CS 4260 and CS 5260
Vanderbilt University

Lecture on AI Story Telling

- [RB13] "Interactive Narrative: An Intelligent Systems Approach" by Mark Owen Riedl, Vadim Bulitko in *AI Magazine*, Vol. 34, No. 1, 2013
  https://www.aaai.org/ojs/index.php/aimagazine/article/view/2449 (read for final!)

- [RY10] "Narrative Planning: Balancing Plot and Character" by Mark Reidl and R. Michael Young in *Journal of Artificial Intelligence Research*, Vol. 39, 2010

Douglas H. Fisher

Talespin: An Example Story
James Meehan (1976)

"*John Bear is somewhat hungry. John Bear wants to get some berries. John Bear wants to get near the blueberries. John Bear walks from a cave entrance to the bush by going through a pass through a valley through a meadow. John Bear takes the blueberries. John Bear eats the blueberries. The blueberries are gone. John Bear is not very hungry. The end.*"

- One of the simplest stories generated by Tale-Spin (the first AI storyteller)
  - Single character
- "Bear world" is a simplified world
  - intrapersonal behavior in this example story
  - Akin to simple children's stories
  - Talking animals escape expectations of human behavior and relationships (like Paro, the therapeutic seal)

`Tale-Spin Comments (from reading The Idea of Tale-Spin and The Specifics of Tale-Spin by E. Charniak, C. Riesbeck, and D. McDermott (1980). Artificial Intelligence Programming (pp. 283 - 300). https://my.vanderbilt.edu/douglasfisher/files/2015/08/TaleSpin.pdf)

Talespin: An Example Story
James Meehan (1976)

"*Joe Bear was hungry. He asked Irving Bird if he knew where soe food was. Irving didn't but he decided to trick Joe. He said he would tell Joe where some bees were lived if Joe brought Irving a worm. Joe did and Irving ate the worm and flew off laughing at Joe. Joe never trusted Irving again. The End.*"

- Interpersonal behavior in this example story
- Even in this simple world, complex behaviors can be included, like deceit among multiple characters
- Not unlike "blocks world" in planning and natural language systems

`Tale-Spin Comments (from reading The Idea of Tale-Spin and The Specifics of Tale-Spin by E. Charniak, C. Riesbeck, and D. McDermott (1980). Artificial Intelligence Programming (pp. 283 - 300). https://my.vanderbilt.edu/douglasfisher/files/2015/08/TaleSpin.pdf)

Talespin: Generating Stories

"*Joe Bear wanted some honey. He went to his cave but there wasn't any there. Then he asked Irving Bird for some. Irving asked Joe for a worm. Joe went to his cave and got a worm. He gave it to Irving, and Irving gave Joe some honey. The End.*"

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

Scripts (skeletal story fragments)

TRADE
Pattern: (GOAL ?W (HAS ?W ?X))
Filter: (GOAL ?Y (HAS ?Y ?Z))
where ?Y <> ?W and ?Z <> ?X
Event Tree: (GOAL ?W (ASK ?W ?Y ?X))
…

ASK-FOR
Pattern: (GOAL ?W (ASK ?W ?X ?Y))
/* note that ?W, ?X, ?Y have same names as ?W,?X,?Y variables in TRADE story fragment, but different variables */
Filter: (AT ?W ?Z) and (AT ?X ?Z)

Talespin: Generating Stories (for acquiring goals)

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

Generating stories is a search;
e.g., could have had Irving be the proactive character

Scripts (skeletal story fragments)

TRADE
Pattern: (GOAL ?W (HAS ?W ?X))
Filter: (GOAL ?Y (HAS ?Y ?Z))
where ?Y <> ?W and ?Z <> ?X
Event Tree: (GOAL ?W (ASK ?W ?Y ?X))
…

(GOAL JOE (HAS JOE HONEY)) matches the pattern of TRADE
(GOAL ?W (HAS ?W ?X)) where ?W is bound to JOE and ?X is bound to HONEY. Applying substitutions yields

TRADE
Pattern: (GOAL JOE (HAS JOE HONEY))
Filter: (GOAL Y (HAS ?Y ?Z)) <this unifies with (GOAL IRVING (HAS IRVING WORM))>
where ?Y <> ?W and ?Z <> ?X
Event Tree: (GOAL JOE (ASK JOE IRVING HONEY))
(JOE asks IRVING for HONEY)

Talespin: Generating Stories (for acquiring goals)

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

Scripts (skeletal story fragments)

TRADE
Pattern: (GOAL ?W (HAS ?W ?X))
Filter: (GOAL ?Y (HAS ?Y ?Z))
where ?Y <> ?W and ?Z <> ?X
Event Tree: (GOAL ?W (ASK ?W ?Y ?X))
…

(GOAL IRVING (HAS IRVING WORM)) matches the pattern of TRADE
(GOAL ?W (HAS ?W ?X)) where ?W is bound to IRVING and ?X is bound to WORM. Applying substitutions yields

TRADE
Pattern: (GOAL IRVING (HAS IRVING WORM))
Filter: (GOAL Y (HAS ?Y ?Z)) <this unifies with (GOAL JOE (HAS JOE HONEY))>
where ?Y <> ?W and ?Z <> ?X
Event Tree: (GOAL IRVING (ASK IRVING JOE WORM)) (IRVING asks JOE for WORM)

Generating stories is a search; e.g., could have had Irving be the proactive character

Talespin: Generating Stories

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

TRADE
Pattern: (GOAL ?W (HAS ?W ?X))
Filter: (GOAL ?Y (HAS ?Y ?Z))
Event Tree: (GOAL ?W (ASK ?W ?Y ?X))
SUCCESS
(GOAL ?Y (ASK ?Y ?W ?Z))
SUCCESS
(GOAL ?W (GIVE ?W ?Y ?Z))
SUCCESS
(GOAL ?Y (GIVE ?Y ?W ?X))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

Talespin: Generating Stories

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

TRADE
Pattern: (GOAL Joe (HAS Joe Honey))
Filter: (GOAL ?Y (HAS ?Y ?Z))
Event Tree: (GOAL Joe (ASK Joe ?Y Honey))
SUCCESS
(GOAL ?Y (ASK ?Y Joe ?Z))
SUCCESS
(GOAL Joe (GIVE Joe ?Y ?Z))
SUCCESS
(GOAL ?Y (GIVE ?Y Joe Honey))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

Talespin: Generating Stories

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

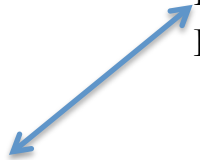

TRADE
Pattern: (GOAL Joe (HAS Joe Honey))
Filter: (GOAL Irving (HAS Irving Worm))
Event Tree: (GOAL Joe (ASK Joe Irving Honey))
SUCCESS
(GOAL Irving (ASK Irving Joe Worm))
SUCCESS
(GOAL Joe (GIVE Joe Irving Worm))
SUCCESS
(GOAL Irving (GIVE Irving Joe Honey))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

Talespin: Generating Stories

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

Scripts trigger other scripts

TRADE
Pattern: (GOAL Joe (HAS Joe Honey))
Filter: (GOAL Irving (HAS Irving Worm))
Event Tree: (GOAL Joe (ASK Joe Irving Honey))
SUCCESS
(GOAL Irving (ASK Irving Joe Worm))
SUCCESS
(GOAL Joe (GIVE Joe Irving Worm))
SUCCESS
(GOAL Irving (GIVE Irving Joe Honey))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

ASK-FOR
Pattern: (GOAL ?W (ASK ?W ?X ?Y))
/* note that variable binding/substitution did NOT happen across story fragments; ?W, ?X, ?Y have same names as ?W,?X,?Y variables in TRADE story fragment, but they are different variables */
Filter: (AT ?W ?Z) and (AT ?X ?Z)

Talespin: Generating Stories

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

TRADE
Pattern: (GOAL Joe (HAS Joe Honey))
Filter: (GOAL Irving (HAS Irving Worm))
Event Tree: (GOAL Joe (ASK Joe Irving Honey))
SUCCESS
(GOAL Irving (ASK Irving Joe Worm))
SUCCESS
(GOAL Joe (GIVE Joe Irving Worm))
SUCCESS
(GOAL Irving (GIVE Irving Joe Honey))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

Piecing together
and instantiating story fragments
To create a cohesive whole

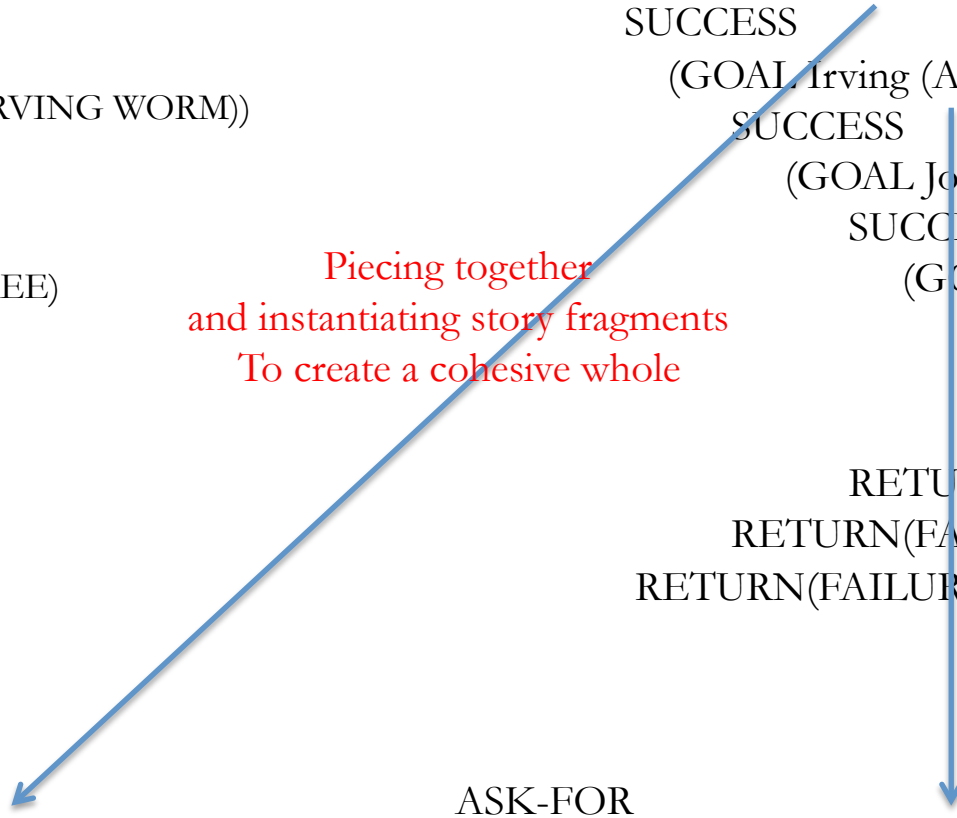

ASK-FOR
Pattern: (GOAL Joe (ASK Joe Irving Honey))
Filter: (AT Joe ?Z) and (AT Irving ?Z)

ASK-FOR
Pattern: (GOAL Irving (ASK Irving Joe Worm))
Filter: (AT Irving ?Z) and (AT Joe ?Z)

Talespin: Generating Stories

Filter is analogous to OP preconditions in planning

Start state of the world

(GOAL JOE (HAS JOE HONEY))

(GOAL IRVING (HAS IRVING WORM))

(HAS IRVING HONEY)

(HOME-OF IRVING TREE)

(HOME-OF JOE CAVE)

(AT WORM CAVE)

(AT JOE ROCK)

TRADE
Pattern: (GOAL Joe (HAS Joe Honey))
Filter: (GOAL Irving (HAS Irving Worm))
Event Tree: (GOAL Joe (ASK Joe Irving Honey))
SUCCESS
(GOAL Irving (ASK Irving Joe Worm))
SUCCESS
(GOAL Joe (GIVE Joe Irving Worm))
SUCCESS
(GOAL Irving (GIVE Irving Joe Honey))
RETURN(SUCCESS)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)
RETURN(FAILURE)

Some scripts change (Add, Delete) conditions in the world

You could implement a Tale-Spin like system using a planning system
Story Generation as planning

Story fragments can be viewed as planning operators and macro operators

Templates can be reused and recombined for multiple stories

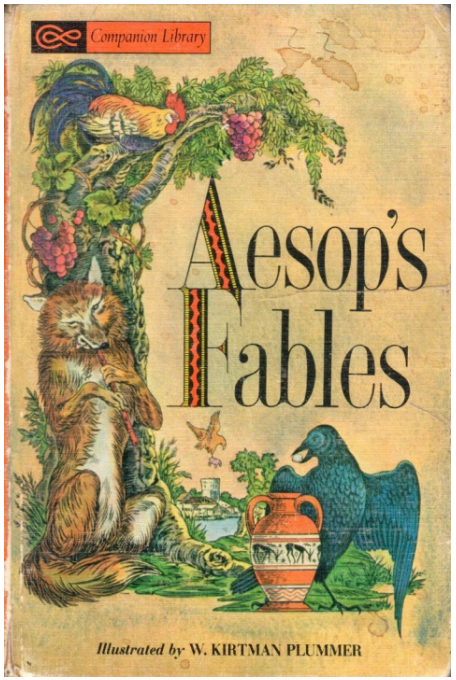

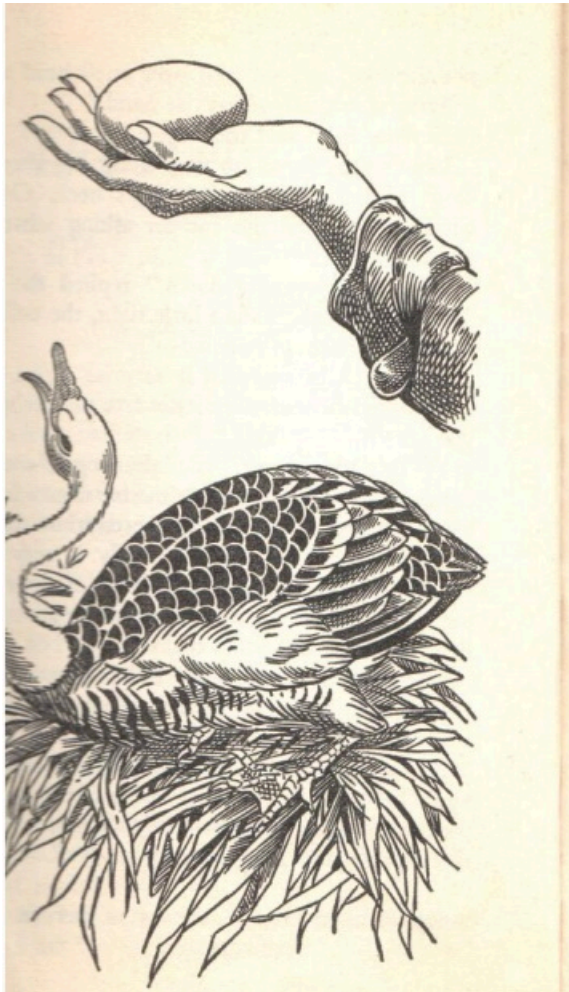

The Goose With the Golden Eggs

A FARMER went to the nest of his goose to see whether she had laid an egg. To his surprise he found, instead of an ordinary goose egg, an egg of solid gold. Seizing the golden egg he rushed to the house in great excitement to show it to his wife.

Every day thereafter the goose laid an egg of pure gold. But as the farmer grew rich he grew greedy. And thinking that if he killed the goose he could have all her treasure at once, he cut her open only to find—nothing at all.

★ *Application:* THE GREEDY WHO WANT MORE LOSE ALL.

[29]

Storytelling as planning

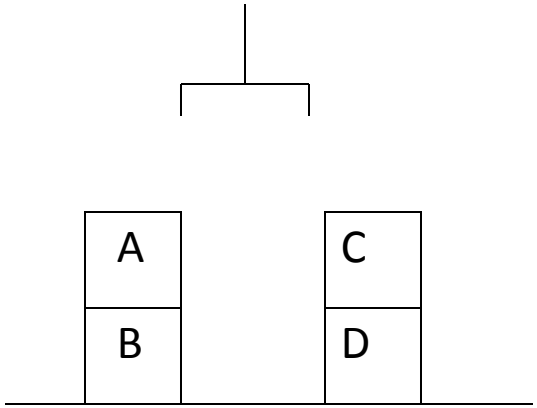

Initial State

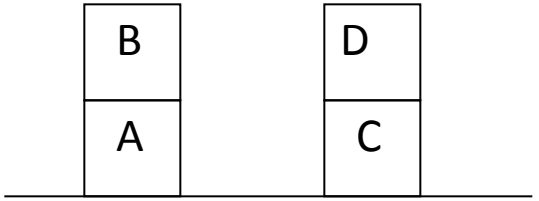

Goal spec

ON(A,B)
ONTAB(B)
CLEAR(A)
ON(C,D)
ONTAB(D)
CLEAR(C)
HANDEMPTY

One could substitute metaphorical characters, places, and things for "blocks" to achieve story-like narratives

"reversal of fortune"
"turning the tables"

ON(B,A)
CLEAR(B)
ON(D,C)
CLEAR(D)

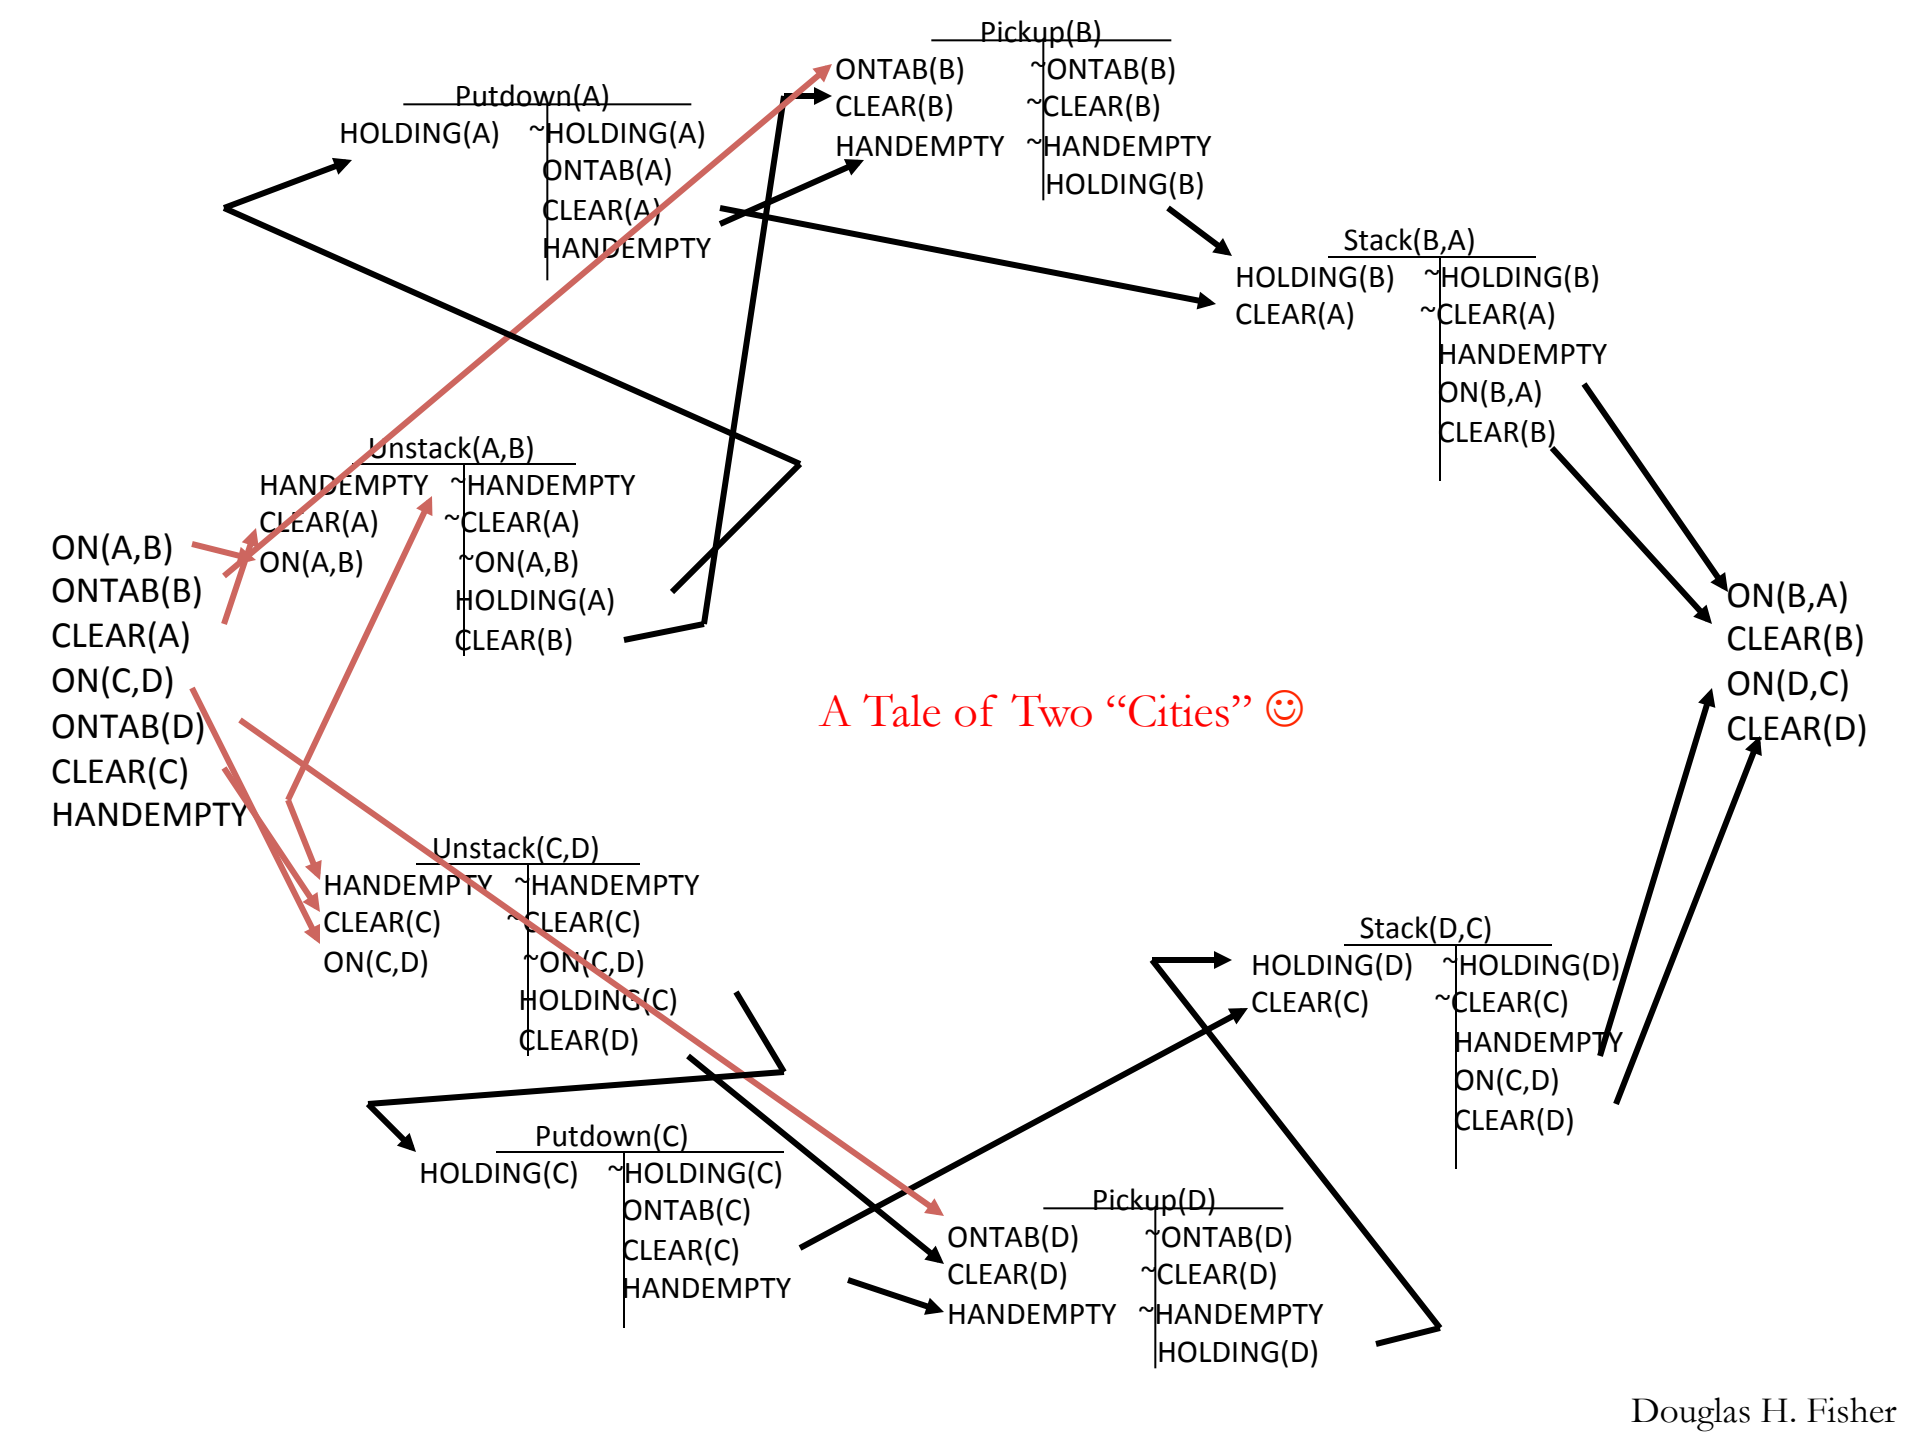

Douglas H. Fisher

Independent Study Possibility

- Using Chess games as skeletal stories
- Utilizing metaphorical characters for pieces, of which Chess has rich possibilities
- Different from other "fantasy chess" games (e.g., http://chessaria.com/) -- we will discuss interactive narrative shortly
- Story teller recounts game (so story generation problem partly taken care of), but …
- … interesting research questions are in developing a story teller that can project on what might have been
- The Bard Project

http://www.radiolab.org/story/153809-rules-set-you-free/
(5:30 – 13:45; 14:50 – 19:00)

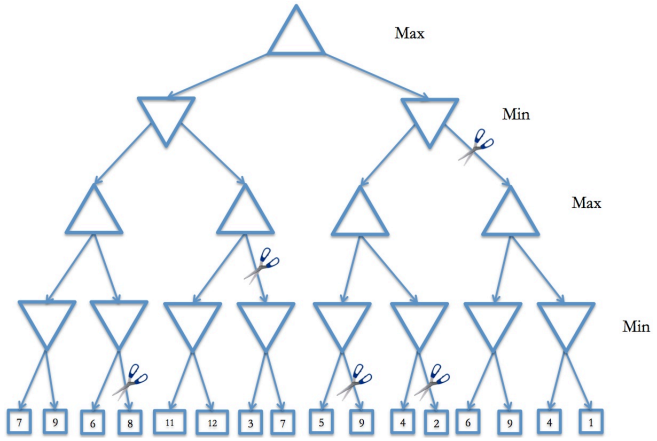

Other story-telling paradigms

- Story-telling with and about data

  For example, https://narrativescience.com/Intelligent-Narratives

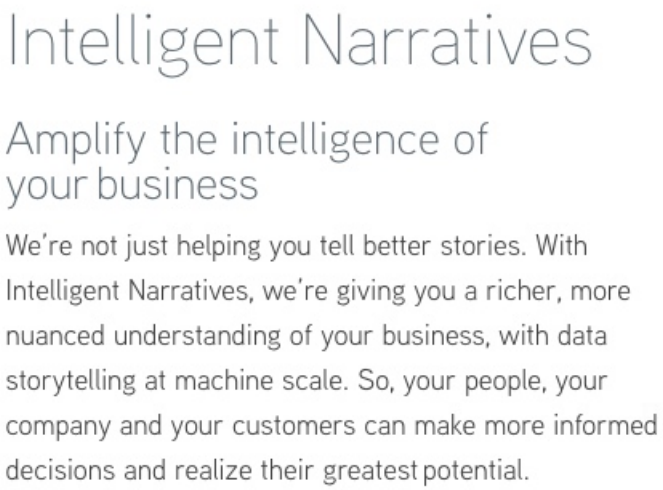

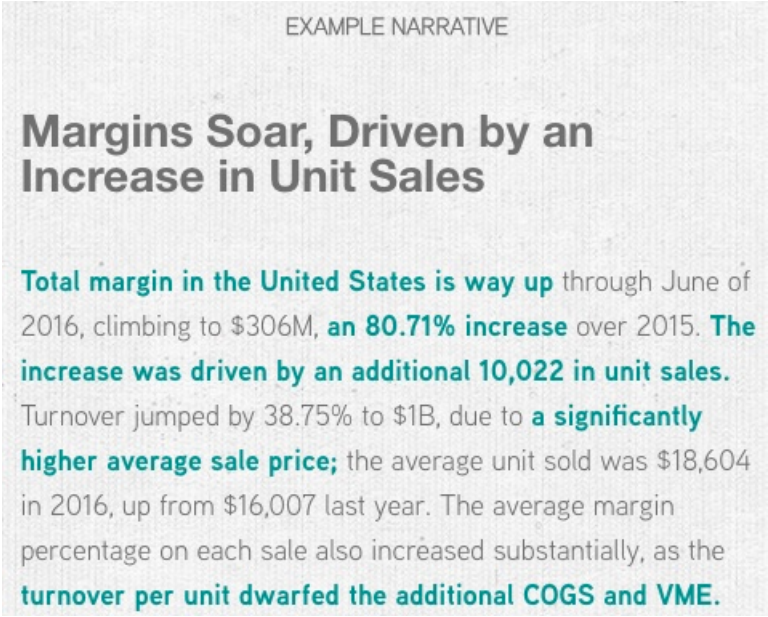

- Tour Guides

  Lim, M. Y, Aylett, R. (2007) Narrative Construction in a Mobile Tour Guide, In Proceedings of the 4th International Conference on Virtual Storytelling
  https://pdfs.semanticscholar.org/c236/d06aae12739510208751f902921bb1696cfe.pdf

  Region Radio (see Sustainability lecture)

Automatic Storytelling: Or, How to Build Your Very Own Data Scientist

by JUSTIN

As a data scientist, one of my primary jobs is to interpret large amounts of data and craft a story about about what I discover. Not everyone is a data geek that wants to wade through large data sets, so it's important to find a way to communicate insights from data that everyone can understand. Unfortunately, the process of analyzing data and compiling interesting results can be very time consuming. Even so, after telling these stories many times, some patterns emerge in the data analysis and communication of the findings. This led the data science team at Chartbeat to ask the question: Can we build an automated data scientist that can seek out interesting stories within our data?

Before we tackle "automated storytelling," as we call it, let's walk through the process I might go through when analyzing some data. Consider the small data set about pageviews, tweets, and average Engaged Time to one article on a completely hypothetical website.

http://blog.chartbeat.com/2014/08/19/automatic-storytelling-build-data-scientist/

Even though "article" had below average engagement for "website.com," readers shared this story 5 times more often than the typical story.

Let's break down where this insight came from. We see that "article" had five tweets, but without context, this does not tell us much. A great way to give context to a number is to compare it to a benchmark. For example, how does this number compare to the typical article on this website or the Internet as a whole? Put into the context of a larger picture, we can judge if a number is good or not. In this case, we are given all we need to know about Twitter shares across the site, so let's compare Twitter activity on "article" to the average article on "website.com." However, since the overall site has much more traffic than "article," comparing the number of tweets for each would be unfair. When comparing numbers, it is important to compare apples to apples. The standard way to deal with this is to normalize your values. In this case, we consider the tweet rate for both. That is the number of tweets per pageview:

http://blog.chartbeat.com/2014/08/19/automatic-storytelling-build-data-scientist/

Rather than writing one template sentence per branch of the decision tree, we can create a collection of templates. This serves to create an illusion of a real data scientist telling you these facts, and will prevent the results from getting stale. We can also use additional data to include related data points. For example, in the case when the story is active on Twitter, we could enhance our original insight in the following way:

Even though "article" had below average engagement, readers shared this story 5 times more often than the typical story. In fact, the tweet from "user" generated 20 pageviews and 100 total seconds of engagement for this story.

Every time a question is asked in the decision tree, if there is additional data available, we can automatically add in extra information to flesh out the narrative.

Other story-telling paradigms

- AI Analysis of Stories

The Six Main Arcs in Storytelling, as Identified by an A.I. - The Atlantic
https://www.theatlantic.com/technology/archive/2016/07/...storytelling.../490733/ ▾
Jul 12, 2016 - The Six Main Arcs in **Storytelling**, as Identified by an **A.I.**. A machine mapped the most frequently .... They collected computer-generated story arcs for nearly 2,000 works of fiction, classifying each into one of six core **types** of narratives (based on what happens to the protagonist):. 1. Rags to Riches (rise). 2.

Vonnegut mapping the narrative arc of popular storylines along a simple graph https://www.youtube.com/watch?v=oP3c1h8v2ZQ

Artificial Intelligence Reveals Popular Storytelling Paths - Glean.info
glean.info/artificial-intelligence-reveals-popular-storytelling-paths/ ▾
Apr 7, 2017 - Using **artificial intelligence**, researchers identifies the most popular emotional paths of stories. ... But **AI** and computer software may be able to help PR and marketing improve their **storytelling** capabilities and develop stories better targeted to ... The results classified stories into six basic **types** of narratives.

- Story Telling Technologies

How Artificial Intelligence Is Changing Storytelling | HuffPost
https://www.huffingtonpost.com/.../how-artificial-intelligence-is-changing-storytelling... ▾
Jul 12, 2017 - Fardinpour says this **kind** of technology can substitute for the lack of mainstream media coverage or misleading coverage to educate kids and even adults on the ... "As **storytellers** it is important to consider that perhaps we are one step closer to creating a truly dynamic story arch with **Artificial intelligence**.

https://futureofstorytelling.org/video/nancy-pearson-watson-you-storytelling-in-the-age-of-ai

How might sentiment analysis be used to characterize narrative trajectories?

Need sentiment analysis be limited to words only?

Interactive Narrative

- Façade: http://www.interactivestory.net/ (play trailer)
- [RB13] "Interactive Narrative: An Intelligent Systems Approach" by Mark Owen Riedl, Vadim Bulitko in *AI Magazine*, Vol. 34, No. 1, 2013 https://www.aaai.org/ojs/index.php/aimagazine/article/view/2449

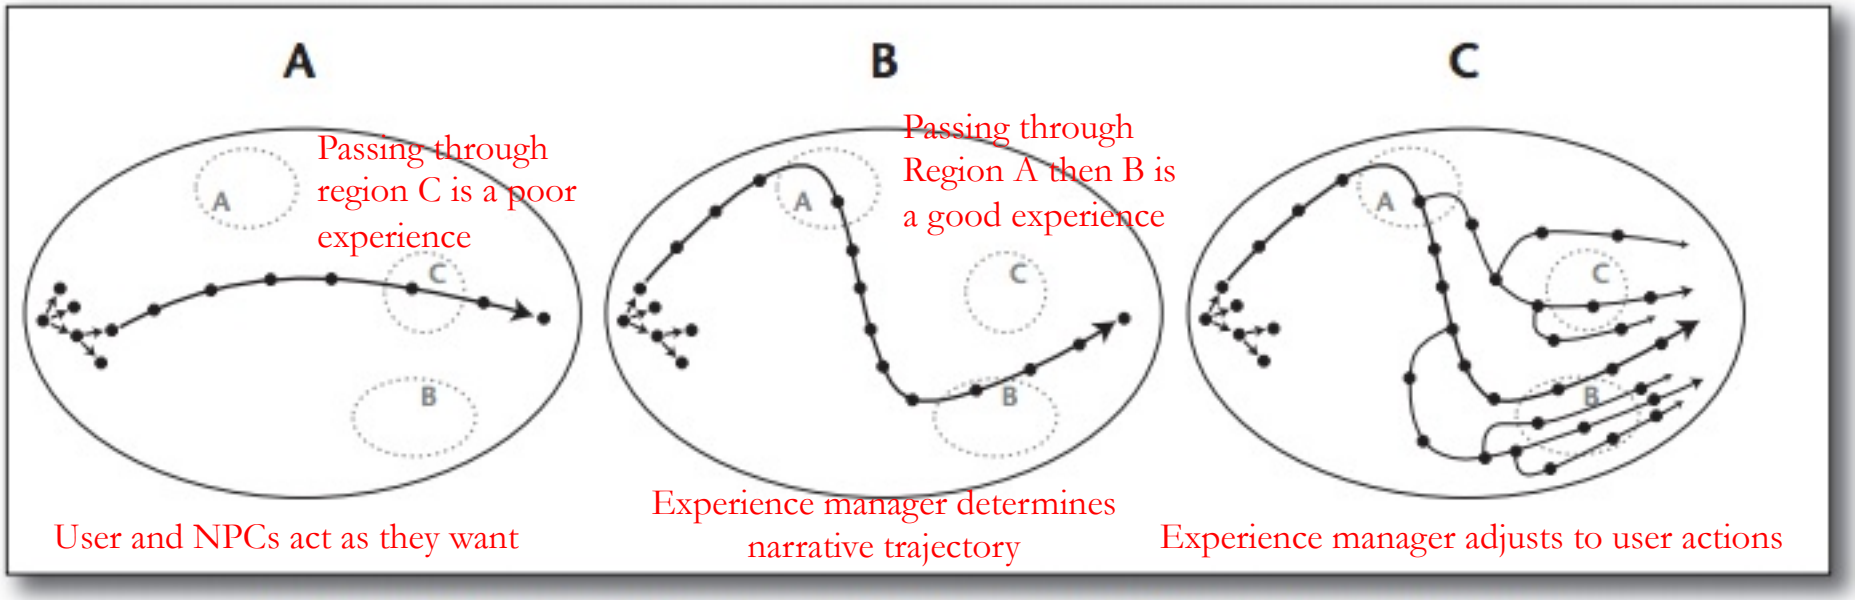

Figure 1. The Experience Management Problem Is to Compute Trajectories through State Space.

a. A possible narrative trajectory through state space. b. A possible narrative trajectory that visits states deemed favorable and avoids states deemed unfavorable. c. Accounting for player interaction.

Strong Autonomy ⟷ Strong Story

Combination/Tradeoff

Interactive Narrative

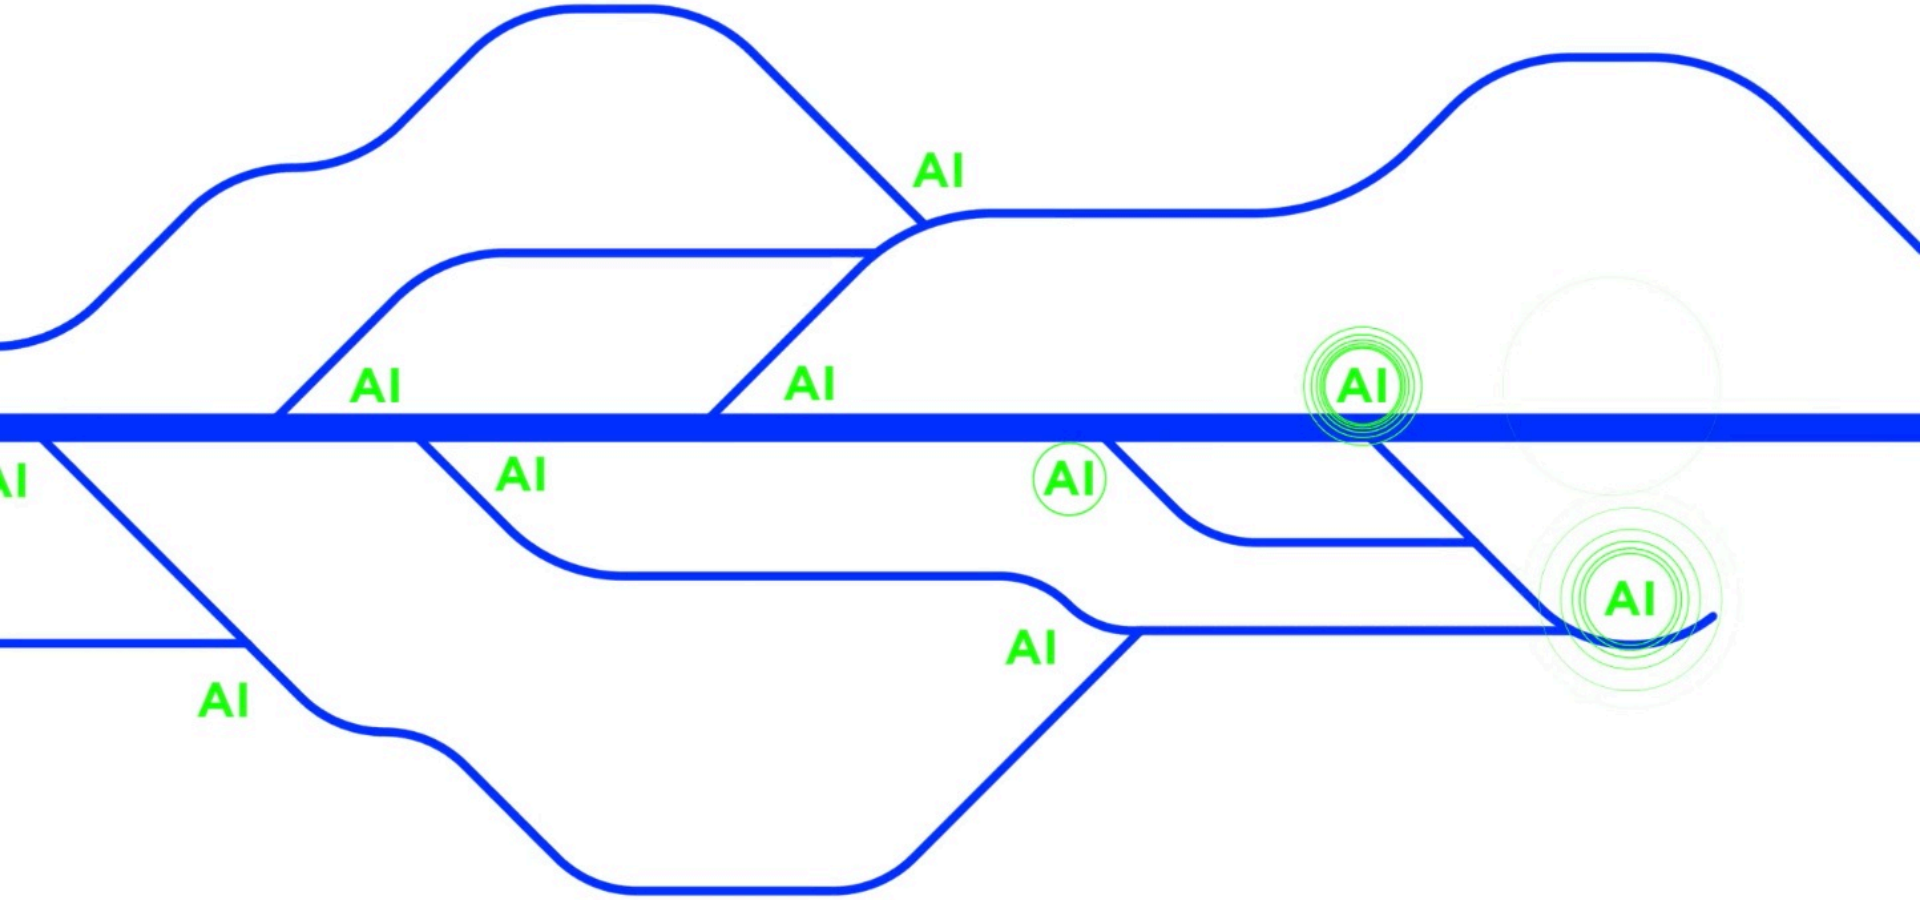

https://futureofstorytelling.org/video/andrew-stern-the-ground-rules-for-ai-storytelling

Interactive Narrative

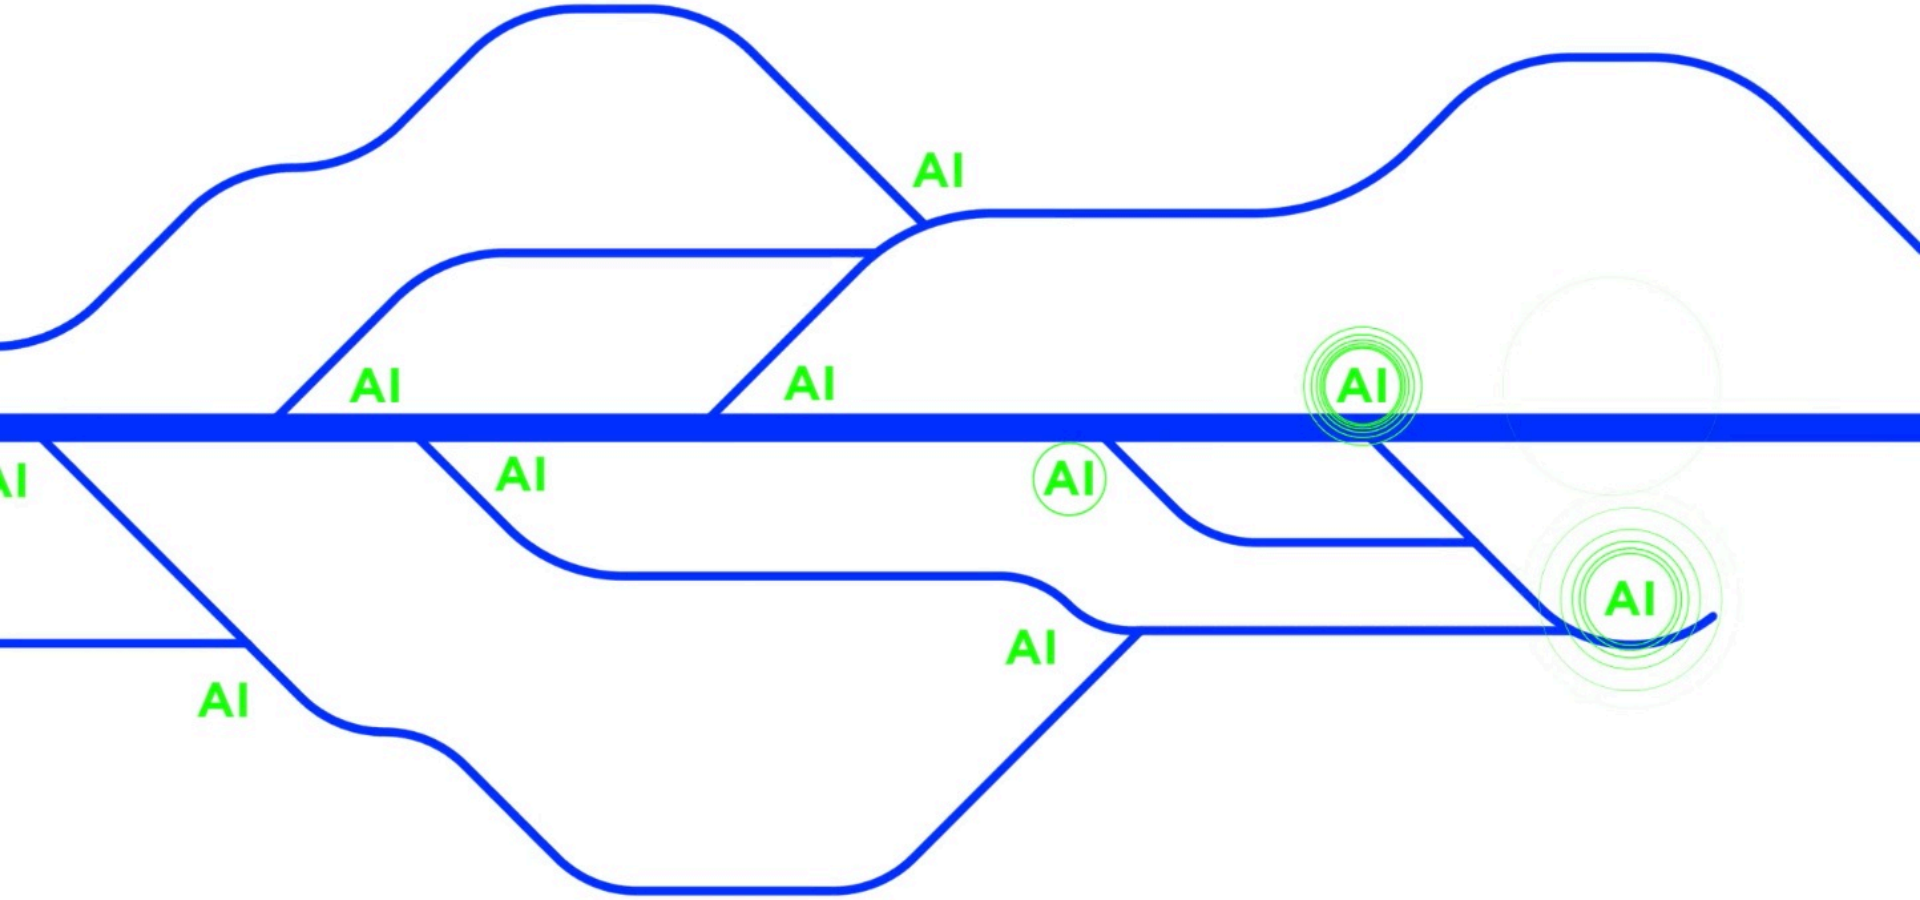

https://futureofstorytelling.org/video/andrew-stern-the-ground-rules-for-ai-storytelling

- How might the Exploitation/Exploration tradeoff be used in interactive AI storytelling?
- How can AI's help manage the tradeoff?
- How might reinforcement learning be used by an experience manager?

Interactive Narrative

- [RB13] "Interactive Narrative: An Intelligent Systems Approach" by Mark Owen Riedl, Vadim Bulitko in *AI Magazine*, Vol. 34, No. 1, 2013
  https://www.aaai.org/ojs/index.php/aimagazine/article/view/2449

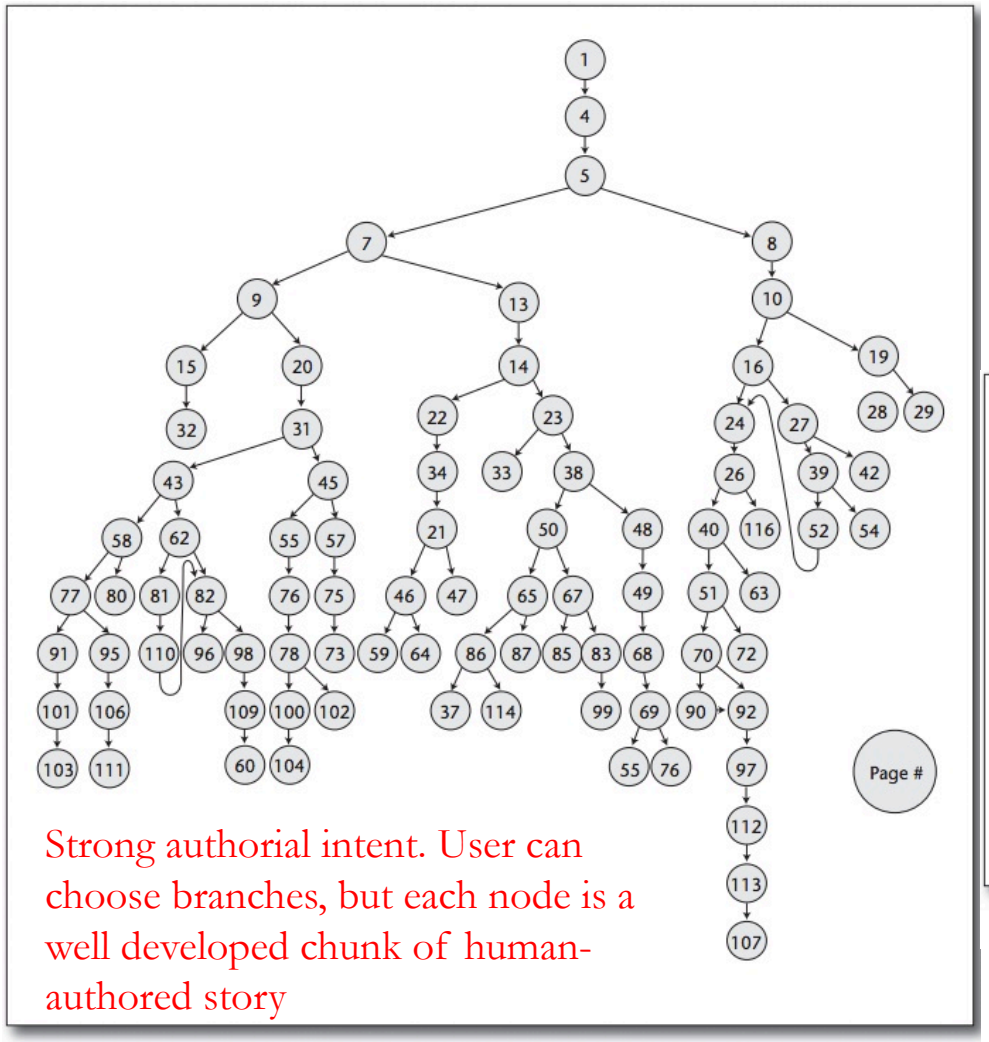

Strong authorial intent. User can choose branches, but each node is a well developed chunk of human-authored story

Figure 3. A Branching Story Graph from The Abominable Snowman.

Greater user options increase the space of possible states beyond what can be reasonably human-created in advance but constraints can be used create realistic scenarios

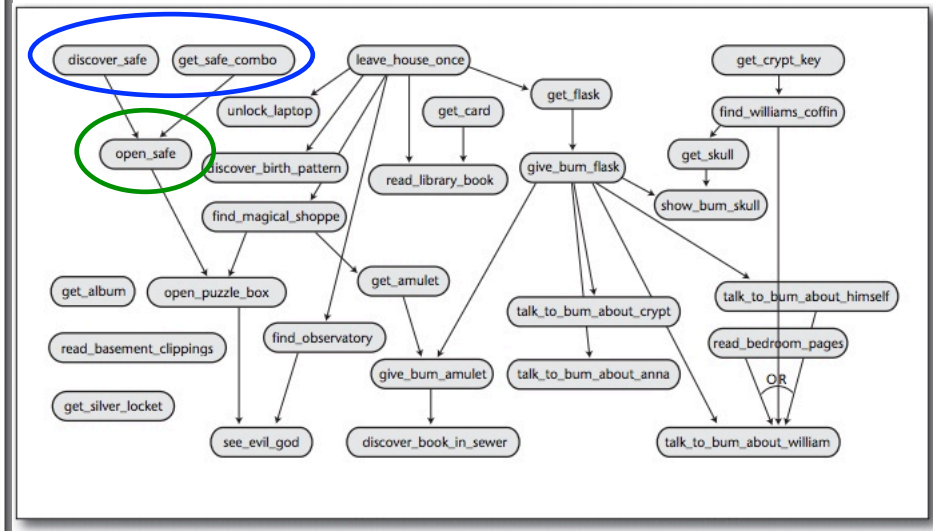

Figure 4. A Portion of the Plot Graph for the Interactive Fiction, Anchorhead (Nelson and Mateas 2005) (reproduced with permission).

"preconditions" must be achieved before selected actions

Architectures for Story Generation

Callaway, Charles B., & James C. Lester (2002). "Narrative prose generation." Artificial Intelligence 139.2: 213–52. http://www.intellimedia.ncsu.edu/wp-content/uploads/npg-ijcai01.pdf

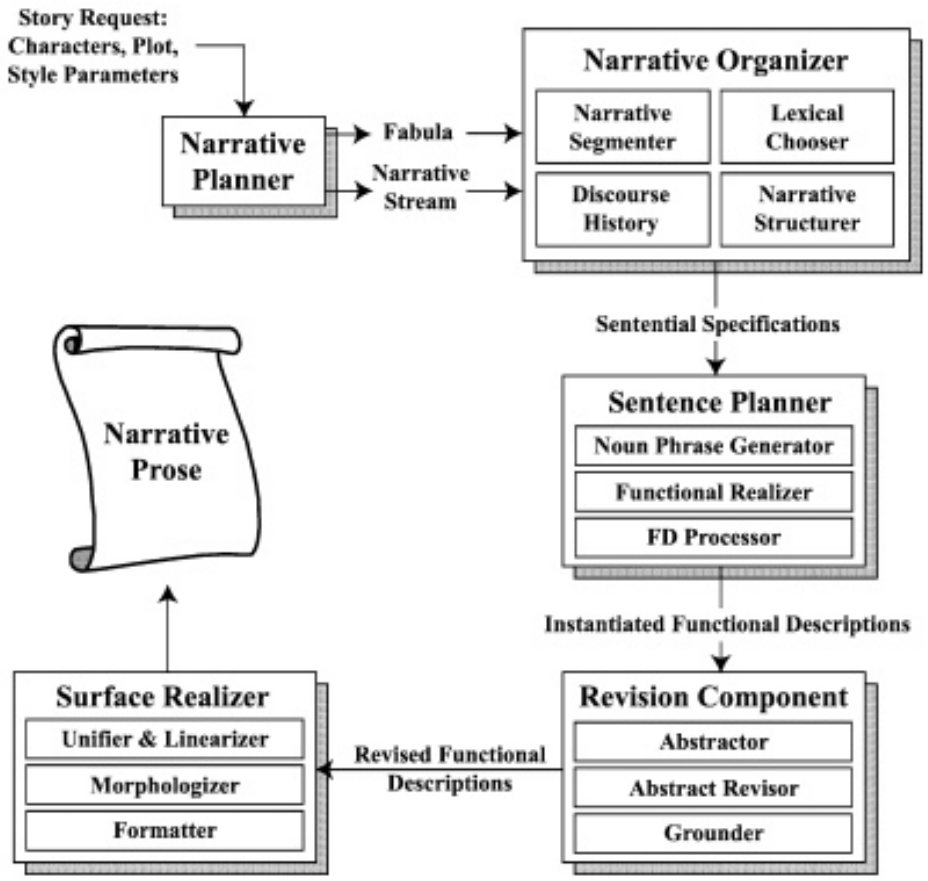

Fig. 7. A narrative prose generation architecture.

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

Fig. 1. Prose generated by TALE-SPIN, 1977.

Once upon a time a woodman and his wife lived in a pretty cottage on the borders of a great forest. They had one little daughter, a sweet child, who was a favorite with every one. She was the joy of her mother's heart. To please her, the good woman made her a little scarlet cloak and hood. She looked so pretty in it that everybody called her Little Red Riding Hood.

Fig. 2. Prose generated by AUTHOR, 2001.

Once upon a time there lived in a pretty cottage, on the borders of a great forest, a woodman and his wife who had one little daughter, a sweet child, and a favorite with every one. She was the joy of her mother's heart, and to please her, the good woman made her a little scarlet cloak and hood, in which she looked so pretty, that everybody called her Little Red Riding-Hood.

Fig. 3. Prose from *Little Red Riding Hood* [61].

How to create "back stories" (Fabula)

Hullett, K., Mateas, M. (2009), GamesScenario generation for emergency rescue training games, Proceedings of the 4th International Conference on Foundations of Digital Games http://dl.acm.org/citation.cfm?id=1536538 , 2007.

The scenario generation paper is concerned with creating simulated physical scenarios (buildings) in which a user can practice rescue skills (after earthquakes, etc). That is, the scenario will support an interactive narrative, in which a human user is making autonomous choices, but the experience manager is making others (e.g., how the building deforms, the fire spreads, etc). INTERNAL CONSISTENCY is an important principle that is highlighted in the paper -- the scenario should be generated so as to avoid inconsistent conditions (e.g., a fire should NOT erupt in a wall that does not include a fire source).

Another interesting aspect of the Scenario generation paper is that it uses an AI planner to "plan" the simulated building and simulated catastrophic event in such a way that user training goals will be satisfied. We have previously talked about AI planners being used to generate the stories (and by experience managers in interactive narratives) to create a story. The common mechanism underlying both context (for story) creation and story creation itself speaks to the fuzzy boundary between story and its context. This in turn suggested a possible research question: If someone asks "Doug, what's your story?", where do I start, and what do I assume as the "jumping off points" of the story? This is a non-trivial question!

The scenario generation paper addressed the creation of a (simulated) physical context. We asked whether the same mechanisms could be used to create social context (e.g., just as there are rules used to talk about how a catastrophic event deforms/collapses a building, can we define rules for how catastrophic rules deform (or reform) persons, psyches, personalities, relationships?