

# COMP 430: Language Design and Compilers

## Spring 2026

**Instructor:** Kyle Dewey ([kyle.dewey@csun.edu](mailto:kyle.dewey@csun.edu))

**Course Web Page:** <https://kyledewey.github.io/comp430-spring26/>

**Office:** JD 4419

### Course Description (From the Catalog)

Examination of the issues involved in the design and subsequent implementation of programming languages. Considerations of implementation difficulties, including various features in a programming language. Tools and techniques to facilitate both the processing of programming languages and the building of programming processors.

### Learning Objectives

Successful students will be able to:

- Design a programming language with (at least):
  - Concrete and abstract syntax
  - Statically-checked types
  - Expressions
  - Subroutines
  - Mechanisms for computation abstraction
- Implement at least the majority of a compiler for the designed language, with:
  - A tokenizer
  - A parser
  - At least one of:
    - A typechecker / static semantic analyzer
    - A code generator

### Course Motivation

A common question with compilers courses: *When am I ever going to need to implement a compiler?* I'll answer that upfront: probably never. However, consider some related questions:

- *When am I going to need to reuse my own code?* - All the time.
- *When will I need to understand how a language works?* - All the time.
- *When will I need to work in a team?* - All the time.
- *When will I need to understand why a language was designed in a certain way?* - At least as often as you evaluate new programming languages. Languages rise and fall frequently, and it's important to know what's worth your time and what isn't.

In this class, you'll incrementally build on your own compiler, likely starting from scratch. You'll live with your coding decisions (good or bad), and will likely have to revisit them. You'll gain a better understanding of how languages work, and understand what is involved with language design. Perhaps most importantly, you'll learn that **languages**

**aren't magic:** you can implement your own, and they are surprisingly straightforward to implement once you understand the basics.

## What this Course Is and Is Not

This is a project-based, implementation-oriented course. Our focus will be on modern compiler development. Grading is based on a series of assignments which will allow you to incrementally implement your compiler. While these assignments accomplish the same high-level tasks between all students, exactly what these assignments are depends on your particular language (e.g., the assignment implementing typechecking will differ for Teams A and B, since both teams implement different languages).

Because this is a project-based course, there are no exams, though the final exam slot may be used for team presentations. Your compiler code serves as evidence of understanding class concepts. Additionally, since the focus is on modern compilers, I'm intentionally **not** planning to cover certain topics in-depth, namely:

- Efficient parsing (e.g., LL, LR, flex, bison, etc.). We won't need very efficient parsers for what we plan to do. Some modern compilers forego efficient parsing entirely. Many of these algorithms and tools were developed at a time when memory and computing resources were scarce, but this is no longer true.
- Very low-level concerns like register allocation or machine-specific optimizations. It's rare nowadays to compile directly to machine code, given existing general-purpose low-level targets like LLVM, the JVM, and CIL. Additionally, these topics are completely irrelevant if everyone selects high-level compilation targets for their projects.

## Textbook

No textbooks are required. That said, the following books may be of interest to you:

- Programming Language Pragmatics, by Michael Scott. Discusses a variety of programming language features from a variety of paradigms, and certain key interactions. Good for language design, though light on compilation.
- Compilers: Principles, Techniques, and Tools, by Alfred Aho, Monica Lam, Ravi Sethi, and Jeffery Ullman. AKA the "Dragon Book" (there is a dragon on the cover). Classic, commonly-used textbook, which discusses low-level compiler details. Heavy focus on parsing, and it's better-suited as a reference than an introduction.
- Modern Compiler Implementation in Java/ML/C, by Andrew Appel. These are each separate books, but they all cover the same content, just with a different implementation language. Good introduction to compilers, with lots of example code. Not very general, but will get you jumping in head-first.
- Engineering a Compiler, by Keith Cooper and Linda Torczon. Great resource on optimizations and low-level concerns. Not a gentle introduction to compilers, and limited information about things like types.
- Types and Programming Languages, by Benjamin Pierce. Specifically about typechecking and *type theory*, specifically as it relates to a number of common programming language features. Highly specialized and math-heavy; can be intimidating, and likely not useful unless your language has some advanced type features.

## Graded Components

All graded components tie directly to your project. Because of this, exactly what you need to do for each component is dependent on your language design. Similarly, how difficult each component will be is very dependent on your language design. As a result, when you submit your language design proposal, you will suggest how much each component should be worth, for a total of 100%. I'll provide feedback on this and potentially adjust these percentages in a way that reflects how much work each one is expected to be. This is why there is a minimum and maximum range for project-specific components.

Component	(Typical) Minimum Percentage	(Typical) Maximum Percentage
Language Design Proposal		5%
Lexer	5%	10%
Parser	10%	20%
Typechecker	15%	50%
Code Generator	15%	50%
Project Report		10%

For each of the lexer, parser, typechecker, and code generator, you will submit each component **twice** - once for an initial round of feedback and a tentative grade, and a second time for a final grade. **If the tentative grade for the initial round of feedback is below 25%, your final grade for the component will automatically be deducted by 50%.** This means you cannot simply skip the initial feedback round and wait until the last minute to submit everything.

For example, say your parser proposal had your parser worth 10% of your final grade. The deadline for the initial round of parser feedback comes up, but you only implement a tiny portion of the parser, which receives only a 15%. Since  $15 < 25$ , your final submission for the parser would be capped at 50% of what it should be worth, meaning you'll only get 5% towards your final class grade (50% of 10 is 5).

## How Grading Is Performed

For the lexer, parser, typechecker, and code generator, students will effectively grade themselves using industry-standard techniques. Specifically, for these components, students are expected to:

- Implement all necessary code
- Write unit tests testing the written code, including relevant assertions

- Run the unit tests, and generate code coverage information showing what parts of the code are covered by the tests. This coverage report will be provided along with your code submission.

On my end, I will:

- Check that the code coverage is relatively high (at least 80%, ideally 90% or above)
- Check that the assertions in the test suite are present, logical, and actually check meaningful things
- Check that all your tests collectively handle all the behavior planned from your language design proposal

With this in mind, the tests you write are just as important as your implementation code.

Students are not expected to understand exactly how to do this ahead of time; we will cover unit testing and code coverage in class.

Code cleanliness is not explicitly part of the grade. However, if I see an issue with a component (e.g., a failing test, an unhelpful assertion, a missing feature), cleaner code will be easier for me to read through and understand what the root cause of the problem is. If I can understand the problem, partial credit is more likely to be awarded.

## Final Grades

Plus/minus grading is used, according to the scale below:

If your score is >=...	...you will receive...
92.5	A
89.5	A-
86.5	B+
82.5	B
79.5	B-
76.5	C+
72.5	C
69.5	C-
66.5	D+
62.5	D
59.5	D-
0	F

## **Late Policy**

Unless prior arrangements have been made or there are extenuating circumstances, for each day a component is late, it will be deducted by 10%. The determination of extenuating circumstances is at my discretion, and is intended for major life events outside of your control. Assignments that are submitted more than 10 days late will not receive any credit. The only exception is the final deadline for everything, where **everything must be submitted by the deadline (don't miss this one!).**

## **Special Note About Teams**

The class is designed for projects to be done in teams. Occasionally I will permit students to work on a project alone, but only if they have prior experience working with programming language implementations. Historically, those working alone without prior experience struggle significantly, and usually deliver low-quality results.

To ensure a fair distribution of work among teams, peer evaluations will be used. Additionally, I will count the number of edits made to code per student per graded component (we will use revision control via GitHub, making this easy to do). If your number of edits is significantly lower than everyone else's without explanation, I will penalize you (and only you, not your team), accordingly. If you think a component is completely implemented already, code can **always** be added in the form of tests. If you are worried that you might not be contributing enough, talk to me and we can try to figure out a way to better distribute work.

## **Plagiarism and Academic Honesty**

You are permitted to collaborate as much as you'd like. If, however, you use code from somewhere else, you must **cite** it. Any violations can result in a failing grade for the assignment, or potentially failing the course for egregious cases. A report will also be made to the Dean of Academic Affairs. Students who repeatedly violate this policy across multiple courses may be suspended or even expelled. From an industry perspective, blindly taking code from other sources can have severe consequences, as this could violate licensing agreements and open a company up to a lawsuit.

## **LLM Policy**

LLM-based tools like ChatGPT are discouraged, though not outright disallowed. It can be easy to over-rely on such tools without realizing it, to the point where the tool does all the work, with zero understanding of what it produces.

Keep in mind that LLMs fundamentally just generate some text that statistically is likely to follow whatever the prompt is based on the training set, with no understanding of what the text says. Any output from an LLM **must** therefore be verified to ensure it is correct. This makes them useful for pointing you in some direction if you are lost, but there is no guarantee that the direction is even relevant.

Specific to COMP 430, LLMs can be good for getting an initial start on things, but they have a tendency to produce code that isn't quite right, and structured in weird ways that

are difficult to build upon. These issues start to compound on each other as more code is added, to the point where everything becomes a total mess requiring practically the entire project to be rewritten.

### **Communication**

In general, any questions should be made through Canvas. You can also email me, though I'm usually much faster to respond to Canvas than my general email. Office hours or talking with me after class are usually the fastest and most detailed way to communicate with me.

### **Class Feedback**

I am open to any questions / comments / concerns / complaints you have about the class. If there is something relevant you want covered, I can push to make this happen. I operate off of your feedback, and no feedback tells me “everything is ok”.

### **Class Structure and Content**

The initial course content is relevant to everyone, so this will be delivered via typical synchronous lectures. You'll then submit a language design proposal defining your project. Soon afterwards, exactly which parts of course content are relevant to you will depend on your language design proposal. For example, someone implementing an object-oriented language probably won't find content on higher-order functions to be relevant, and someone implementing a language that compiles to JavaScript won't find content specific to an assembly-level target relevant.

For these reasons, content that could be project-specific may be delivered in an asynchronous format, based on prerecorded lectures. We will still meet synchronously, but the class time may be used like a lab wherein you can collaborate or ask me questions. This structure was adopted based on student feedback from early versions of the course.

### **Class Schedule and Component Due Dates (Subject to Change):**

Items in **bold** are tentative deadlines.

<b>Week</b>	<b>Monday</b>	<b>Wednesday</b>	<b>Friday</b>
1	1/19: MLK Jr. Day; Campus Closed	1/21: Introduction, project information	1/23:
2	1/26: Project information, team formation	1/28: Project information, grammars, tokenization	1/30:
3	2/2: Grammars, tokeniation	2/4: Tokenization	2/6:

Week	Monday	Wednesday	Friday
4	2/9: Tokenization, ASTs, architecting around ASTs	2/11: Parsing	2/13: <b>Language Design Proposal</b>
5	2/16: Parsing	2/18: Parsing	2/20:
6	2/23: Parsing	2/25: Type theory basics, typechecking	2/27: <b>Lexer - Initial Feedback</b>
7	3/2: Typechecking	3/4: Typechecking	3/6:
8	3/9: Typechecking	3/11: Typechecking	3/13: <b>Parser - Initial Feedback</b>
9	3/16: Spring Recess (no class)	3/18: Spring Recess (no class)	3/20:
10	3/23: Code generation	3/25: Code generation	3/27:
11	3/30: code generation	4/1: Code generation	4/3:
12	4/6: Open lab	4/8: Open lab	4/10: <b>Typechecker - Initial Feedback</b>
13	4/13: Open lab	4/15: Open lab	4/17:
14	4/20: Open lab	4/22: Open lab	4/24:
15	4/27: open lab	4/29: Open lab	4/1:
16	5/4: Open lab	5/6: Open lab	5/8: <b>Code Generator - Initial Feedback</b>
17	5/11: Classes over	5/13: Classes over	5/15: Classes over; <b>Final Version of Whole Compiler; Project Report</b>