

## CS106X Course Syllabus

---

I don't even try to promise a day-by-day lecture schedule, since even the most disciplined and organized of instructors have a difficult time staying on track, and I'm far outside the set of disciplined and organized instructors. I do, however, think it's reasonable to give a sense of what topics we'll be covering, and how much time we'll be dedicating to each of them.

Here goes:

- Week 1: **Basic C++ syntax, control idioms, program structures, strings, and libraries.** *Beyond the introductory remarks, the week will feel like a transition course designed to get Java programmers to emulate their craft using C++ instead. The overarching lesson here is that C++, like Java and most other programming languages, has **ints**, **floats**, **for** loops, **switch** statements, **return** values, classes, and methods.*
- Week 2: **Templates, abstract data types, containers, vectors, stacks, queues, sets, maps, scanners, and lexicons.** *You're already familiar with templates and containers, even if you didn't call them that when you took AP Java or CS106A. Containers are data structures that are designed to store other data structures, and you already have plenty of practice with them—specifically, the Java **ArrayList** and the **HashMap**. C++ affords its own versions of the **ArrayList** and the **HashMap** (the **Vector** and the **Map**, respectively). We'll invest a lot of energy teaching the new container classes and the metaphors you should subscribe to when you code with them (stacks of cafeteria plates, a queue of customers at Harrods in London, etc.)*
- Week 3: **Recursion, drawing examples from mathematics, graphics, and human language.** *Recursion didn't originate with computer science or programming—simple mathematics gets all the credit there. But virtually all modern programming languages support programmatic recursion, which is a function's ability to either directly or eventually call itself. Many very practical programming problems exhibit recursive structure, and the ability to code using recursion makes it easier to exploit that structure.*
- Week 4: **Advanced recursion, recursive backtracking, memoization.** *Recursion is a difficult enough topic for newcomers—even those as talented and motivated as the typical CS106X student—that I want to spend a good stretch of time providing increasingly more sophisticated examples from a variety of application domains. Recursive backtracking is a form of recursion where you recognize one particular recursive call turned out to be a dead end of sorts, but*

that some other recursive call could pan out with something useful.

Memoization is a technique used when there are a finite number of recursive sub-problems that will otherwise be repeatedly called an exponential number of times unless some form of caching is used.

Week 5: **Memory, Memory Addresses, Pointers, and Dynamic Memory Allocation.**

One of the most powerful features of the C++ language—and admittedly, the feature that makes C++ a terribly difficult language to master—is that it affords you the ability to share the physical memory locations of your data with other functions and to manually manage dynamically allocated computer memory. We'll discuss a new type of variable—the pointer—that's designed to store the location of other figures in memory, whether those figures are integers, arrays, **ArrayLists**, records, or even other pointers. For the next several weeks, we focus less on application-focused programming and algorithms and more so on the use of pointers and advanced memory management to implement all of the ADTs you've come to use and appreciate since Week 2.

Week 6: **Linked Lists and Their Variations.** You're all very familiar with the arrays and the **Vector**—so familiar, in fact, that you should recognize that insertion and deletion from the front can be a very time consuming operation if the size of the array or **Vector** is large. The **Queue's dequeue** operation, which is supposed to be fast regardless of queue length, couldn't possibly be backed by anything that's truly array-like. The linked list makes a few sacrifices in the name of fast insertion at and deletion from both the front and the back of the sequence. We'll see the linked list and a few of its variations as the most basic linked structure in the series of linked structures we learn about over the remainder of the course, and we'll understand why it (and not anything array-like) backs the **Queue**.

Week 7: **Hash Tables and Trees.** Once you're fluent in the construction and manipulation of the basic linked list, we'll be in a position to build and talk about more advanced linked structures like hash tables, binary search trees, tries, and skip lists. The hash table is the backbone of your **Map** container, the binary search tree is more or less the core of your **Set**, and the trie is a simplified version of what backs your **Lexicon**. (The skip list is a fairly recent randomized data structure that could back the **Set** if we opted for it over the BST.) We could spend 40 lectures talking about data structures. I'll try to get as much of those 40 lectures into the two or three I have.

Week 8: **Graphs and Fundamental Graph Algorithms.** The graph is the Holy Grail of all linked data structures, allowing us to model generic peer-to-peer systems like road, rail, and airline systems, computer intranets, social networks, and the WWW. There are a number of fascinating and fairly well-understood graph algorithms (Dijkstra's shortest path algorithm is the most important we'll study), as well as a number of other algorithms that we're not 100%

*convinced are the most efficient ones possible. We'll study as many of them as time permits, and without stealing the thunder of later theory classes, explain why some algorithms appear to be the best we can do even though they're exceptionally slow for large graphs.*

Week 9: **C++ Interfaces, Inheritance, and Class Hierarchies.** *There are a good number of scenarios where multiple classes are intentionally implemented to the same interface. Inheritance is a unique form of sub-typing and code sharing that recognizes common implementation patterns across multiple classes and works to unify them so their public interfaces overlap as much as possible. You've already seen and benefited from inheritance to some degree if you've done any significant coding in Java, as all Java classes extend the core `Object` class, and therefore respond to a small core of methods like `equals` and `hashCode`. (And `KarelProgram`, `GraphicsProgram`, and `ConsoleProgram` all extended `Program` in CS106A.) We'll extend that basic understanding and construct collections of related classes that exhibit even more aggressive sharing of interface and implementation, and I'll demonstrate how fundamental inheritance is to large, scalable, object-oriented systems.*

Week 10: **Modern Programming Languages.** C++ is a great systems language, and it's arguably the de facto standard for implementing core OS devices and services like device drivers, process control, synchronization, and compilers. But more and more applications are being built in younger languages like Python, PHP, JavaScript, Ruby, and JavaScript. I'll select one or two of these and present a cursory introduction to them over the course of three lightweight lectures, if for no other reason than to familiarize you with the syntax and demonstrate that programming and software development is very much the same regardless of the programming language you speak. (In all likelihood, I'll have at least one guest lecturer this week: a colleague of mine at Facebook is a JavaScript expert and has guest lectured in the past, and he understands JavaScript better than most browsers do.)