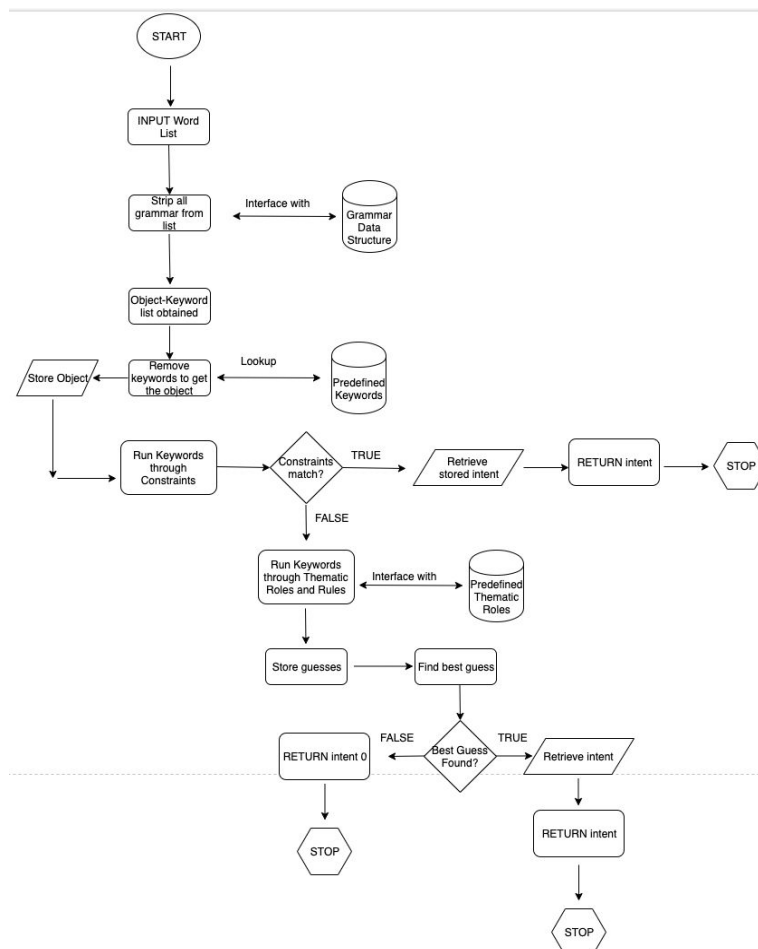
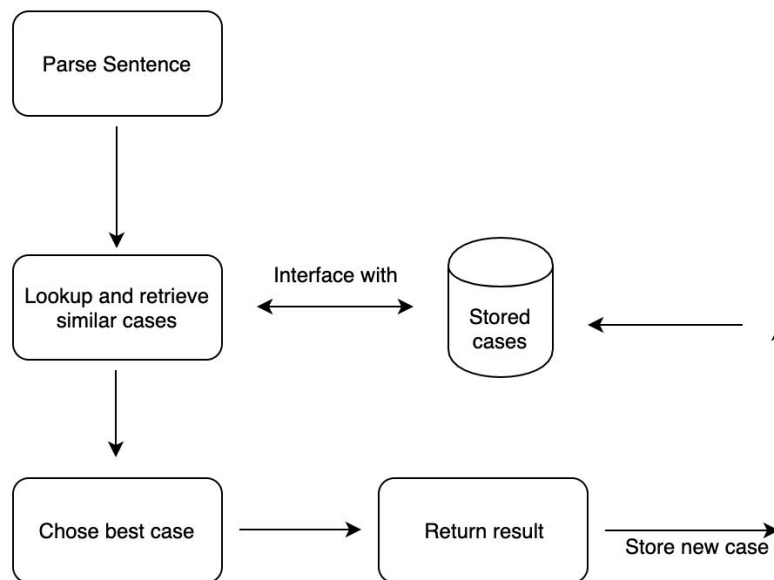


Project 2 Reflection

Throughout project 2, I followed a strict and thoughtful algorithm in order to design an agent that could answer questions as accurately and intelligently as possible. In my algorithm, I first began by creating a class for each project/ assignment and assigning its respective attributes (ie 'DUE DATE', 'RELEASE DATE', etc). I then examined each user question and removed all unnecessary grammar in efforts to parse the sentence as accurately as possible. Once the grammar was removed, the object and keywords of the sentence were identified. Then the object was used to reference a specific class and then the keywords were used to reference a specific attribute of that class. For achieving that, I assigned weights to each attribute to help with fine-tuning the agent. However, a certain kind of question wouldn't be able to be answered using this approach, specifically, the questions which would ask about objects pertaining to a week. So I designed a separate helper function to answer these questions. The helper function would use a syllabus class which had each week with its respective assignments released and due stored as lists. The helper function would then find the specific week being reference and then answer the question to the best of its ability. Upon parsing the sentence and a high confidence count which is affected by the weight of each keyword, the respective intent number (which is stored in the class) would be returned. The thought process that the agent followed was very similar to my own. Below is a diagram explaining the architecture of my agent:



For designing the agent in project 2, I created a best-fit analogy of how the human brain would answer the questions. For example, when asked ‘When is project 1 due?’, my cognitive process as a human would be to first find the project 1 frame in my mind and then look up the ‘DUEDATE’ attribute and return the data associated with that attribute. My overall classification process was to first analyze each question to find certain characteristics which would help in the classification of the type of sentence and what the proper response is. One insight that I programmed into the agent, which stems from human cognition, is that upon encountering a question beginning with ‘Where’, that question must always be classified as a ‘PROCESS’ question, as humans would recognize that ‘where’ refers to a location or process of accomplishing a task. If it is a question beginning with ‘When’, it must only be about ‘RELEASEDATE’ or ‘DUEDATE’, since the human brain associates the word ‘when’ with a time connotation. The overarching logic behind the sentence classification process that humans often use also involves case-based reasoning, a concept we learned in lecture. Case-based reasoning was utilized so that humans can learn from cases and then when they see a case which is not what they typically come across, they can draw similarities and guess the best answer to that case. The questions in this project were fairly easy for me as a human to classify. I did not find myself struggling at all with any of the questions. However, I can see how an agent could get potentially get confused by some questions. For example, the question ‘What project can we start during week 11?’ stumped my agent. This sentence would require the agent to have a very deep understanding of the entire syllabus in addition to being able to extract what the question was asking, which in this case was a bit more complicated than for other simpler questions. The agent became confused particularly on the part of the process where it must understand that the question is asking to find the explicit project that would be released in week 11, however, since there could be an assignment and a project released during that time, the agent could not decipher that question correctly. Below is a diagram mapping the mental process for classifying sentences with KBAI techniques:



The two diagrams are similar and different in several ways. First off, the human brain obviously works very high level as opposed to the very binary conditions that an agent operates with. Humans can

use analogies, verbal reasoning, logic, and several other linguistic techniques that an agent simply does not have. This is why the diagram I made for the human cognitive process is much more condensed and dynamic than the agent's diagram--it is because humans can use a combination of semantic networks, thematic roles, and learning by correcting mistakes all together quite easily, which is not existent for the agent since it mostly uses production systems with minor case-based reasoning and a small level of semantic networks. This is why the architecture of the agent is much more complex and thus its diagram is much more elaborate. I chose to design my agent in this way because the agent is based on the use case of a syllabus chat bot, and its design need not be very generalized for other use cases. This is why i used production systems, case-based reasoning, and semantic networks. Using these in such a specific use case largely yields accurate results as the agent just does not need to be dynamic in nature.

The agent could very easily answer simple and straight-forward questions. However, I designed test questions that were more complicated to challenge the agent and assess its accuracy. Below is a list of 5 questions the agent correctly parsed, and 5 questions it could not properly understand:

Correct:

1. How long is there to complete project one?
2. What assignment can we start during week 1?
3. What assignment can we work on during week four?
4. How long do we have for the final project?
5. How much time is remaining for the 1st project?

Incorrect:

1. Which assignment can I work on in week three?
2. When should we have project 2 completed by?
3. How long will the final final be open?
4. What week is the midterm?
5. What is due in week 16?

The agent answered these questions incorrectly because these questions had more than one possible reply where the agent got confused and instead of replying with intent 0, replied with a specific intent. Other reasons include errors in sentence parsing where the placement of certain keywords, especially numbers, would confuse the agent. Other questions such as number 4 failed because the agent did not know whether to answer with a 'RELEASEDATE' intent of 'DUE DATE' intent.

This project taught me a lot about human cognition. Throughout the development process, I was challenged to delve deeper into how I think and how the human mind operates in order to continuously improve the intelligence of my agent upon it encountering failures in sentence parsing. However, I do feel that I discovered how different and more complicated human cognition is than it seems in day-to-day life. Integrating the way I think into how the agent processes data proved to be a much bigger challenge than I expected. However, it was incredibly rewarding for me to learn about how to teach an agent to interpret human language using concepts that I make use of every day. I am looking forward to using this new framework of artificial intelligence thought processes towards progress on project 3.