
Note: It is advisable that you hand in to my mailbox by Nov 5 a brief (1 or 2 paragraph) project proposal, outlining your proposed project topic. I will comment and hand this back to you.

For the final project, each student should select a topic that is relevant to the course. Focus on a problem being addressed by current researchers. Include a description of current research on the problem and an analysis of this research. Your analysis should comment both on strong points of the research and areas where the research falls short. You may choose to focus on one main researcher or to compare and contrast various approaches. If focusing on only one researcher, include reference to related work, demonstrating some understanding of the background.

Once you have described existing work, move forward to include some original ideas on how to advance research on the problem, how to extend or improve on the current approaches. Try to specify your ideas with some precision, rather than simply stating a problem that is worth addressing. I am looking for some depth of understanding and some creativity to emerge, here. Include some general conclusions at the end of the project.

The size of the project is not fixed. Aiming for about 10 (500-word) pages seems reasonable. A project which is too long will be difficult for someone to survive reading. A project which is too short will have problems convincing the marker that sufficient time, energy and thought was devoted to its cause.

It is expected that each student will spend time reading research papers, understanding the research, analyzing the research to see the benefits and drawbacks, glancing at related work (for instance, other papers cited in the bibliography of the paper(s) in focus), thinking of some new ideas.

If the focus of the project is a complex theory, some effort will be devoted to summarizing the research and presenting it in the project. Nonetheless, there must still be some time spent in the project describing the value of the work and its relationship to previous work in the field. Each student must decide the appropriate amount of effort and space to devote to describing the solution(s), analyzing the solution(s) and proposing new ideas.

The topic selected must fall under the general area of intelligent computer interfaces. Any of the subtopics we have addressed in the course (in lectures or suggested by the papers proposed for presentation) can be used as the starting point for finding a project topic. The list of papers available for presentation in the course is intentionally quite broad in scope, to provide you with some insight into a nice range of interesting relevant subtopics. You are not required to use one of these papers as a starting point; any topic that you feel is relevant to the course can be proposed. Many of you may choose to find a topic that is of use towards your own thesis research. I will explicitly approve your topic, if you submit a project proposal by the November 5 deadline. Make sure that your topic is brought into the context of artificial intelligence and that the papers you read and analyze come from this field.

If you prefer, it is possible to do an implementation-oriented final project. For this option, the report will include details on the implementation itself: the design decisions made, some annotated sample output, etc. But, there must also be some component which discusses the significance of the work, mentioning some related work, which convinces me that the student has also done some deeper thinking/reading for the project. In some sense, the description of the implementation takes the place of the original work component of the project and this will likely be a fairly lengthy piece, but there still needs to be some discussion of the problem that motivated the creation of the system.