

Conversational Artificial Intelligence

Principles and Practice of Virtual Assistant AI

EECS 498, WINTER 2019

(prereqs: 280, 281)

Summary

The science and art of creating conversational AI spans multiple areas in computer science. Learning about and leveraging advances in these areas to create state-of-the-art conversational virtual assistants is the central focus of this course. Throughout the course, students will learn principles of deep learning applied to natural language processing, dialogue management, response generation, and other applications. Students will also build end to end virtual assistants using tools that span both traditional techniques as well as cutting edge techniques for these AI with the goal of creating experiences that surpass existing technologies such as Google Assist, Siri, and Alexa. Students will also be tasked with integrating the AI experiences they create with a set real APIs of their choosing (e.g., Spotify, Fitbit, Unity, etc.). The course culminates in a demo day where creations are shared for others to try, and the best such creations will have the opportunity to be featured at a technology conference such as SXSW, CES, or Techcrunch Disrupt.

In the creation of these virtual assistants, students will form groups of around 5 students to select a use case, design the virtual assistant, train AI models, implement logic, and execute the end to end build out and integration into an existing product API. Throughout this process important software engineering and design practices will be exercised. Students will learn about the process of data collection, curation, and crowdsourcing to build good clean datasets. Students will also learn state of the art techniques for the iterative process of debugging and quality assurance for AI. Students will do at least 2 rounds of user testing and refinement. In parallel throughout the course, readings on the underlying deep learning techniques the enable state-of-the-art conversational AI will be selected and discussed in class. This course is going to be awesome!

Instructor: Jason Mars
(<http://web.eecs.umich.edu/~profmars> or <http://www.jasonmars.org>)

Lecture: MW 1:30-3, EECS 3427

Discussion Section: 1hr per week coaching session with industry practitioner (per group, times determined by each group)

Credits: 4

Prerequisites: EECS 280, EECS 281 (optional: EECS 485)

Office Hours: Wednesdays 3-4

Course Website: Coming soon.

Tools Used (provided): Google's Dialogflow, Amazon Lex, Microsoft Luis, IBM Watson (bluemix), Clinc Platform, other free API

Schedule by Week *(subject to slight changes)

1/9	Course Introduction <ul style="list-style-type: none">- Syllabus overview- Project examples- First reading assigned
1/14	Introduction to AI and Deep Learning <ul style="list-style-type: none">- What is machine learning?- Whats a DNN, CNN, RNN, xNN?- Why is data so important?
1/21 <small>*(no class Monday, MLK Day)</small>	Impact of Deep Learning on Industry <ul style="list-style-type: none">- How the scientific breakthroughs in ML and data is revolutionizing industry.- The story of Siri, Google Now, Alexa, Facebook's M and what it means for all business
1/28	Project Pitches and Feedback <ul style="list-style-type: none">- One page write-up due- First project assigned using tools from Google, Microsoft, Amazon, and IBM
2/4	Components of a Virtual Assistant <ul style="list-style-type: none">- Slot value extraction

	<ul style="list-style-type: none"> - Intent classification - Dialog Management - AI Reasoning Subsystem - Response Generation
2/11	Tools of the Trade Presentations and Demos <ul style="list-style-type: none"> - Project 1 due - Student presentations of how they built project in Google's Dialogflow, Microsoft's Luis, Amazon's Lex, and IBM Watson - Compare and contrast discussion
2/18	Clinic Platform Introduction and Overview <ul style="list-style-type: none"> - Special guests join the class - Project planning working group - The science and art of dialog management
2/25	Data Collection, Intro to crowdsourcing <ul style="list-style-type: none"> - Classic techniques for data collection - Using the crowd as a tool - The creativity of getting good data
3/4 *(Spring Break)	No Class
3/11	Getting Clean Data, Advanced crowdsourcing for NLU <ul style="list-style-type: none"> - Data iteration and refinement - Debugging bad data - NAACL paper reading assigned
3/18	Science of Good Data (special guest) <ul style="list-style-type: none"> - Paper author presentation - Scoping and building conversational flows
3/25	The power of RNNs <ul style="list-style-type: none"> - DNNs that capture seq-to-seq features - fastText and Elmo, good stuff for production
4/1	Special Presentation from Industry Practitioner
4/8	QUIZ and Special Topics
4/15	Final Project Presentations
4/22	Last Day Wrap-Up

- | | |
|--|---|
| | <ul style="list-style-type: none">- All projects final version Due- All writeups due |
|--|---|

Grading

This is a very 'do heavy' course. A significant portion of the grade is allocated to the project and demos. Lets build some amazing stuff!

Quiz: 10%

Project write-ups: 15%

Presentations: 20%

Demos: 25%

Final Project Presentation, Code, and Report: 30%

Participation (extra credit): Up to 3% Extra Credit (some/good/excellent participation levels)

Logistics and Details

Late Policy

Assignments are expected to be completed on time. However, for unavoidable situations where this is not possible, you may use up to 1 late day (24 hours) per assignment for a penalty of 10%. Anything submitted more than 1 day late will receive 0 credit.

Assignment Submission

Projects should all be submitted in as a .zip file containing project code, and a PDF containing README content (setup / usage instructions) and a link to a hosted version of your code. For the final project we will also collect a link to your hosted AI and we will perform testing. We'll also use Github in the class.

MDE Project

This course will have a large team-based project that will require designing and building a virtual assistant. After group formation, an initial 'pitch' document and in-class presentation will help teams get feedback on their ideas. After that, there will be a few milestones on the way to a final project document, presentation, and demos. The objective of this project is to build a system from the ground up that work really well in practice. Be creative!

Honor Code

All students (including LS&A and Engineering) are required to observe the Engineering Honor Code in all assignments and exams. A copy of the honor code can be found at <http://ossa.engin.umich.edu/honor-council/>. Please make sure that you clearly understand what constitutes cheating. If you are not sure in any specific case, you should ask the teaching staff. The University takes honor code violations seriously, and penalties can be severe. You are not allowed to share your code with anyone other than your partners. You are not allowed to make use of project or homework solutions by others, including solutions from previous semesters. Make sure that you do not upload your code on github public repositories, as this also constitutes violation of the honor code.