

Project notes

20260827-1145

These are some rules and guidelines that apply to all of the projects in this course.

Project domains

Every project handout starts with an overview that sets up the problem. It gives you information about the problem domain that you might not be familiar with, and it sets the goals and motivations for the project. In general, the handout will go out before I've covered all the material you'll need to *complete* the project, in part because I will use the ongoing project as a running example in class—which means you'll need to read it as soon as it's assigned, since I'll be talking about it right away.

The Two Biggest Problems

I'm foregrounding this so you are certain to see it. These will be explained in more detail below, but if you submit

- a program that fails to compile, or immediately crashes on all inputs;
or
- with only zero or one `git` commits for a final handin with a nontrivial amount of work;

then your **entire project submission grade will be zero**. This fate is easy to avoid! Exercise care!

Checkpoints

After the overview of the project topic, the next portions of each project handout will be checkpoint assignments. The first, which I refer to as “prep work”, generally won't work as much with the core conceptual content of the project, but will give you a head start on things like input formats and data structures that you'll need. Which means you can (i.e. should) start

on it right away, even before we’ve finished covering the related content in class. If you happen to finish the prep work before its deadline and want to work forward, that’s ok but *commit it* (and tag the commit as “**prep**”) before you start the other work.

The second checkpoint is the “design work”, which is to be written on paper (or the equivalent) and brought to class: diagrams, test cases, and other planning work that will help you understand how the program you’re writing will need to fit together. I am looking for good-faith best-effort work here, not perfection; we’ll talk about everybody’s design in class and improve them, so that you’ll be well-prepared to really forge ahead on the main body of the work.

Self-analysis: commit messages and correctness justifications

As you move from the prep and design work into the main body of the project, you should have a plan for how to move from one version that works a little bit to another version that works a little more (rather than trying to type the whole thing in all at once and then debug it). To that end, I’m requiring you to `git commit` your work regularly, at points where you’ve completed an identifiable chunk of the work. That doesn’t mean the new piece has to be fully working, just that you’ve made some progress (though it’s a good idea to at least make sure it compiles/runs before committing).

Each time you commit, it is important (both for me to see and to develop good professional practice) to include a **descriptive commit message**. The message can be brief and doesn’t need to be “full sentences” but should identify what that commit’s changes implemented or fixed or documented or whatever.

The other part of the self-analysis here—and also both for me and to develop good professional practice—is to **justify why you think the program works**. This is not a “proof” in the mathematical sense, but some justification is important, whether code is written from scratch, pulled from a website, borrowed from an older assignment, or even (especially!) generated by AI, to know why you believe that it accomplishes the goal. Preferably, it should be organised by rubric line and point to what part of the code, or what test case, or what other evidence you have that shows you’ve met the criteria for the rubric point. If your project manager or boss asked why you

thought you were done, what would you say?

I recommend you keep the justifications more-or-less in sync with the code development itself, although they can be in separate commits. If you follow an agile development style (my preference), you'll write test cases before the associated code is developed; but you can also follow a code-first style, and write test cases and justifications after the fact, and still get full credit on that front.

This all does a few things. First of all, this kind of self-reflection seems to make learning “stick” better; second, it's structured in tandem with my AI policies to make it possible to use and experiment with AI assistance but also make it hard to just blindly hand off all the work to the AI.

Note, though: in order to make all of this work, I need you to commit fairly frequently. If you **only commit a couple times** for work that is substantially beyond the prep work, **you will get zero points** for the commit messages, the justifications, and the code itself.

Generative AI

You have the option of not using a large-language-model (LLM) generative AI system on the projects.

However, at least *some* use of generative AI will be permitted on each project, and one of them will strongly encourage it. If you're at all interested in playing with the LLMs, I generally encourage it—now's the time—but there are a few constraints, laid out below. (Check each project handout to see if there are any adjustments to these rules; this is still somewhat experimental and they may evolve a little.)

- You **MUST NOT** just paste in the project handout to get a final program to hand in. (That would cause problems with the multiple-commit rule anyway.)
- In fact, you **SHOULD NOT** paste in text from the handout at all, though you can paraphrase it when you're writing your AI prompt. Pasting someone else's work into an LLM is unethical and a form of copyright violation, since you're giving stuff you don't own to the AI company (which they'll then use in training their models).

- Any code that is entirely, largely, or even just partially AI-generated MUST be marked off with comments that indicate what parts were AI generated, which AI (e.g. Claude, Gemini, etc) provided them, and the prompts used, via the “share” link on the chat or edit session. If your AI use is a single long chat window, a single valid share link in the readme would be fine, but if your interactions are more targeted you’ll want the share link in a code comment in the vicinity of the code that was AI-influenced.
- Remember that you should be working one piece at a time, so your AI requests SHOULD be sort of focused, and talking about how the AI helped you SHOULD be part of your self-analysis in the commit messages.

Code submission

The program code you submit can in general be in any language that runs on the lab systems (although Project 1 is restricted to C++, Java, or Python due to the starter code). It’s a good idea to make sure the code compiles/runs for each of the intermediate submissions, but it is vitally important that the final submission [compiles and] runs.

If your **code fails to compile** or immediately crashes on all inputs, the **submitted code will receive a zero**.

That means if you test your code and then “just add a quick comment” or some other quick change... you should test it again before handing in.

Documentation

Your project must have documentation. It’s nice if it’s in a single file called `README.md` or `README.txt` but if it’s spread across multiple files, just make sure I can find it all. The things that absolutely must be in the documentation are:

- how to compile/make/setup your program (if relevant)
- how to run your program
- how to run your tests (if relevant)

with real, concrete things that can be typed on the command line (and preferably, with those things on lines by themselves in the documentation, for easy pasting).

Once you’ve got documentation, though, that’s a good place to put anything else you want me to know about. The required documentation and the correctness justifications can all go in the same file (but again, if they’re in different files, just make sure I can find them).

Project followup

A short time after each project is due, and I’ve given you at least initial feedback, I will give an in-class followup assessment on the project. The point of the project is to learn some things; and the followup will then assess if you actually did. Is it an exam? Kind of, but with a very practical style and very narrow focus; my intent is that simply *doing the project* is all you should need to do to “study for” the followup. I aim to calibrate it so that if you legitimately work through the project, even with a bit of human or AI help, you’ll get roughly the same grade on the followup as on the submitted code (but if someone were to blindly submit code produced by AI or another human, without understanding it, they’d have a hard time on the followup).

The followups for Projects 1 and 3 will be standalone and take up only a portion of the class day; the followups for Project 2 and 4 will be given together with the midterm and final exam respectively.

Timeline

As a general rule, projects will be initially assigned on a Wednesday and due three weeks later. (Project 1 is slightly shorter.) The prep work will be due after the first week; they’re due **by 8pm** on the due date. (This gives you a chance to ask questions in class and go back and fix things.)

The design work will generally be due on the following Monday, sometimes alongside other homework. Like the homework, the design work is due at the start of class, and should be brought with you to class.

The final handin for a project is due also **by 8pm** on the due date.

Grading

Project grades are out of 150 total possible points.

The **prep assignment** is worth 10 points. You get 10 points if it [compiles,] runs, and it does at least approximately what it's supposed to. You get 5 points if it [compiles,] runs, and does something relevant to the task. You get 0 points if it doesn't compile, if it immediately crashes, or if I can't easily figure out how to make it work. (So, you should include some documentation.) If your prep work gets a zero, work with me ASAP and you can get back up to 5 of the 10 points. (The purpose is to get you started, not to create a penalty!)

The **design work** is worth 10 points, and is graded for completion rather than correctness. If you're absent on the design day and haven't made some sort of advance arrangements, that will normally be a zero on the design work.

The **documentation** is worth 10 points if there's enough documentation that I can *easily* see what to do with your code and what your code does. The more I have to work to figure out what your code even does, the lower this gets. If you have no documentation that I can find, you get zero points for documentation.

The **commit messages** and **correctness justifications** are each worth 10 points, graded on the same 10/7/4/0 scale as homework problems; submissions that are even slightly beyond the level prep work are eligible for full credit on progress reports and "proof", as long as they're clear and sensible for the project work that was done. (For these points, you're not evaluated on whether you have made *good* design decisions (though I hope you do), but rather on how you explain why you made the choices.)

The **code submission** is worth 20 project-specific points, and the **project followup** is worth 80.

That adds up as follows:

10	prep work
10	design work
10	documentation
10	commit messages
10	correctness justifications
20	submitted code
80	project followup
150	

The goal will be to calibrate the rubrics so that if you get all of the first 50 points and do roughly as well on the followup as on your submitted code, you'd be getting

- a low D for not making it past the checkpoint
- a middle C for a program that does some relevant thing
- a high B for a program that avoids major pitfalls and correctly solves at least a simple version of the project problem
- a high A or full credit for a program that really gets into the conceptually interesting parts of the problem.

I'll put rubric specifics in each project handout. Note that I'm not planning to, in general, read your code or comment on it after you hand in—if you want substantive feedback on your code, see me while you're working on it. I will assign scores based on reading your docs and running your code, but if you're serious about the self-evaluation in your correctness justifications, you should be able to predict your score within a few points based on the rubric I give you in advance.

Each project will be worth 15% of the final grade.

Collaboration and citation

The projects in this course are collaborative, meaning you can (and should!) discuss your ideas with other students, but the code you write needs to be your own. See my collaboration policy for many examples of what is and isn't acceptable.

The nature of this course and these projects also makes it more reasonable to look on the internet for explanations, algorithms, and even code examples. That's great! But it also makes it more important than ever that you cite your sources. Make sure you read my citation policy for when and how to do this for programs.

Handing in

There is a `handin` command on the lab machines that you will use to hand in all your work. (You still need to use `git` to start your projects—either using `clone` or `init` as appropriate—but submitting will involve handing in your entire repo, which will let me see your commit history.) If you use your own computer for development, transfer your files to the lab machines (including the `.git` directory!), *verify that they still work*, and then hand them in by typing, if you're outside the directory,

```
handin cmsc262 proj1 dirname/
```

or if you're inside the directory, just

```
handin cmsc262 proj1 .
```

to hand in an entire directory of files, replacing `proj1` with the actual name of the assignment. It's important to hand in the whole directory so I get the git history.

Don't forget to include your documentation, and make sure it's easy for me to find it. `README.md` is a great name for a file of documentation that you want me to be able to find.

Use the same project name for both the prep work and the final handin.

These instructions have been updated