

594 Special Topics: The art and science of Dialogue-Based Systems

January 12, 2016

1 Course Goals

Dialogue-based systems are a hot topic right now, with the popularity of Siri on the iPhone; a variety of personal assistants for Android based on Google Voice, like Robin, Jeannie, or Andy; Cortana, the Windows personal assistant; etc. These systems are grounded in research areas such as speech processing, language processing, and human computer interaction. This course will provide students with scientific foundations in all these areas as they relate to dialogue systems, insight into the many remaining research issues, and hands-on experience in building their own dialogue systems.

2 Course Outline (Tentative)

Week	Topic	Readings
	Weeks 1-7: Foundations	
Week 1	Introduction to dialogue systems	Ch. 1 [Jokinen and McTear, 2009]
Week 2	Tools to support building dialogue systems	
Week 3	Speech recognition	Ch. 9 [Jurafsky and Martin, 2009]
Weeks 4-5	Dialogue State, Dialogue Modeling, Dialogue Acts	Ch. 2 [Jokinen and McTear, 2009]; Sec 24.2.4, 24.2.3 [Jurafsky and Martin, 2009]; Research Papers
Week 6	Speech synthesis	Ch. 8 [Jurafsky and Martin, 2009];
Week 7	Evaluation of dialogue systems	Research Papers
Week 8	Multimodality in dialogue systems	Research Papers
	Weeks 8-14: Applications	Research papers throughout
Week 9	Affect and emotion	
Week 10	Dialogue systems for information access	
Week 11	Human-robot interaction	
Week 12	Dialogue interfaces to educational technology	
Week 13	Applications for Health and Medicine	
Week 14	Commercial dialogue systems	
Week 15	Project Presentations	

3 Course Work

1. Reading, discussions, presentations.
2. Two-three small assignments to try out existing platforms: for example, Luis from Microsoft <https://www.luis.ai/>, Virtual toolkit from USC <https://vhtoolkit.ict.usc.edu/>, etc
3. Course project: build a (toy) dialogue system using one of the available platforms; or, focus on one module of a dialogue system and develop it in depth. This is a **group** project, to be done in groups of 2-3 students

4 Textbook and readings

Recommended: [Jokinen and McTear, 2009, Jurafsky and Martin, 2009]. 80% research papers taken from a variety of sources, including: the Proceedings of SigDIAL, the ACL / ISCA Special Interest group on Discourse and Dialogue, <http://www.sigdial.org/> (ACL= Association for Computational Linguistics; ISCA = International Speech Communication Association)

5 Prerequisite / Corequisite

421 or 521 (or another equivalent NLP class). Corequisite: you can register in this class if you are taking 421 at the same time.

References

- [Jokinen and McTear, 2009] Kristiina Jokinen and Michael McTear. Spoken dialogue systems. *Synthesis Lectures on Human Language Technologies*, 2(1):1–151, 2009.
- [Jurafsky and Martin, 2009] Daniel Jurafsky and James H. Martin. *Speech and Language Processing*. Pearson Education – Prentice Hall, second edition, 2009.