

CS570 Artificial Intelligence

Individual Research Project Outline

Overview:

An important part of graduate studies is learning how to learn independently, by “working your way into” a new area of intellectual interest to you. In particular, this involves:

1. Identifying an intellectual area of interest. Early discover to learn a bit more about the area beyond its name, to understand if/how it might be relevant to your current research or learning goals.
2. Finding and digesting key projects and publications in that area. This will help you understand the current state of the art – the “research terrain” – in the area. A strong understanding will include not just what has been done, but also how well relevant questions have been answered...and what questions have *not* yet been answered.
3. Demonstrating your understanding of the area. This generally involves a presentation and a survey paper on the area. Again, the key is not just to list what has been done, but to digest and summarize the state of the art, including gaps or areas with promise for further research.
4. Hands-on mastery. For this course, demonstrating understanding will often also require a small software project demonstrating some of the behaviors studied.
5. Novel research in the area. The next step in this “research development” sequence would then involve focusing on some promising subproblem discovered in the area, in an effort to make novel contributions.

Of course, the last step, performing novel research in the area, is beyond the goals of this course; it is something that would normally be reserved for an MS or PhD thesis. Indeed, while the level of exploration we are doing here would likely suffice as a basis for an MS project, PhD-level studies would typically require much more thoughtful and in-depth exploration of a potential area (guided by your thesis advisor!) than we’ll have time for in one semester. Still, the steps remain the same; the goal of this project is to teach you a thoughtful, measured process of intellectual discovery that you will apply again and again during your career.

Individual Research Project: Elements and deliverables

Because every student will choose and pursue a different project, it is expected that the exact nature and content of the deliverables will vary slightly to fit the project. Moreover, as a graduate level project that extends over a semester, there is a bit more leeway in the timeline than might be expected for undergraduate studies. Still, it will be important to pace yourself conscientiously throughout the semester, and we have established a rough timeline of due date to help with that.

The milestones and deliverables that have been set for this project roughly correspond to the steps in the individual research process outlined above. The following sections describe these milestones and associated deliverables in more detail.

[Milestone 1: Selecting a research topic. Your project proposal](#)

The project proposal is the result of some initial investigation of one or more areas of potential interest, to determine what topic you are truly interested in pursuing. Often this will be driven by your broader research interests, e.g., what might be most complementary to your broader thesis topic. The deliverable for this stage is a 1-2 page “project proposal” that roughly corresponds to the following outline:

1. Introduction. An opening paragraph where you introduce the area of interest in general. Your goal here is to introduce the area, describe what it is (goals, how it works), mention some exciting projects, and sketch a picture of what it could bring in the future. You are trying to make your area sound attractive, exciting, worthwhile here.
2. Next, talk about some of the specific interesting aspect in the area. This could be key challenges that have been tackled...or remain as open problem. Or it could be simply some interesting aspects of the area, e.g., underlying computational approaches or solutions that have been used to get ahead. You don't need great detail here (that comes later!), your aim is just to convince us that there are some interesting things to investigate here.
3. Finally, you want to outline your planned work on this project. Obviously, it will center around building up the knowledge for a decent survey paper on the area, but try to be a bit more specific based on what you've learned in your preliminary explorations, e.g., “I'd like to compare and contrast this approach/orientation to that one”, etc. Or maybe bullet out a list of questions that you'd like to be able to answer by the end of the term.
4. Your software project. Outline what you'd like to do as the software element of your project. This could be writing a small demo program the implements a technique found in the research, it could be downloading and configuring an existing system and then doing some explorations with it, whatever. This goal can change as you learn more; the point here is just to indicate that you've thought about it.
5. Conclusion. Every document your ever write needs a conclusion! Just a few sentences that rehash the intro to remind us that the area is hot and exciting, then summary what you plan to do, and what you hope it will yield.
6. Reading list: Append an annotated list of potential sources of information. We're not talking about specific papers/book/websites here (though you could include those), we're talking about places you intend to search for relevant information. “Annotated” means that you list the place, and then add a sentence or so explaining what it is and what you expect to find there.

[Milestone 2: Literature Review and Analysis.](#)

This milestone occurs when you have completed your background research, i.e., you've located, read, and digested enough literature to understand the “contours” of your research area. What has been done, what are the main projects, and (the digestion part) how these could be organized into differing themes or approaches within the area. The deliverable is

a document (probably 4-8 pages) in which you present your findings. As a survey of existing work, this document will become the heart of your final paper for the class, so it's worth doing a good job here. The outline usually goes something like this:

- **Intro.** Here you talk about the area in general; you can likely borrow some text from your project proposal to start things off. After a half-page of introductory material, you will want to introduce your analytic framework, i.e., how you've decided that organize/group the work in the area. So something like "While all work in computer vision is, of course, fundamentally concerned with identifying objects represented within flat digital images, the approaches taken vary considerably, and can be differentiated along several dimensions."...and then you go on to introduce and carefully describe the dimensions. Metaphorically, think of yourself as organizing and giving a guided tour of a museum. You wouldn't just start by saying "ok, and here in room one we have a pile of random stuff...there's this and this and this...ohh look, isn't this nice?". Right? That would be useless! Instead, you would want to have thought about the whole collection, and decided on how it could usefully be organized. Just organizing it by chronological time is usually the least interesting approach. Instead, you'd want to find dimensions of difference that somehow give a new and interesting perspective on it all, e.g., "Here are works produced by process X versus by process Y", or "works influenced by religion contrasted with secular works from the same period". See what I mean? Try to discern dimensions of difference within your targeted area of research that are interesting and meaningful.
- **Ok, now that you've done the hard work of introducing the dimensions of difference that will organize your discussion, you can dive in:** in a sequence of subsections, you discuss each dimension by introducing/describing/citing key projects that characterize that dimension, often including some graphics or screenshots as needed. Note that a single project could be cited several times along the way: it falls in this group on this dimension, and in this other group on that dimension. Your plan here is not to be absolutely exhaustive in just dumping out 50 project across the dimensions. Instead, you are using the projects to illustrate and characterize each of these dimensions and the differences along it. Thus, you might cite five projects, but only actually describe one or two in detail: "Having just examined a set of projects that take pre-defined feature sets as input to their processes, we might ask whether a more dynamic approach to discovering feature sets could provide additional flexibility in novel scenarios. Indeed, a handful of recent projects (five citations) have explored this approach with varying degrees of success. For example, Project X <intro and describe one in detail>. Similarly, project Y <intro and describe another>. Although Project X and Y both takes the same fundamental approach, <comment on detailed differences between the projects>." See what I mean? You taking us on a tour of each dimension, showing us highlight projects, and giving your running analysis.
- **Discussion.** This is where you bring it all together, tell us what it all means in the end. This is "your summary take" on the status of the area...the real value in this paper! In the end, you'd like to end up in a place where your N dimensions of difference have nicely characterized the existing research place and...tada!...it's obvious that there is this one corner of that research space that hasn't received a lot

of attention. And that would then motivate your own research. Of course, we won't be going quite that far in this course project, but understanding this end goal should help you towards a more interesting and effective literature analysis. What you should therefore shoot for in this wrap-up discussion is just to summarize what's been done, what has been successful and what questions appear to remain open. This is your opinion. There is no right or wrong here...you give your analysis of the state of the area – where it's been, where it's at, and where it's going – based, of course, on the discussion of existing work you've presented.

At the end of this document you should, of course, have your bibliography, generated based on the citations that you've made in the document. You are **required** (as a vital learning experience) to use a bibliography management program to organize your sources and citations for this project. You may use any of the many systems that exist; if you have no experience or no preference, **I can highly recommend Zotero** (zotero.org), which works as a web app, a stand-alone application for Mac, and has helpful plugins for several popular word processing applications. And it's free! Please attach a screenshot of the main library window of whatever bibliography manager you choose as the last page of your document.

Milestone 3: Project proposal and presentation

This milestone is just a small step from your literature review...the next natural step, really. You use what you have learned about the area in your lit review to propose a project for yourself. In particular, this is some hands-on project where you write up a short demo program that embodies/explores/demo's the techniques you've learned about. Or you propose to download an existing program or system and propose some experiments that you might do with it to understand it better. You get the picture: something interesting that will make this more than you just reading a bunch of papers.

The deliverables here are:

- a short “addendum” to your lit review, where you build on the discussion to propose the work/experiments that you have in mind.
- A presentation on your work so far to the class at large. Shoot for 10-15 minutes maximum. You intro the area, outline it along the lines of your lit review (throw in some cool graphics and screen shots to illustrate and keep the audience stimulated), and then tell us what you plan to do for the hands-on part of your project.

Milestone 4: Draft paper.

This milestone is near the very end of the term. You will have completed your hands-on part and will now bring all of your work together into your final paper. Fortunately you have produced many of the pieces already: a good intro to the area, a detailed review of current state of the field, proposal for an interesting hands-on exploration...and now you add on the results of that exploration. Thus, the outline for your final paper is something like:

1. Title/Head: Title, your name, the course, date
2. Abstract: about 300-words that succinctly summarize your project and its outcomes. Goes right below the title/header.

3. Introduction: Starts several spaces below the abstract. Intros the area, intros the key questions/issues that you'll explore, hints at the outcomes. Ends by outlining what the upcoming sections will discuss.
4. Review of existing work: basically your lit review.
5. Discussion: basically the discussion section of your lit review, modified so that it gets around to motivating and introducing your hands-on project
6. Project work: Here you intro and describe your project in great detail, probably across several subsections: intro, methods (where you describe exactly what you did), outcomes (where you present the outcomes of your tests), analysis and discussion (where you discuss your findings and what they showed).
7. Discussion: Here you bring your project and its exploration together with the existing work somehow. What did you learn that was new, how does it change how you see existing work and the area in general.
8. Conclusions and Future Work: Here you close the whole thing off. Start by drawing on your intro to remind us of the area and how exciting it is. Then go on to say some summary items about the existing work, followed by a summary of what your project aimed at and achieved. End by speculating on implications of what you have learned, the most productive directions for future work, predictions of where work in this area will lead (exciting future visions) and when those might arrive.

You will present your draft paper *at least two weeks* before the end of the term, leaving plenty of time for comments and corrections. It is absolutely vital that your draft be a *complete, penultimate draft!!* That means a highly-polished and completely formatted document, complete to the best of your ability. Not "an outline". Not a "rough draft". This is an important lesson to learn! It is difficult to give high-quality meaningful feedback on partially complete products; these might be worth an oral discussion in office hours, but you are wasting your time and theirs by giving any reviewer a rough product!

[Milestone 5: Final paper and presentation](#)

This speaks for itself and is the culmination of the course. You will need to turn in your professionally bound and presented final paper, and give a second presentation on your project. The presentation will be similar to the first one: about 15 minutes, but this time compressing the review of existing work to leave more time to discuss your project and its outcomes.

Timeline for Milestones:

As stated earlier, the timeline for this graduate project is somewhat flexible, without hard down-to-the-minute deadlines. They are roughly:

- Milestone 1: Project proposal. By end of week three
- Milestone 2: Completed literature review. Before Spring Break.
- Milestone 3: Project Proposal and Presentation: Week after Spring break
- Milestone 4: Draft paper. About mid-April, two weeks before finals week.

- Milestone 5: Final presentation and paper. The presentation will be scheduled as shown on the syllabus course schedule in the final two weeks of the term. The final paper is due at the scheduled time of the final exam, at the latest.