# Model AI Assignments 2025

**Todd W. Neller[1], Rasika Bhalerao[2], Eun Kyung Ko[3], Vishodana Thamotharan[3], Lisa Zhang[4], Sonya Allin[5], Mahdi Haghifam[2], Michael Pawliuk[4], Rutwa Engineer[4], Florian Shkurti[4], Cunyan Ma[7], Daniella DiPaola[8], Cynthia Breazeal[8], Loreto Alonzi[9], Brian Wright[9], Ali Rivera[9], Kristin Fasiang[10], Duri Long[10], Shruthi Chockkalingam[11], Giulia Toti[11], Evan Shieh[12], Princewill Okoroafor[13], Thema Monroe-White[14], Mustafa Haiderbhai[4], Carolyn Quinlan[4], Ashwin R. Bharadwaj[2], Anio Zhang[2], Rajagopal Venkatesaramani[2], Sarah Wharton[8], John Masla[8], Lydia Guterman, Mary Cate Gustafson-Quiett[8], Christina Bosch[8], Samar Abu Hegley[8], Calvin Macatantan[8], Eric Klopfer[8], Hal Abelson[8], Shira Wein[15], Mercy Wairimu Gachoka[16], Li-Hsin Chang[17], Maryam Mirzaei[18], and Mohammad Mahdi Ajallooeian[18]**

[1]Gettysburg College, [2]Northeastern University, [3]National Louis University, [4]University of Toronto, [5]York University, [6]Northeastern University, [7]Brown University, [8]Massachusetts Institute of Technology, [9]University of Virginia, [10]Northwestern University, [11]University of British Columbia, [12]Young Data Scientists League, [13]Cornell University, [14]George Mason University, [15]Amherst College, [16]University of Eastern Finland, [17]University of Turku, [18]NorQuest College.

**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 thirteen AI assignments from the 2025 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.

## Spelling Fixer - Rasika Bhalerao

This is a series of two assignments to build an algorithm to fix spelling errors in user-inputted text.

1. Part 1: Using a Hidden Markov Model (HMM) and the Viterbi algorithm to "correct" words which are not in the dictionary. This part includes instructions to calculate emission and transition probabilities and use them in the Viterbi algorithm. (The HMM "emissions" are the typed letters, and the "states" are the ground-truth characters.)
2. Part 2: Improving the algorithm by using BERT to take context into account. In this assignment, students choose how to design the final algorithm which can take into account context (via the language model), probabilities from the Hidden Markov model, or text edit distance (via Levenshtein distance). There is ample PyTorch starter code to help students with BERT.

Students also answer questions to analyze various types of errors in their algorithm's output and the choice of training dataset.



## Designing Culturally Responsive Lessons with AI: Enhancing Engagement and Inclusion - Eun Kyung Ko and Vishodana Thamotharan

This assignment engages teacher candidates in the integration of artificial intelligence (AI) tools to design culturally responsive teaching (CRT) lessons. The objective is to explore how AI can support the development of inclusive and engaging lesson plans tailored to diverse student populations. The assignment begins with a series of AI reflective journals aimed at building foundational knowledge of AI's capabilities and limitations in the classroom. Through these reflective entries, candidates critically examine how AI tools can be utilized to foster culturally responsive learning environments and consider their potential impact on teaching practices.

Following the reflective journal phase, candidates will apply their newfound AI understanding to develop a CRT lesson plan. Using AI tools, such as adaptive learning platforms, language translation applications, or content generators, candidates are tasked with creating lessons that address the unique cultural and linguistic needs of their students. The lesson plans must demonstrate thoughtful integration of AI tools that enhance differentiation, engagement, and accessibility while promoting an inclusive classroom environment.

The assignment culminates in a reflection on how AI supported the development of culturally responsive instructional strategies, highlighting both the opportunities and challenges encountered. This process prepares future educators to thoughtfully integrate technology in ways that support diverse learners and promote equity in education.

## Differential Privacy with MedMNIST - Lisa Zhang, Sonya Allin, Mahdi Haghifam, Michael Pawliuk, Rutwa Engineer, and Florian Shkurti

This sequence of two advanced deep learning assignments introduces learners to issues surrounding differential privacy and model memorization of training data. The assignments, which are implemented in Jupyter Notebooks, guide learners in implementing an advanced optimization method (DP-SGD, or Differentially Private Stochastic Gradient Descent) in PyTorch.

In the first assignment, learners build a Multi-Layer Perceptron (MLP) model that performs predictions on MedMNIST's PneumoniaMNIST data set. This assignment introduces students to MLPs, and issues surrounding building and evaluating neural networks.

In the second assignment, learners analyze the model's predictions on the training vs test data and notice that the model's prediction pattern differs between the two, making it possible for end users to identify the training data. The assignment then provides scaffolding for learners to implement DP-SGD. Finally, learners compare the results of models trained using SGD vs DP-SGD.

Emphasis is placed on analyzing and understanding emergent issues that can arise with neural network training that are of active interest in the research community, including where it is possible for a model's training data to be reconstructed by advanced adversaries, and where such data leaks could produce harm. Emphasis is also placed on developing skills to implement novel algorithms from descriptions written in machine learning papers.

## AI & Employment Lesson Plan - Cunyan Ma, Daniella DiPaola, and Cynthia Breazeal

AI is an evolving technology with immense social impact. As narratives of AI outperforming and replacing humans put conventional career paths at risk, many young people perceive AI's impact as anxiety about future employment prospects. Notably, part of the anxiety is manufactured by an education that teaches AI as a technology with a de facto impact rather than a socio-technical system that can change and be changed by societal factors. This lesson plan targets young people's understanding of future employment by addressing the current shortcomings in AI education. In this 60-minute lesson plan, students use the Writers' Guild of America Strike in 2023 as the primary case study to explore four ways in which AI can impact future employment. Students are then encouraged as active agents by adopting the roles of workers or decision-makers in drafting AI-related policies. As this lesson is less technical and more reflective, it is suitable for students of all subjects in high school and teachers without an AI background, lowering the barriers to learning about AI.

## Mastermind Has Entropy - Loreto Alonzi, Brian Wright, and Ali Rivera

Nothing motivates a learner like the desire to win. Our assignment design philosophy is to introduce a game that the students want to win and thus trick them into learning how to use AI techniques for victory. This assignment's game is the timeless classic Mastermind. The learning objective is quantifying the amount of missing information by applying the concept of Entropy. This game can be easily purchased and played by students in physical form. The introduction of the game physically reduces the cognitive load on the students as they grapple with the tricky concept of Entropy and how to apply it. As the students search for the correct code, they compute the Entropy and gain an understanding of missing information. As they learn the concept they start winning. A suggested reinforcement assignment would be coding up the game and exploring it at scale. The popularity of the game also has a few beneficial features including many online tutorials and math explainer videos that students may use to enhance their experience. Also of note is a 1976 paper published in the Journal of Recreational Mathematics by Donald Knuth called "The Computer As Master Mind".

## Large Language MadLibs - Kristin Fasiang and Duri Long

Large Language MadLibs is an unplugged, two-day lesson sequence introducing high school students to the mathematical concepts and ethical issues behind Large Language Models (LLMs). In the first lesson, students learn about independent probability and why LLMs need to learn from patterns in data. To do this, they take on the role of a LLM that only knows one distribution of word probabilities, and roll dice to randomly choose words from that distribution to fill in blanks independently of context in a "MadLib"-style story. In the second lesson, students learn about conditional probability and how context-dependent word prediction improves story quality but may introduce biases. Students use a similar dice-rolling and coin-flipping mechanic to generate a set of 5-word sentences, though this time they encounter different word distributions based on their previous rolls or flips. Students reflect on how gender bias is encoded in the word distributions. As an unplugged activity, this lesson is designed to engage students with limited computer science backgrounds and encourage more critical and reflective engagement with common LLM tools by making abstract concepts tangible. It is also aligned with high school Common Core State Standards for Math related to probability and statistics.

## Gender Bias in Word Embeddings - Shruthi Chockkalingam and Giulia Toti

In this assignment, students gain first-hand experience with gender bias in word embeddings. By using pre-trained word embeddings (GloVe), they learn how to evaluate and quantify the presence of gender bias. Additionally, students are introduced to recent debiasing techniques and how to assess the extent to which bias persists after debiasing. De-

spite covering advanced topics, the assignment is accessible to undergraduate students with a basic understanding of machine learning concepts, such as training, classification, and clustering, and does not require prior experience in Natural Language Processing. It is currently part of a course focused on equitable and responsible use of data science techniques, which is popular among students from various majors who possess basic machine learning knowledge. The assignment is implemented as a Jupyter Notebook using Python. A suggested grading rubric is provided, and the solution is available upon request for interested instructors, in order to maintain the integrity of the exercise. As the demand for teaching the responsible use of AI applications continues to grow, we hope this assignment will help instructors raise awareness of gender bias in language-related AI systems.

## "Guess My Passcode": A K-12 Game-Based Introduction to AI and Backpropagation - Evan Shieh, Princewill Okoroafor, and Thema Monroe-White

"Guess My Passcode" is a social math game tested for grades 6 and above that demystifies what is commonly referred to as "artificial intelligence". Learning sciences research on the role of language in scientific knowledge acquisition shows that highly technical terms like those that are frequently used in Intro to AI courses (e.g. "machine learning", "AI", or "backpropagation") may widen learning gaps by presenting additional cultural barriers for diverse students. However, "AI" is far from inscrutable, and contrary to common knowledge it is possible to gain a deep understanding of how AI algorithms "learn" from data using mathematics commonly taught prior to middle school. Indeed, perhaps the most central algorithm in modern AI - backpropagation - is often implemented using only the four arithmetic operations. In "Guess My Passcode", students are challenged to learn the intuition behind backpropagation through a mathematical guessing game that can be played collaboratively or individually. The game uses a low-floor, high-ceiling approach to challenge learners ranging from middle school to the undergraduate level. No prior programming experience is required, although supplementary activities are included for computer science classrooms. Participants finish with a concrete understanding of how computer algorithms learn from feedback using elementary mathematics.

## Model Building and Risk Analysis with Health Survey Data - Sonya Allin, Lisa Zhang, Mustafa Haiderbhai, Carolyn Quinlan, Rutwa Engineer, and Michael Pawliuk

In this sequence of three programming labs, students build and analyze machine learning models to predict the presence of heart disease using the NHANES survey responses. Students build a decision tree, a logistic regression, and a neural network model. More importantly, these labs integrate discussions surrounding model evaluation, group-level fairness analysis, and bias-variance decomposition.

The first lab scaffolds the exploratory data analysis process to develop an intuition of the data distribution and data limitations. Learners then build decision tree classifiers and perform a grid search for hyperparameters. Learners choose a final model and report the test accuracy, but reflect on the limitations of accuracy as a measurement.

The second lab asks learners to implement stochastic gradient descent and train a logistic regression model. Learners analyze the errors made by the model, with consideration for error types and their differences between sensitive subgroups (e.g. men vs. women). The lab explores the use of different thresholds for men and women to predict heart disease and asks learners to critically consider the safety of this approach.

The final lab begins with a demonstration of bias-variance decomposition using synthetic data. Learners return to using the NHANES data set and empirically explore sources of error. The lab introduces model averaging techniques, then asks learners to discuss the impact of data on sources of error and risk.

## Shapeshifting Coloring Problems - Ashwin R. Bharadwaj, Anio Zhang, and Rajagopal Venkatesaramani

In this Model AI Assignment, we use coloring problems to teach students about search techniques. Specifically, we randomly initialize a grid-environment with some cells colored in and ask students to color the remainder of the grid such that no two cells sharing an edge share the same color. We provide a set of colors and list of brushes to choose from – each of which takes on one of nine shapes and spans up to four cells on the grid. The objective is to also minimize the number of brush strokes and colors used. We model the environment as a black box, providing an interface that mirrors real-world applications, such as the popular Gymnasium environments for reinforcement learning. Thus, the assignment naturally lends itself to a range of algorithms from basic searches to reinforcement learning, allowing instructors to control difficulty.

The open-ended nature of the assignment helps students build multiple vital skills:

a) correctly and efficiently modeling a search space,
b) working with a black box model through a limited set of interactions,
c) optimizing code for runtime and memory,
d) testing for edge-cases and generalizability.

This assignment was well-received in an undergraduate AI class at a large R1 university.

## Act Out An LLM - Sarah Wharton, John Masla, Lydia Guterman, Mary Cate Gustafson-Quiett, Christina Bosch, Samar Abu Hegley, Calvin Macatantan, Eric Klopfer, Cynthia Breazeal, and Hal Abelson

In this activity for AI novices, students learn how large language models (LLMs) generate text by acting out the pro-

cess. First, the concept of a token is defined with examples and counterexamples. Then, the class is given a prompt: "Tell me a fun fact about potatoes," along with the first token generated by the LLM: "Potato." Choosing from a small bank of tokens, students vote on which token should be next generated; the most popular choice is added to the LLM's response. This process is repeated until a full sentence is generated. The activity is followed by a discussion about the accuracy of generated text.

## Lecture Material and Worksheets for AI Applications and Societal Impact - CS Core 2023 - Giulia Toti, Shira Wein, Mercy Wairimu Gachoka, and Li-Hsin Chang

In January 2024, a joint Task Force comprising members from ACM, IEEE-CS, and AAAI introduced revised curriculum recommendations for undergraduate computer science education. The updated guidelines emphasize the growing importance of AI in the curriculum, increasing the recommended dedicated hours and adding three core hours on AI Applications and Societal Impact. Drawing on their collective teaching experience, the authors of this submission developed materials to assist CS instructors in covering these recommended topics. The materials include two sets of slides addressing key issues, such as defining fair and equitable AI applications, dataset bias, algorithmic and evaluation bias, environmental considerations, and a high-level overview of large language models. Additionally, two worksheets are provided to facilitate and assess students' understanding, while promoting reflection on these critical subjects. The worksheets are designed to be unplugged, making them accessible to a wide range of students and not requiring any prior AI knowledge beyond what is included in the lectures. To preserve the assignment's integrity, solutions to the worksheets are available to instructors upon request. We believe the material in this submission will help and encourage CS instructors, regardless of their familiarity with AI, to integrate important topics on the responsible use of AI applications into their courses.

## Enhancing Data Science Education through Environmental Impact Prediction and Data Storytelling - Maryam Mirzaei and Mohammad Mahdi Ajallooeian

The "ML-Ready Dataset Preparation for Environmental Impacts Prediction" assignment offers students practical experience in data preparation, machine learning, and data visualization using real-world data from the National Pollutant Release Inventory (NPRI) of Environment and Climate Change Canada. The assignment is structured in phases, guiding students through exploratory data analysis, data cleaning, feature engineering, and machine learning model development to address problems like predicting pollutant releases based on economic and policy changes.

Emphasizing hands-on learning, students detect patterns, resolve data issues, and engage with stakeholders to refine their models, mirroring real-world data science practices. The assignment encourages deep thinking about machine learning problem framing and solution formatting, making it an excellent project for students preparing for industry roles. It requires careful consideration without being overly complex.

An optional task to enhance engagement involves creating interactive dashboards with Plotly Dash that focus on data storytelling. This encourages students to creatively present findings in a narrative format, improving their ability to communicate complex insights to non-technical audiences. Adaptable to various educational contexts, this assignment builds essential data science skills and emphasizes AI's potential to address environmental challenges, particularly climate change, aligning with the EAAI conference's focus on AI for social good.