Note that this curriculum is still being finalized. The activities presented will be included, but the exact order of activities may be rearranged or expanded.

Course Title: Computational Thinking and Machine Learning	
Lesson Date/Day: Day 1	
Daily Objectives	- *Students should have a general understanding of the history of artificial intelligence, including some classical, statistical approaches to machine learning*
Essential/ Debatable Question(s)	- *What properties make an artificial system "intelligent?"* - *Can computers be "intelligent?"*
Concepts	- *What is AI?* - *AI overview and history* - *Turing test* - *"Games" in AI* - *Solved problems/games* - *Minimax*
Materials	- *Powerpoints will be used to cover the initial topics about AI and its history* - *Games will be introduced in a hands-on way by having students act as characters in a simple turn-based game* - *Solved problems will be shown with examples of how they can be played optimally (multimedia presentation)* - *The minimax algorithm will be explained in some short slides.*

Activities	- The first activity will have students get up and act as players in a simple turn-based game (specifics not available as of yet)
	  - Students will think about the rules of the game and act in intelligent ways to try to get higher scores
	  - We will pivot to a metacognitive discussion about their decision processes, which will lead into the AI topic
	- The next activity (after history etc.) will involve students playing Tic-Tac-Toe and/or connect-4 against a custom AI of our creation
	  - This gets into game theory in AI
	  - We will transition to getting students to make a naive AI to play Tic-Tac-Toe and have them compare its performance to the original AI they did not create
	  - We will discuss if/how this could apply to Connect-4 and how larger games become intractable
	  - This will lead into the minimax activity
	- Time permitting, we will attempt to get students to create their own

	implementation of the minimax algorithm for general gameplay optimization. - A starter project will be created and distributed to students to dramatically simplify this process (Parsons problem) - This activity is non-critical and can be left incomplete if time expires; in either case, students will be given a link to a completed version for their own personal study
Assessment	- *Throughout the presentations, we will periodically stop and ask for engagement from students in open discussions* - *During the initial phases of this camp and the occasional ethics components, there are no wrong answers and we are looking specifically to get a large diversity of opinions/ideas in various areas*

Course Title: Computational Thinking and Machine Learning	
Lesson Date/Day: Day 2	
Daily Objectives	- *Students should be able to describe the key steps of deep learning and explain how it works in a general sense (mathematics not required)* - *Students should be able to identify possible sources of bias in machine learning systems and the potential impact on users*
Essential/ Debatable Question(s)	- *What things can or cannot be learned by a machine learning system?* - *How do AI engineers teach computers to perform new tasks?* - *What techniques are used to gather data for training modern machine learning models?* - *What are the ethical ramifications of using individuals' data without permission?*
Concepts	- *Deep learning (basics)* - *Supervised learning* - *Image classification*
Materials	- *The first presentation on the basics of deep learning will involve a powerpoint presentation* - *Projects will be distributed for the interactive activity on supervised learning and function optimization*
Activities	- When teaching the concept of supervised learning, students will be

	given a project that has several clusters of points (blue and red) - students must manually fiddle with parameters to achieve a good fit (visually) - We will then introduce the concept of gradient descent as a means of automatically tuning these parameters. - This will also be a hands on activity, but only on 1-dimensional data - Notably we will not mention any calculus terms like derivatives/gradients and will instead have students repeatedly look at both x+e and x-e and take whichever is better - Image classification will be introduced at a high level in presentations - An interactive activity will have students create their own image classifier for a personally-interesting set of labeled classes - Students will be given 30 mins to gather images online to represent each of their classes - The training portion of their data will be fed into Teachable Machine, which will train on their collected data - Students will then test the model on some of their testing data and see how the model performs - Presumably with a small amount of data the model will perform poorly - they will be given more time to collect more data, retrain the model, retest, and compare the new results - This should also give some insights into the problems faced in real world data collection - To teach some other important concepts, we will also have students intentionally compromise the quality of their labels by introducing some samples from incorrect classes (small percentage) and see how dramatically that can affect the performance of the model - Other concepts that may be covered (time permitting) include activities where students share models in pairs and attempt to fool them with adversarial samples (e.g., a model might learn to classify Marissa vs Mark, but do so by the fact that every example of Mark had a window in the background, rather than by actually looking at his face to distinguish the two as we had intended) - Other activities include exploring Quick Draw, which attempts to guess what you are drawing in real time, as well as an interactive project where students will create a custom hand gesture recognizer and use it to create a game in NetsBlox - This hand gesture recognizer project will consist of a starter template and all the code needed to gather samples of

	custom labeled gestures. - It will then be up to students to create a game that incorporates hand gestures as the control mechanism and test the reliability of such a system - This project will likely not be completed in this day, and students will be allowed to spend some of the following day completing their projects. - Some time will also be included for voluntary sharing of the projects that students have created.
Assessment	- *Throughout the presentations, we will periodically stop and ask for engagement from students in open discussions* - *Each time we reach some difficulty in a project, we will ask for student suggestions on what might be contributing to the problem and/or how to solve it (e.g., the issue of function optimization or the poor performance of image classifiers on small amounts of training data)*

Course Title: Computational Thinking and Machine Learning	
Lesson Date/Day: Day 3	
Daily Objectives	- *Students should have a general understanding of the history of artificial intelligence, including some classical, statistical approaches to machine learning* - *Students should be able to describe the key steps of deep learning and explain how it works in a general sense (mathematics not required)*
Essential/ Debatable Question(s)	- *What techniques are used to gather data for training modern machine learning models?* - *What things can or cannot be learned by a machine learning system?* - *What are the ethical ramifications of using individuals' data without permission?*
Concepts	- *Unsupervised learning* - *Clustering*
Materials	- *The introduction to unsupervised learning will initially include a powerpoint presentation* - *The explanation of clustering will be provided by a multimedia presentation (likely a pre-recorded video) and/or a whiteboard*

	demonstration of the concept

Activities	- The day will begin by providing students some more time to complete their hand gesture classification games from the previous day. Voluntary sharing of completed projects will also be included.
	- After introducing the concepts of unsupervised learning and clustering, students will be given some interactive projects
	  - The first project has students construct a simple clustering algorithm (distributed as a Parsons problem) for attempting the same task of distinguishing red vs blue points
	  - This is a supervised problem, but is easier to understand
	  - The follow-up to this project will be truly unsupervised clustering on drawn shapes.
	  - For this unsupervised version, the project students use will be provided for them already completed.
	  - Students will use this project and create their own (hidden) classes of drawings (e.g., “heart”, “circle”, “square”, etc.)
	  - The code will then attempt to cluster them by similarity and display the contents of each class/family that it discovered
	  - Students will then reflect on which classes were most accurate or inaccurate and form hypotheses about what makes these classes easier or harder to distinguish
	- We will also include an exploratory exercise looking at semantic word embeddings from popular language model pre-processors
	  - This includes things like seeing what words are spatially/semantically close (according to the model), as well as the model’s understanding of analogies between words represented as offset deltas followed by nearest neighbors (this part is hidden from students)

Assessment	- *Throughout the presentations, we will periodically stop and ask for engagement from students in open discussions* - *For this day, the primarily metric of success is engagement in both activities and discussion*

Course Title: Computational Thinking and Machine Learning	
Lesson Date/Day: Day 4	
Daily Objectives	- *Students should be able to explain the high-level differences between supervised learning, unsupervised learning, and reinforcement learning and know when each could be used*

	- *Students should be able to describe the key steps of deep learning and explain how it works in a general sense (mathematics not required)*
Essential/ Debatable Question(s)	- *What things can or cannot be learned by a machine learning system?* - *How do AI engineers teach computers to perform new tasks?* - *What techniques are used to gather data for training modern machine learning models?*
Concepts	- *Reinforcement learning* - *Reward functions* - *Exploration*
Materials	- *The initial presentation on reinforcement learning will be presented in slideshow form, which will lead into the first hands-on activity*

Activities	- Because reinforcement learning is very much tied to game theory, the first activity of the day will again take students back to pretending to play characters in a simple turn based game
	  - This has many analogies to minimax, but we will focus instead on local rewards rather than optimality of global game state
	  - We will discuss this difference and how it helps with the intractability problem
	  - We will also discuss specific reward functions and have students act them out
	  - We will intentionally introduce some seemingly sensible reward functions that have unintended consequences such as stalling the game forever
	  - Before we address this more, we will lead into the coding activity so that students can get faster feedback from attempted fixes to their reward functions.
	- The coding activity will be a virtual version of the same game that students acted out before (possibly with a larger game space)
	  - The activity will be fully-built out of the box, but students will have to create the specific reward function it uses out of a collection of high-level game-specific blocks that we provide
	  - Students will spend some time trying to tweak the reward function and achieve good results and avoid unforeseen consequences like were encountered in the initial exercise
	  - Hints will be given along the way for specific issues that students face

Assessment	- *For this activity, we primarily are interested in making sure students are engaged and continue to make incremental improvements to their reward function* - *Students should understand and be able to explain (theories for) why some of their earlier reward functions were bad and how they managed to improve them.* - *We will emphasize that students should copy and save their good reward functions and keep the best around to share at the end of the day*

Course Title: Computational Thinking and Machine Learning	
Lesson Date/Day: Day 5	
Daily Objectives	- *Students should be able to explain how tools such as ChatGPT generate text* - *Students should be able to identify possible sources of bias in machine learning systems and the potential impact on users*
Essential/ Debatable Question(s)	- *What are the ethical ramifications of using individuals' data without permission?* - *What might be some of the societal impacts of the latest advances in AI?*
Concepts	- *Natural language processing* - *Probability models* - *Large language models*
Materials	- *The initial presentation on natural language processing will be in powerpoint form* - *This will be followed by two coding activities, for which we will distributed starter projects and/or completed versions* - *Finally, we will use the main ChatGPT website to conclude with several problematic aspects of the technology and ways that we might defend against them*

Activities	- *We will begin with a presentation on natural language processing*
	  - *We will present a single task that users might want to accomplish (e.g., via alexa) and pose a diverse collection of various ways that same request could be phrased*
	  - *This is the fundamental problem of natural language processing (even after solving the problem of speech to text)*
	  - *We will ask students if they have any ideas on how this*

	might be accomplished or first steps that might be helpful (e.g., transforming words to semantic embeddings to gain some initial understanding of word relations) - *Likely, students will have no idea, so not much time will be spent here* - *We then introduce the concept of probability models for next-words and use this to explain how text generation works* - *The first coding activity of the day involves playing with a simple text generation program based on short-window markov chains* - *This project will by provided to students in a completed form, but we will explain the process of how it works with a simpler version that students can assemble as a separate Parsons problem (possibly as a take home assignment for those who are personally interested after camp is completed)* - *From this, students will see that the text generation abilities of the model are rather limited and can result in run-on, repetitive, or otherwise incoherent results* - *This will lead into a conversation on the context available to these models* - *This will be treated as a separate activity where students are shown many next-word prediction tasks with various context sizes and asked to predict the next word* - *This will give students insight into the problem* - *From there, we will lead into large language models, which are able to have dramatically expanded context sizes, but are still fundamentally restricted by the same problem* - *We will related this to GPT models such as ChatGPT* - *The final activity of the day will have students connect to a virtual chat room where they will be randomly paired with either another student or an instance of ChatGPT* - *Students are tasked to perform a Turing Test (discussed on day 1) and determine if they are talking to a human or a computer* - *This will be played in several rounds - human/computer identities will be revealed at the conclusion of each round* - *This wraps up the whole class and ties back to very early AI theory* - *This also ends the course with state-of-the-art AI technology and poses a serious dilemma in what forms of media can be trusted as human-sourced vs machine generated and mass distributed (possibly for malicious purposes)* - *Time permitting, we will also explore several problematic aspects of ChatGPT and other language models, such as how to circumvent*

	established response fencing and ways that ChatGPT could be weaponized if put in the wrong hands (e.g., one could use ChatGPT to act as a virtual email scam artist targeting millions of people with zero human oversight) - *Naturally, we will also discuss ways that these attacks could be mitigated.*
Assessment	- *For this activity, we primarily are interested in making sure students are engaged in conversations and ethical issues*