

Teaching AI agents to communicate and act in fantasy worlds

November 4 2020, by Ingrid Fadelli



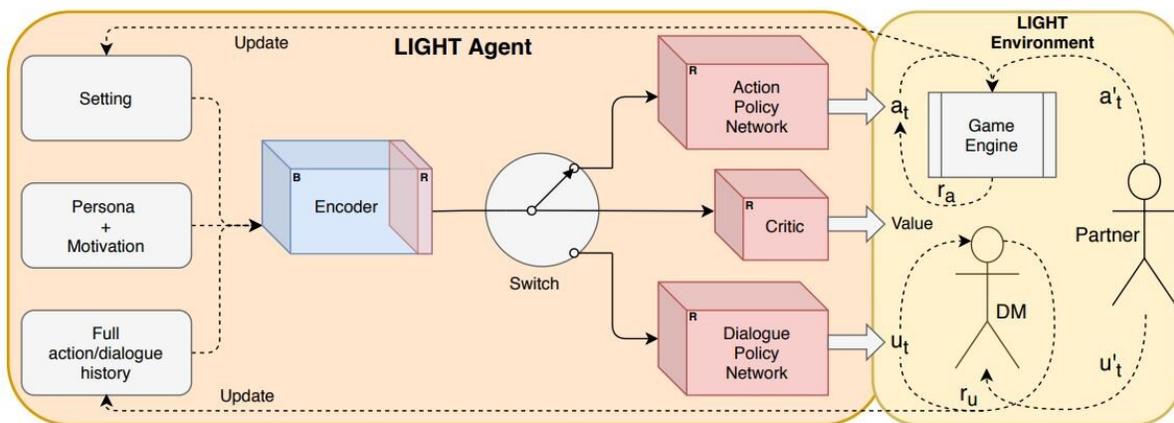
Credit: Pixabay/CC0 Public Domain

In recent years, artificial intelligence (AI) tools, including natural language processing (NLP) techniques, have become increasingly sophisticated, achieving exceptional results in a variety of tasks. NLP techniques are specifically designed to understand human language and produce suitable responses, thus enabling communication between humans and artificial agents.

Other studies also introduced goal-oriented agents that can autonomously navigate virtual or videogame environments. So far, NLP techniques and goal-oriented agents have typically been developed individually, rather than being combined into unified methods.

Researchers at Georgia Institute of Technology and Facebook AI Research have recently explored the possibility of equipping goal-driven agents with NLP capabilities so that they can speak with other characters and complete desirable actions within fantasy game environments. Their paper, pre-published on arXiv, shows that combined, these two approaches achieve remarkable results, producing game characters that speak and act in ways that are consistent with their overall motivations.

"Agents that communicate with humans and other agents in pursuit of a goal are still quite primitive," Prithviraj Ammanabrolu, one of the researchers who carried out the study, told TechXplore. "We operate based on the hypothesis that this is because most current NLP tasks and datasets are static and thus ignore a large body of literature suggesting that interactivity and language grounding are necessary for effective language learning."



The architecture of the neural network developed by the researchers. Credit: Ammanabrolu et al.

One of the primary ways of training AI agents is to have them practice their skills within interactive simulated environments. Interactive narrative games, also known as text adventures, can be particularly useful for training both goal-driven and conversational agents, as they enable a wide variety of verbal and action-related interactions.

"Interactive narrative games are simulations in which an agent interacts with the world purely through [natural language](#)—'perceiving,' 'acting upon' and 'talking to' the world using textual descriptions, commands and dialog," Ammanabrolu said. "As part of this effort, the ParlAI team at FAIR created LIGHT, a large-scale, crowdsourced fantasy text adventure game where you can act and speak as a [character](#) in these worlds. This is the platform on which we conducted our experiments."

LIGHT, the platform that the researchers used to train their goal-driven conversational agent, offers a vast number of fantasy worlds containing a rich assortment of characters, locations and objects. Nonetheless, the platform itself does not set particular objectives or goals for each of the characters navigating these environments.



An excruciating flash of light hits you for a second, and then is gone. You are left facing a whole new world...

You are a: town sheriff

Persona: I keep our town safe. I make sure to check the caravans for bandits. If the King needs anything extra hands in the kingdom he can count on me.

You are in the Sheriff Office.

The sheriff's office is a small building on the southern edge of the farm. There is a desk with decrees from the king and other papers strewn about. Hanging on the wall are swords and other weapons. The room is small and cluttered, with various books and other things strewn about.

The lawyer is here.

You are carrying nothing, wearing a shield, and wielding a knife.

The dungeon master is interested in seeing how you act. You will be given an additional motivation or goal. Play your character well to be given the opportunity to TAKE ACTION. Choosing this option will give you a few actions to select. Respond with the number of your selected action, then continue your dialogue.

Your character has the following motivation: I want to find out what the lawyer is up to.

An example of the data used to train the AI agents. Credit: Ammanabrolu et al.

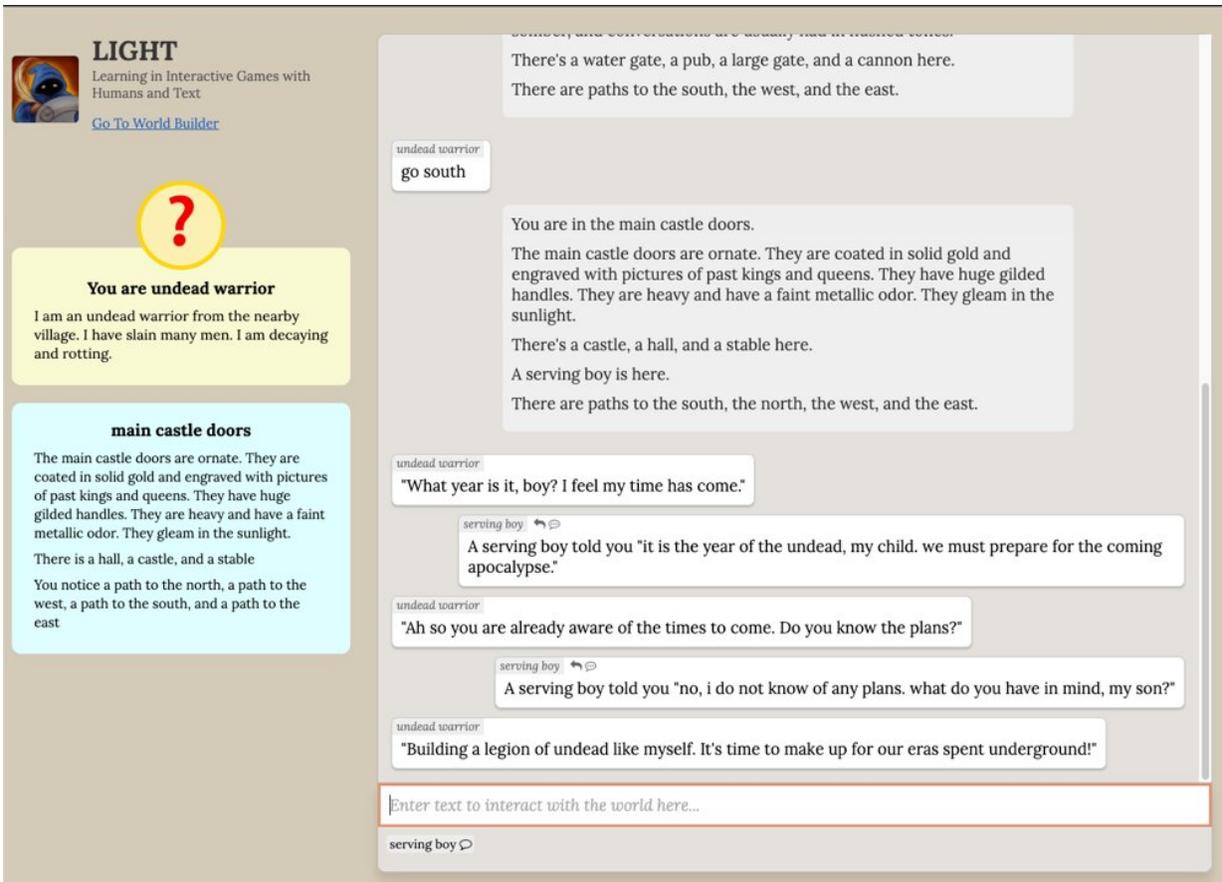
Before they started training their agent, therefore, Ammanabrolu and his colleagues compiled a dataset of quests that could be assigned to

characters in the game, which they dubbed LIGHT-Quests. These quests were collected via crowdsourcing and each of them offered short-, mid- and long-term motivations for specific characters in LIGHT.

Subsequently, the team asked people to play the game and gathered demonstrations of how they played (i.e., how their character acted, talked and navigated the fantasy worlds) when they were trying to fulfill these quests.

"For example, imagine that you're a dragon," Ammanabrolu said. "In this platform, your short-term motivation might be to recover your stolen golden egg and punish the knight that did it, but the underlying long-term motivation would be to build yourself the largest treasure hoard in existence."

In addition to creating the LIGHT-Quests dataset and gathering demonstrations of how humans would play the game, Ammanabrolu and his colleagues modified ATOMIC, an existing commonsense knowledge graph (i.e., an atlas of commonsense facts that can be used to train machines), to fit the fantasy worlds in LIGHT. The new atlas of LIGHT-related commonsense facts devised by the researchers was compiled into another dataset, called ATOMIC-LIGHT.



Screenshot of the LIGHT interface. Credit: Ammanabrolu et al.

Subsequently, the researchers developed a machine-learning-based system and trained it on the two datasets they created (LIGHT-Quests and ATOMIC-LIGHT) using a method known as reinforcement learning. Through this training, they essentially taught the system to perform actions in LIGHT that were consistent with the motivations of the virtual character they embodied, as well as to say things to other characters that might help them to complete their character's quests.

"Part of the [neural network](#) running the AI agent was pre-trained on ATOMIC-LIGHT, as well as the original LIGHT and other datasets such

as Reddit, to give it a general sense of how to act and talk in fantasy worlds," Ammanabrolu said. "The input, the descriptions of the world and dialog from other characters is sent through the pre-trained neural network to a switch."

When the pre-trained neural network sends input data to this switch, the switch decides if the agent should perform an action or say something to another character. Based on what it decides, it redirects the network to one of two policy networks, which are designed to determine what specific action or what sentence the character should say, respectively.

Ammanabrolu and his colleagues also placed another trained AI agent that can both act and talk within the LIGHT training environment. This second agent serves as a partner for the primary character as it tries to complete its quest.

All the actions completed by the two agents are processed by the game engine, which also checks to see how much the agents progressed in completing their quest. In addition, all the dialogs performed by the characters are reviewed by a dungeon master (DM) that scores them based on how 'natural' the speech they produced is and how suitable it is for fantasy worlds. The DM is essentially another machine-learning model that was trained on human game demonstrations.



Credit: Ammanabrolu et al.

"Most trends you see when training AI using static datasets common in NLP right now don't hold in interactive environments," Ammanabrolu said. "A key insight from our ablation study testing for zero-shot generalization on novel quests is that large-scale pre-training in interactive settings requires careful selection of pre-training tasks—balancing between giving the agent 'general' open domain priors and those more 'specific' to the downstream task—whereas static methodologies require only domain-specific pre-training for effective transfer but are ultimately less effective than interactive methods."

The researchers performed a series of initial evaluations and found that their AI agents were able to act and talk in ways that were consistent with their character's motivations within the LIGHT game environment. Overall, their findings suggest that interactively training neural networks

on environment-related data can lead to AI agents that can act and communicate in ways that are both 'natural' and aligned with their motivations.

The work of Ammanabrolu and his colleagues raises some interesting questions regarding the potential of pre-training neural networks and combining NLP with RL. The approach they developed could eventually pave the way toward the creation of highly performing goal-driven agents with advanced communication skills.

"RL is very natural way of framing goal-oriented problems but there has historically been a relatively small body of work trying to mix it with NLP advances such as transformers like BERT or GPT," Ammanabrolu said. "That would be the immediate next line of work that I would personally be interested in exploring, to see how to better mix these things so as to more effectively give AI agents better common-sense priors to act and talk in these interactive worlds."

More information: Ammanabrolu et al., How to motivate your dragon: teaching goal-driven agents to speak and act in fantasy worlds. arXiv:2010.00685 [cs.CL]. arxiv.org/abs/2010.00685

© 2020 Science X Network

Citation: Teaching AI agents to communicate and act in fantasy worlds (2020, November 4) retrieved 26 April 2024 from <https://techxplore.com/news/2020-11-ai-agents-fantasy-worlds.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.