Big Idea #4: Natural Interaction

Key Insights	Explanation
Computers can use natural language to communicate factual information, but struggle to understand non-literal modes of expression such as metaphor, imagery, humor, and sarcasm.	Language is an essential part of what makes us human. Metaphor, imagery, and humor are parts of human culture, but we lack precise and detailed theories of how people make sense of these things, or how computers should address them.
Computers can recognize but not experience emotions. Appropriate responses to emotions must be programmed by humans.	AI systems use algorithms to identify human emotions. But there is no heart to ache for us when the machine expresses empathy for our suffering.
Current AI systems are narrow reasoners specialized for well-defined problems. Flexible human-like reasoning, known as "broad AI" or "Artificial General Intelligence", has not yet been achieved.	It is unclear when we'll be able to build Artificial General Intelligence systems that can mimic or exceed human reasoning capabilities.
Current AI systems lack consciousness and self-awareness. A self-aware computer would require representations of its own existence and thoughts, and memories of its past experiences.	Consciousness is understanding that one exists, i.e., awareness of one's body and one's environment. Self-awareness is the recognition that one is a conscious, thinking being with the ability to reason about one's thoughts. See Jabr (2012) for a brief discussion of these ideas.

Big Idea #4: Natural Interaction	*Intelligent agents require many types of knowledge to interact naturally with humans.*	LO = Learning Objective: What students should be able to do. EU = Enduring Understanding: What students should know. Unpacked descriptions are included when necessary to illustrate the LO or EU	

Concept	K-2	3-5	6-8	9-12
Natural Language *(Structure of language)* **4-A-i**	**LO:** Demonstrate knowledge of the structure of language through tasks such as (a) generating plausible and implausible novel words, or (b) reordering the words in a scrambled sentence so that it makes sense. **EU:** Human languages have rules for how words and sentences are constructed, and computers can use these rules to help them figure out what people are saying. **Unpacked:** A plausible novel word is "flurg"; an implausible one is "fnurg" because English doesn't allow words to begin with "fn". Turning to word order, the words "ate pancakes John" only make sense in the order "John ate pancakes". **Activity:** Use this SpeechDemo to see how Google responds to *"I ordered a large **flurg** with whipped cream."* vs. *"I ordered a large **fnurg** with whipped cream."*	**LO:** Demonstrate that human language is infinite by showing how any sentence can be repeatedly extended to form a more complex sentence. **EU:** Human language can express an infinite number of ideas and form an infinite number of sentences. This property makes it impossible to pre-program a computer with a response to every sentence. Thus, to understand a new sentence a computer must recognize how the words combine into phrases and clauses to communicate complex ideas. **Unpacked:** Sentences can be lengthened by adding new phrases or clauses,with no limit. Repeatedly extending a sentence will yield something like, "John said that Mary knew that Peter saw that Lisa gave Megan the book about hamsters from outer space that Harry recommended." **Activity:** Have the class take turns extending a sentence until it is exceptionally long. Then run it through the Berkeley Neural Parser to see if it can correctly recognize the phrase structure.	**LO:** Demonstrate a computer's grasp of grammar by using a parser program to display the syntactic structure of a sentence, and explain what the nodes represent. **EU:** Parse trees are a way of representing the syntactic structure of a sentence, showing the relationships between words. Computers can use syntax trees to both analyze and generate sentences. **Unpacked:** The grammatical structure of a sentence is key to understanding its meaning. For example, if the task is question answering, then we need to understand the meaning of the question to perform the appropriate query. If the task is to get a robot to do something, the robot must understand what is being requested. Parsers produce syntax trees whose nonterminal nodes are grammatical categories such as NP (Noun Phrase), VP (Verb Phrase), and PP (Prepositional Phrase). "Put the cup on the saucer" vs. "put the saucer on the cup" illustrates the importance of syntax. Both sentences contain the same words, but the words have different syntactic relationships. That's why changing the word order changes the meaning. In addition, "cup" can be either a noun or a verb, but in this context the syntax indicates it's being used as a noun. **Resource:** The Berkeley Neural Parser demonstrates both POS (part of speech) tagging and parse tree generation. The tags come from the Penn Treebank project.	**LO:** Identify portions of a text that would be difficult for a computer to understand, and explain why. **EU:** Computers have difficulty understanding text that makes use of metaphor, imagery, hyperbole, sarcasm, or humor, or word play. **Unpacked:** Currently, we don't have satisfactory formal explanations of metaphor, imagery, hyperbole, sarcasm, humor, or word play. This is the focus of current AI, linguistics, and cognitive science research.

Big Idea #4: Natural Interaction		*Intelligent agents require many types of knowledge to interact naturally with humans.*	LO = Learning Objective: What students should be able to do. EU = Enduring Understanding: What students should know. Unpacked descriptions are included when necessary to illustrate the LO or EU	
Concept	**K-2**	**3-5**	**6-8**	**9-12**
Natural Language *(Ambiguity of language)* **4-A-ii**	**LO:** Illustrate the ambiguity of language by giving examples of homophones and homonyms and showing how the correct word can be determined using context. **EU:** Different words may sound the same, but people and computers can use the surrounding words to determine the correct one. **Unpacked:** homonyms (e.g., "bear" can be either a noun or a verb) and homophones (bear/bare, there/their/they're, or which/witch). This is called *lexical ambiguity* **Activity:** Use this SpeechDemo to see alternative interpretations of homophones.	**LO:** Illustrate how understanding a sentence could be challenging for a computer by giving sentences where a pronoun could refer to either of two nouns. **EU:** Sentences can have multiple meanings depending on which noun a pronoun is thought to be referencing. People and (to a lesser extent) computers can use context and world knowledge to select the most likely meaning. **Unpacked:** An example of *reference ambiguity* is "John handed Pedro his cellphone", because "his" could refer either to John or to Pedro. Often world knowledge can resolve this type of ambiguity, as in this example: "The trophy would not fit in the suitcase because it was too large/small." The choice of "large" vs. "small" changes our decision about what "it" refers to. Reference ambiguity is an issue of *semantics (meaning). The* parse tree (grammar) is the same for both sentences, but the meanings are different. Compare this with attachment ambiguity discussed in 6-8.	**LO:** Illustrate how understanding a sentence could be challenging for a computer by giving sentences where a prepositional phrase could attach in either of two paces, and show how this ambiguity can sometimes be resolved based on meaning. **EU:** Language is often ambiguous, but some possible meanings can be ruled out if they don't make sense. Computers are not yet as good as people at using context and world knowledge to make these judgments. **Unpacked:** "John saw the man with the telescope" is ambiguous because it's unclear who had the telescope. But "John saw the man with the violin" is unambiguous. *Attachment ambiguity* is an example of *syntactic* ambiguity because the parse trees are different: "with the telescope" could attach to either "John" or "the man". Semantic knowledge (meaning) can help humans choose the most plausible parse. Compare this with reference ambiguity discussed in 3-5. **Activity**: Create ambiguious sentences and run it through the Berkeley Neural Parser to see if it can attach the the prepositional phrase to the correct word.	**LO:** Illustrate how understanding a sentence could be challenging for a computer by describing multiple senses of a given word. **EU:** A single word can have multiple senses. People use context and world knowledge to determine the correct sense, but computers don't typically represent meaning with this degree of subtlety. **Unpacked:** Multiple senses is called "polysemy". For example, "book" can refer either to a commercial product ("the book was a bestseller"), a collection of words ("the book's language is appropriate for middle school students"), or a physical instantiation of the text ("the book was heavy in my backpack"). Similarly, "class" can refer either to an academic course ("a math class"), a specific instance ("Tuesday morning math class at Riverdale High"), or the persons attending it ("the teacher addressed the class"). Compare this with lexical ambiguity due to homonyms and homophones discussed in K-2. Modern AI systems using neural nets and huge training sets take a statistical approach to resolving word sense ambiguity, but statistics alone aren't as powerful as humans' common sense. **Activity:** Use Google Translate to explore how computers resolve multiple senses of a word when selecting appropriate words for the translation.
Natural Language *(Reasoning about text)* **4-A-iii**	**LO:** Demonstrate how a computer can produce different forms of a verb, such as present or past tense. **EU:** In order for a computer to speak naturally with humans it must be able to understand how words are constructed and put words in the proper form. **Unpacked:** Verbs can take different forms in first, second, and third person, singular or plural, and present or past tense. For a computer to respond correctly to the question "When did dinosaurs live?" it should say "Dinosaurs lived over 65 million years ago." **Activity 1:** List all the different forms a verb can take, then check your results using this online verb conjugator. **Activity 2:** This past tense converter is not very good. Experiment to see where it succeeds and where it fails. **Activity 3:** Set up Google Translate for English -> Spanish and enter a sentence like "I walk to school". Then try different forms of the sentence ("I walked", "I am walking", "I was walking" "I had been walking", etc.) and observe how the translation changes.	**LO:** Experiment with a speech to text system to see if it resolves alternative word choices correctly based on context. **EU:** Speech recognition systems can use grammar and context to resolve ambiguous words but they don't always get it right. **Activity:** Use this SpeechDemo to see alternative interpretations of *"the dishes' size and weight were impressive"*, or homophones such as "which/witch" in *"I couldn't tell which of the witches was the witch with the broomstick."*	**LO:** Illustrate how word embeddings can be used to reason about the meaning of words. **EU:** Word embeddings represent words with similar meanings as nearby points in a semantic feature space, and allow us to reason about words by doing arithmetic. **Unpacked:** An example of reasoning about word meanings is **outlier detection**, e.g., which word in the list "breakfast", "lunch", "banana", "dinner" does not fit with the others? Similarity of words can be measured as distance in feature space; the outlier will be farther from the other words. Another example of reasoning about meaning is **analogy completion**, e.g., "man is to king as woman is to *queen*". This can be solved by calculating the feature vector "king" - "man" + "woman" and finding the word closest to the result. The real benefit of word embeddings is that they are used as inputs to transformer networks which perform much more complex reasoning operations on text. *See 2-A-iv.6-8 For more on word embeddings and 2-A-iv.9-12 for transformer networks.* **Activity 1:** Experiment with WordEmbeddingsDemo to explore realistic word embeddings. **Activity 2:** Use knowledge of word embeddings to explain how the Semantris game works (Blocks activity).	**LO:** Demonstrate how a small context-free grammar can be used to parse or generate simple sentences. **EU:** Context-free grammars describe how words are combined into phrases and clauses, and can represent much of the syntactic structure of natural language, but don't handle things like subject-verb agreement well. **Unpacked:** An example of an agreement constraint is that subjects and verbs must agree on number, e.g., "he says" but "they say". Expressing this in a context-free grammar would require separate rules for singular and plural subjects. This gets cumbersome when sentences become more complex, e.g., coordinating conjunctions may require agreement on tense as well as number. See 4-A-ii.6-8 for examples of syntactic ambiguity, which leads to multiple syntactic parse trees that must be disambiguated using semantic knowledge. Even a simple grammar, if it's recursive, can generate an infinite number of sentences; see 4-A-i.3-5. **Resource:** This Stanford Context Free Grammar tool can be used to construct parse trees, generate sentences, and verify the grammaticality of a sentence. See this reference for creating a context free grammar.

Big Idea #4: Natural Interaction		*Intelligent agents require many types of knowledge to interact naturally with humans.*	LO = Learning Objective: What students should be able to do. EU = Enduring Understanding: What students should know. Unpacked descriptions are included when necessary to illustrate the LO or EU	
Concept	**K-2**	**3-5**	**6-8**	**9-12**
Natural Language *(Applications)* **4-A-iv**	**LO:** Demonstrate the kinds of tasks an intelligent assistant can and cannot perform. **EU:** Intelligent assistants (e.g., Siri, Alexa) are computers designed to respond to a limited set of requests. They cannot engage in a conversation like a human. **Unpacked:** Siri and Alexa can answer questions, play music, set alarms, and make lists. But intelligent agents are not people. They have limited conversational abilities despite their ability to recognize spoken language. One reason is that they have trouble maintaining context, i.e., remembering what was said before to help them understand what is being said now. **Activities:** Talk with an intelligent agent such as Alexa. What are the kinds of things they can do well or not so well? What are the limits of its understanding?	**LO:** Demonstrate some types of questions that a search engine or intelligent assistant can answer, and some types that it cannot answer. **EU:** Search engines (e.g., Google) and intelligent assistants (e.g., Siri, Alexa) have a collection of specialized and general purpose modules they draw upon to answer different types of questions. **Unpacked:** Examples of queries handled by specialized modules include: definitions of words; unit conversions (e.g., inches to millimeters, dollars to euros); current time and weather anywhere in the world; biographical and geographical facts (birth date of Abraham Lincoln, capital of Belgium); store hours and driving directions; and airline, train, and bus schedules. Current intelligent agents do less well when asked to reason about relationships between entities (e.g., is an alligator bigger than an ostrich), or when asked to produce explanations rather than simple facts (e.g.,Why didn't Shakespeare write about airplanes?). When no specialized module can handle a query, search engines fall back on keyword search, but the results are often unsatisfying.	**LO:** Describe some NLP (Natural Language Processing) tasks computers can perform, and explain how they work. **EU:** NLP (Natural Language Processing) tasks include text summarization, text generation, sentiment analysis, question answering, machine translation, and conversational interaction.	**LO:** Describe several approaches to Natural Language Processing, ranging from simple to more sophisticated. **EU:** Simple NLP approaches include keyword matching, dictionary lookup, and template matching, while newer, more sophisticated but less transparent approaches use deep neural networks and machine learning. **Unpacked:** Taking chatbots as an example, the simplest approach looks for keywords in the user input to decide what response to give. A slightly more sophisticated approach uses templates to describe all the variations a phrase might take rather than looking at single phrases. More sophisticated chatbots use deep neural networks for "intent recognition", which detects when the *meaning* of an input matches a template, rather than looking for specific words or phrases. **Activities:** Students should be able to describe when these approaches might be helpful or not. Explain the limitations.
Commonsense Reasoning **4-B-i**	**LO:** Give commonsense explanations of human behavior or events that might be useful for a robot to know. **EU:** Choosing the most likely ending for a story is an example of commonsense reasoning, which requires knowledge about people and things, and how they behave, that computers may not have. **Activities:** (1) Choose the most likely ending to a story given a list of alternatives, and explain your choice; (2) Explain why a character in a story took a particular action; (3) Suggest an appropriate action for a character in a particular situation, and explain why it is appropriate.	**LO:** Explain what knowledge would be required for a computer to understand a story. **EU:** AI has trouble understanding stories because it doesn't have the knowledge that humans have about everyday life. **Unpacked:** "Everyday life" includes both cultural knowledge (what is an umbrella for) and naive physics (dropped objects will fall due to gravity). **Activity:** Compose a story that may be hard for a computer to understand, and explain what makes it hard.	**LO:** Explain the knowledge a computer would need to solve a naive physics reasoning problem. **EU:** Computers can reason about physical phenomena using *naive physics* inference rules, which formalize our intuitive understanding of concepts such as mass, volume, forces, and motion. At present, computers are not very good at this. **Unpacked:** Naive physics includes knowledge of conservation laws. Pouring a large container of water into a smaller container will cause an overflow because small containers have less volume and the amount of water doesn't change. By grades 6-8 students should have mastered conservation laws; see this article on [Piaget conservation tasks](). Naive physics also includes inference rules for reasoning about forces, such as "When a heavy thing hits a light thing, the light thing moves." **Activity:** Propose a reasoning problem that requires knowledge of naive physics to solve, and say what that knowledge is.	**LO:** Explain the cultural and naive physics knowledge required for a computer to correctly interpret a fable or fairytale. **EU:** For humans or computers to infer the meaning of a story requires understanding of cultural knowledge, naive physics, and folk psychology. This is still difficult for computers. **Unpacked:** Fables and fairy tales are challenging for computers because they can incorporate cultural knowledge, complex human motivations, imagery, humor, and metaphor. Folk psychology refers to our everyday ability to attribute mental states to other people, including their beliefs, desires, and intentions. **Activity:** Explain what knowledge is required to correctly resolve the referent of the pronoun in a Winograd sentence. An example Winograd sentence is "Anna did a lot [better/worse] than Lucy on the test because she had studied so hard." Here, the required knowledge is that more studying for a test leads to better performance.

Big Idea #4: Natural Interaction		*Intelligent agents require many types of knowledge to interact naturally with humans.*	LO = Learning Objective: What students should be able to do. EU = Enduring Understanding: What students should know. Unpacked descriptions are included when necessary to illustrate the LO or EU	
Concept	**K-2**	**3-5**	**6-8**	**9-12**
Understanding Emotion **4-C-i**	**LO:** Demonstrate how computers recognize emotions in faces. **EU:** Computers recognize emotions in faces by looking at the shapes of the mouth, eyes, and eyebrows. **Activity:** Use a face demo that recognizes facial landmarks (e.g., eyebrows, eyes, and mouth) to infer emotional state. The activity can be introduced using icons to demonstrate different emotions.	**LO:** Illustrate how computers can judge the emotional tone of text. **EU:** Computers can recognize positive and negative statements about a topic using natural language processing techniques known as "sentiment analysis". **Unpacked:** Sentiment analysis is used for a variety of tasks, such as: - **Movie reviews:** Analysing online movie reviews to get insights from the audience about the movie. - **Restaurant reviews**: analyzing reviews of restaurants to measure customer satisfaction. - **News coverage:** Analyzing news coverage of an event, person, or company to assess media opinion. - **Social media analysis:** analyzing the sentiments of Facebook, Twitter, or Instagram posts to assess public opinion about events, persons, or products. **Activities**: Use appropriate AI services (e.g., Scratch plugins) to create artifacts that demonstrate sentiment analysis. Use a web-based sentiment analysis tool to analyze the polarity of text that is marked subjective.	**LO:** Describe how computers use different types of cues to recognize human emotional states. **EU:** Computers can recognize human emotional states by looking at facial expressions, gaze, gestures, body language, tone of voice, and choice of words. **Unpacked:** For example, body language can indicate engagement, hostility, anxiety, or boredom. Body language includes such factors as how a person holds their arms, how they position their feet, whether they are leaning in or leaning back, and how they are tilting their head. Computers can extract this pose information from webcam images.	**LO:** Identify ways AI applications can modify their behavior to respond to people's emotional states. **EU:** Computers can respond appropriately to human emotions by acknowledging what the human is feeling and responding in a way that the human finds supportive and socially appropriate. **Unpacked:** For example, AI-based tutoring systems monitor student behavior to identify frustration, boredom, and tiredness so they can adapt the instruction or prompt the student to take a break. Automated customer service agents that detect human emotions could adapt their responses accordingly, e.g., by adjusting their speech rate and tone of voice. **Activity:** Given that computers can use AI to identify human emotions, discuss how should they respond to people in a way that we think is supportive and socially appropriate.
Philosophy of Mind **4-D-i**	**LO:** Demonstrate some ways that Intelligent agents don't understand things the way people do. **EU:** Artificial Intelligence devices are not people, they are computer programs designed by people. **Unpacked:** Today's intelligent agents cannot answer most common sense questions, such as "Would smelly socks make a good birthday present?" They can read stories to children but cannot answer questions about the story. They are also unaware of repetitive questions, e.g., if asked the same question five times in a row they would not become frustrated or even acknowledge that the questions are repetitive.	**LO:** Demonstrate what a broad AI system could do that current narrow AI systems cannot do. **EU:** "Narrow AI" has given us a collection of algorithms that are good at specific things. "Broad AI" or "General AI" promises to give us general reasoning ability comparable to a human. **Unpacked:** At present, only people exhibit general intelligence. They can do lots of different things, and learn to do new types of things. Current AI programs are narrowly focused on a single kind of problem. For example, programs that are good at playing chess or driving a car are not good at answering questions about a story. We don't currently have algorithms or machine learning technologies that can mimic all the facets of human intelligence.	**LO:** Define criteria for consciousness and evaluate AI systems or fictional AI characters according to those criteria. **EU:** The nature of consciousness is an unsolved problem, but two important criteria seem to be self-awareness and personal history. **Unpacked:** Questions about consciousness: (1) is self-awareness (or "awareness of awareness") required; (2) is a personal history (memories of past experiences) required; (3) is a mental language required; (4) are there levels of consciousness or is it all or nothing; (5) what kinds of consciousness do animals exhibit. **Activity:** Research discussions about whether animals can be conscious, and how animal consciousness differs from human consciousness. If consciousness exists on a scale, do you think a computer could ever meet or exceed humans on this scale?	**LO:** Debate alternative perspectives on human vs. artificial intelligence **EU:** Some AI experts and philosophers have argued that computers can never achieve fully human-like intelligence, while others believe that AI will eventually *exceed* human intelligence. **Unpacked:** One problem with arguing that a computer possesses human-like intelligence is that we don't have a clear idea of how human intelligence is achieved. A stronger claim against computer intelligence is that because computers lack physical bodies they cannot experience the world the way humans do. On the other hand, people have limited memories and reasoning abilities that a computer could potentially surpass, leading to an intelligence that was more capable than humans (artificial superintelligence). But we can also have devices that perform complex information processing tasks without possessing human-like intelligence. **Activities**: (1) Explain the Turing test and pull out a snippet of Turing's paper and compare the dialogue proposed by Turing to what modern-day AI systems can do. (2) Compare animal to human intelligence. (3) Compare animal intelligence to AI.