the details | real talk with robots

Students explore the world of machine perception, supervised learning and classification models while increasing their self-efficacy as AI creators.

driving question

Can we make robots that are conscious? Should we?

nutrition facts

[P]Perception and Data Gathering

[P2-a1] Speech recognition
[P2-c1] Sentiment Analysis

[R]Representation and Reasoning

[R2-a1]Rule-based expert systems

[L]Learning

[N]Natural Interaction

[N2-b]Human-AI interaction
[N1-a1]Natural language understanding
[N1-a2]Natural language generation
[N1-b]Dialogue flow
[N1-d]The Turing Test
[N2-a]Human-robot interaction
[N3]Consciousness and philosophy of mind

[S]Social Impact

[S3-a] Impact on work
[S3-c] Impact on human-human interaction
[S5-a] Algorithmic bias
[S5-b] Bias in data collection

learning goals: knowledge

Students will be able to...

- Define Natural Language Processing (NLP) and give a summary of the NLP pipeline.
- Describe how NLP converts word relationships into mathematical relationships.
- Describe tools used for intent recognition in NLP.
- Understand how NLP classifies types of phrases to make it easier to understand patterns in conversation.
- Connect human conversation with logical flows by creating a conversation flowchart in their chatbot development process.

learning goals: ski s

Students will be able to...

- Evaluate the appropriateness of creating an interactive chat-based social agent to address a societal need.
- Develop an interactive chat-based social agent that augments their social-emotional learning by incorporating mindfulness strategies.

learning goals: a itudes

Students will...

- Identify at least one listening strategy that they can use to have productive conversations.
- Practice listening strategies centered on active and sympathetic listening.

suggested tools

- PRG Version of Scratch (link)
- Chatbot builder software (link)

leveling up & down

computational complexity

High: Students train their own models of intent reasoning, and then use them in Scratch
Low: Students use word embeddings in Scratch with dialogue flowcharts to make prescriptive chatbots

programming ability

High: Students use custom blocks and subroutines to create increasingly complex subtopic interactions.
Low: Students use basic Scratch blocks and intent reasoning extension blocks

design thinking

High: Students engage in participatory research to create a dataset of conversation, on which they base dialogue explicitly or via transfer learning.
Low: Students make a character in Scratch that can interact using dialogue based on their own theory of mind.

on the menu | real talk with robots

We are building chatbots that can pretend to be human by having interesting, interactive conversations.

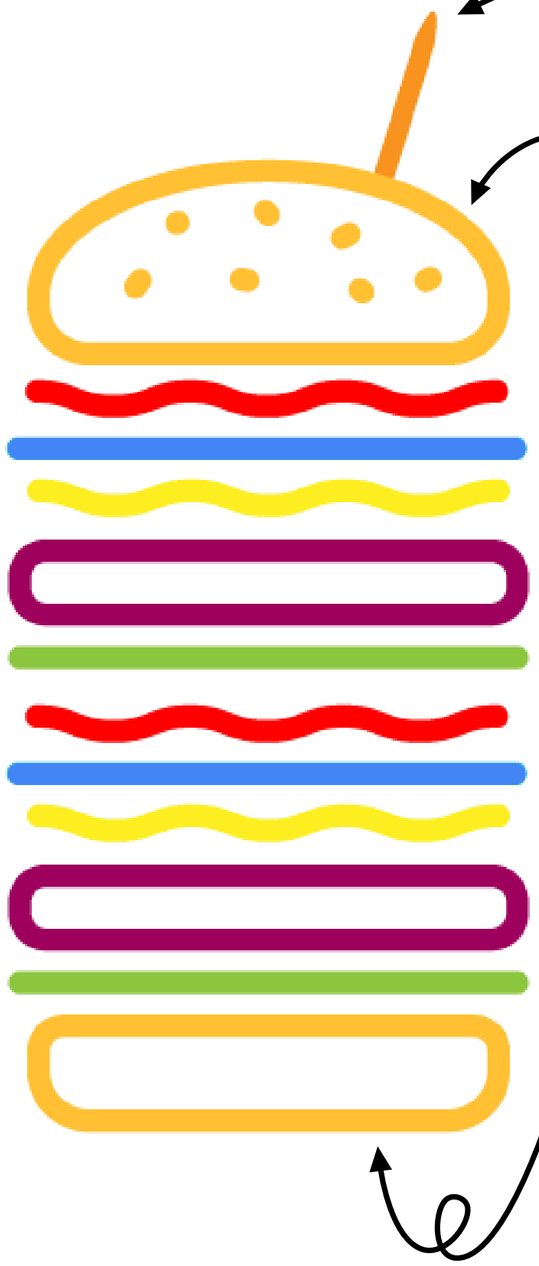

toothpick

The driving question: How can we teach artificially intelligent robots about feelings?

top bun

Students engage school community members in live Turing Tests of their chatbots.

assessment & re ection

Students track their project progress and reflect on their learning using their design journals.

condiments

Student-facing resources aimed to help students solve technical issues and understand how to use text classification tools in Scratch.

project creation

The structure of studio time and the backbone of the module, where students complete the bulk of their projects. Students create a Scratch project with text classification that can have a unique conversation with the user by responding to user inputs.

curated ings

Whole-class activities for teachers looking to deepen students' learning around specific topics, including:

- AI-generated speech
- Artificial neural networks

core ings

Whole-class activities where students learn the principles behind chatbots. Taking a step back from their projects, students learn about:

- natural language processing
- dialogue flows
- intent reasoning
- word embeddings

bo om bun

The exciting project launch: exploring the Turing Test and how robots can mimic humans.

timeline | real talk with robots

Students follow the design thinking process in a hands-on development of a chatbot.

start

Turing Tests 1-2 days	Students talk about how robots can be similar to humans, and if robots have feelings. The class then discusses existing chatbots, ranking them to find the most human- like chatbot.
Exploring projects 1-2 days	Students are presented with the project prompt: We are building chatbots that can pretend to be human by having interesting, interactive conversations. Students then explore example projects.
What is NLP? 1-2 days	Before starting their own projects, students learn about the larger field of Natural Language Processing, modeling a Naive Bayes Classifier.
Assessment point <1 day	[TBD, based on Assessment team]
Empathiz e 2-3 days	Groups think about who will use their chatbots and learn how they might best have kind conversations with users.
Define <1 day	Students define the goal of the chatbot: what kind of conversation will it have with users?

Dialogue acts in NLP 1-2 days	The class learns how to classify parts of a conversation, then considers how to incorporate dialogue acts in creating a chatbot.
Dialogue flows 1-2 days	The class learns how to script an overall conversation between a chatbot and a human, mapping out possible paths based on how the human responds.
Ideate 2-3 days	Based on the goal of their chatbots, students consider what kinds of personalities to give to their chatbots. Students also brainstorm and map possible conversation flows.
Simple intent reasoning 1-2 days	The class learns about rule parsing, a simple way to decode what a user is talking about in a conversation with a chatbot.
Prototyp e 3-5 days	Students create video or paper prototypes of possible chatbots, based on selected ideas from the Ideate stage.
Intent reasoning: emotion and toxicity 1-2 days	Humans aren't always direct when speaking, so this lesson focused on how computers can detect what users want. Students also learn how emotion can be sensed by chatbots.
Using word embeddings 1-2 days	Students learn how synonyms can be mathematically represented, showing the distance between similar words.

Build

3-5 days

Teams complete the bulk of their projects, building a text classification model in Scratch and scripting their chatbots.

Assessment point

<1 day

[TBD, based on Assessment team]

Test & assess

2-3 days

Students iteratively test and improve their chatbots, incorporating peer feedback. They put the final touches on their final projects.

Wrap-up & portfolio

<1 day

[TBD, based on Assessment team]

Presentations

1-2 days

Students present their work to the community in a Turing Test event, where participants must determine if they're speaking with a human or a student-created chatbot.

end

background | real talk with robots

You don't have to be an AI expert to teach this module. Here's what you need to know.

summary

Natural Language Processing (NLP) is the sub-field of AI focused on processing and generating written and spoken human speech. The field of linguistics is a prominent contributor to NLP. Tools that utilize NLP can be speech-based or text-based; examples include autocorrect and autocomplete tools, Google Translate, and voice-activated assistants, including Siri and Alexa. In this module, students will create a chatbot, which is a text-based NLP tool. Chatbots are often employed by websites and apps for simple customer support, deciphering user requests to provide tips and connect the user with another human for more support, if needed. The chatbot students will make will utilize human-like listening skills to make the users feel heard.

A car website chatbot

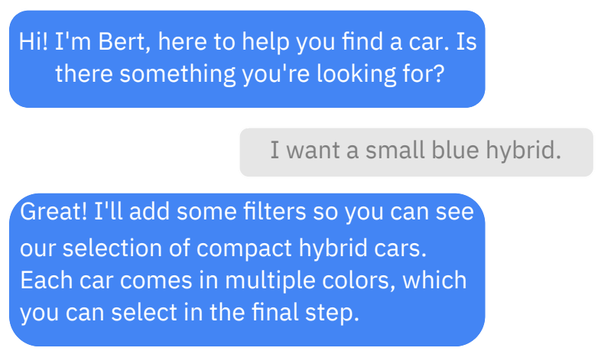

When designing a chatbot, designers start by scripting possible dialogue flows, represented with flowcharts. Points where a user could answer "yes" or "no" cause branches in the possible conversation, which can be represented with two different paths following that conversation point.

Dialogue flows

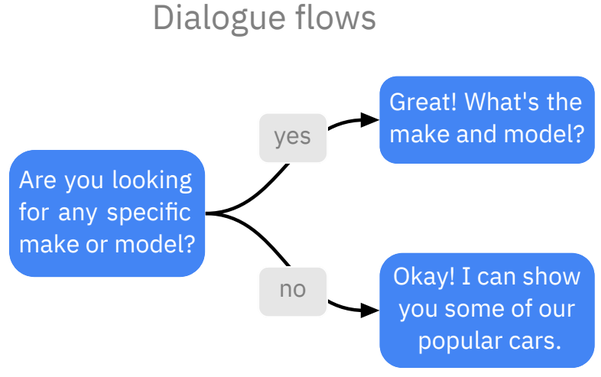

To parse speech inputs, NLP tools use dialogue acts, which are classifications for all parts of a conversation. Questions, "yes" answers, backchannels (acknowledgement of understanding/paying attention; e.g., "mhmm"), and declarative statements all have their own patterns and uses within human speech. By paying attention to the patterns of dialogue acts within human speech, we can create better chatbots that sound more human. The ultimate goal for many NLP tools is to pass the Turing Test, where it becomes indistinguishable from a human. Note that many NLP tools have claimed to pass the Turing Test, but none have reached the point seen by AI tools in movies-- like HAL 9000 from *2001: A Space Odyssey* or J.A.R.V.I.S. from the *Avengers* franchise.

Word embeddings mathematically represent the distance between words' meanings. Practically, word embeddings can be used by NLP tools to detect similar words, allowing users to have more flexibility in their requests. For example, a user could receive similar responses for "where should I go for breakfast" and "where should I go for brunch," since "breakfast" and "brunch" are semantically close.

Word embeddings

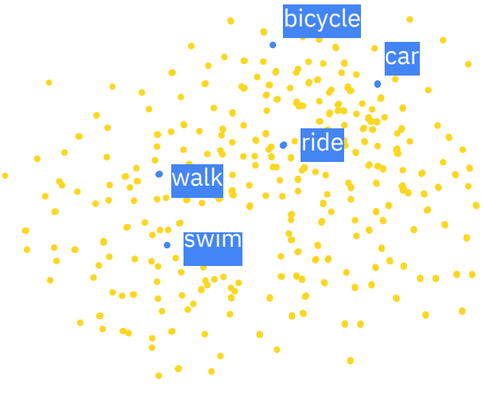

However, synonym detection isn't sufficient for understanding what humans want when speaking. For that, NLP developers look to intent reasoning. Simple intent reasoning relies on rule-based parsing, where the machine follows specific grammatical rules that the designer set out. However, human grammar fluctuates and can be hard to specify, so new AI tools are used to better detect what humans are saying. New AI tools can even understand less obvert communication; for example, some NLP tools can understand "can you pass the salt" as a request, rather than a question asking about another's ability to move a salt shaker. Many intent reasoning algorithms can also detect toxicity and emotion, furthering their overall ability to process human speech.