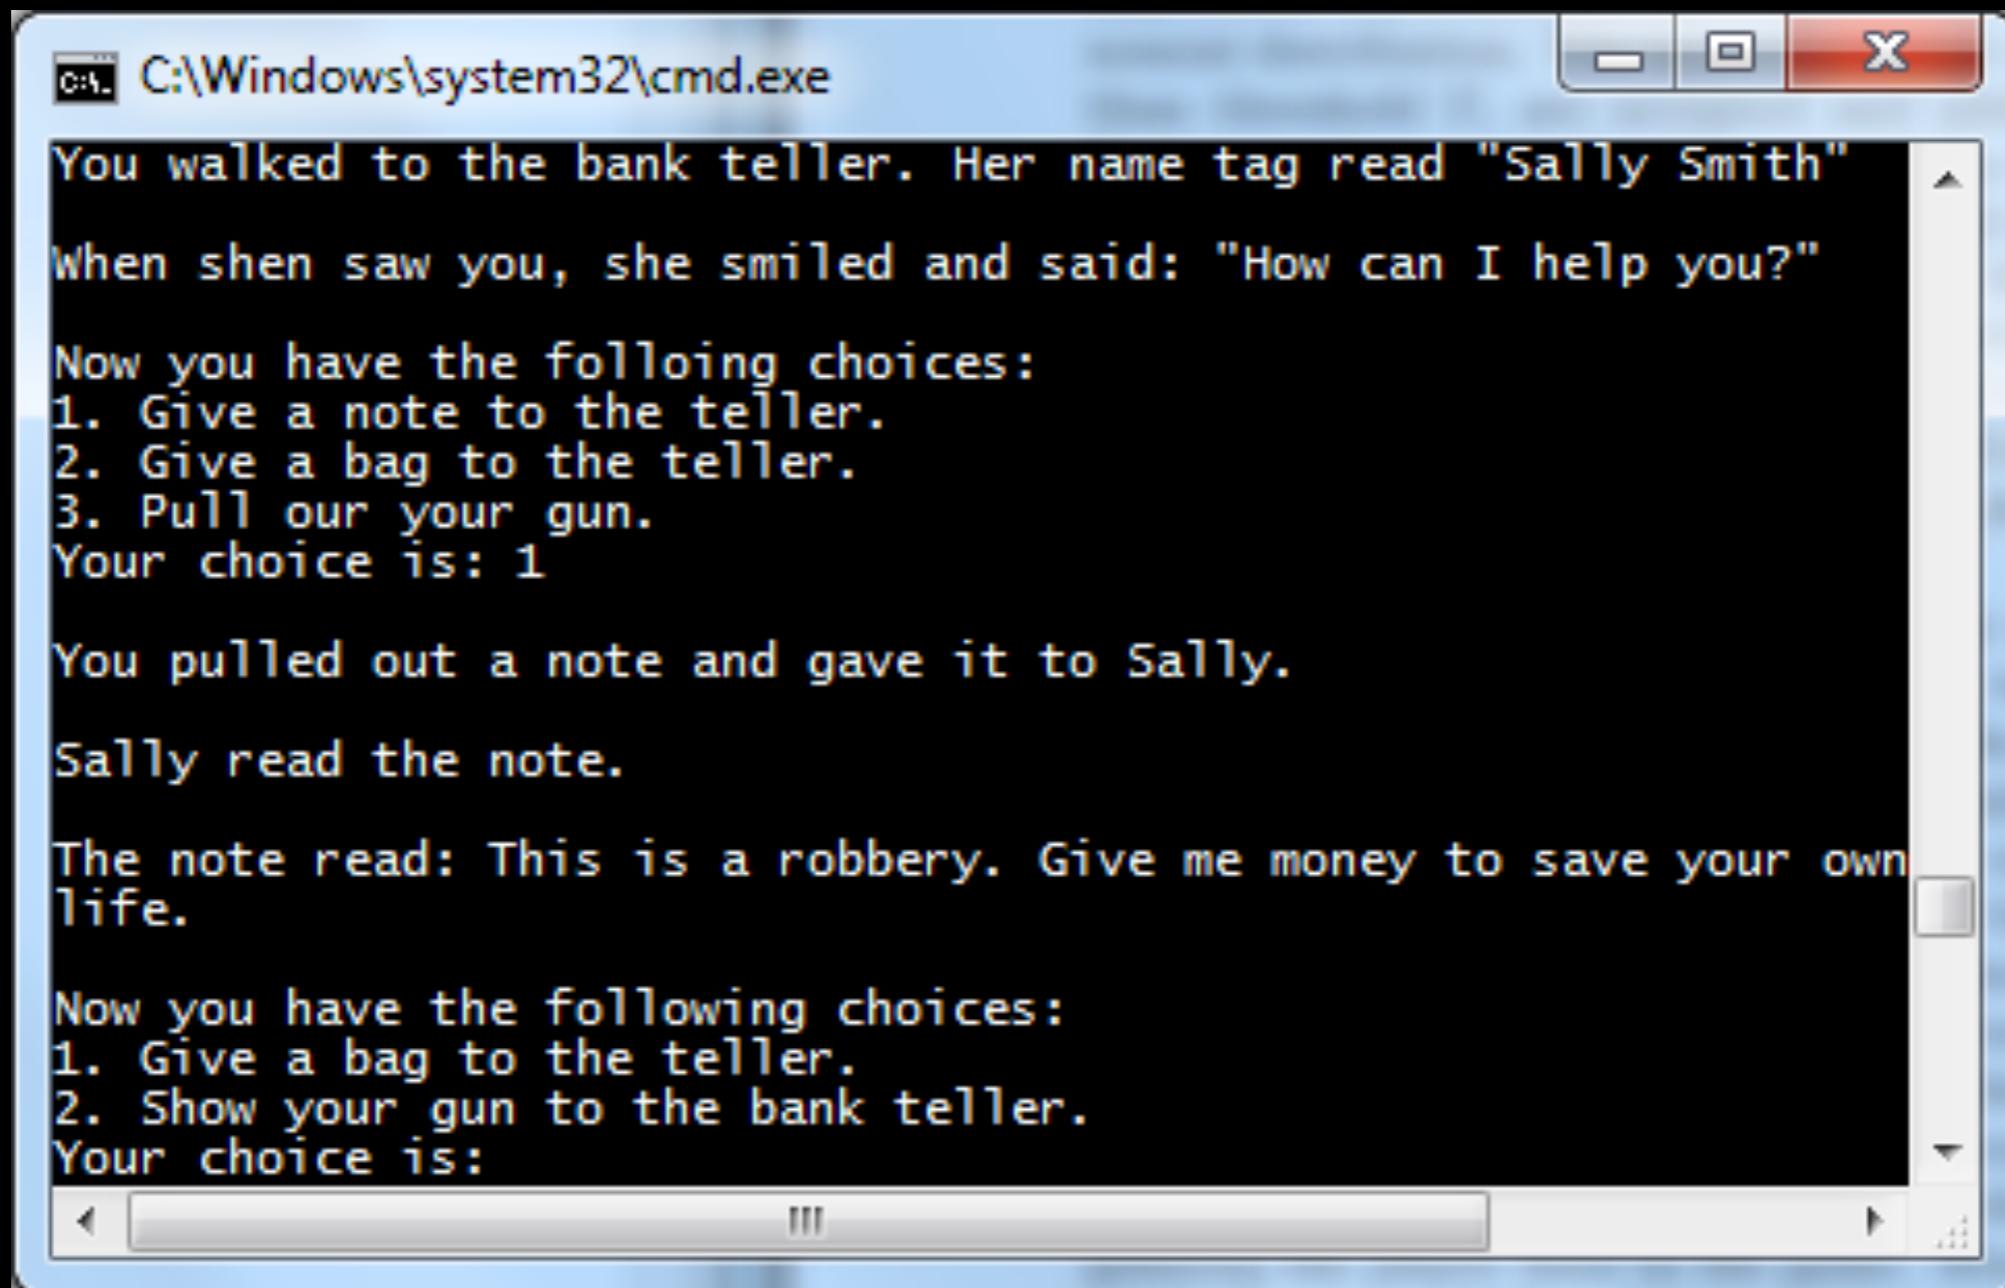


# Narrative generation

- Script defines a space of linear sequences
- Random walk, preserving temporal order and mutual exclusion
- Evaluation shows generated stories statistically indistinguishable from human-authored stories



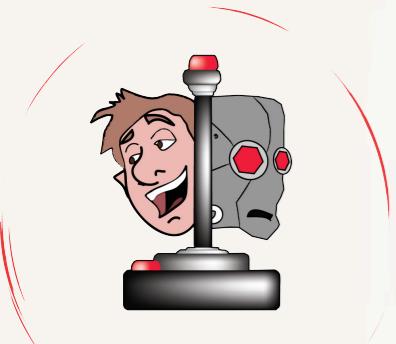
# Interactive narrative generation

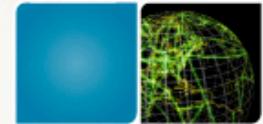




# The road ahead

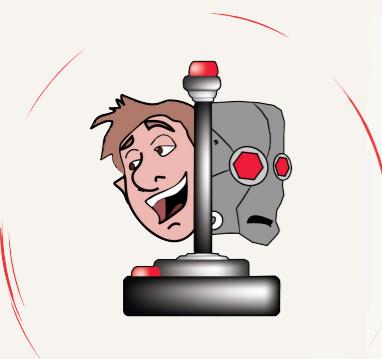
- Starting with basics: causality, character believability, typicality
- Progressively layering complexity
- Better models of story “goodness”



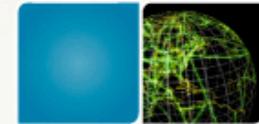


# Observations

- Just-in-time generation of a corpus
- Crowd control circumvents some NLP challenges
- Causal-like structure from data
- Concepts grounded in language
- Points in a direction where cognitive and NLP approaches may come together

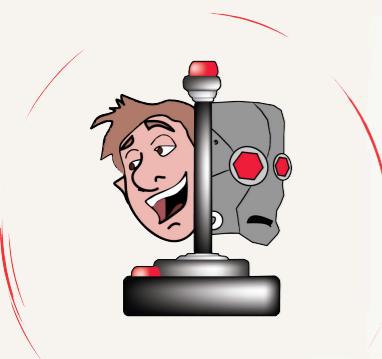


58

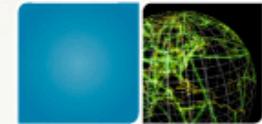


# Conclusions

- Story generation is a key capability that unlocks many practical, real-world applications
  - Create and manage user experiences in virtual worlds
  - Games, interactive narratives, training simulations, narrative learning environments, virtual characters
- Narrative intelligence is a step toward human-level intelligent systems
- What can I do to help?

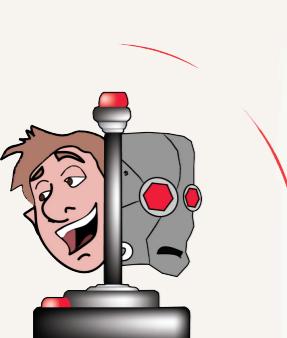


59



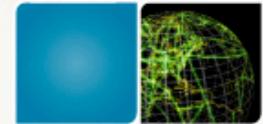
# Acknowledgements

- Boyang Li
- Brian O'Neill
- Alexander Zook
- Hong Yu
- Ken Hartsook
- Michael Drinkwater
- Steve Lee-Urban
- Sauvik Das
- Rania Hodhod
- Mohini Thakkar
- Sanjeet Hajarnis
- George Johnston
- Charles Isbell
- R. Michael Young



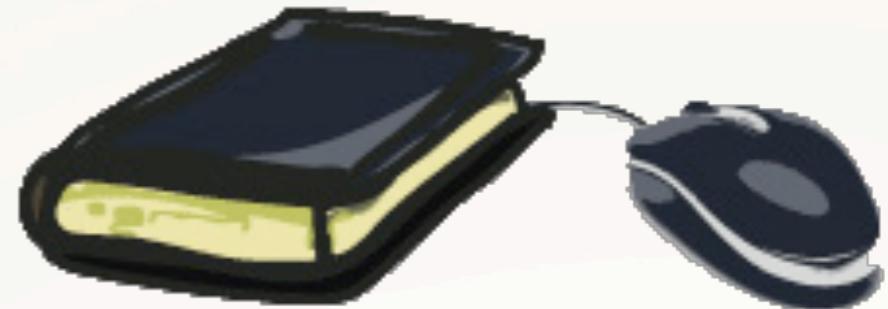
## Thanks! Questions?

riedl@cc.gatech.edu  
@mark\_riedl



# Ninth AAAI Conference on Artificial Intelligence & Interactive Digital Entertainment

<http://aiide.org>



# Sixth Workshop on Intelligent Narrative Technologies

Papers due: July 3, 2013

<http://www.int6.gatech.edu>

October 14-18, 2013, Boston MA

