Model AI Assignments 2021

Todd W. Neller[1], Nathan Sprague[2], John Maraist[3], Lisa Zhang[4], Pouria Fewzee[4], Duri Long[5], Jonathan Moon[5], Brian Magerko[5], Alex Leto[6], Toni Lefton[6], and Tom Williams[6]

[1]Gettysburg College, [2]James Madison University, [3]University of Wisconsin - La Crosse, [4]University of Toronto, [5]Georgia Institute of Technology, [6]Colorado School of Mines

Abstract

The Model AI Assignments session seeks to gather and disseminate the best assignment designs of the Artificial Intelligence (AI) Education community. Recognizing that assignments form the core of student learning experience, we here present abstracts of six AI assignments from the 2021 session that are easily adoptable, playfully engaging, and flexible for a variety of instructor needs. Assignment specifications and supporting resources may be found at http://modelai.gettysburg.edu.

ScalarFlow: Implementing Reverse Mode Automatic Differentiation - Nathan Sprague

Many of the most significant successes in artificial intelligence over the past decade have been based on progress in the area of deep neural networks. That progress has been facilitated by high-quality deep-learning libraries like Theano, Tensorflow and PyTorch. The key algorithm at the heart of all of these libraries is reverse-mode automatic differentiation. Despite its importance, automatic differentiation is rarely covered in undergraduate AI or machine learning courses. Students may learn to use deep-learning packages, but they are often treated as black boxes. The purpose of this assignment is to give students a better understanding of deep learning algorithms and frameworks by providing them the opportunity to program an automatic differentiation engine and to modify machine learning code built on top of that engine.

Rush Hour: Designing and Comparing Heuristics for a Children's Puzzle - John Maraist

Rush Hour is a children's puzzle based on sliding blocks back-and-forth on a grid. One of the blocks represents the family car, which is stuck in traffic: one or more blocks sit in between the family car and an exit at the end of the family car's row of the grid. Solving a puzzle configuration requires moving cars within each one's row or column so that the family car can move to the exit. We used Rush Hour to explore the design of heuristics for A* search. Students constructed different heuristics for this single problem, and applied effective branching factor and other concepts to compare their solutions to each other's and to naive BFS. The provided code for A* and for modeling game gives a low programming expectation, and focuses the assignment on the concepts relating to heuristics.

Text Denoising Autoencoder for News Headlines - Lisa Zhang and Pouria Fewzee

In this assignment, students combine the techniques they learned throughout a deep learning course to build a denoising autoencoder for news headlines, to be able to retrieve similar headlines. The assignment combines students understanding of autoencoders, language modelling via recurrent neural networks (RNN), data augmentation, and working with embeddings. The cumulative nature of this assignment makes it a good choice for a final assignment in an introductory Artificial Intelligence, Machine Learning, or Deep Learning course that covers the prerequisite topics. For a traditional AI course, it is possible to modify the assignment to use various other Information Retrieval techniques on the provided dataset.

"Unplugged" Semantic Networks and Knowledge Representations - Duri Long, Jonathan Moon and Brian Magerko

AI agents store and organize information in their memory using structures known as knowledge representations. One type of knowledge representation is a semantic net-work. Semantic networks are a way of representing relationships between objects and ideas. For example, a net-work might tell a computer the relationship between different animals (e.g. a cat IS-A mammal; a cat HAS whiskers). In this unplugged (i.e. no computer required) assignment, learners can create their own semantic net-works by gluing down printable cards containing concepts and arrows containing relationships. Provided card decks contain concepts related to animals, family, and musical instruments. Blank cards are also provided for networks on custom topics. Learners will be encouraged to reflect on the networks they create and consider the strengths and limitations of the knowledge representation using a provided list of questions that can be used

to foster discussion or as a written activity. This assignment is suitable for young learners (6 and up), but can also serve as a simple, accessible introduction to knowledge representations or knowledge-based AI for learners of all ages. This assignment could be adapted as either a take-home written activity or an in-class group project.

Using Markov Chain Text Generators to Facilitate Found Poetry Creation - Alex Leto, Toni Lefton, and Tom Williams

The combination of AI and artistic or literary techniques provides a unique opportunity for students and practitioners from both communities to simultaneously achieve Learning Outcomes (LOs) from both fields, as well as LOs situated at the intersection of those fields. In this Model AI Assignment, we show students how simple AI techniques can be combined with Found Poetry techniques to create unique poetic works that raise challenging questions regarding the relationship between algorithm and poet, and between technology and society.

Specifically, we present an educational activity, suitable for complete AI novices, in which students are (1) provided with code for a Markov Chain based text generation algorithm, (2) instructed in how to create novel blackout poems using the text generated by this algorithm, and (3) asked to reflect on the creative relationship between themselves and the algorithm. Furthermore, we provide supplementary guidelines for adapting this activity to tailor it to different communities and use cases, allowing it to be used with either coders or poets, and as either an outreach activity or a formal class assignment.

Introducing AI Worksheet Activity - Duri Long, Jonathan Moon, and Brian Magerko

What is artificial intelligence? Where have you used it before? How do you feel about it? How does it work? In this assignment, students will explore questions like these through an interactive worksheet activity. This is intended as a high-level introduction to AI and can be used to kick off an AI-related class, unit, or workshop. This can be completed as a worksheet activity or the worksheet can be used as a guide to lead an in-class activity. The activity requires the use of a printable deck containing cards with examples of AI technologies and possible inputs, algorithms, and outputs for AI devices. The worksheet activity begins by prompting students to consider where they have seen AI before and how they feel about AI. Then, they are asked to look at the examples of AI in the card deck and select cards they have interacted with previously. Finally, students are walked through a high-level explanation of how AI works and are guided by the worksheet to create an imaginary AI using the input, algorithm, and output cards. This assignment is suitable for young learners but can also serve as an introduction for students of all ages. No technology is required.