

CS2001: SPOKEN DIALOGUE SYSTEMS

HOMework: DUE IN 1 WEEK, PUT HARDCOPY IN
PROF. LITMAN'S MAILBOX

I: Talk to a SDS

Call two spoken dialogue systems, and discuss their limitations.

- MIT (www.sls.lcs.mit.edu/applications/jupiter.shtml) (research)
 - 1-888-573-TALK
- Amtrak (commercial)
 - 1-800-USA-RAIL

II: Examine a Dialogue

Transcribe the first 40 seconds of:

- www.cs.pitt.edu/~litman/wav/NORMALIZED_genius.wav

“Transcribe” means that you should make a record on paper of what you hear – so that when someone else reads the transcript, they know what went on. Please indicate and explain any special annotations you use to indicate turn taking, speaker overlap, disfluencies, pausing, intonation, etc.

Discuss the issue of what makes an adequate transcription, and what challenges a computer might have in interacting in our example dialogue. For example:

- What makes a sufficient on-paper record of a spoken interaction? What parameters need to be transcribed (the words, the intonation, etc.)?
- Suppose that you had perfect word recognition, what are the most challenging issues in processing the dialogue? Are some of these challenges specific to the domain? What complexity does speech recognition noise introduce? What type of higher-level constructs would be useful to infer and annotate?

Turn in a printed transcription and your discussion.