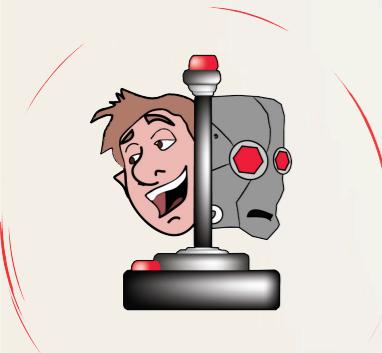
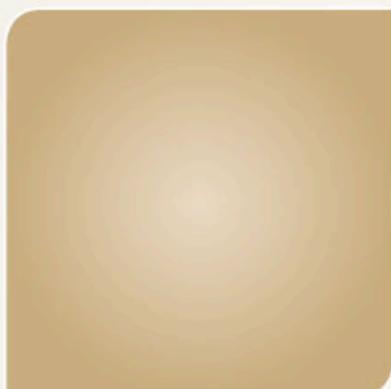


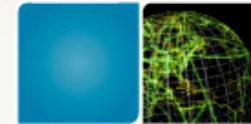
# Intelligent Narrative Generation: Creativity, Engagement, and Cognition

Mark Riedl

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

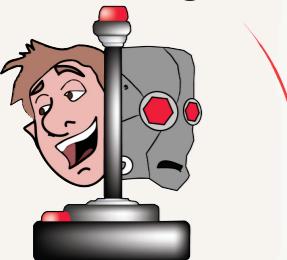


**Tell me a story**



# Storytelling

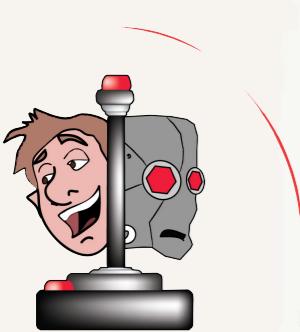
- Storytelling is pervasive part of the human experience
  - Books, movies, computer games, training scenarios, education, every-day communication, etc.
- **Narrative Intelligence:** Narrative is a fundamental means by which we organize, understand, and explain the world
- Instill computational systems with the ability to craft and tell stories in order to be better entertainers, educators, trainers, communicators, and generally more capable of relating to humans





# Why study story generation?

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

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.

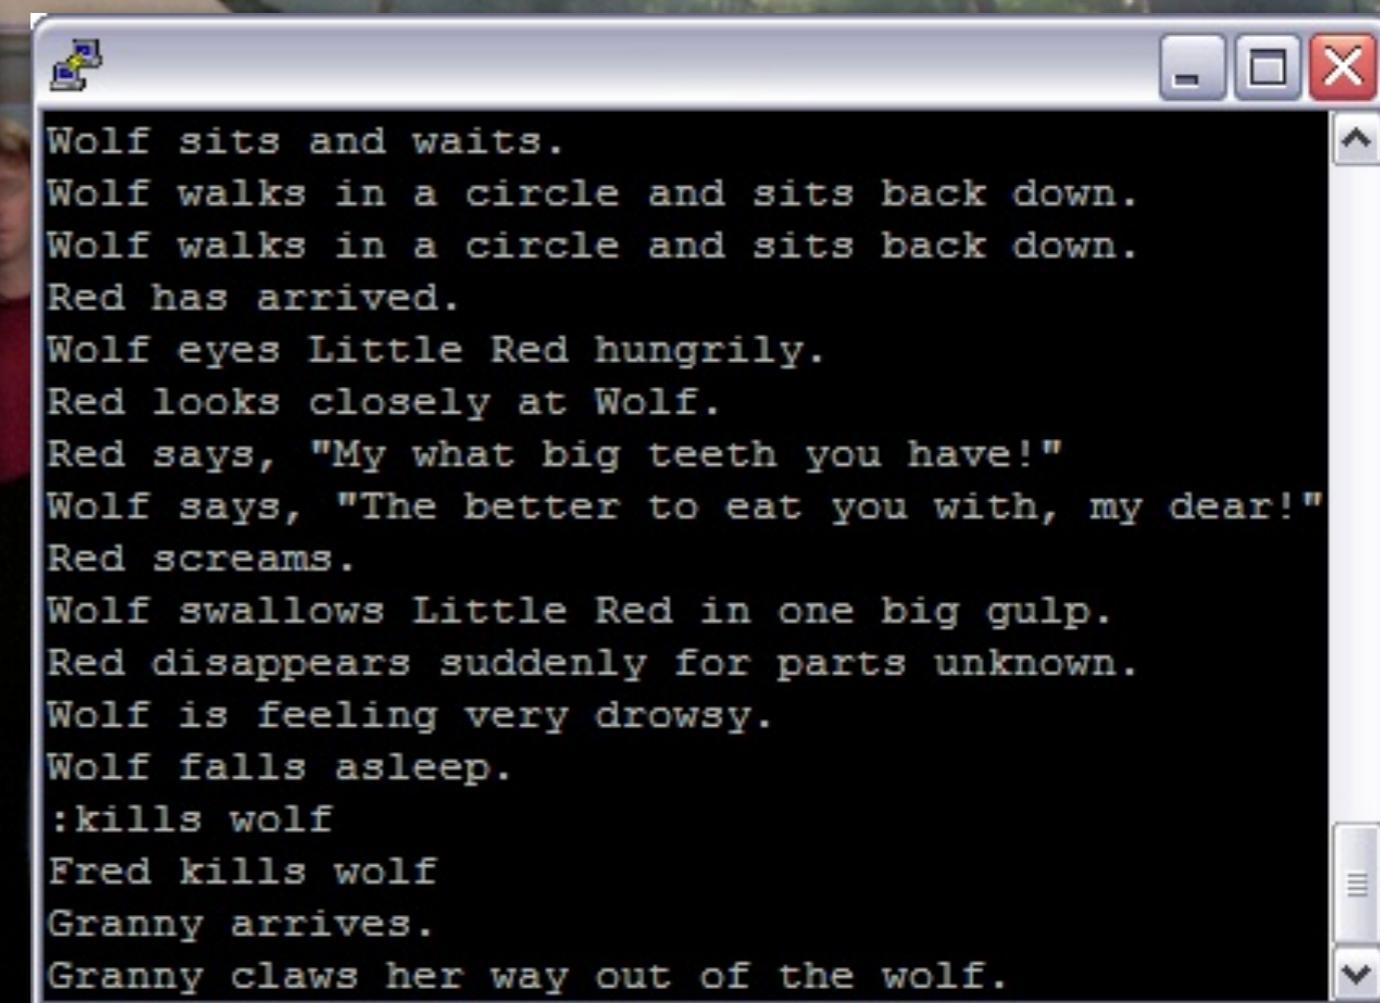
King Jafar is not married. Jasmine is very beautiful. King Jafar sees Jasmine and instantly falls in love with her. King Jafar wants to marry Jasmine. There is a brave knight named Aladdin. Aladdin is loyal to the death to King Jafar. King Jafar orders Aladdin to get the magic lamp for him. Aladdin wants 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're taking customers  
away from me, Hassan!

Leave me alone, Saleh!



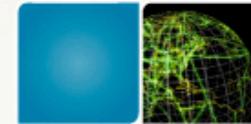
Wolf sits and waits.  
Wolf walks in a circle and sits back down.  
Wolf walks in a circle and sits back down.  
Red has arrived.  
Wolf eyes Little Red hungrily.  
Red looks closely at Wolf.  
Red says, "My what big teeth you have!"  
Wolf says, "The better to eat you with, my dear!"  
Red screams.  
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

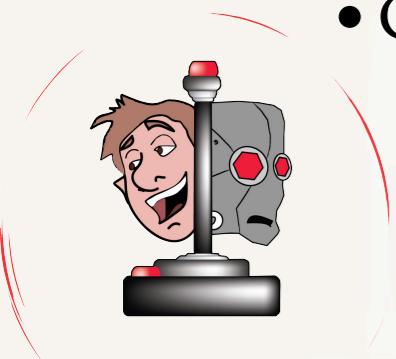
Computer  
Games

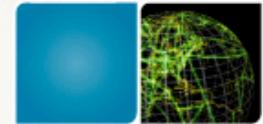
Narrative  
Intelligence



# Automated story generation

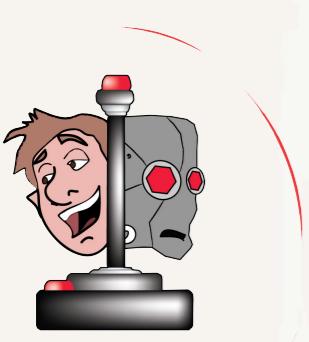
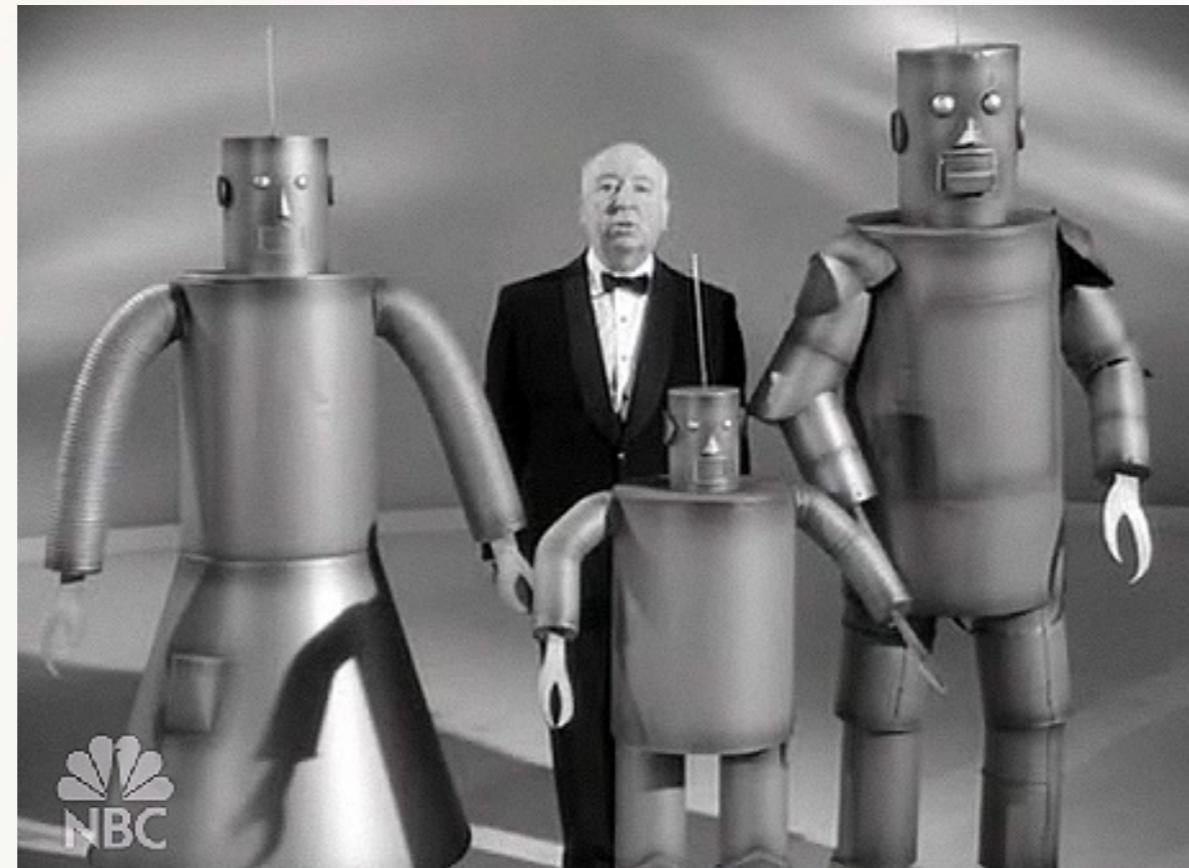
- Automatic creation of meaningful fictional sequences is hard
  - Complexity, subtlety, nuance, mimesis
  - Focus on plot
- Two nearly-universal properties of story
  - Causal progression
    - Perception that the main events of the story make up causal chains that terminate in the outcome
  - Character believability
    - Perception that characters act in a manner that does not distract from one's suspension of disbelief
    - Characters are perceived to be intentional agents

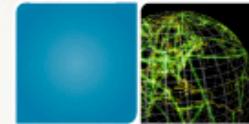




# Computer as author

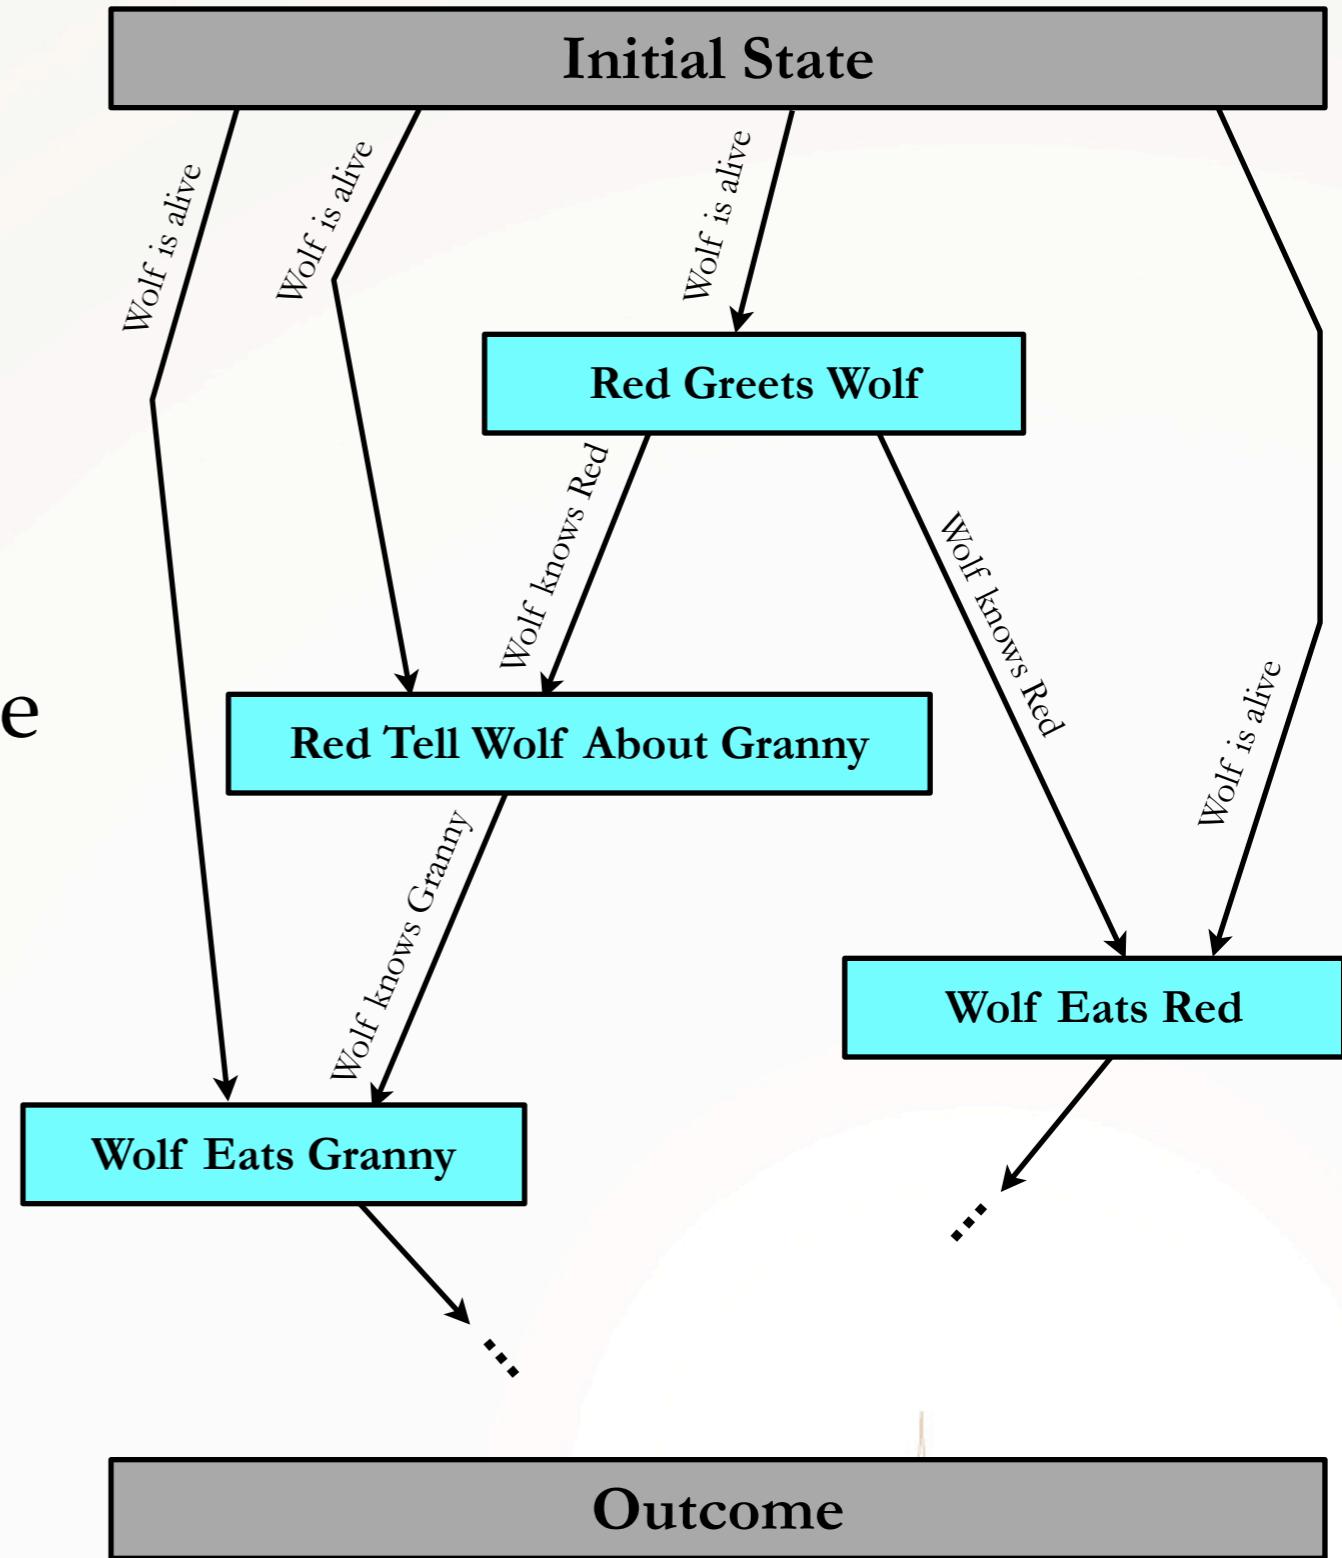
- Creative writing is a problem-solving activity
- Author goals vs. character goals
- Plan out the events that should occur in the narrative





# Narratives as plans

- Partial-order plan is a good representation of plot
  - Action, temporality, causality
- 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?

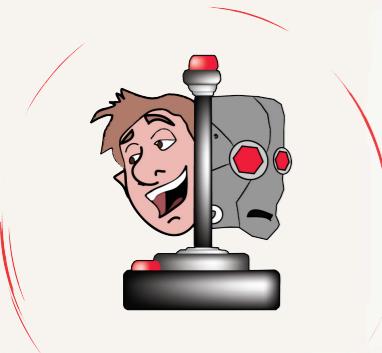
- Conventional Planning

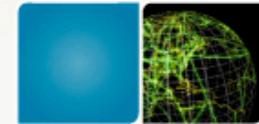
- Single agent
- Goal state is agent's desired world state
- Goal state is intended by the agent

- Narrative Planning

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

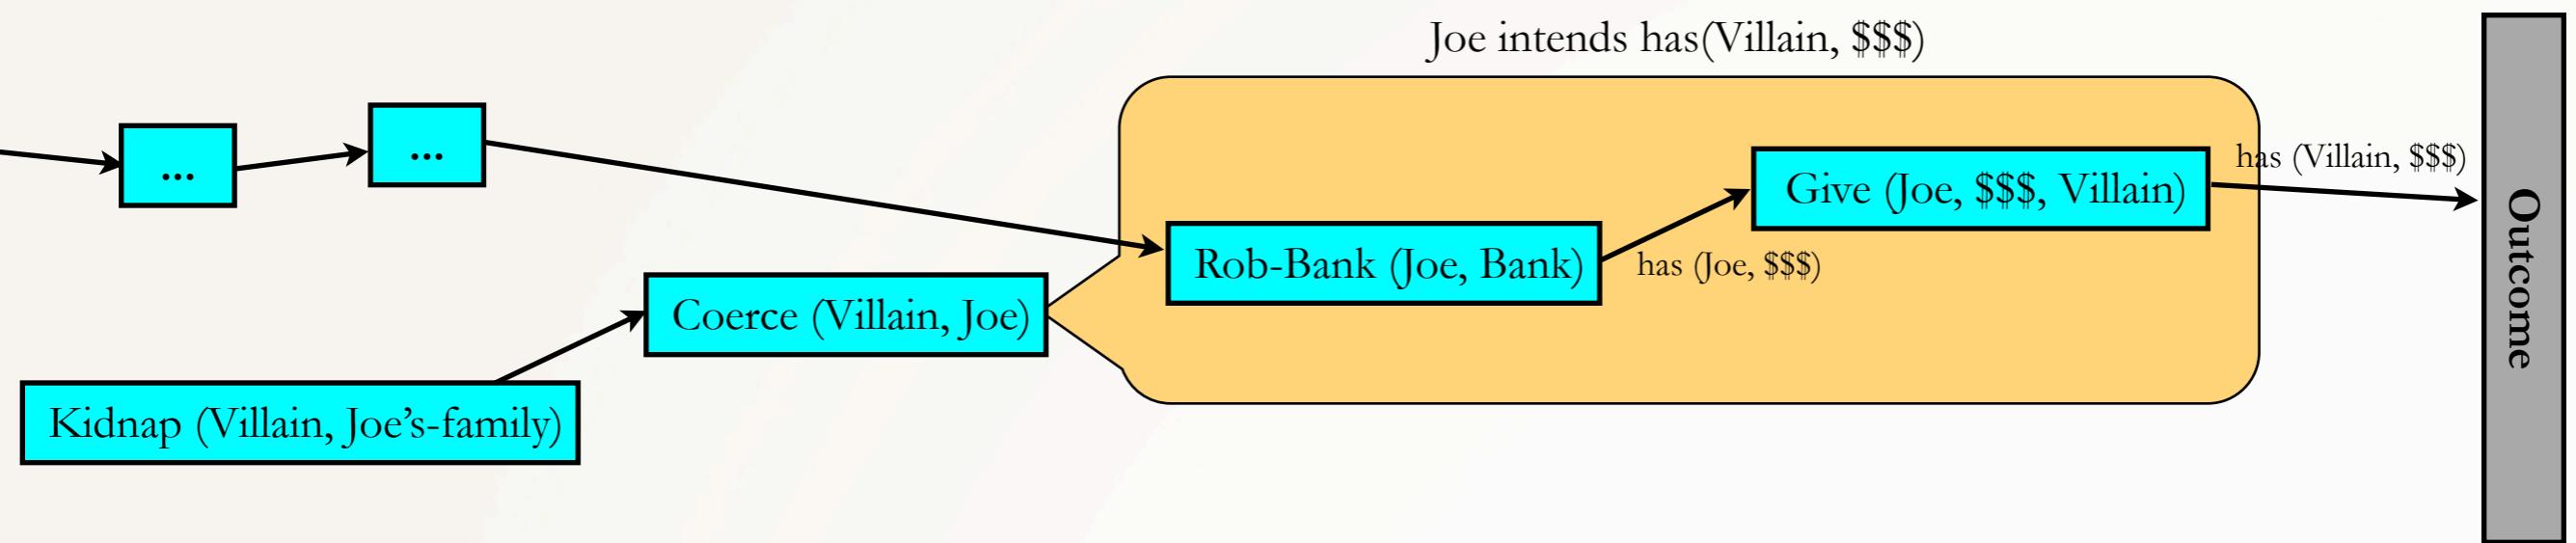
- Augment planning algorithm to reason about author goals **and** character goals



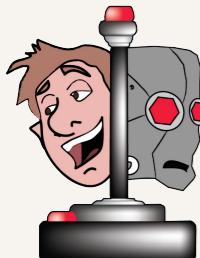


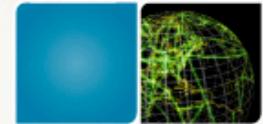
# Fabulist

- 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





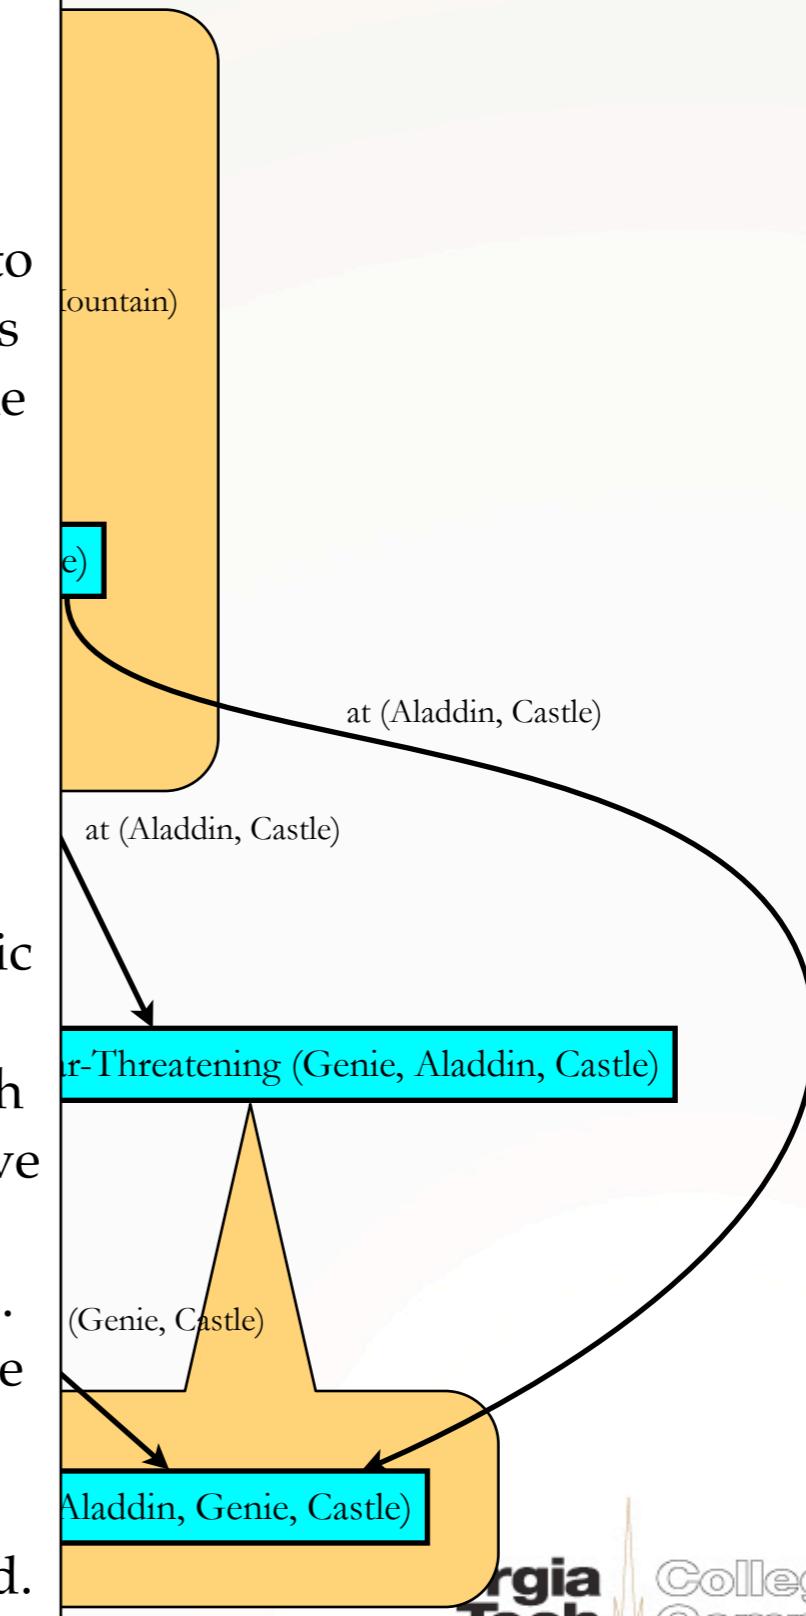
There is a woman named Jasmine. There is a king named Jafar. This is a story about how King Jafar becomes married to Jasmine. There is a magic genie. This is also a story about how the genie dies.

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.

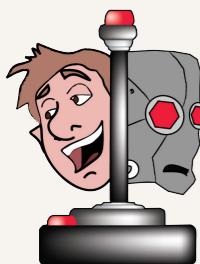
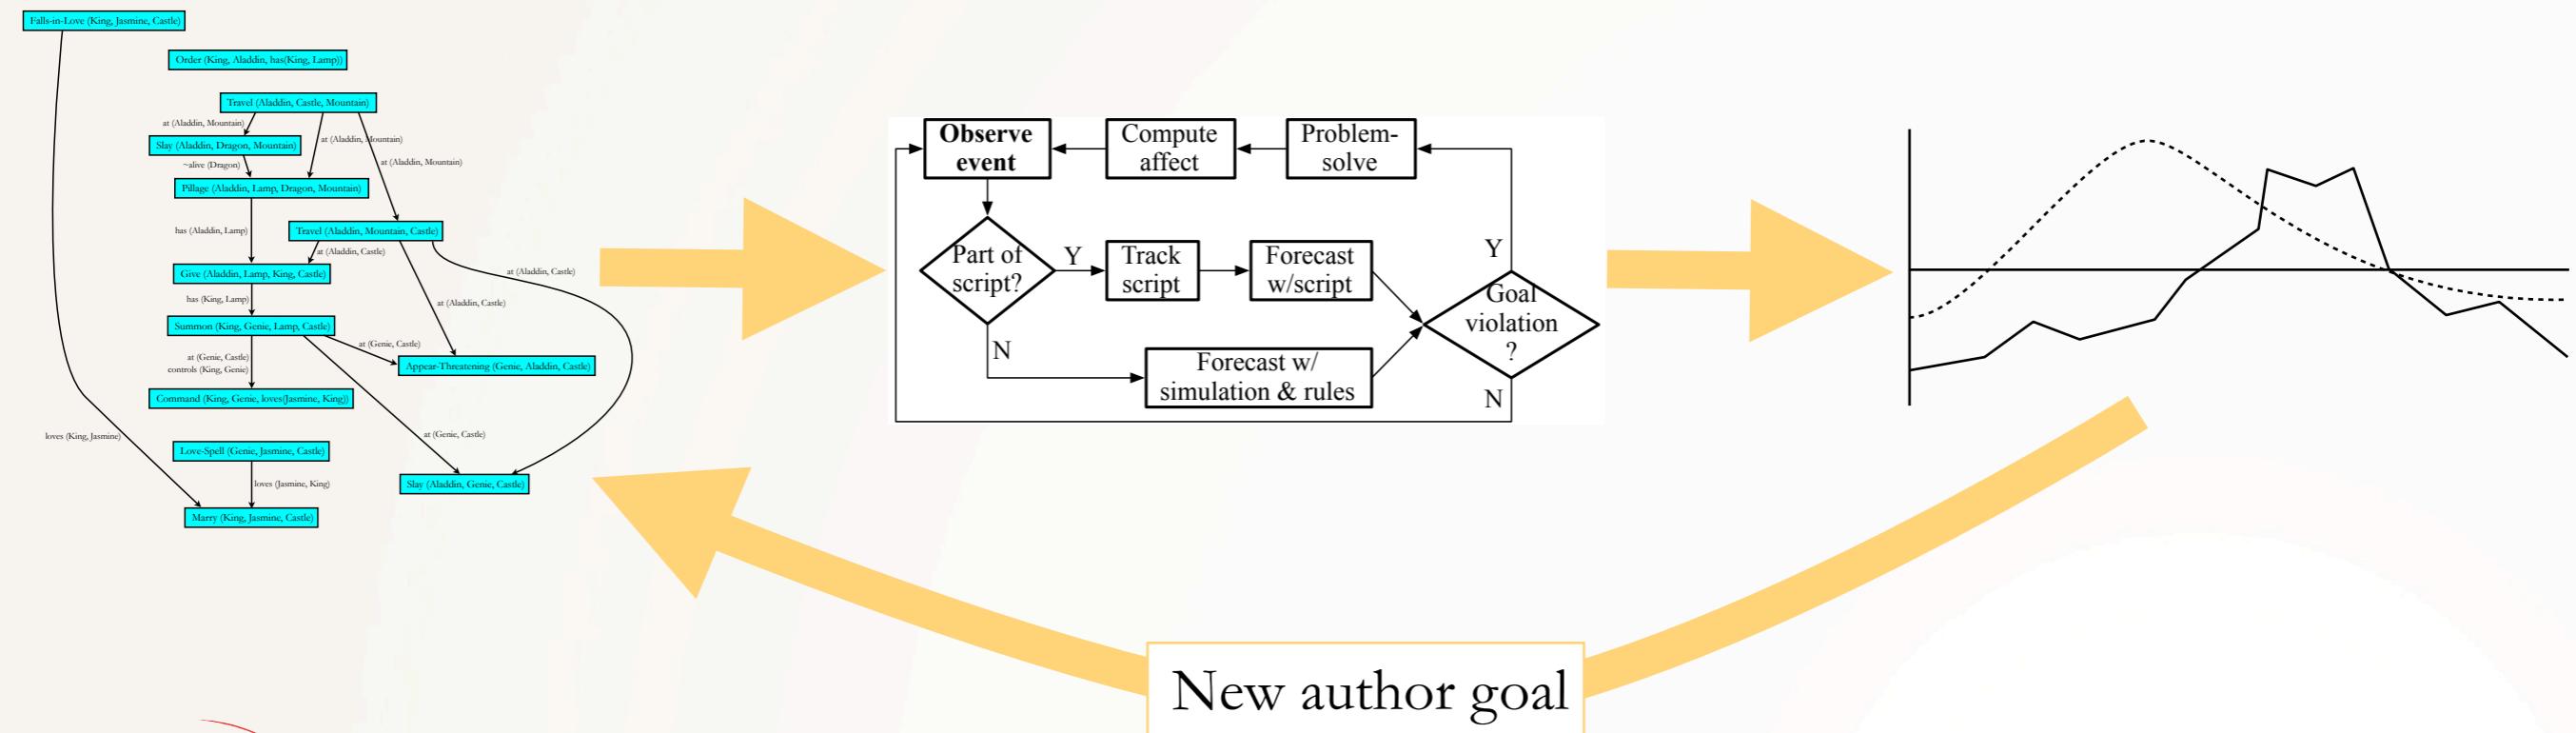
Marry (King, Jasmine, Castle)





# The “goodness” question

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



Story  
Generation

# Computer Games

Narrative  
Intelligence

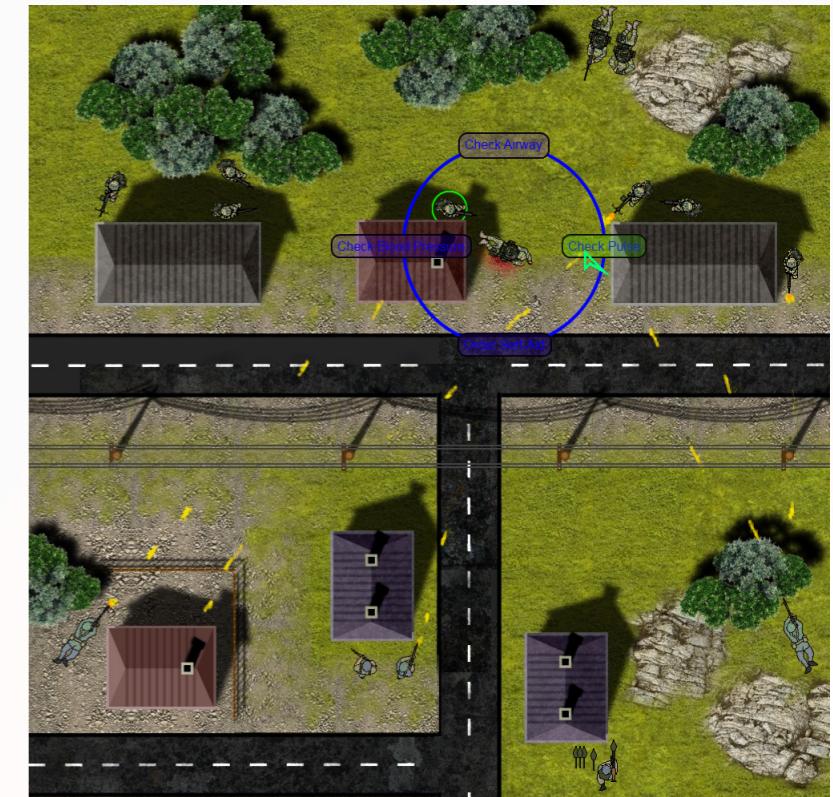
# Replayability

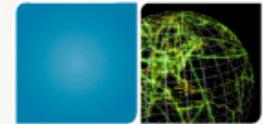
- The replay value of a game diminishes as game affordances and story content are exhausted
- Can we adapt an existing game plot so that the player never experiences the same plot progression twice?
- Can we personalize the game to an individual player?



# Game Personalization

- Role playing games (and training scenarios) interleave challenges and narrative connectives
- Narrative authored by humans and thus one-size-fits-all
- Can we learn a model of player (learner) preferences, skills, desires, needs?
- Solution: narrative adaptation + data-driven player modeling



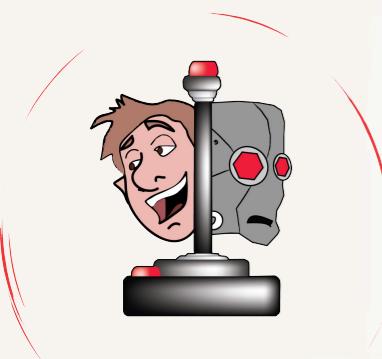


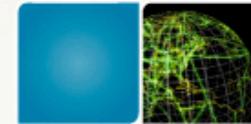
## Challenge tailoring

Choose and order the challenges the player should experience

## Challenge Contextualization

Provide motivating story context between challenges

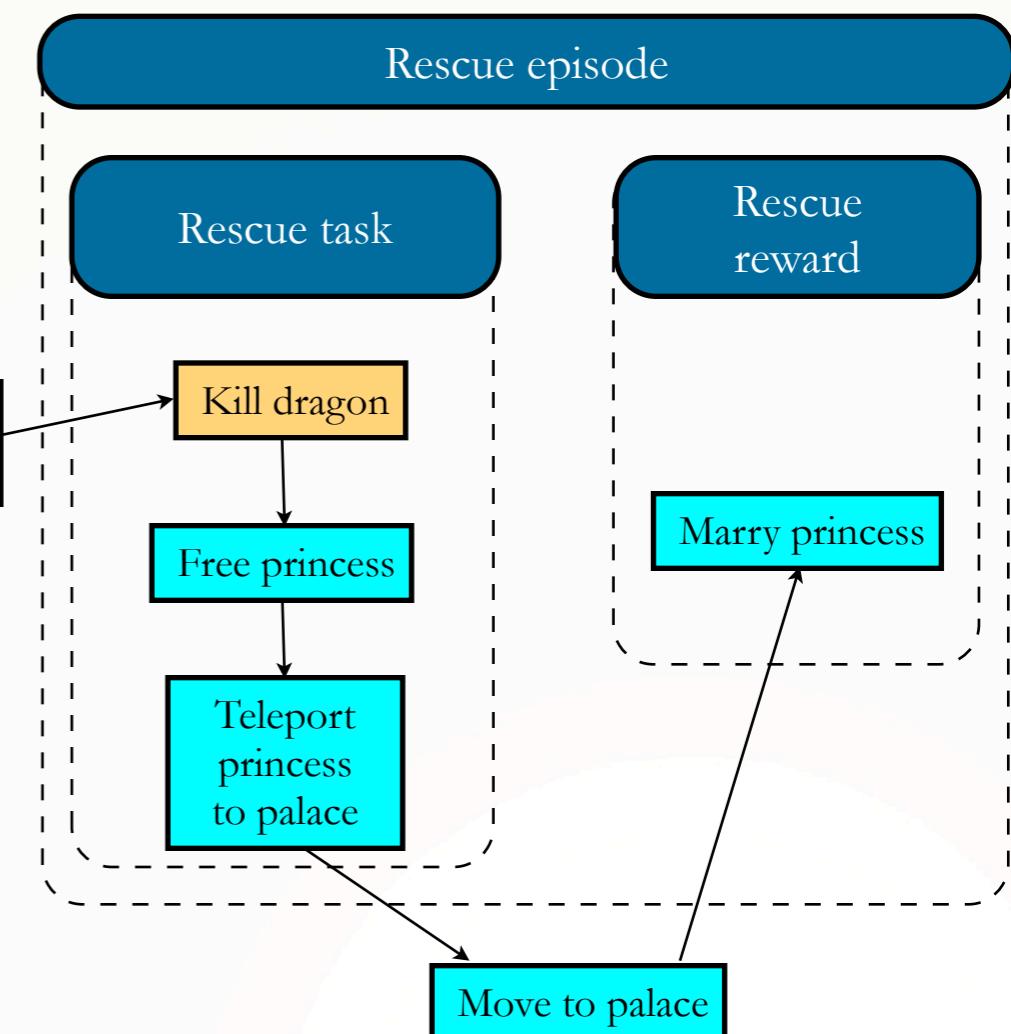


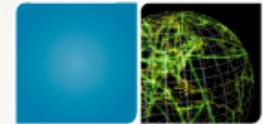


# Game plot representation

- Partial-order plan of events
- Skill-based events

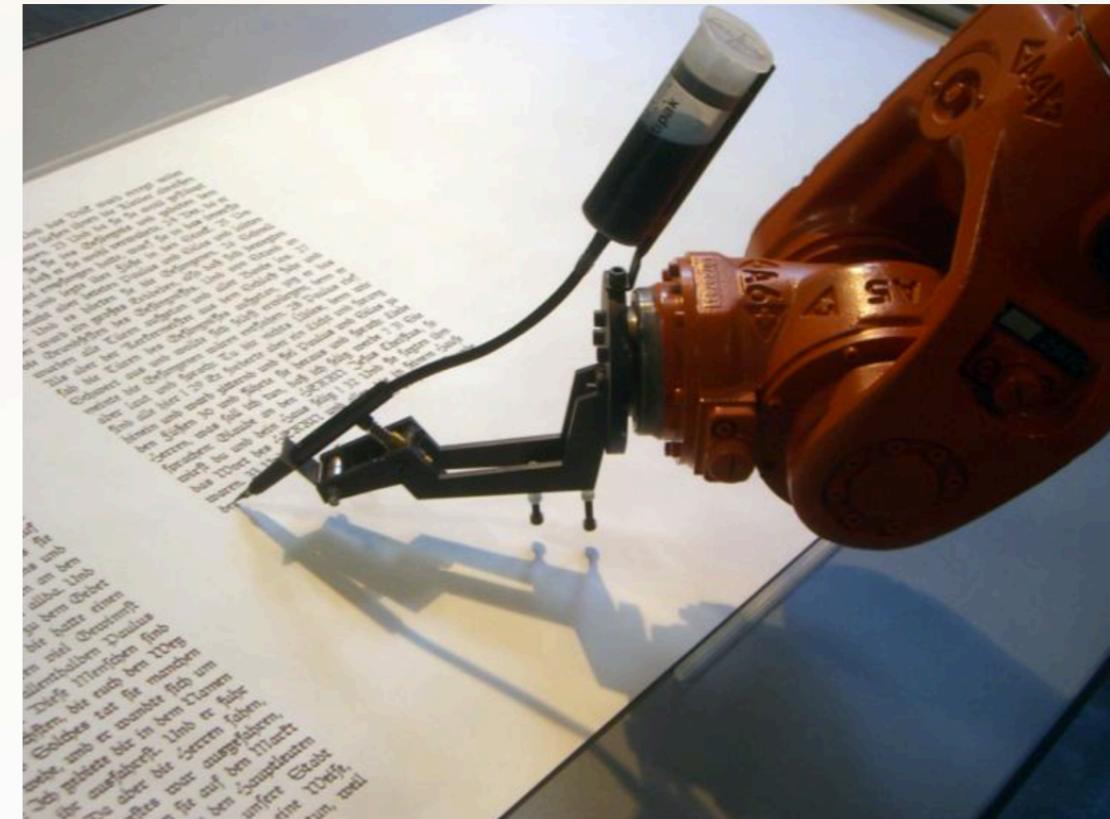
- Hierarchical relationships





# Challenge Contextualization

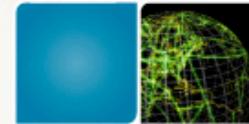
- Start with hand-authored narrative
- Given: which skill events and episodes should be present (or not)
- Plot adaptation
  - Search for a set of modifications to the plot structure such that the required episodes are present and plan is free of flaws
  - Add, remove, re-order, and causally-rewire events



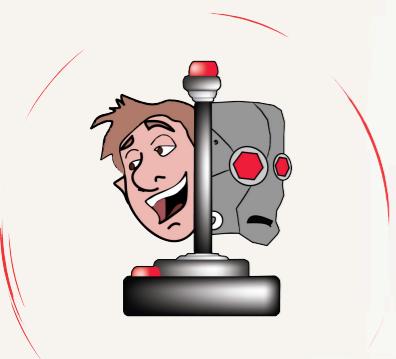
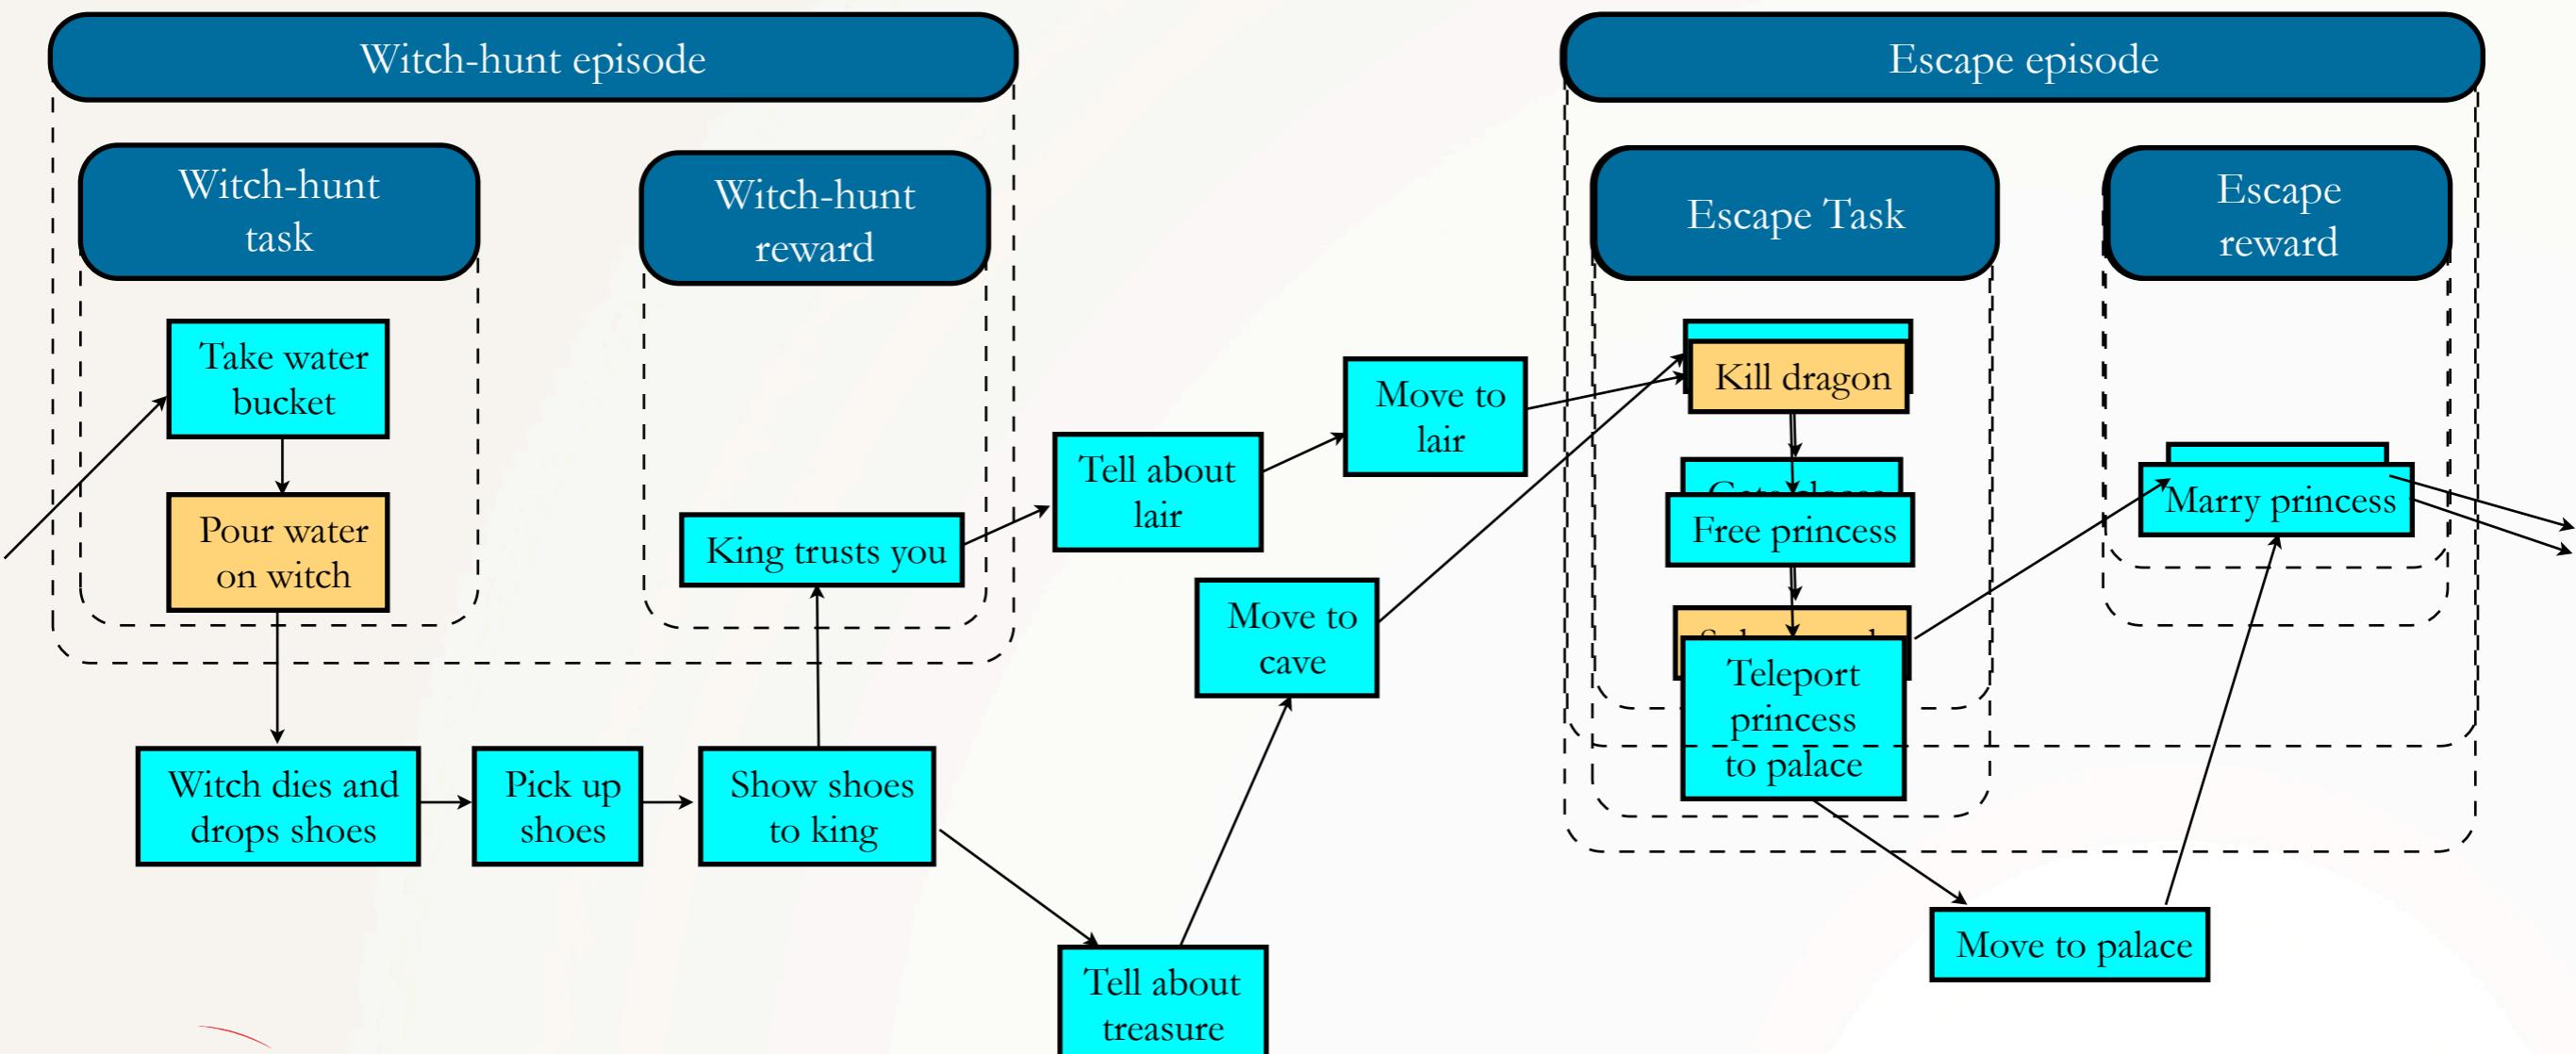


# Flaws and repairs

Flaw	Description	Repair Strategies
Open condition	Event $e$ has a precondition $p$ that is not satisfied by a causal link	<ol style="list-style-type: none"> <li>Instantiate new event <math>e_{\text{new}}</math> that has an effect that unifies with <math>p</math>. Extend a causal link from <math>e_{\text{new}}</math> to <math>e</math>.</li> <li>Select an existing event <math>e_{\text{old}}</math> that has an effect that unifies with <math>p</math>. Extend a causal link from <math>e_{\text{old}}</math> to <math>e</math>.</li> <li><b>Delete <math>e</math>.</b></li> </ol>
Causal threat	Event $e_k$ has an effect that negates a causal link between events $e_i$ and $e_j$	<ol style="list-style-type: none"> <li>Promotion: temporally constrain <math>e_k</math> before <math>e_i</math>.</li> <li>Demotion: temporally constrain <math>e_k</math> after <math>e_j</math>.</li> <li><b>Delete <math>e_k</math>.</b></li> </ol>
Un-decomposed abstract event	Event $e$ is abstract but has no children	<ol style="list-style-type: none"> <li>Select and apply a decomposition rule, instantiating new events or reusing existing events as children.</li> </ol>
Dead end	Event $e$ is a dead end	<ol style="list-style-type: none"> <li>Select an existing event <math>e_{\text{old}}</math> that has a precondition that is unsatisfied and that unifies with an effect of <math>e</math>. Extend a causal link from <math>e</math> to <math>e_{\text{old}}</math>.</li> <li>Select an existing event <math>e_{\text{old}}</math> that has a precondition that is satisfied by causal link <math>c</math> and unifies with an effect of <math>e</math>. Transfer the starting point of <math>c</math> to <math>e</math>.</li> <li>Delete the flaw.</li> <li>Ignore the flaw.</li> </ol>
Superfluous event	Event $e$ is superfluous	<ol style="list-style-type: none"> <li>Link effects of earlier steps to effects of <math>e</math>.</li> <li>Ignore the flaw</li> </ol>

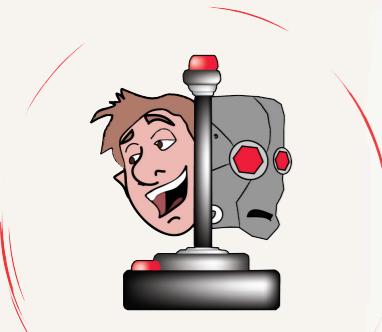
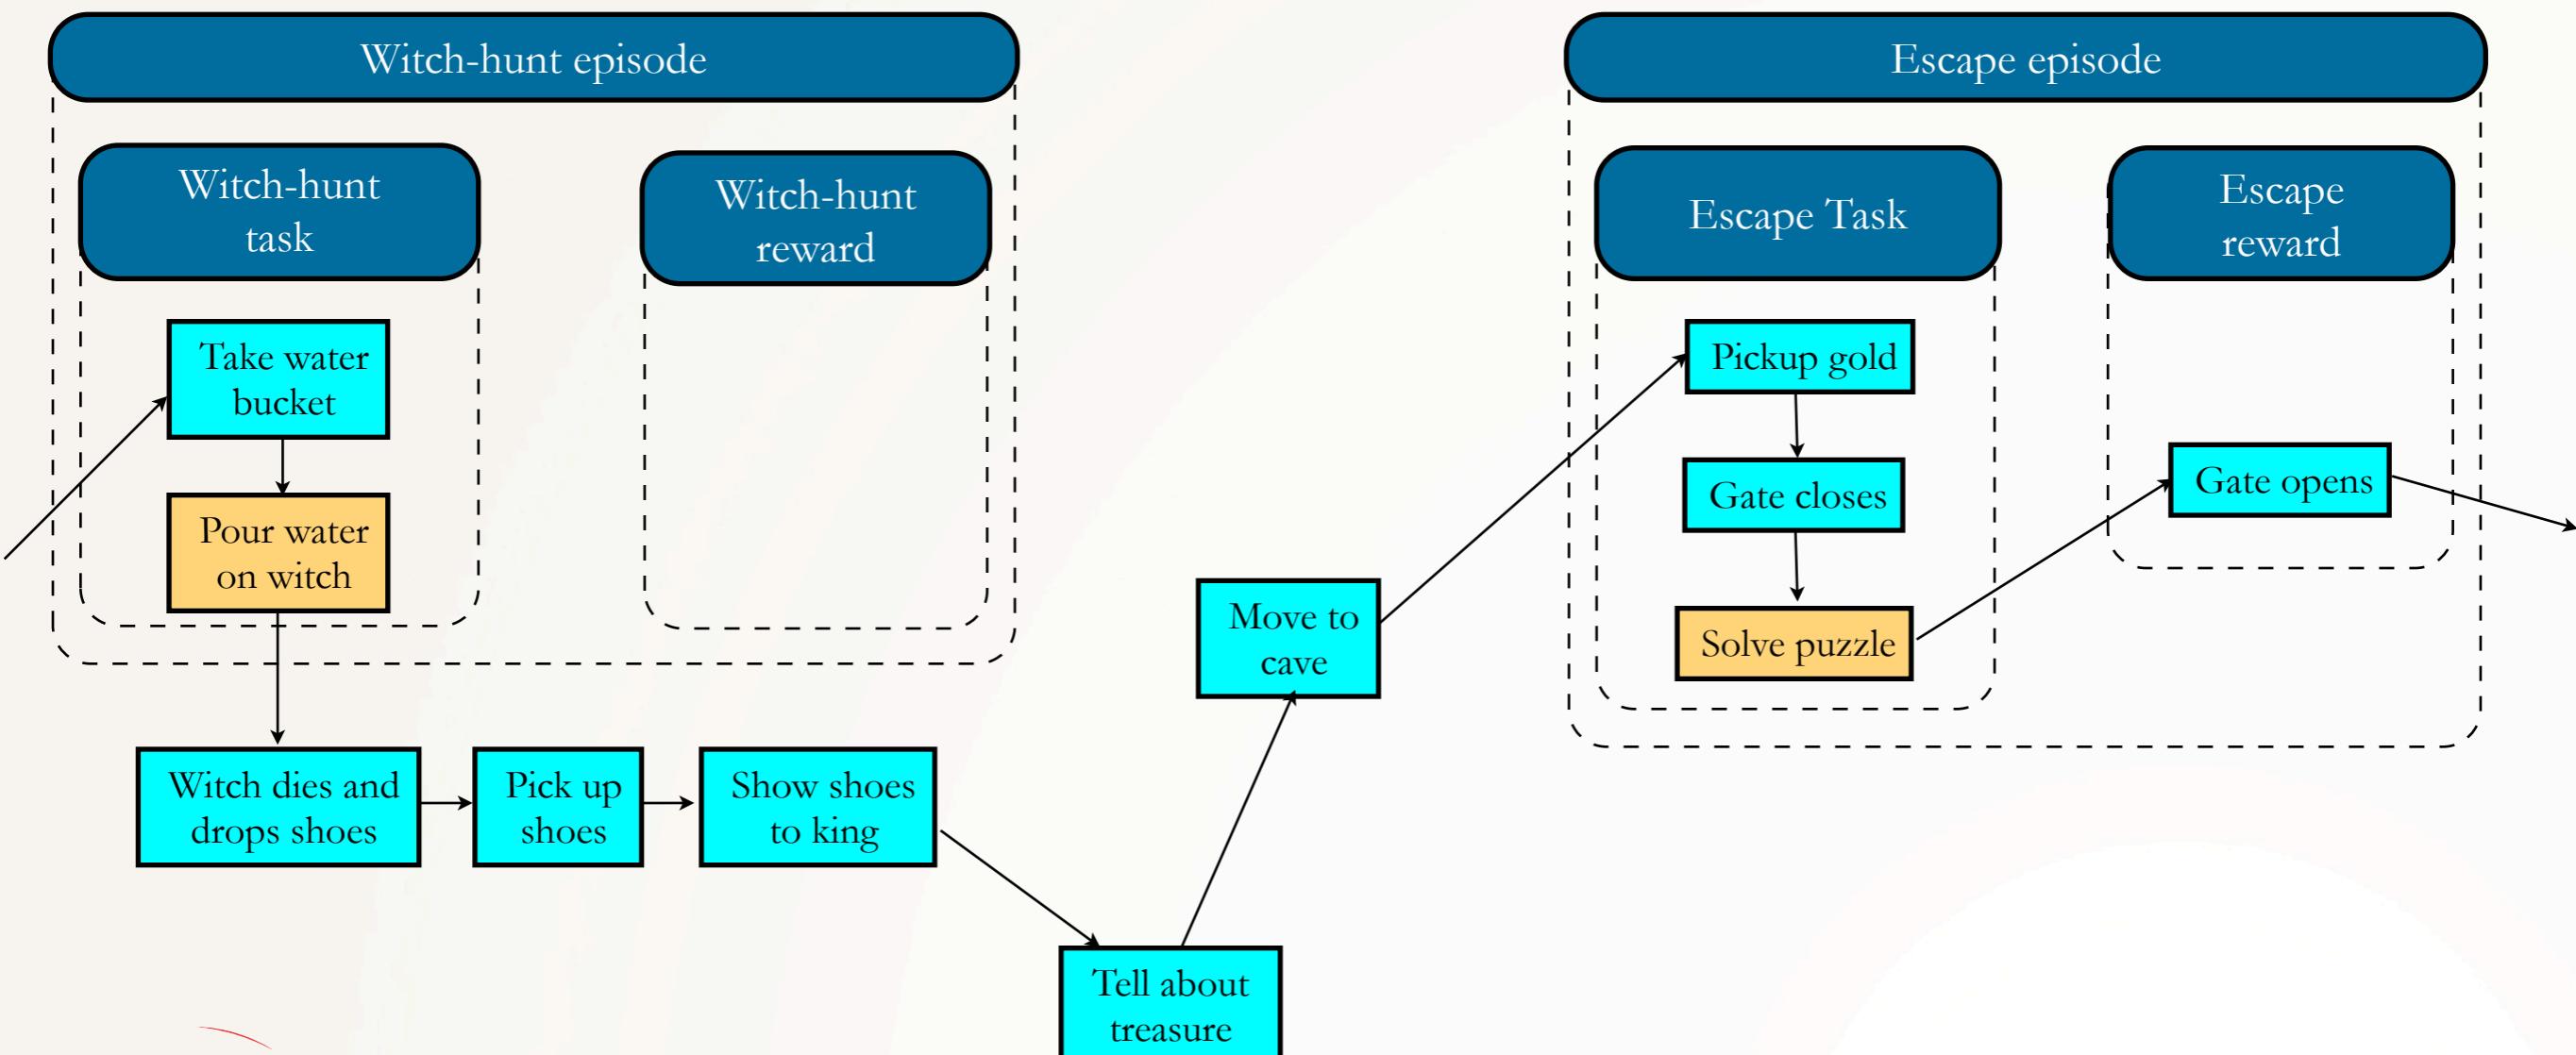


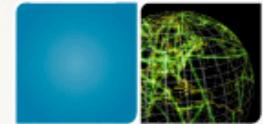
# Challenge contextualization example





# Challenge contextualization example





# Challenge tailoring

- Which skill-based events and in what order?
- Simple example: dynamic difficulty adjustment
- Data-driven player model predicts skill performance

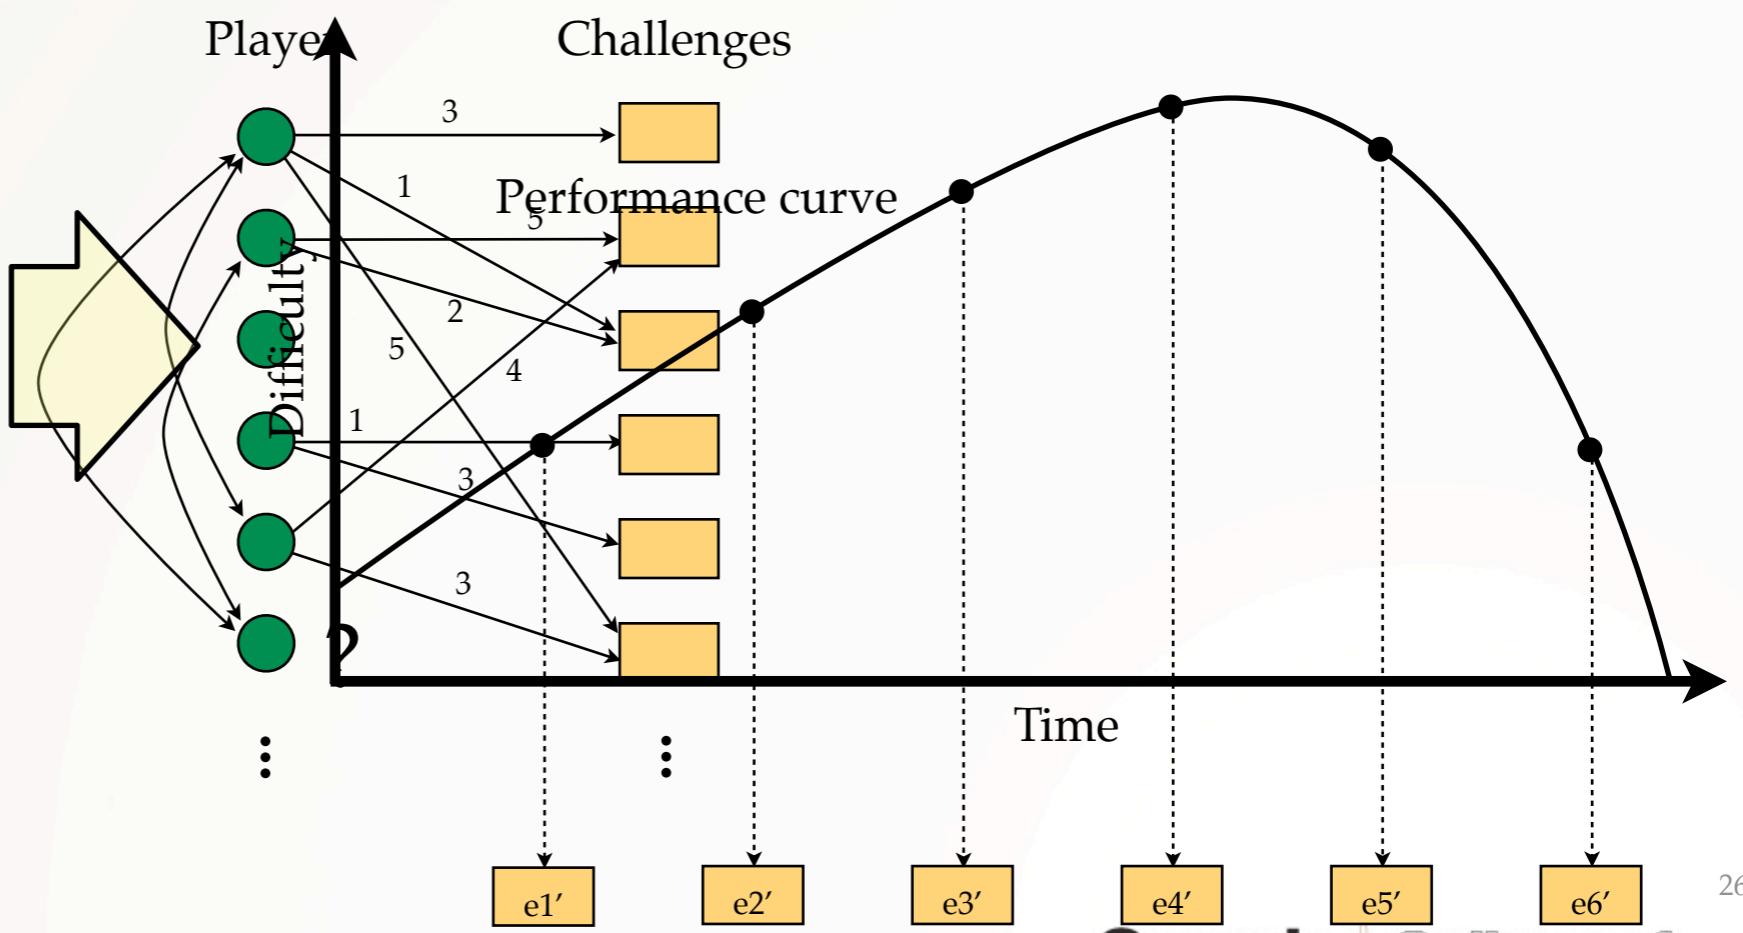


Zook & Riedl. Proc. AIIDE 2012 Conference.



# Player modeling

- Observe many players' proficiency at various challenges
- Exploit statistical correlations between players' performances
- Temporal matrix factorization



26



# Game world generation



Story  
Generation

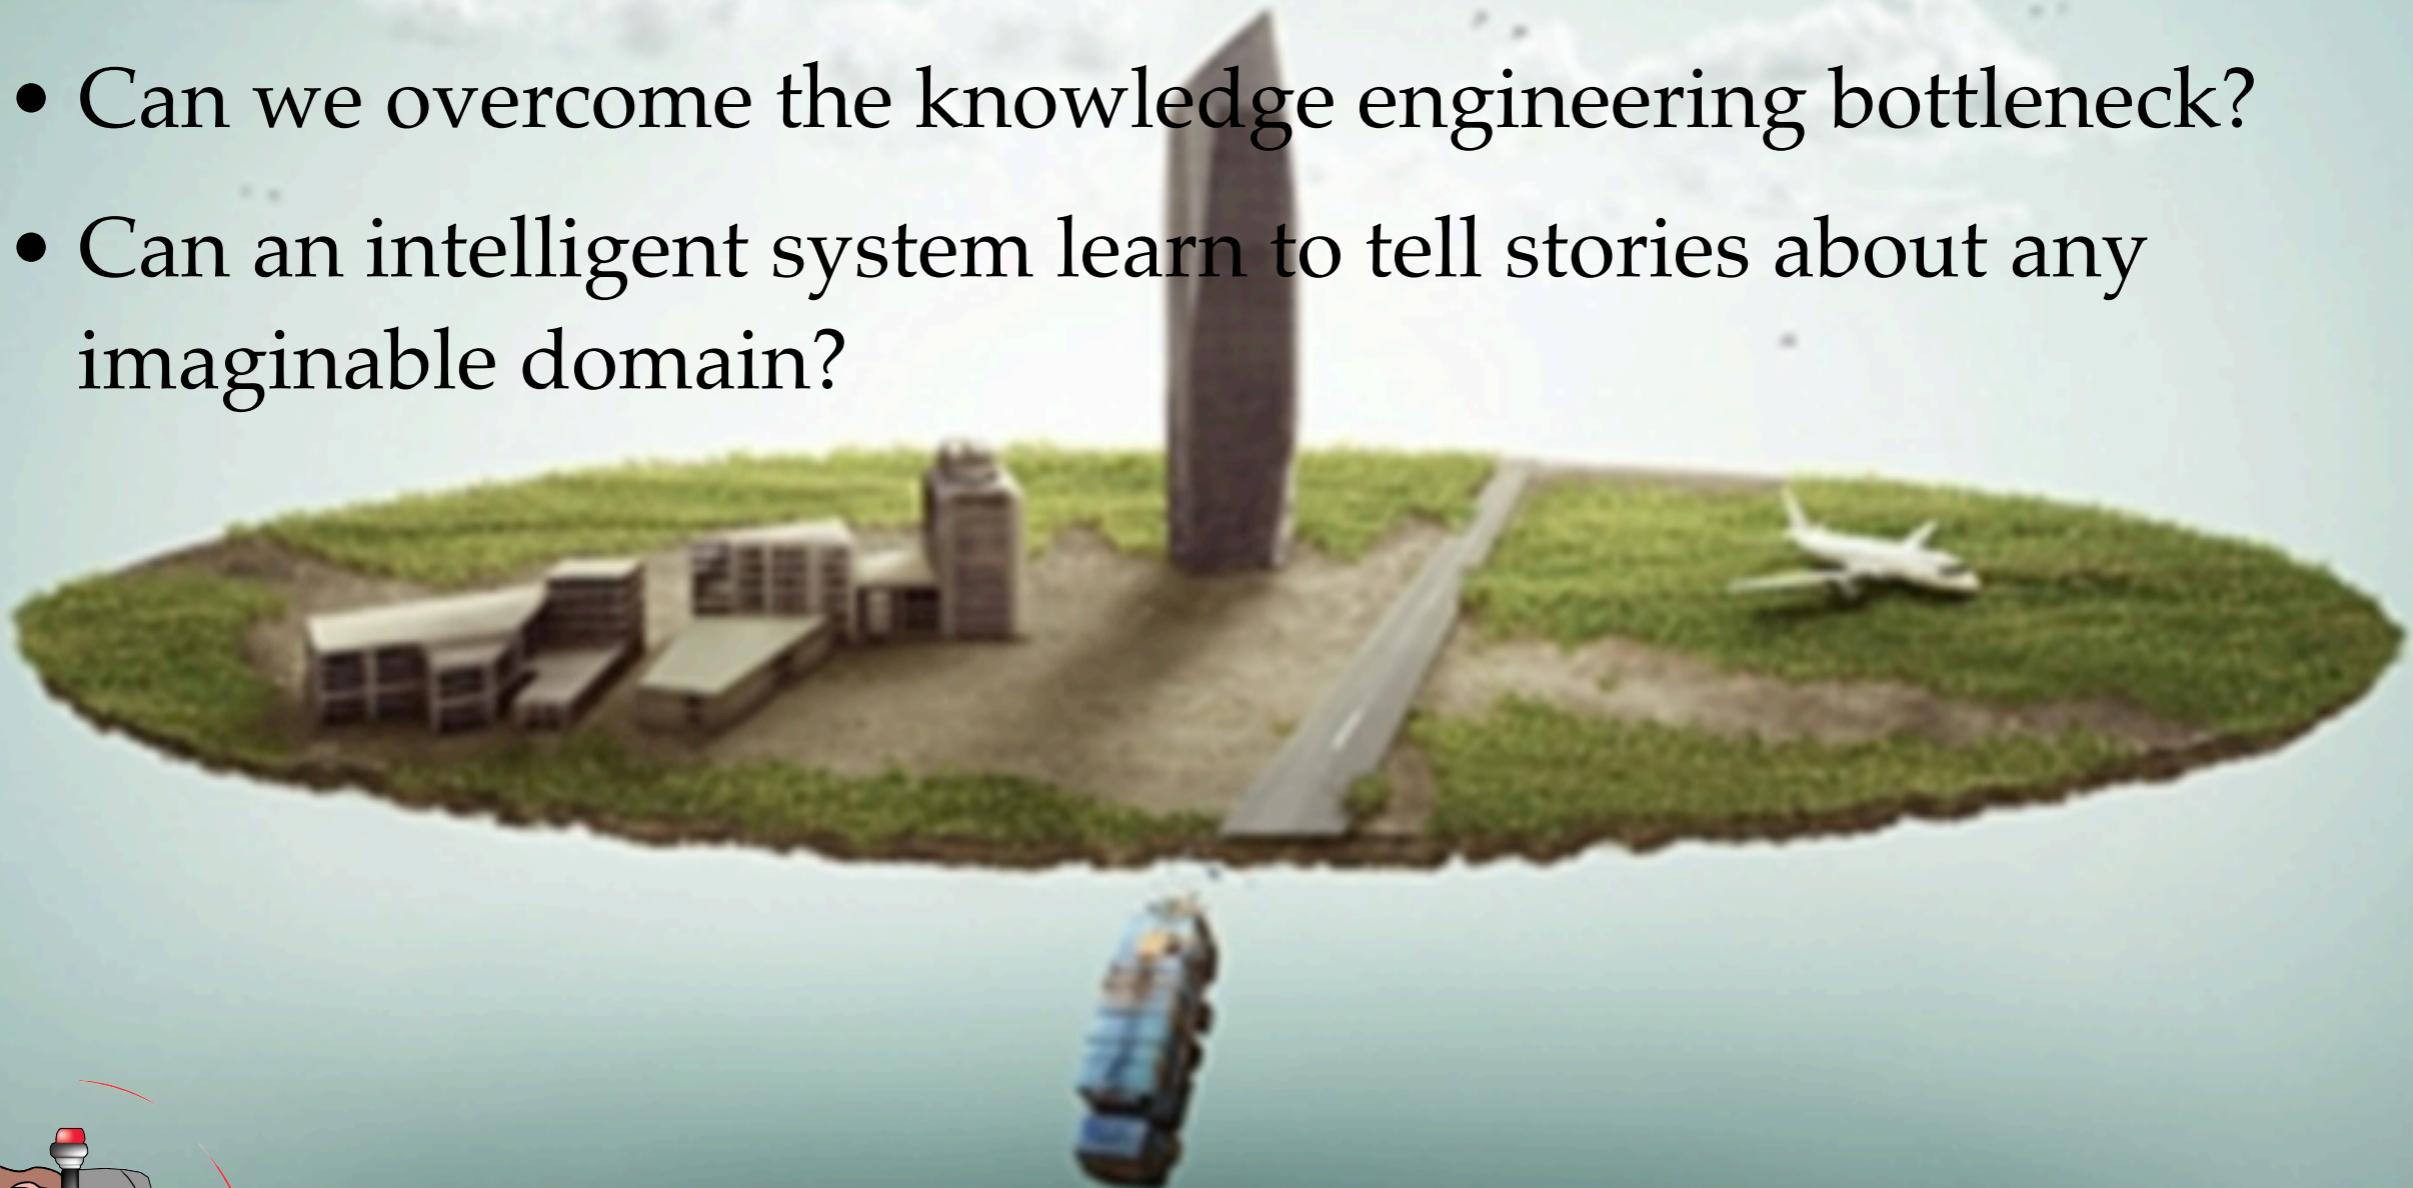
Computer  
Games

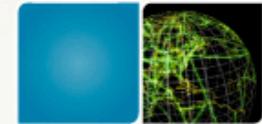
**Narrative  
Intelligence**



# Open story generation

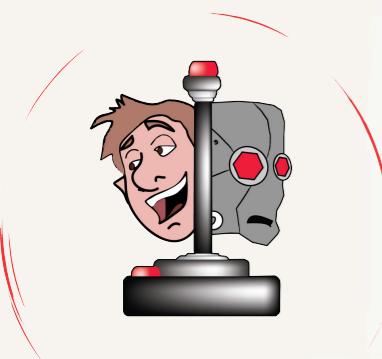
- Narrative intelligence is knowledge-intensive, resulting in micro-worlds
- Can we overcome the knowledge engineering bottleneck?
- Can an intelligent system learn to tell stories about any imaginable domain?





# Learning sociocultural knowledge

- Humans rely on a lifetime of experiences from which to explain about stories, tell stories, or act in the real-world
- Computational systems can now live in a rich information ecosphere, including the Web, other agents, and humans
- Possible approaches:
  - Read natural language corpora & websites
  - Mine commonsense knowledge bases
  - Learn from humans

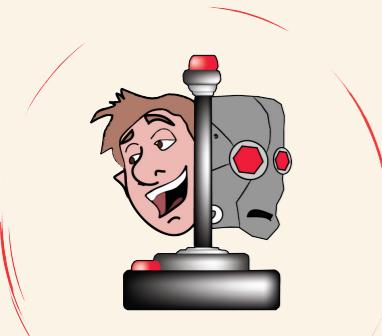
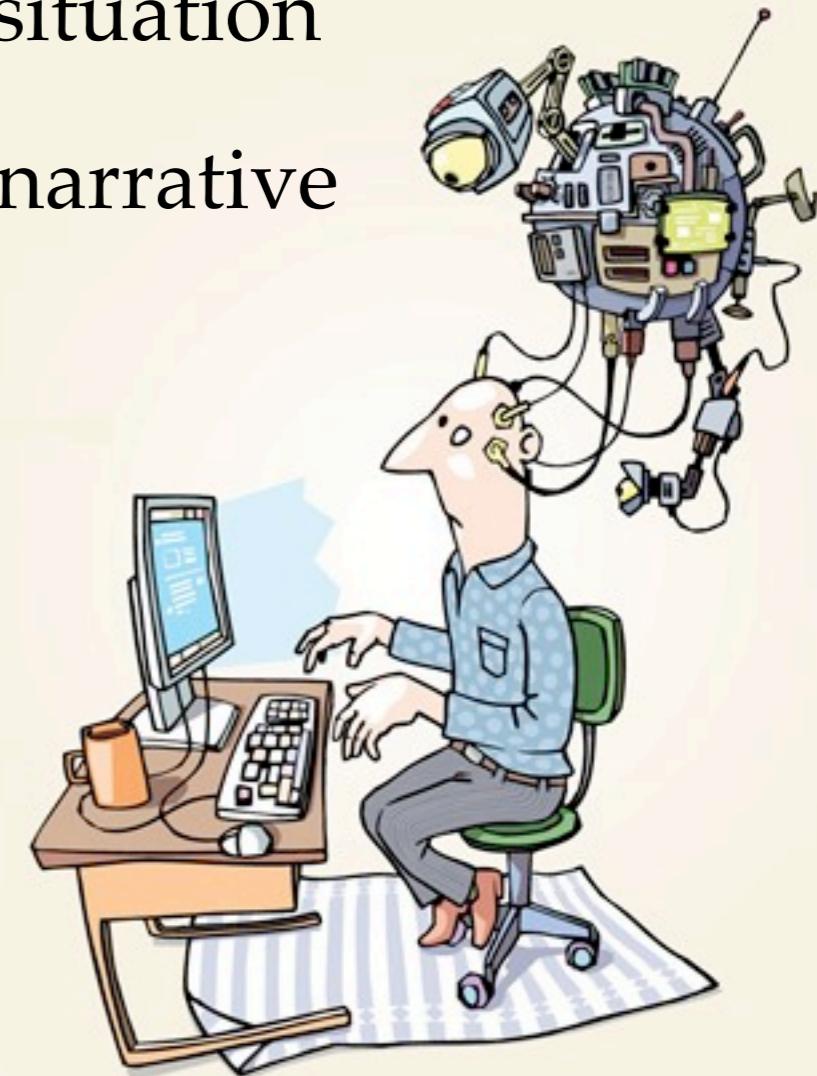


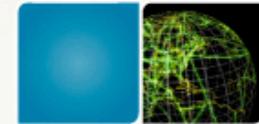
31

# Crowdsourcing narrative intelligence

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

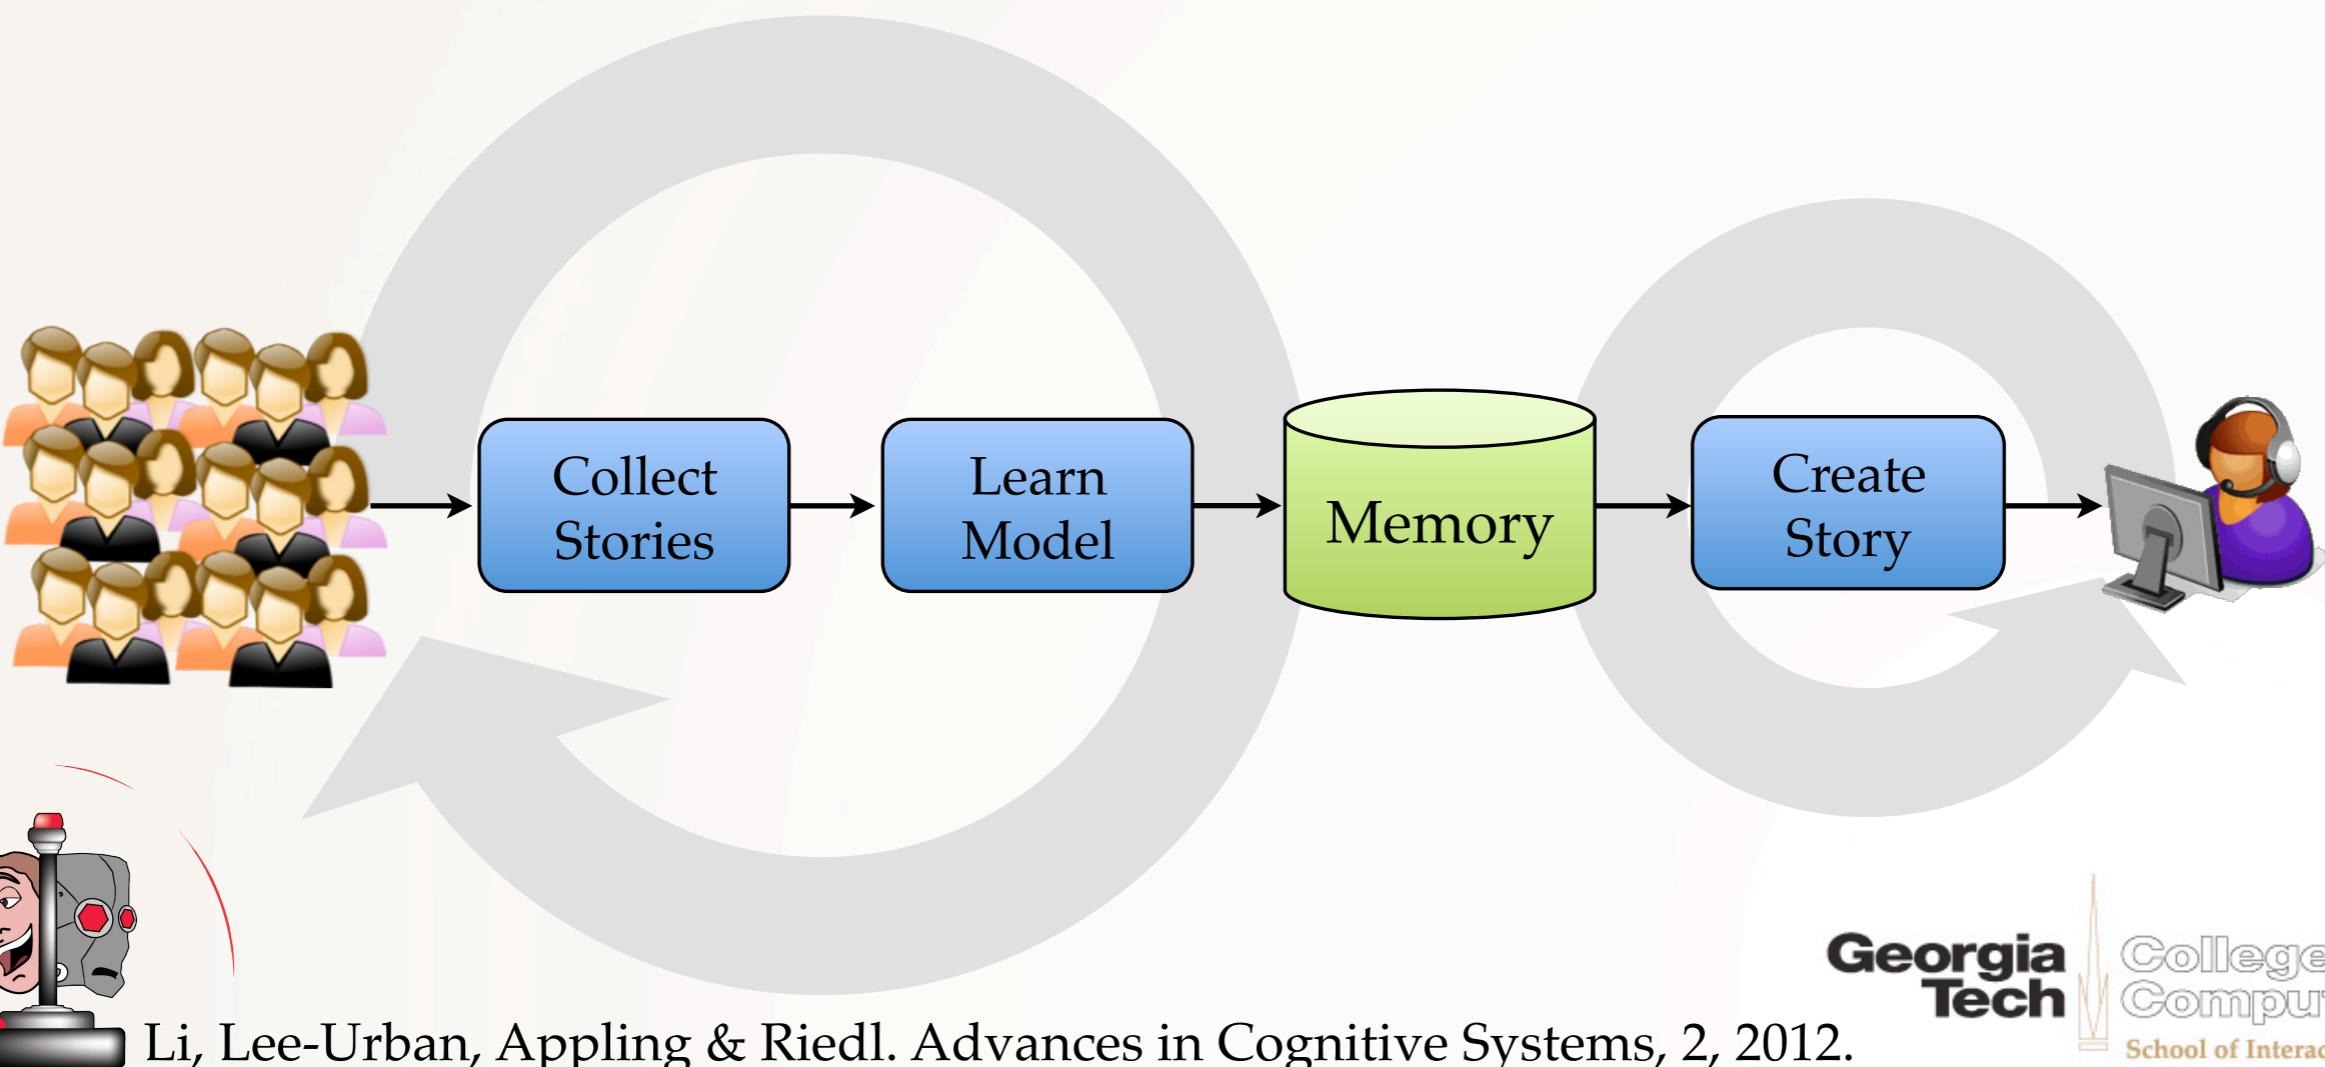
Automatically generate stories and interactive experiences **without** a priori domain knowledge authoring





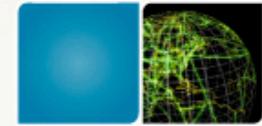
# Scheherazade

- Just-in-time model learning
  - Semantic parsing, pattern mining, global optimization
  - Barthes: human narratives implicitly encode causality
  - Grice: maxims of quality and relevance (and quantity)



33



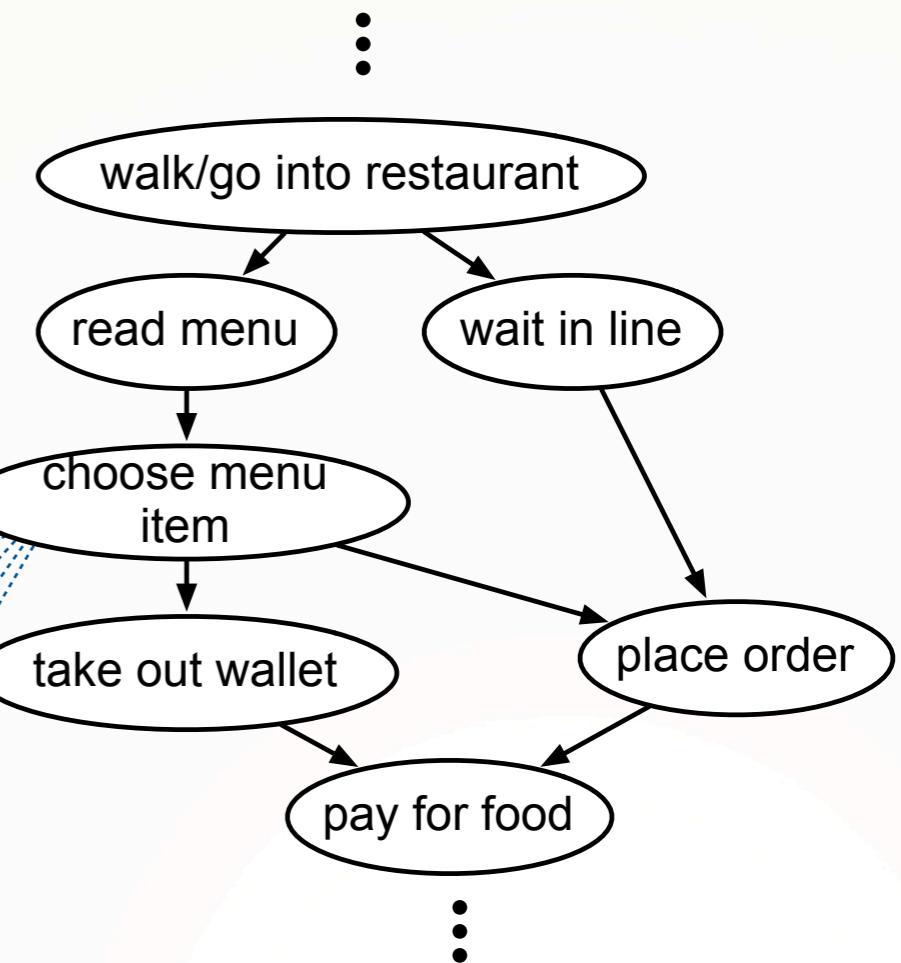


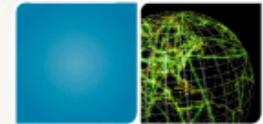
# Sociocultural knowledge representation

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

- Set of temporally ordered events
- Grounded in language

Sally chose item from the menu  
Sally thought about what Sally wanted  
Sally decided what to order  
Sally chose something from the menu  
John made a selection  
...



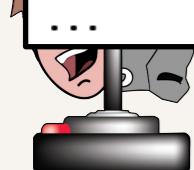


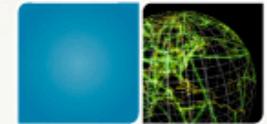
# Learning to tell stories

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

- 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.	...
...	

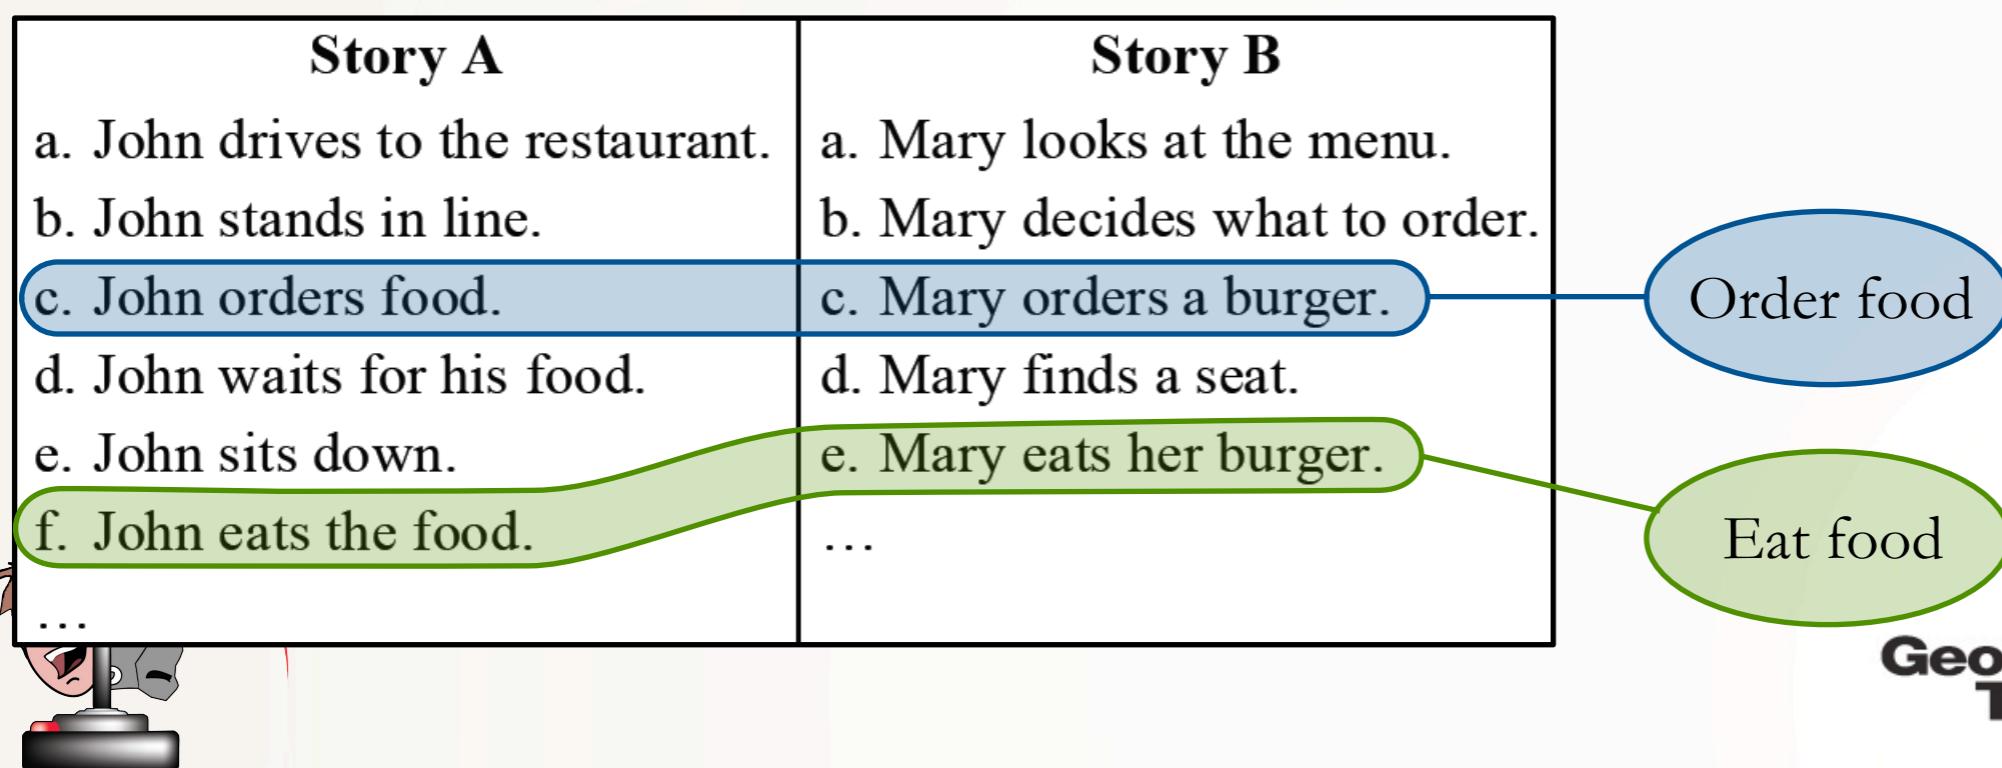




# Learning to tell stories

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

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

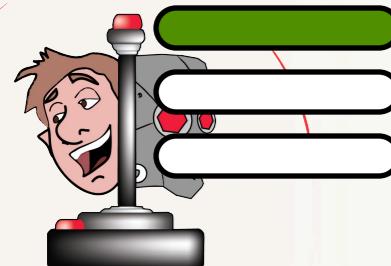
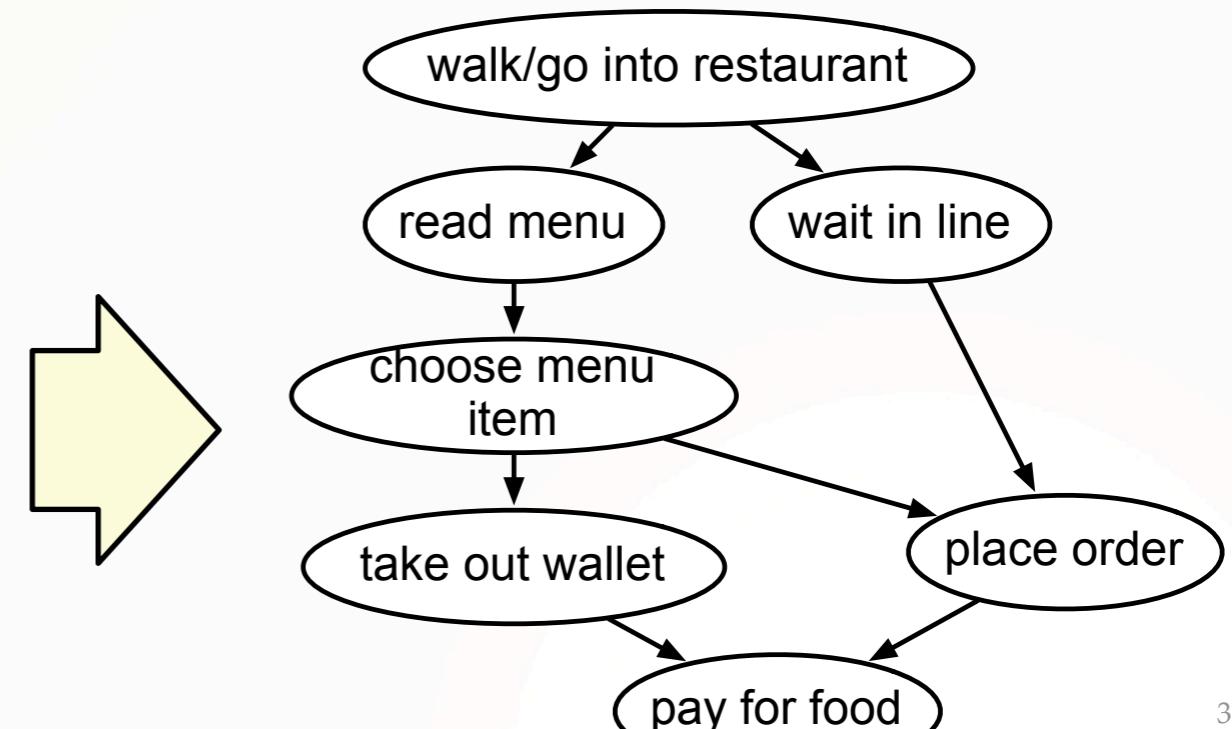
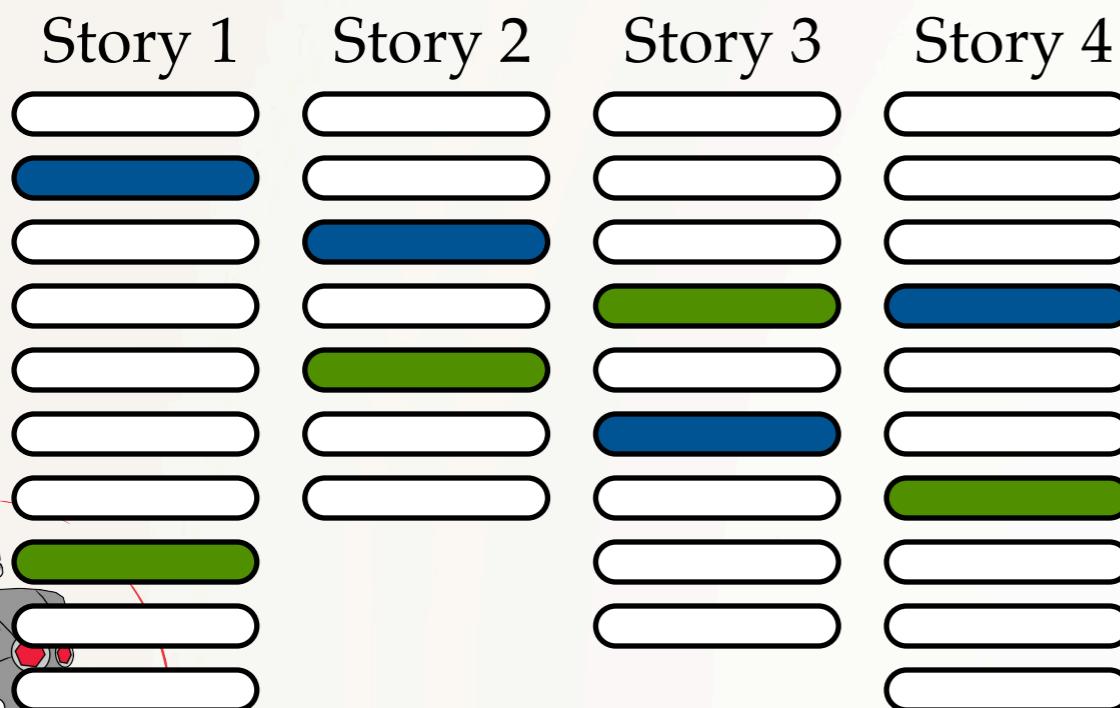


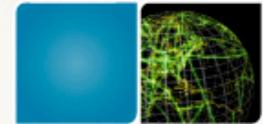


# Learning to tell stories

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

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

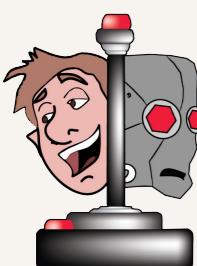
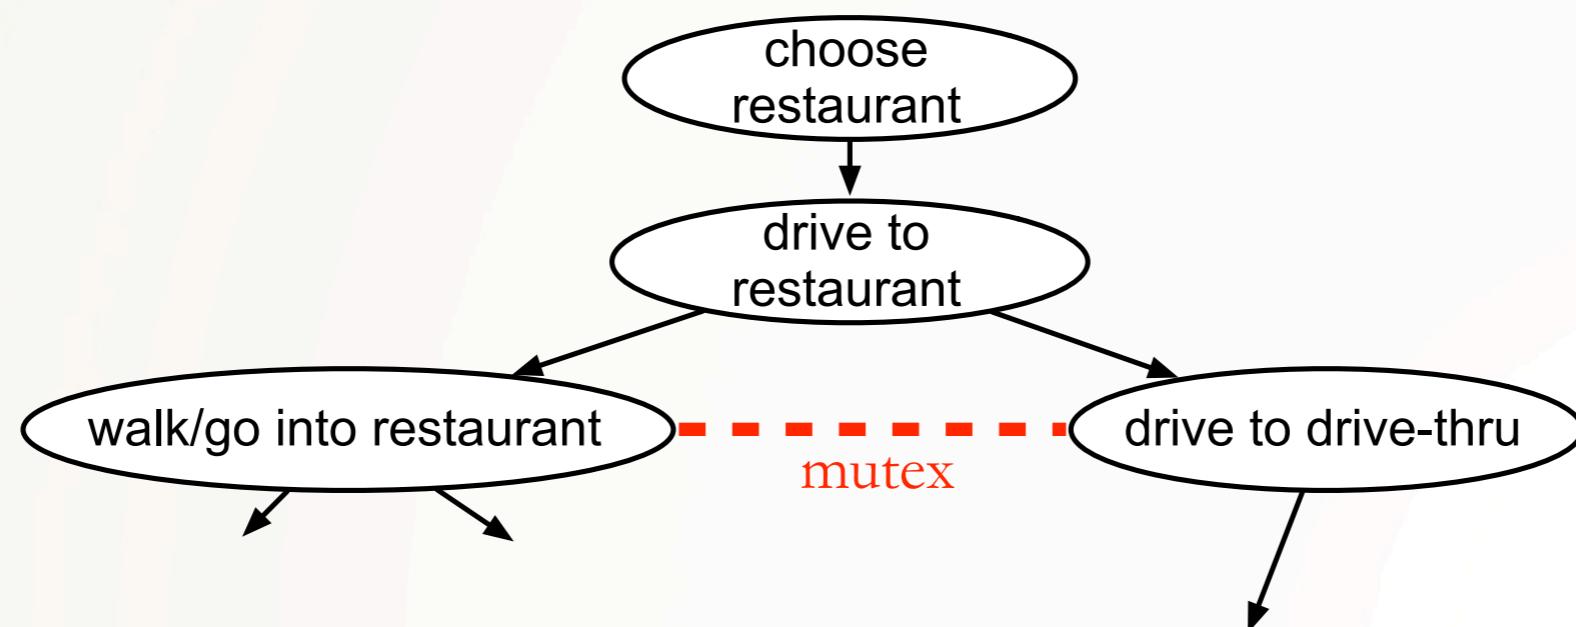




# Learning to tell stories

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

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

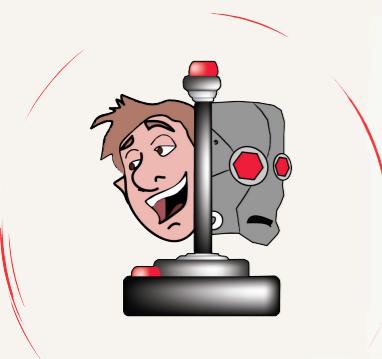




# Learning to tell stories

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

- Crowd can improve results at every stage of process
  - Re-cluster
  - Verify relations
  - Acquire additional examples



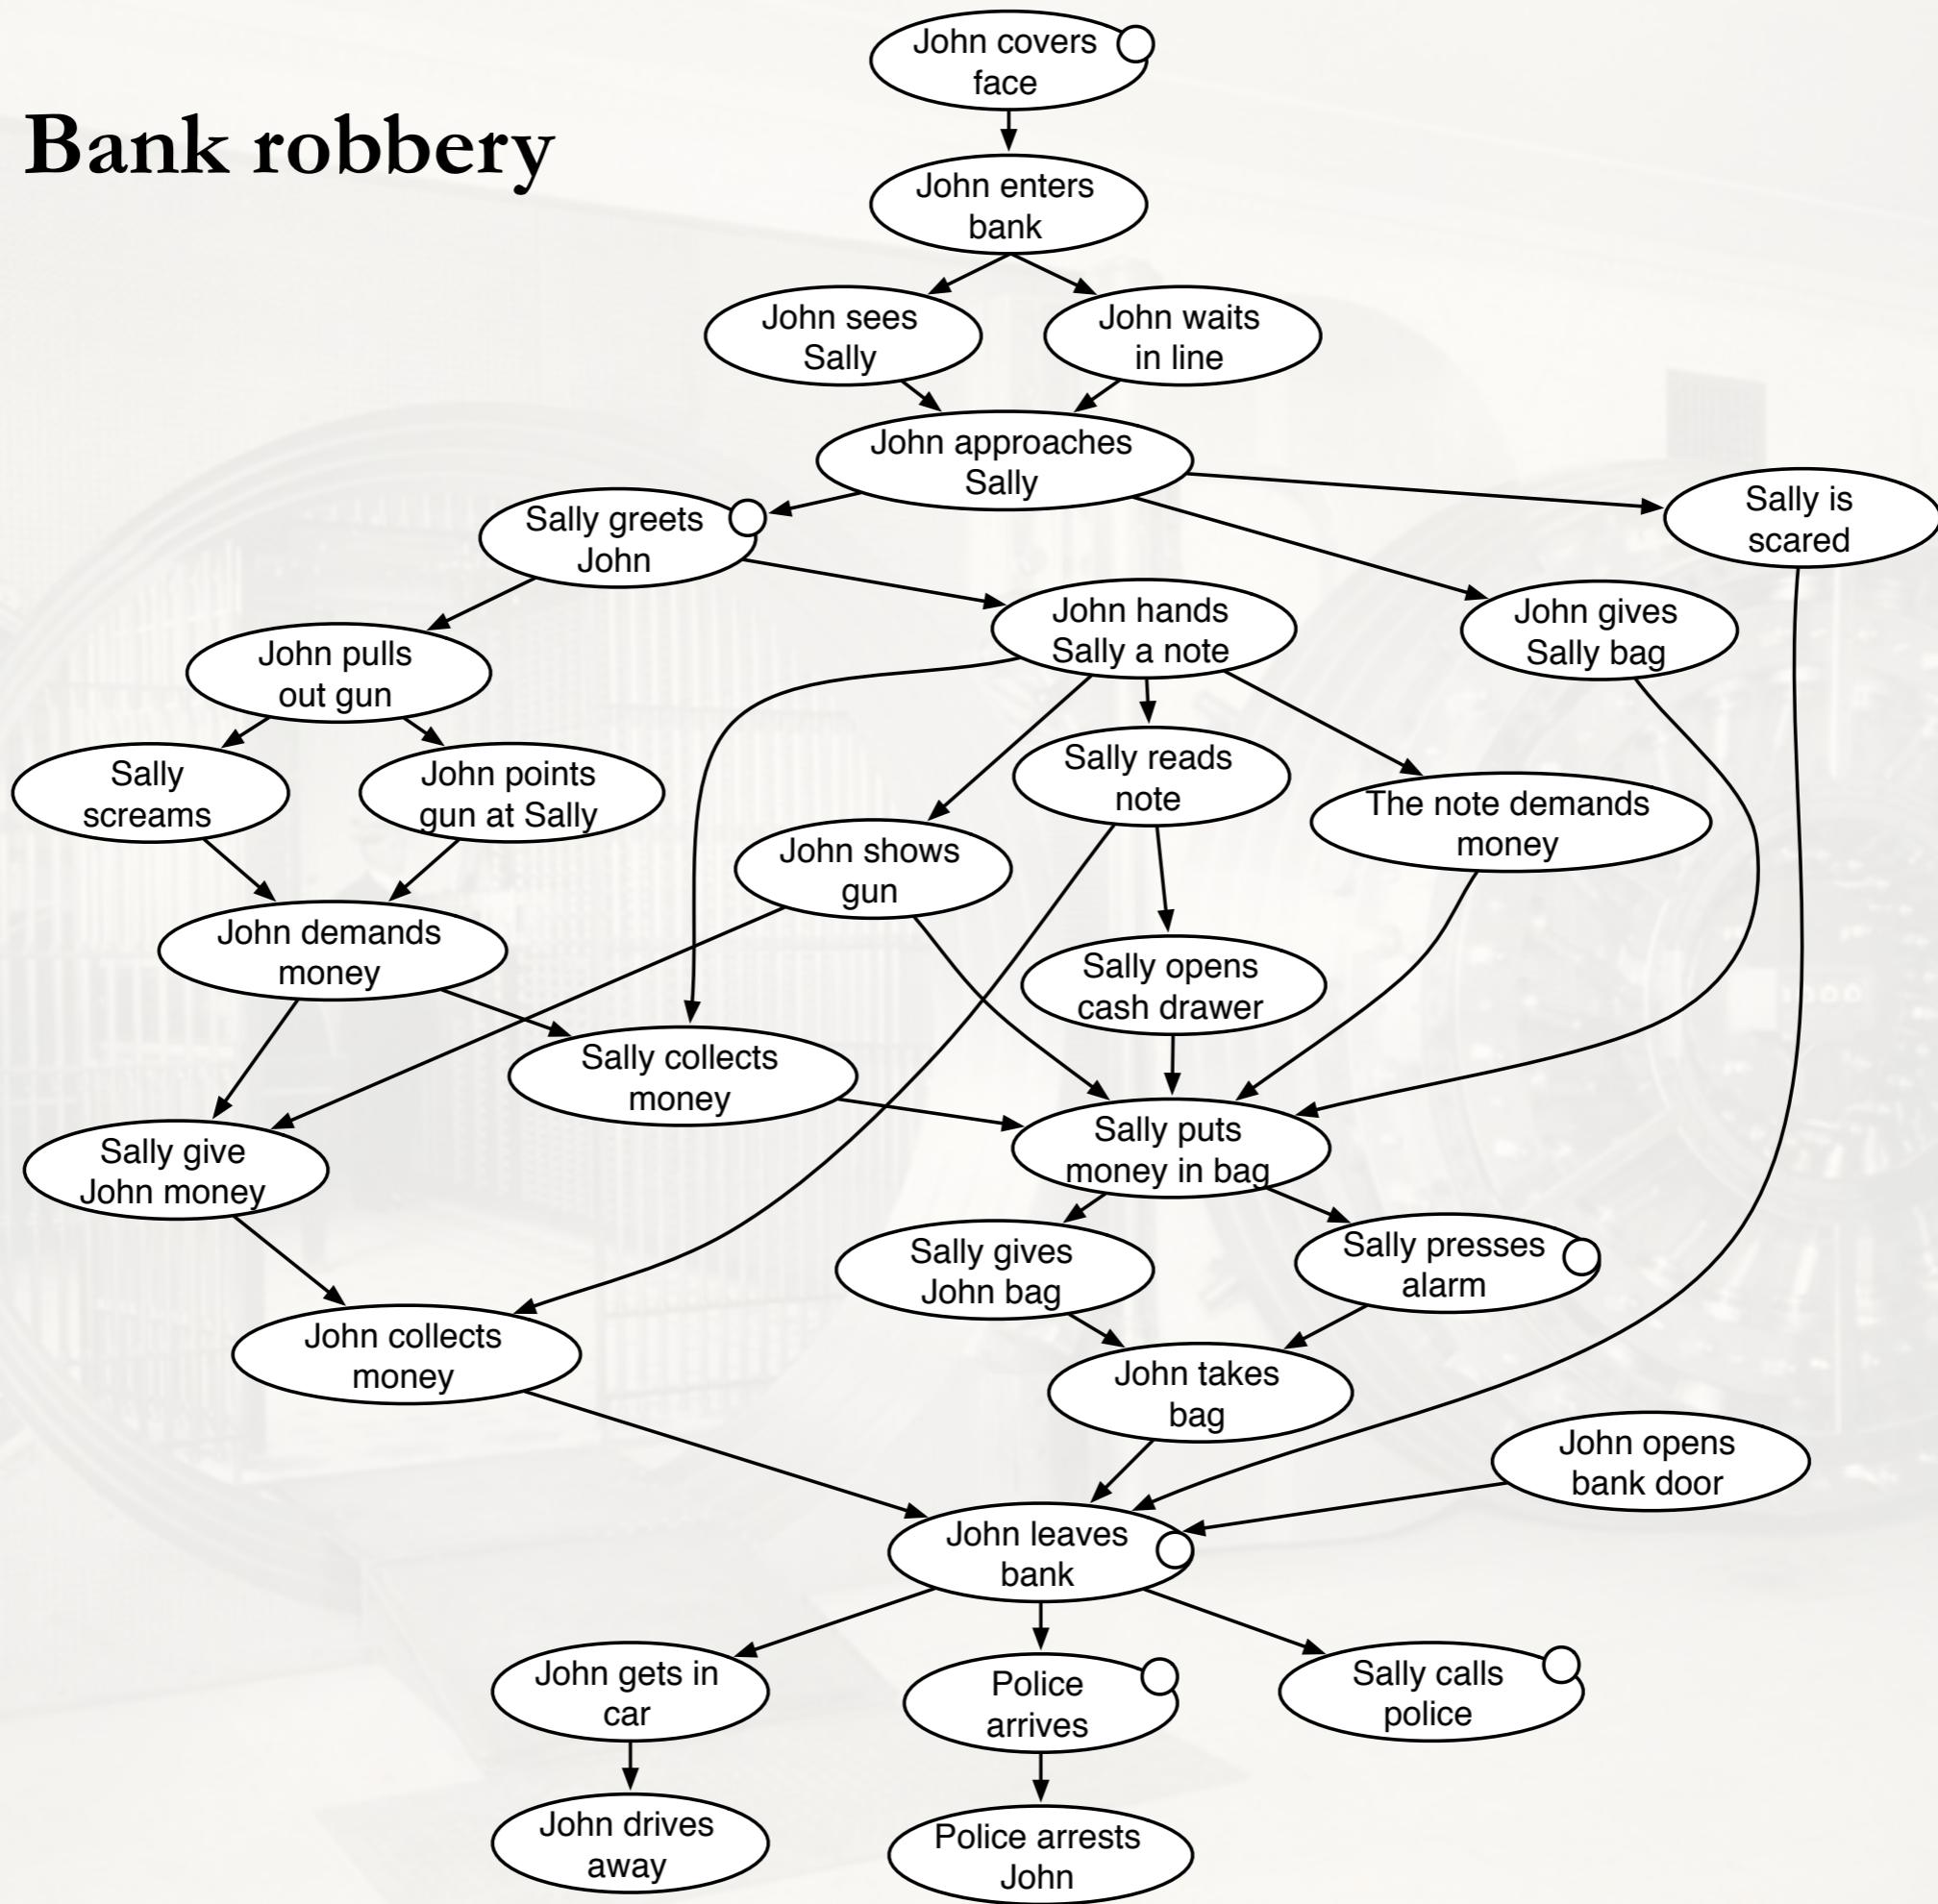
# Fast food restaurant

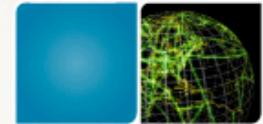




## Going on a date to the movies

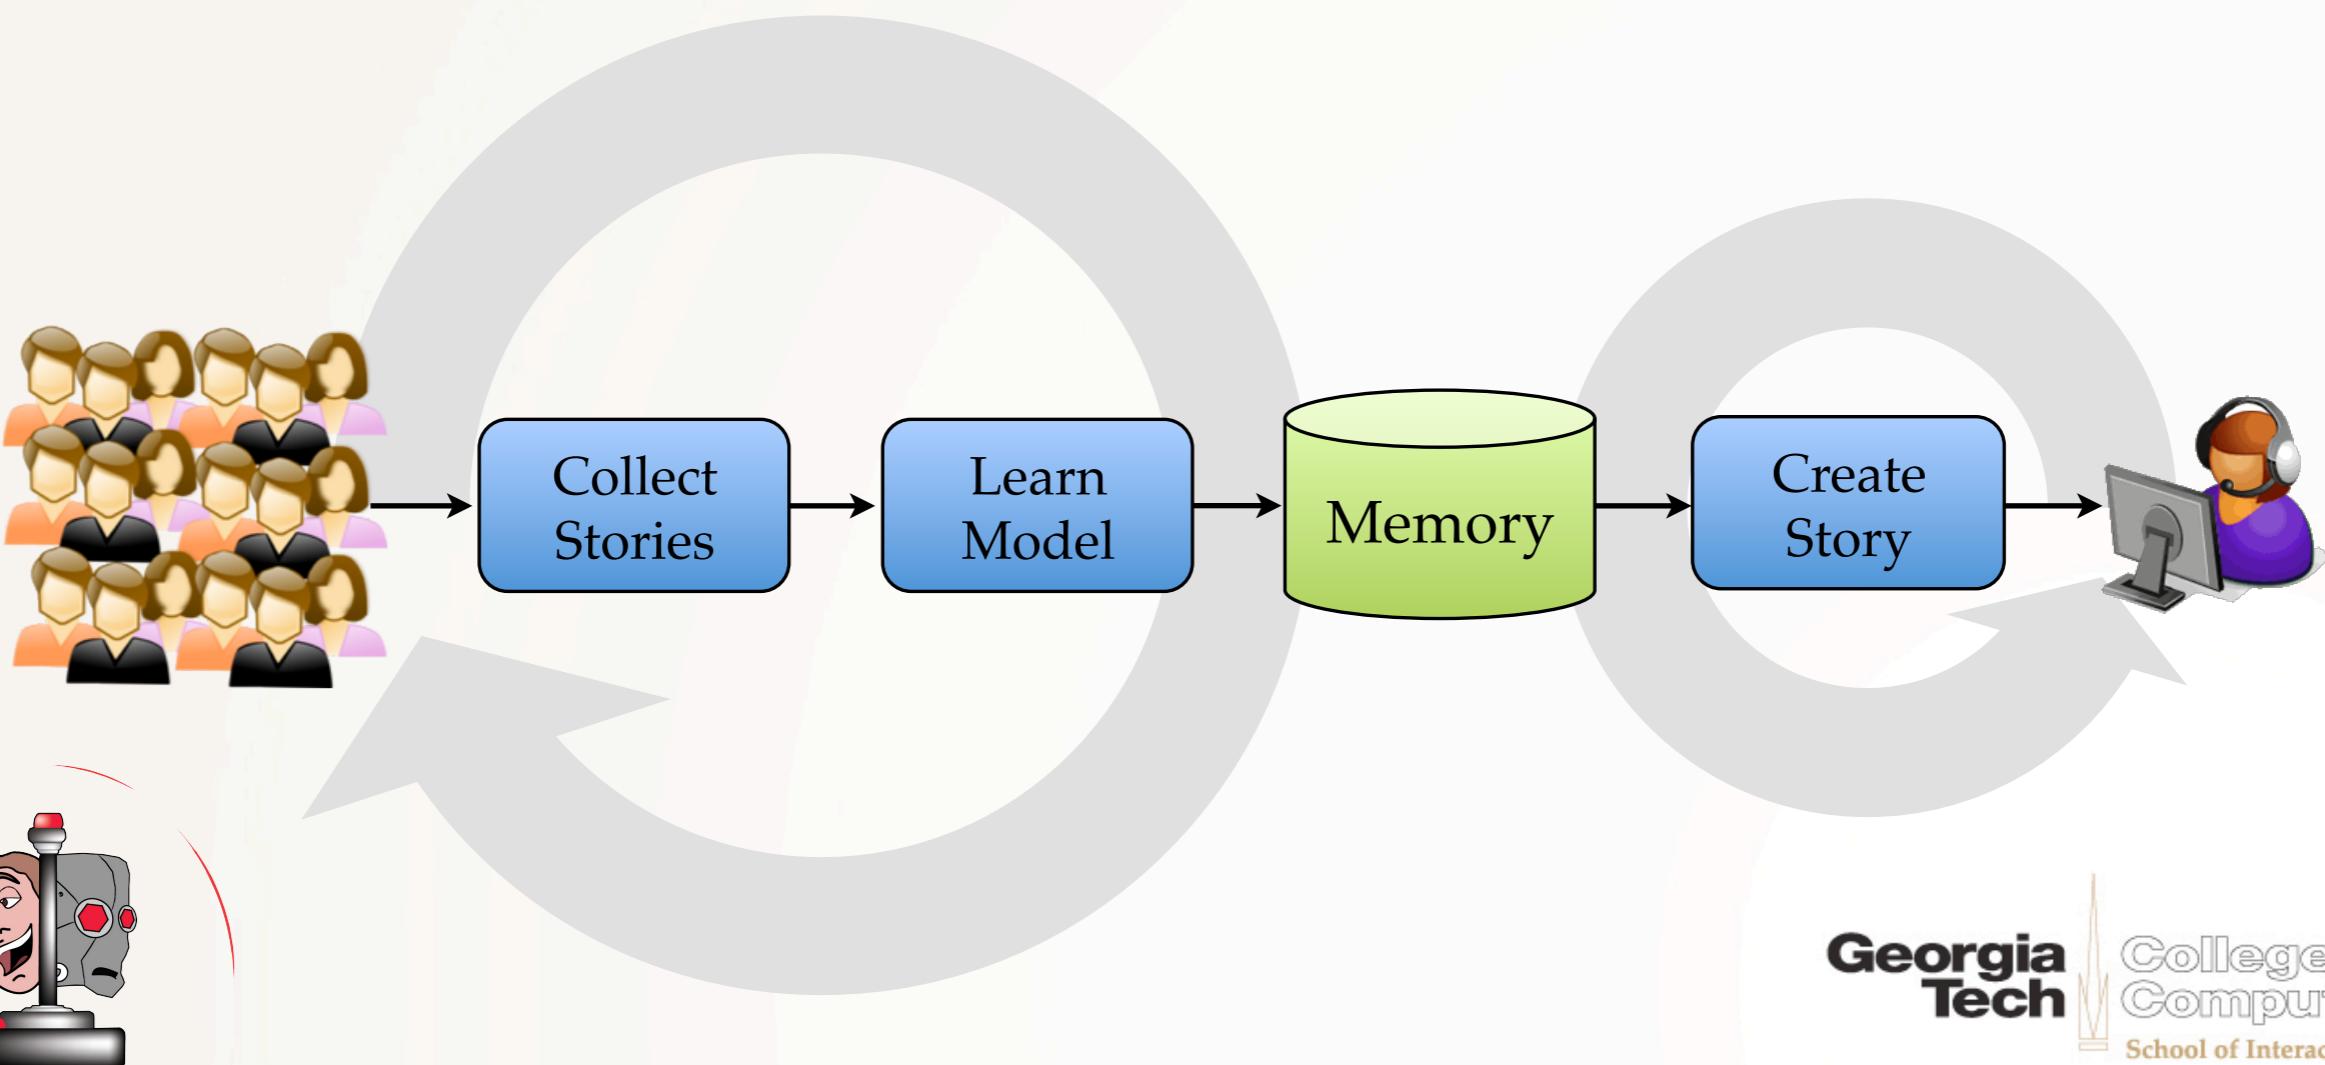
# Bank robbery



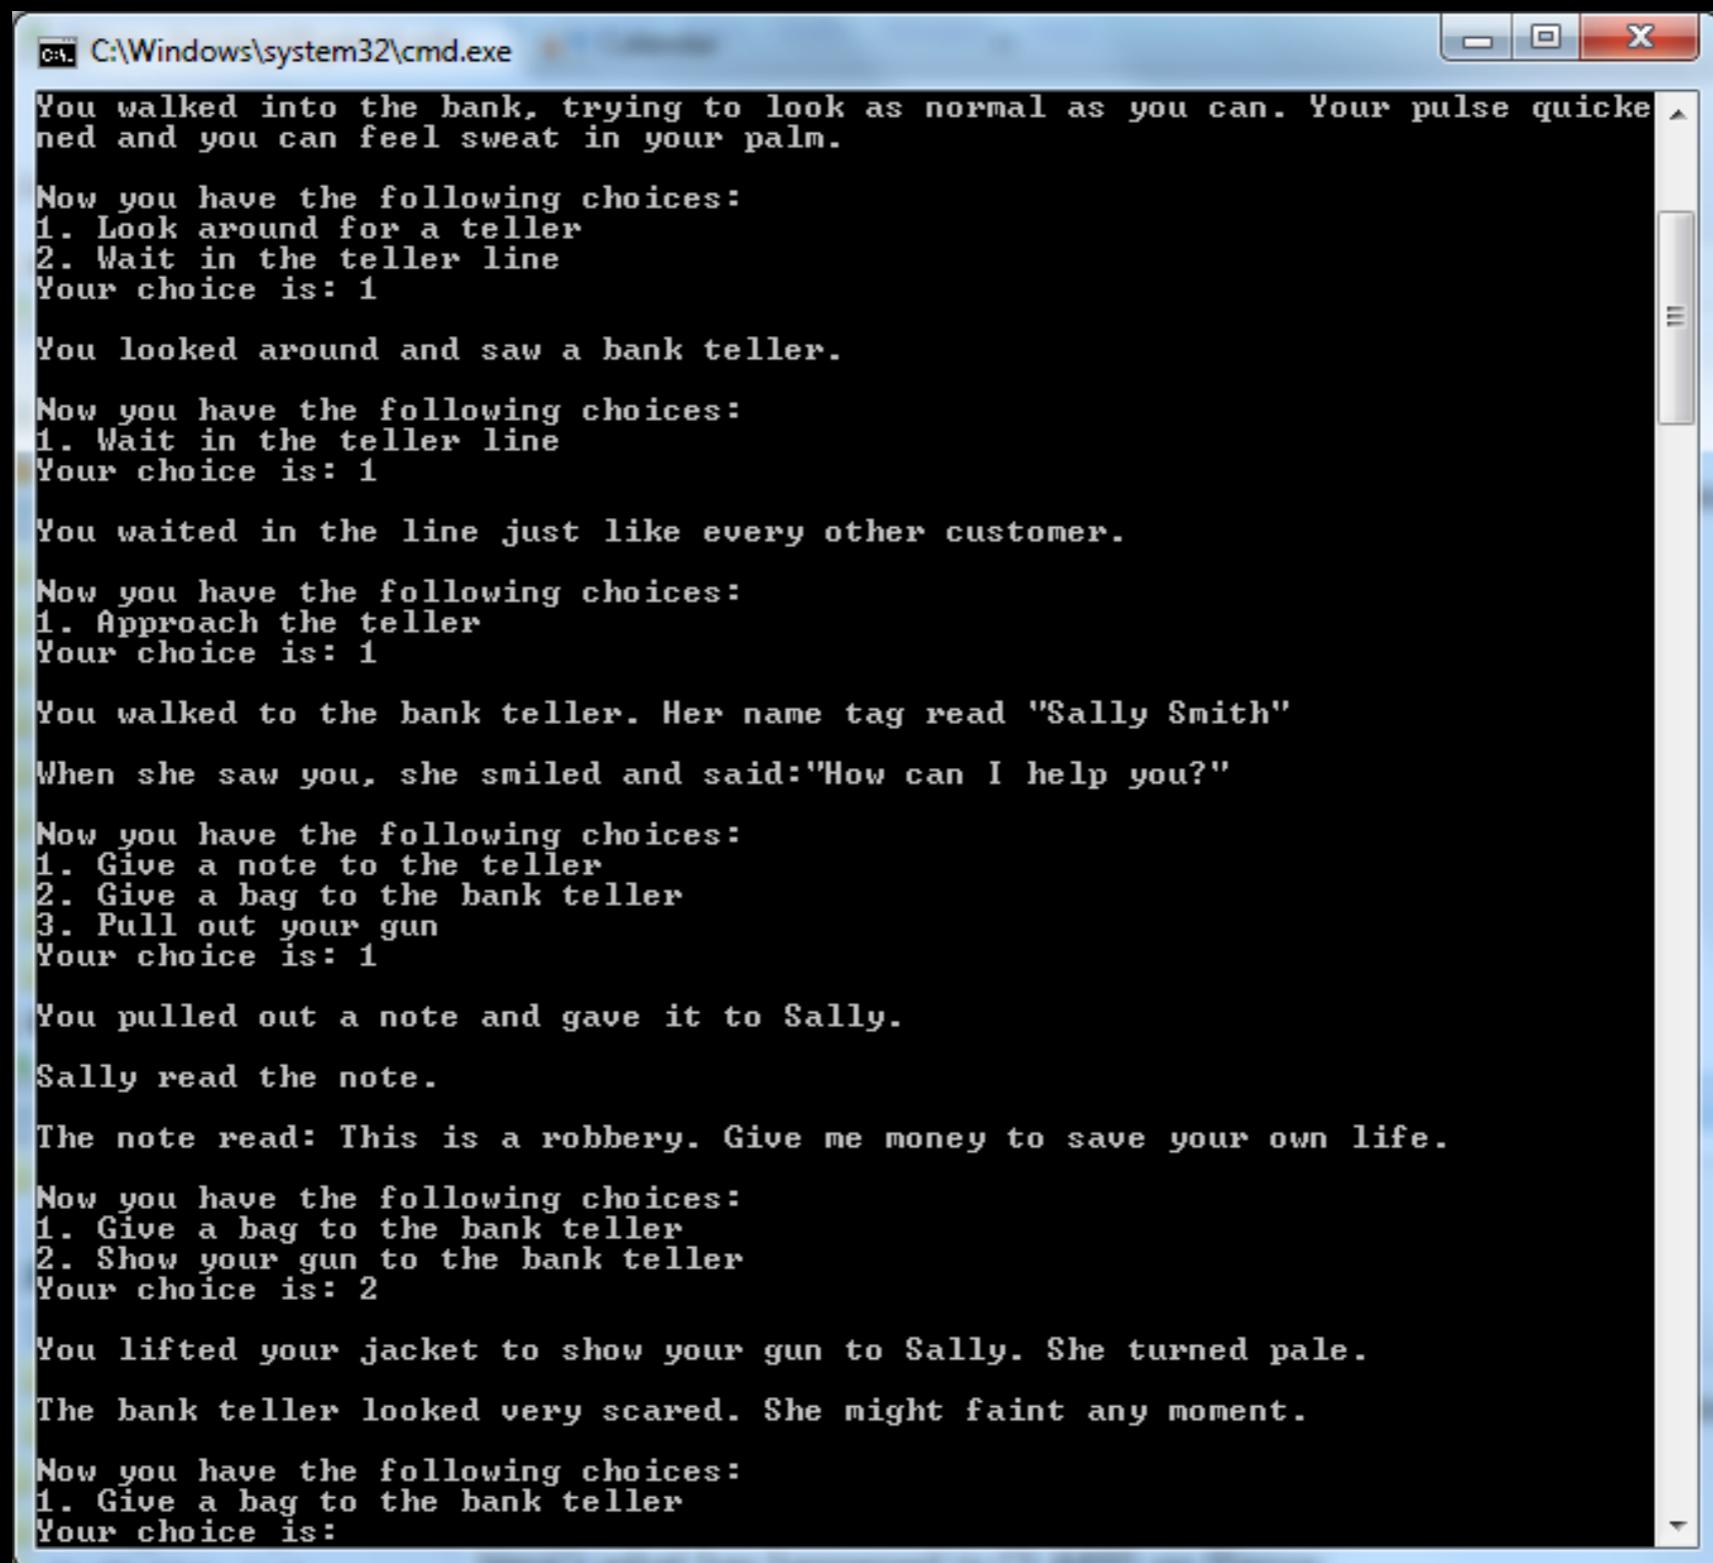


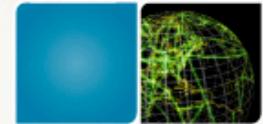
# Narrative generation

- Script defines a space of linear sequences
  - Search for most typical, least typical, specific events, surprising, etc.
  - Comparable to human-authored stories



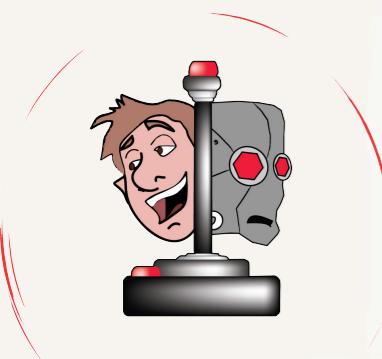
# Interactive narrative generation





# 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
- Creative, expressive computational systems

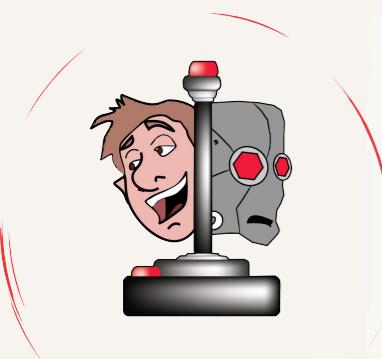


45



# Acknowledgements

- Boyang Li
- Brian O'Neill
- Alexander Zook
- Hong Yu
- Ken Hartsook
- Michael Drinkwater
- Steve Lee-Urban
- Sauvik Das
- Andrew Stern
- James Niehaus
- Don Dini
- Jason Alderman
- R. Michael Young



## Thanks! Questions?

riedl@cc.gatech.edu  
@mark\_riedl