

The date, time, and location are set (and have not been altered): **Wednesday, May 17, 7-9:30pm, Bard 140** (same room as always). I will be holding my regular Wednesday 10:30am-noon office hours during the next two weeks, and will also be available by appointment.

You may bring to the exam up to five sheets of notes (8"×11", both sides OK), but no other reference material. Calculators are not permitted (but won't be needed, either).

*Please send in the revised versions of pending lecture guides soon* so that they may be posted to the course website, <http://www.cs.cornell.edu/courses/CS630/2006sp/>, as study aids.

Note that the exam is comprehensive, in the following sense: *all material from lectures 1 (1/26/06) onwards, including lectures for which no scribe notes are available, is fair game*<sup>1</sup>. On the other hand, the exam will also be held to *approximately four to six questions* (anticipated to be roughly of the same nature as those on the midterm) to allay time pressure; hence, it is easily inferred that the *coverage will be very heavily (but not completely) weighted towards material not tested on the midterm*.

It bears repeating that what is most of interest is whether one understands how the various methods and models we have discussed were developed. This is because we have introduced fundamental concepts and tools that one simply must be familiar with to understand current research in human-language technologies, and because such understanding should help enable one to develop new methods and models. As has been previously mentioned, an argument can be made that one really understands a concept when one can understand the implications when some assumption or other aspect of the setting is changed. My favorite kind of question is based on this principle.

---

<sup>1</sup>Recall that no scribes were assigned for the last round of lectures to prevent students from having formal course-work during study week.