

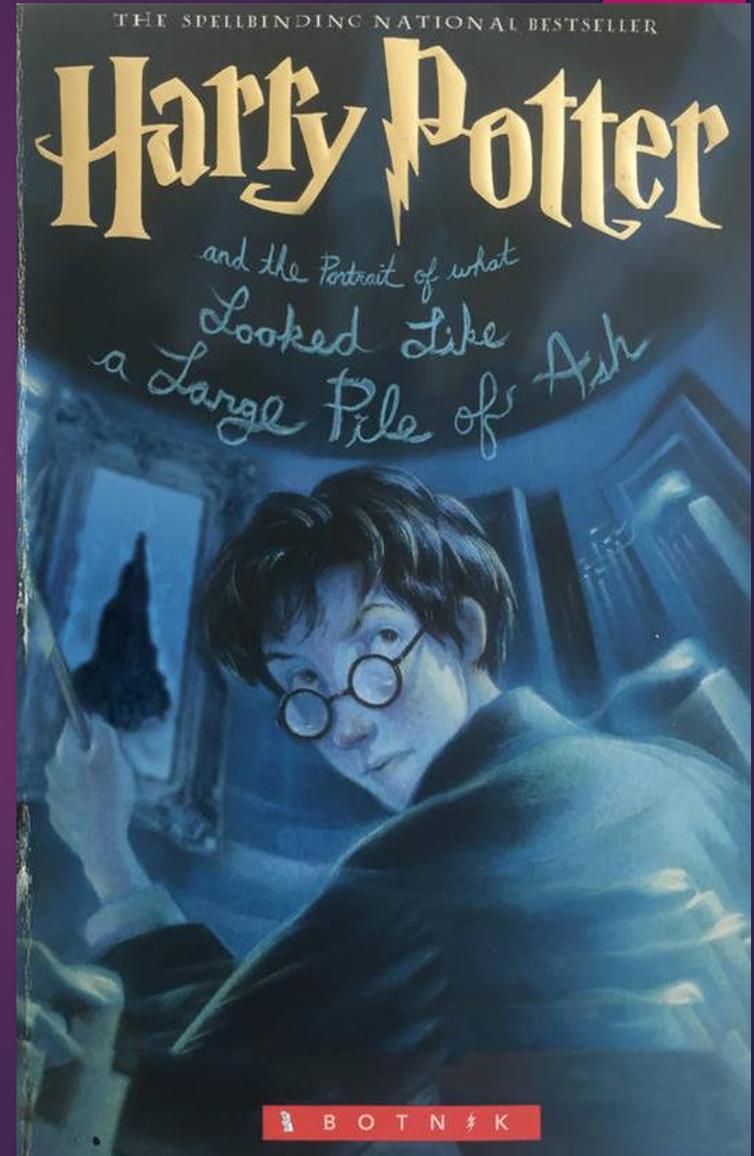
# Dungeons and DQNs

Toward Reinforcement Learning  
Agents that Play Tabletop Roleplaying  
Games

LARA J. MARTIN, SRIJAN SOOD, MARK O. RIEDL

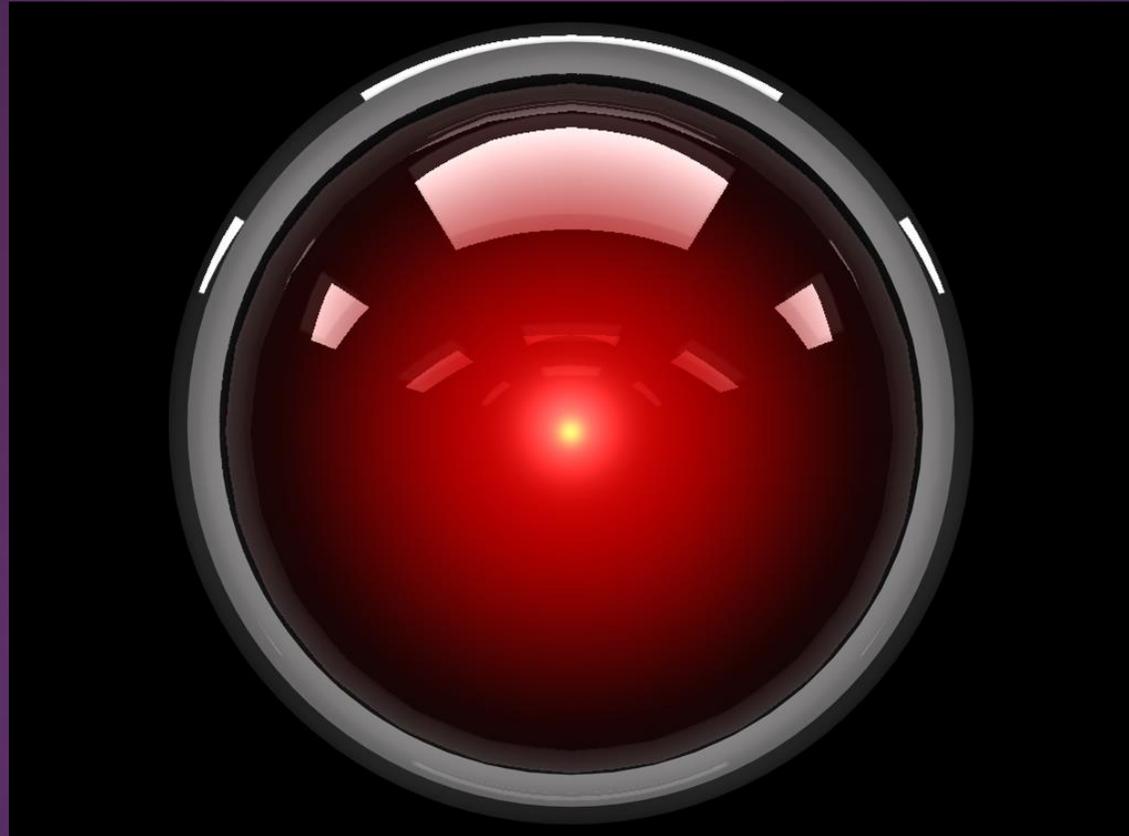


It's an exciting time for AI



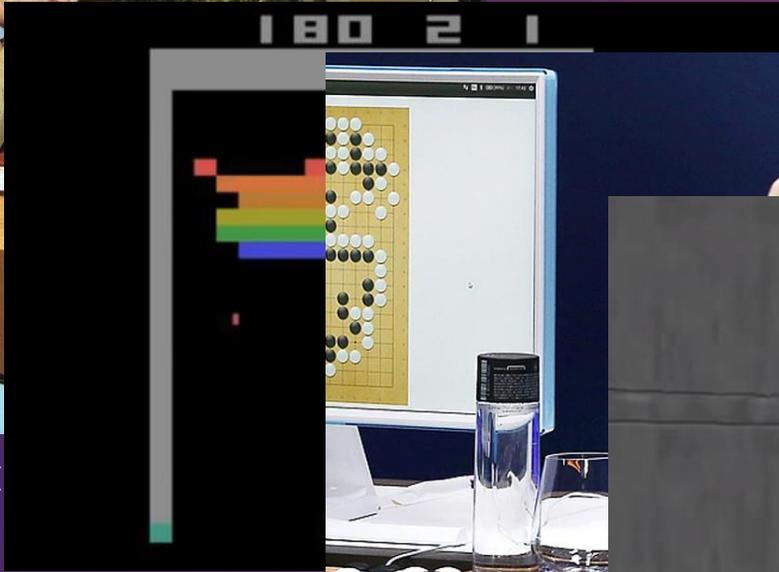
# How do we push the limits of AI?

3

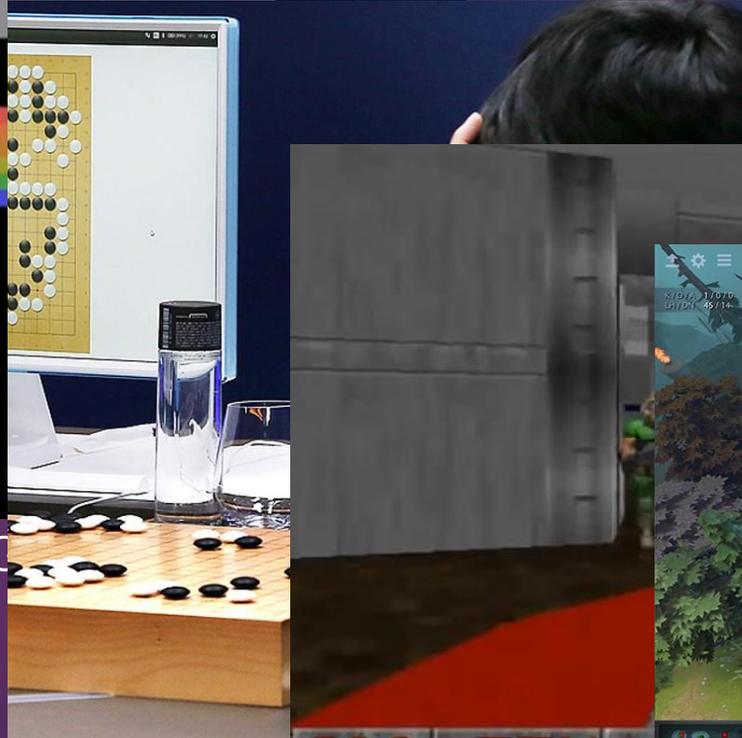


I'm sorry, Dave. I'm afraid I can't do that.

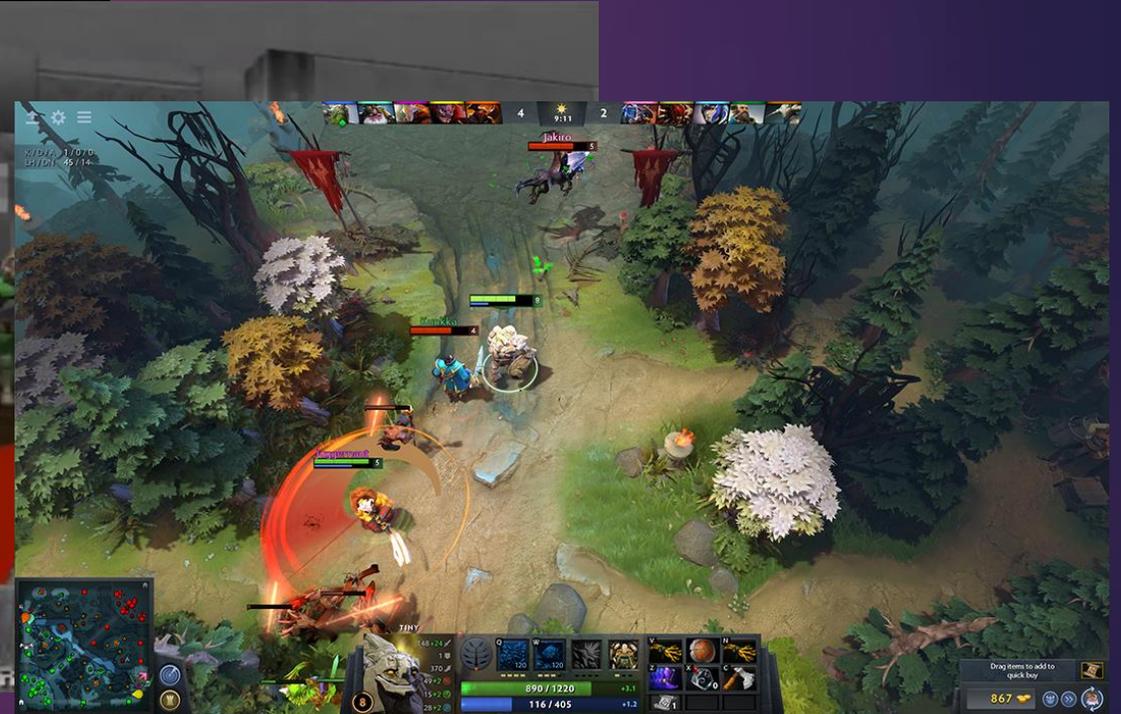
# Games!



Atari C



100 AMMO 55% HEALTH FR



Doom - 2016

DOTA 2 - 2018

# What about Dungeons & Dragons?

5

- ▶ Players create characters to play & describe their character's actions
- ▶ Characters exist in a shared imaginary world
- ▶ Game/Dungeon Master (GM/DM) mediates and sets up scenarios—or campaigns

# Why Dungeons and Dragons?

- ▶ Unlimited actions (discourse)
- ▶ Actions can have unexpected consequences and/or DM can get unexpected player actions
- ▶ Actions cannot cleanly map to states (model of the world changes as game progresses)
- ▶ Distributed game world (across players and DM)
- ▶ Players receive intrinsic reward for actions (unclear win condition)
- ▶ Collaborative

# Outline

- ▶ TRPGs compared to:
  - ▶ Interactive Fiction
  - ▶ Experience Management
  - ▶ Automated Story Generation
- ▶ Our starting point:
  - ▶ Genre Expectation Model + Commonsense Rules Model
  - ▶ Deep Q-Learning

# TRPGs vs the World

Medium	Comparison to TRPGs
Interactive Fiction (IF) Playing	<ul style="list-style-type: none"><li>• Use puzzles to uncovers pre-existing story</li><li>• Often simplified grammar</li></ul>
Experience Management (Used in Interactive Narrative)	<ul style="list-style-type: none"><li>• Intervenes in storyline to keep things “on track” for quality</li><li>• Often fixed set of actions</li></ul>
Automatic Story Generation	<ul style="list-style-type: none"><li>• Generates new story</li><li>• Uses planners to create actions for characters for well-defined domains</li></ul>

# Outline

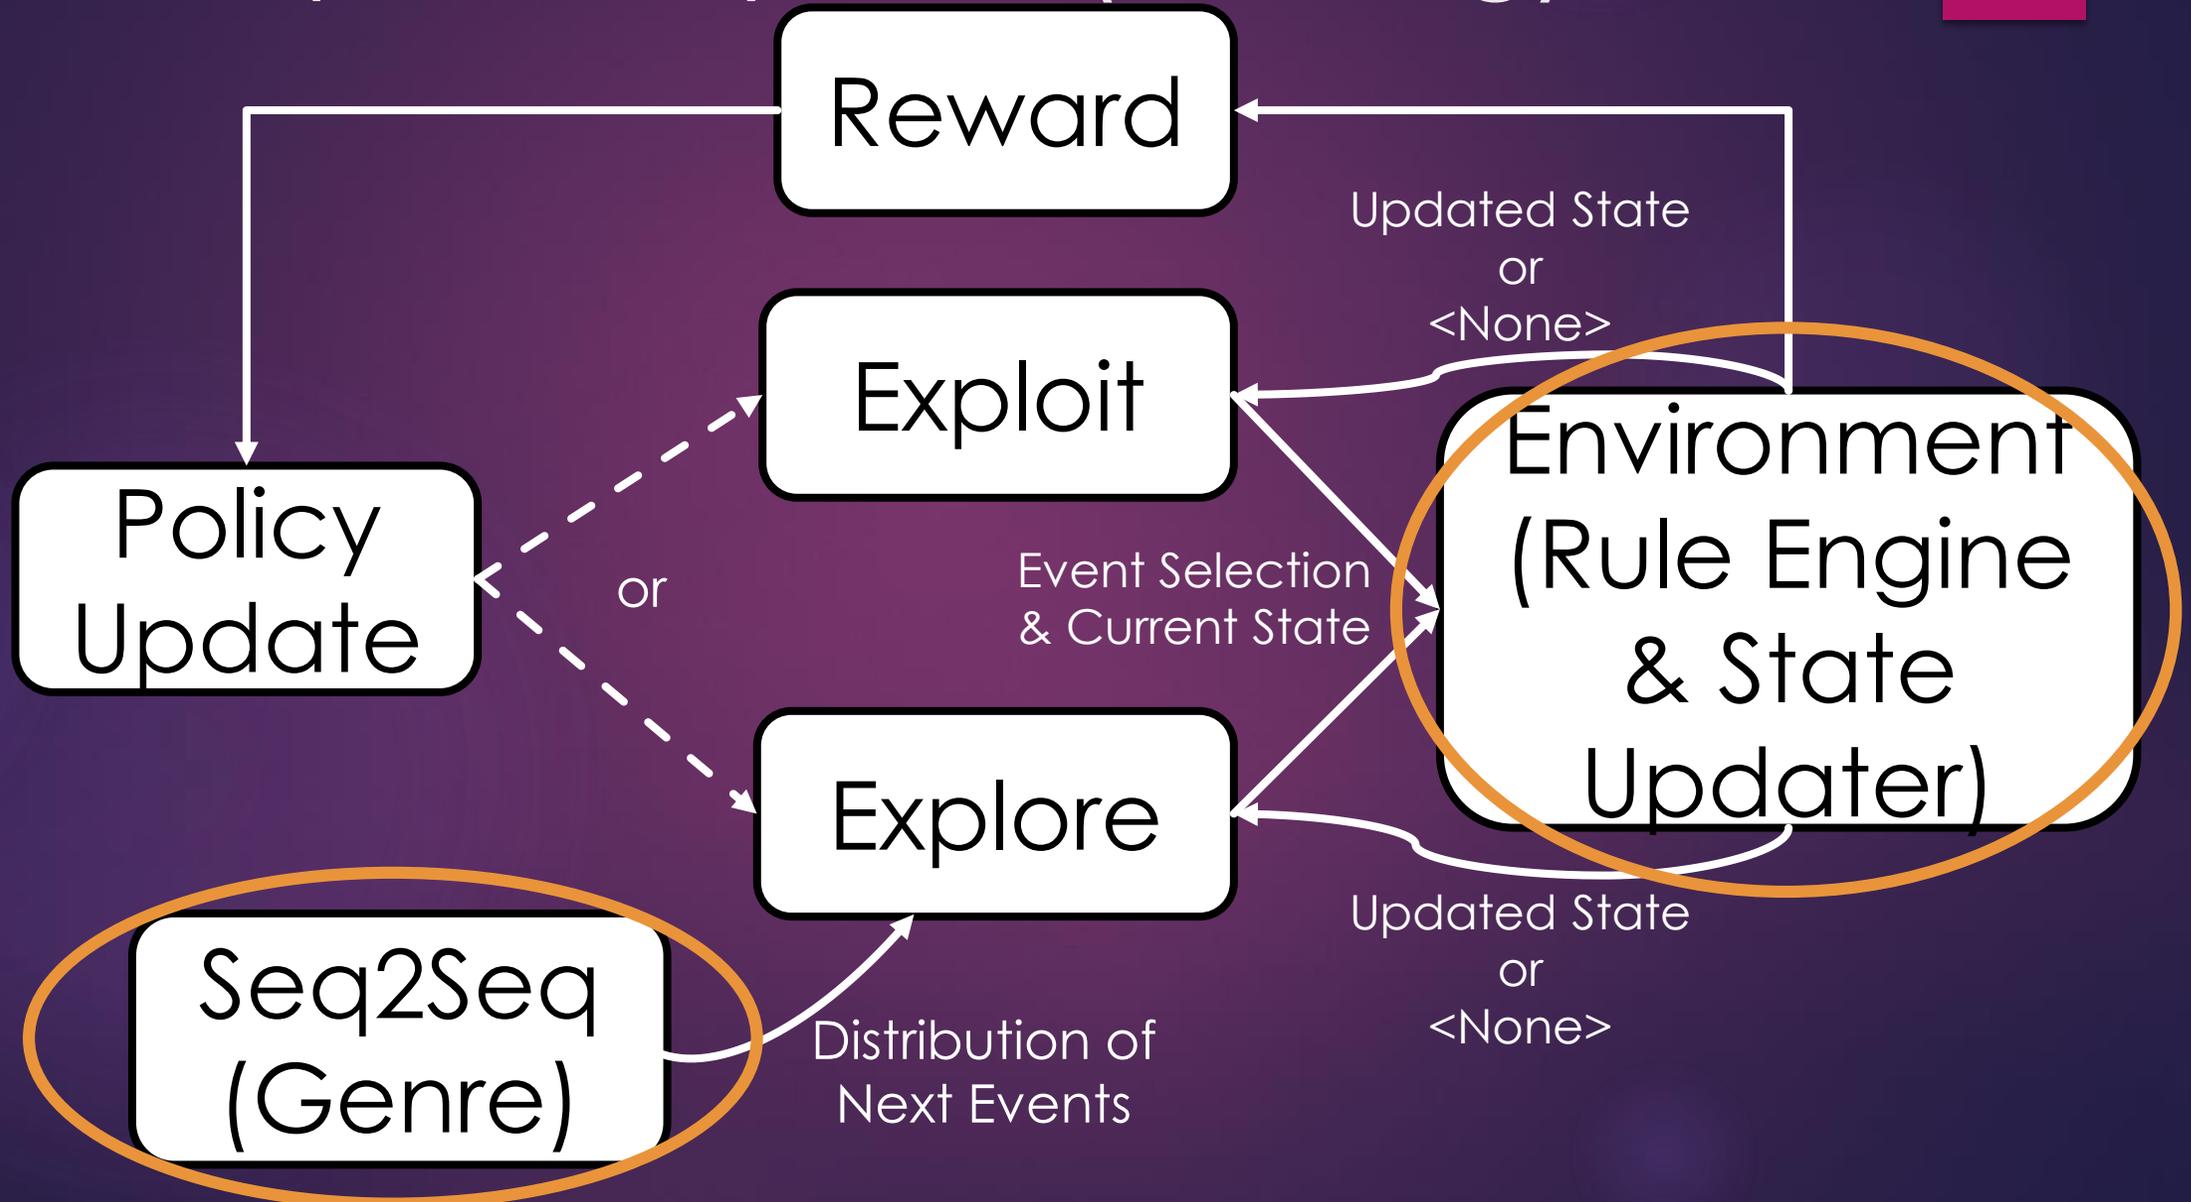
- ▶ TRPGs compared to:
  - ▶ Interactive Fiction
  - ▶ Experience Management
  - ▶ Automated Story Generation
- ▶ Our starting point:
  - ▶ Genre Expectation Model + Commonsense Rules Model
  - ▶ Deep Q-Learning

# Assumptions

- ▶ No dice rolling (i.e. no combat, etc.)
- ▶ Agent is always in character
- ▶ GMs aren't refereeing

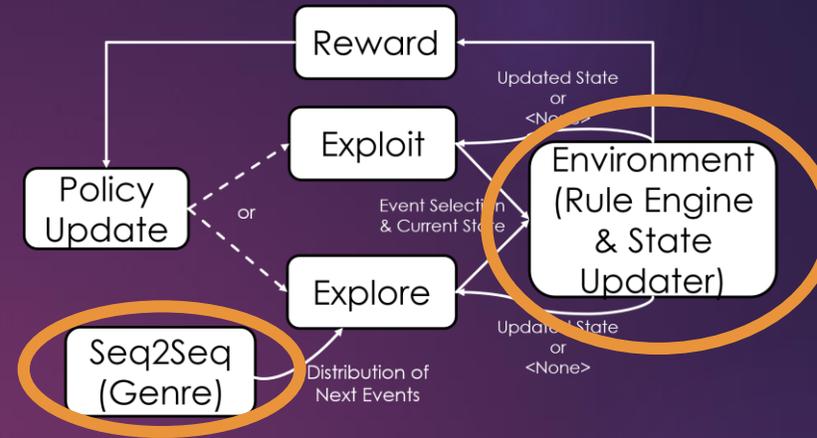
# The Proposed System (Training)

11



# World Model

12



## 1. Genre Expectation Model

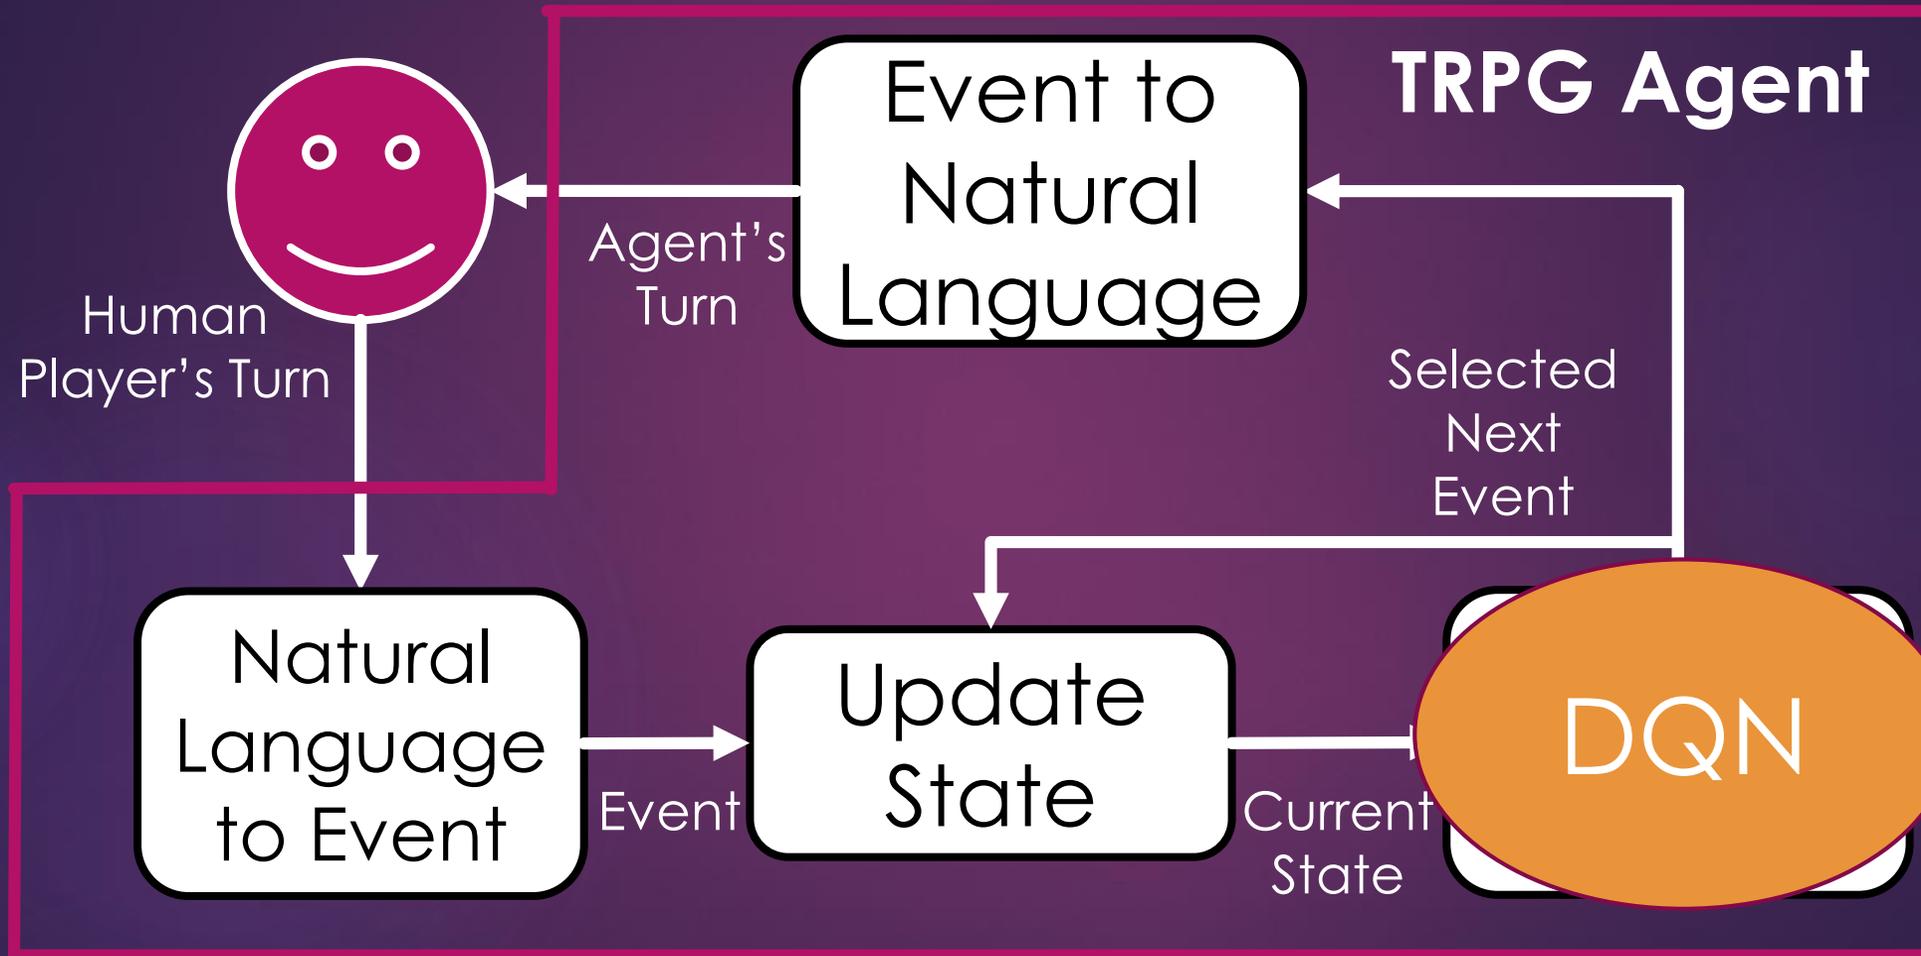
- ▶ Seq2Seq network generates next event in the story
- ▶ Trained on relevant genre

## 2. Commonsense Rules Model

- ▶ Things that aren't mentioned in stories (see: Principle of Minimal Departure)
- ▶ Temporal & physical rules

# The Proposed System Pipeline

13



# Back to Games!



Atari Games - 2015



Go - 2016



Doom - 2016



DOTA 2 - 2018

# Conclusion

15

- ▶ TRPGs are the next AlphaGo

# Thank you!

LARA MARTIN

LJMARTIN@GATECH.EDU