

syllabus for computing for poets comp 131

Instructor: **Mark LeBlanc** SC-1322 508.286.3970
mleblanc@wheatoncollege.edu Hours: T 11-12,
<http://cs.wheatoncollege.edu/mleblanc> F 10:30-11:30 or by appt

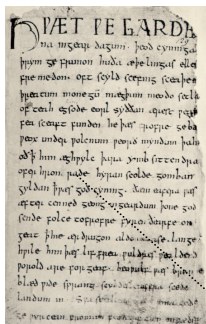
Meeting Times: Mon-Wed 2:00-3:20pm, Room SC 1349, Science Center

The use of computers to manage the storage and retrieval of written texts creates new opportunities for scholars of ancient and other written works. Recent advances in computer software, hypertext, and database methodologies have made it possible to ask novel questions about a poem, a story, a trilogy, or an entire corpus. This course exposes you to leading markup languages (HTML, CSS, XML) and teaches computer programming as a vehicle to explore and “data mine” digitized texts. Programming facilitates top-down thinking and practice with computational thinking skills such as problem decomposition, algorithmic thinking, and experimental design. Programming on and with texts introduces students to rich new areas of scholarship including stylometry and authorship attribution. Prerequisites: A love of the written (and digital) word; no previous computer programming experience is required.



<poem>Three Rings for the Elven-kings under the sky,
Seven for the Dwarf-lords in their halls of stone,
Nine for Mortal Men doomed to die,
One for the Dark Lord on his dark throne
In the Land of Mordor where the Shadows lie.
One Ring to rule them all, One Ring to find them,
One Ring to bring them all
and in the darkness bind them
In the Land of Mordor where the Shadows lie.</poem>
The Fellowship of the Ring, JRRT

Using computers to analyze digitized texts is an exciting new area of research; it enables new forms of both a “close reading” of a single text and a “distance reading” of many texts. Although many tools exist for working *with* texts, what do you do when the tools you have at your disposal cannot answer *your* questions or expect data in another format? In this course, you will learn to write computer programs (also called “software” or “scripts”) that can morph your data and answer your original questions.



The programming language we will learn is named Python – a modern, accessible, yet powerful language when computing with and on texts.



```
array = line.split()
for word in array:
    # CASE 1: working on poetry
    if word in dictionary:
        counts[word] = counts[word]+1
```


poetry in the Anglo-Saxon period (Cædmon, Cynewulf and King Alfred), and there are various problems with linking these names (much less their biographies) with more than a very few specific poems. Although a few prose texts are by known authors (particularly the homilists Ælfric and Wulfstan), even the majority of the prose is anonymous. Thus for years scholars of Old English have struggled to divine relationships between texts based on vocabulary, meter, and style. These results have been at best contentious and at worst completely unsuccessful. As we proceed to learn more and more scripting in this course, we will set up and run experiments using the entire Anglo-Saxon corpus. Some of the questions you could ask may never have been asked before. For J.R.R. Tolkien, we will design and execute an experiment across multiple texts, e.g., the Lord of the Rings trilogy ... but what about the *Silmarillion*? and *The Hobbit*? and ...

We will learn to combine Python programs, comma-separated output from those programs, and Excel spreadsheets in a manner very similar to the way your instructors do when in their Lexomics Research Group (see <http://lexomics.wheatoncollege.edu>). **A major “take home story” for this course is to learn ways to view and analyze texts in an entirely new way.** By leaning on computing, new methods of analysis (that you could *not* do by hand) will now be at your disposal.


In addition to programming in Python and working in Excel, HTML and XML, we will also study how computers store individual characters, including the traditional (English-only) ASCII character code and the international standard called Unicode.

At this point, you may be asking: “Why Python (as a choice of programming language)?” Good question (and it is one your instructor has considered at length). Python is an excellent choice for your first programming language. The interpreter environment encourages you to experiment. Note however, that Python is an industrial-strength programming environment. It is used increasingly in industry in a broad range of areas including but not limited to web sites, data and text mining, and scientific computing. We recently used Python in our Artificial Intelligence (AI) course.

NOTE: This course is but an introduction to using computing to explore digitized texts. Computers allow us to study texts in exciting new ways that we could not otherwise do; however, as we'll discuss at length, **we are wise if we keep in mind what computers cannot do.** The following quotes can help us (1) stay humble and (2) stay focused.

*“As students of a powerful new form of scholarship, we have much to offer.
We do ourselves no justice when we forget that the quantifiable features
we deal in are but the shadow of a shadow.”*

John Burrows, *Computers and the Humanities*, v37, 2003, p30.



*“The onus of competency, clarity, and completeness is on the practitioner.
The researcher must document and make clear every step of the way.
No smoke and mirrors, no hocus-pocus, no ‘trust me on this.’ ”*

Joseph Rudman, *Computers and the Humanities*, v31, 1998, p353.

In computer science, if you are almost correct you are a liability.

Fred Kollett (1941-1997), *MathCS*, Wheaton College, Norton, MA

Digital Humanities

Computing and the Humanities, eh? If you are reading this syllabus, I'll wager you are beyond the "yeah right" stage of thinking when hearing of the connection. The truth is, the Humanities as a discipline is meeting the challenge of a world of digitized texts and scholarship with a vengeance. You don't have to look too hard to find outstanding examples of the range of work in the digital humanities. Scholars from around the globe are marking-up, morphing, mashing, mapping, and mining. So truth in advertising: this course cannot do justice to all of the great work in the digital humanities. Rather, as I hope you are gleaning from the syllabus at this point, **we will focus on an introduction to "text mining", in particular learning to write programs to explore digitized texts.** So, if you were hoping to learn the details of the Text Encoding Initiative (TEI) or how to mash Graphical Information System (GIS) data with Asian author birthplaces to make an interactive map, sorry ... we just won't touch those topics. To further explain the "take home" story for this particular course, read on about computational thinking.

Computational Thinking

There is much misconception about computer science these days. For many, computing means using an iPhone. Yeah, that's cool ... but that is not computing. Computer Science is the study of computation – what can be computed and how to compute it. The discipline encompasses "computational thinking" (Wing, 2006)¹, a universal metaphor of reasoning that defines how creative and imaginative humans use computation to facilitate communication, model complex systems, and visualize content.

Given the tangible, ubiquitous, embedded, and rapidly evolving nature of computing in our lives, the discipline of computer science faces the challenge of how to attract students to study within a discipline where some perceive they are already "experts". Clearly, relative to the previous generations that defined technology and computing by the electronic technologies at hand, the thumb-wielding, wireless generation of new students appears undaunted as they master and demand new hardware. But hardware is not computer science; a power-user is not a computational thinker. This course is about becoming a computational thinker.

Computational thinking is:

- a move away from "literacy" and toward "fluency," a broader concept including contemporary skills and intellectual capacities
- using abstraction and decomposition when attacking a large complex task
- choosing an appropriate representation for a problem
- modeling relevant aspects of a problem to make it tractable
- using invariants to describe a system's behavior succinctly
- developing heuristics to posit if and when an approximate solution is good enough
- using randomization to our advantage
- planning and scheduling in the presence of uncertainty
- search, search, and more search ... in our case "search texts!!"
- about ideas, not artifacts – it's not just the hardware and software that will be physically present everywhere, it will be the computational concepts we use to approach and solve problems, manage our lives, and interact with other people.

¹ Wing, J. (2006). A Vision for the 21st Century: [Computational Thinking](#), *CACM* vol. 49, no. 3, pp. 33-35.

Online (free) Text:

Interactive Python: How to Think Like a Computer Scientist
<http://interactivepython.org> by Brad Miller and David Ranum

How to Think Like a Computer Scientist



I also *strongly recommend* that you buy a 3-ring binder. I'll pass out lots of handouts.

We'll be using lots of online websites as reference, including those from Scott Kleinman, Associate Professor of [English](#) at [California State University, Northridge](#). Scott is working with the Lexomics Research Group on some research.

He works in Old and Middle English literature (<http://scottkleinman.net/>).

Your Grade:

Things to do	Grading Percents	Due Dates
Labs (lecture and lab are “blurred”)	5% overall	in class as needed
MOOCs	10% overall	TBA
5 Assignments	45% overall	
a1: Set up a website -- Ngrams	5%	Wed., Feb. 05
a2: RPB: Reading Poetry Backwards	10%	Mon., Feb 24
a3: Regex Play	10%	Mon., March 5
a4: “tall” ... “elf” in Tolkien	10%	Wed., Mar 26
a5: Anglo-Saxon poetry	10%	TBA
Quizzes	20% overall	
Quiz I	5%	Mon, Feb 10
Quiz II	5%	Wed, Feb 26
Quiz III	5%	TBA
Quiz IV	5%	TBA
<i>No Final Exam</i>		
Final Project: Text Mining Experiment	20% overall	
Presentation	5%	last week of class
Paper: Methods, Results, Discussion	15%	Wed., April 30

Late Submissions:

Due is due. Always turn in whatever you have on time. Something turned in on time is much better than not having it accepted because it is late. Late is not an option. (Good, glad we can all agree with this).

Note: **Python Programs** are due on various dates (see detailed syllabus in moodle (onCourse)); however, since I know from experience that many students like to use the last night for testing, I will allow you to submit your assignments until 4am the following day. For example, assignment a1 is due Wed., Feb. 5th (see above), but you can submit it electronically until 4am Thur., Feb. 6th – but be careful! The course website (onCourse) makes it appear as if the program is due on Thursday, but remember, that means Thursday at 4am!

Honor Code Revisited:

It goes without saying that all submitted work will be the student's own, in keeping with the Wheaton Honor Code, unless the assignment has assigned groups. For labs, you may get “help” from fellow classmates, but remember that all completed work must be your own. Use discretion; don't ask your colleague for “the” answer or for lines of code. However, I do encourage you to discuss the problem in general, such as the type of statements or functions

one might use. For homework, your answers and software must be your own from beginning to end. Here is an analogy. Almost no one would every “use/steal” a line or two from another person’s poem. Consider it the same with your programs. Don’t “borrow/use” lines or sections of your program from another classmate. Your program is (like) your poem; everyone’s program should be unique. Be wise. If a colleague is asking you for too much help, be honest and remind them your program is just that, *your* program.

MOOCs (Massive Open Online Courses): We won’t specifically be using a MOOC in this class, however, 10% of your grade will be based on you completing online materials at `codeAcademy.com`. To earn these points, you must *show me that you have* completed the Badges for each section assigned. If you complete the assigned sections that will be graded as “above average (B).” If you do more sections than assigned, this will be considered Superior (A) effort.

Tips for working on your own

- (0) It is expected that you spend at least 3+ hours on reading, study and preparation for every 100 minutes of lecture and lab.
- (1) It is expected that you spend at least 6-10 hours per week on your current programming assignment.
WARNING: Programmers typically underestimate the time it takes to complete a software project; 6-10 hours per week on your programming assignment may be one of those “underestimations.”

In classroom “LABS”

- (0) The computer work in class (labs) are a critical part of the course. Appropriately so, it may be hard to distinguish if you are in lecture or “lab”. In a way, hands-on class time is your time to “hack”, solve unique problems, and show that you can focus on the problem at hand. Typically, your lab time will prepare you to work on your next programming assignment. You must be in class to get credit for the session. If you happen to miss a class, you are strongly encouraged to complete the in-class work on your own time, but please do not ask for credit.
- (1) In order to best grasp the material presented in lab, I strongly suggest that you completely redo any labs that you find difficult. (Read that last sentence again, unless of course you've already reread it once).

Final Projects

As you can see, the final projects are a significant part of your final grade (20% total: 5% for your final talk, 15% for your final report on your experiments). You will work with another person on a two-person team. When working on a team, the collective team will be given a certain number of points and those points will be, by default, divided and assigned equally. Separate grades will be given for the oral presentation and for the final report. For example, if a team is awarded 170 points for their presentation, each person will earn a grade of 85. If that team is awarded 158 for their final report, each person will earn a grade of 79. Note: if a team agrees that one member should be assigned more of the points, it will be up to the team to let me know how to split the grade. Again, unless I hear from a team, the default grading system will be to equally divide the grades.

Quizzes

Your four quizzes will test your comprehension and ability to write your own Python, regular expressions, and HTML. During lecture, I will give “hints” of what a typical quiz question might be; take good notes! There will be no make-ups, nor will the lowest quiz be dropped. If you are an athlete and/or you have a conflict with a quiz date, please see me within the first week of classes.

HELP

I have listed my office hours on the syllabus at the top. But I'm usually near a keyboard so we can schedule alternate times to meet. Study, study, study and talk about it with me and others as often as you can. *Please don't wait too long before you see me; a quick chat in my office can often clear things up. I'm here a lot...*