

Vectors for CS Majors - Fall 2010

THESE RULES APPLY TO STUDENTS ENTERING CORNELL AY 2010/11 OR BEFORE.

Choosing a Coherent Set of Electives

The grade point average is but one way to measure the quality of an undergraduate record. An equally important attribute concerns depth-of-education, something that is best measured by the choice of elective courses. While required core CS courses ensure breadth-of-education and set the stage for more specialized work, it is through the choice of electives that you can communicate an appealing level of curiosity and a readiness for the future.

The newly redesigned CS major has a streamlined set of core courses to allow for more choices in electives, and a collection of *vectors* that promote coherent, in-depth study. A vector has a magnitude (typically four to five courses, which may be internal or external to CS, and which are allowed to “double-count” towards other requirements of the major or College), and a direction representing a "line of inquiry" within CS.

[Renaissance \(Basis\)](#)

[Network Science](#)

[Artificial Intelligence](#)

[Programming Languages](#)

[Computational Science and Engineering](#)

[Security and Trustworthy Systems](#)

[Data-Intensive Computing](#)

[Software Engineering / Code Warrior](#)

[Graphics](#)

[Systems](#)

[Human-Language Technologies](#)

[Theory](#)

The faculty may introduce or modify vectors as they see new directions for the department and the field; however, changes will be introduced no more frequently than once a semester.

Completion of at least one vector is required for graduation. (Here, for reference, is [a sample checklist showing completion of multiple vectors simultaneously](#) according to the Fall 2009 rules.) You should begin thinking about possible vectors of study when you affiliate, and, since your interests or perspectives may change, we recommend you start out following at least two vectors and consult regularly with your advisor about your choices. Because there is considerable overlap between certain pairs of vectors and because the CS major has a large elective component overall, if you follow this advice, there is usually no problem should you decide to change vectors (or launch another) along the way.

Your choice of vector(s) should be driven by academic interest alone. It should not be based on risky predictions of where the field is going. The best way to ensure a successful career after you graduate as a CS major is to practice *now* the art of being intellectually nimble. Tumultuous times call for mobility. Cultivate a broad knowledge of the field and show the world through your transcript that it empowers you to go deep and to learn new areas. Choose your electives as if your life depended on it. It does.

Important deadlines for declaring and verifying vectors:

- *Declaring your vector* - Students should inform the CS department, no later than the last day of classes in junior year, of the vector(s) they are planning to pursue by submitting the 'Vector Declaration Form' to the CS Undergraduate Office.
- *Verifying completion of vector* - During senior year, CS majors must complete the 'Vector Verification Form' which indicates the vector(s) that will be completed by their graduation date. The deadline for submission of the 'Vector Verification Form' is October 15th of senior year for January candidates, and March 15th of senior year for May and August candidates.

Both of these forms are available online (see ['Useful Forms'](#)) and/or from the CS Ugrad Office in 303 Upson Hall.

Vector Descriptions

*NOTE: The "F" notation Fxxx (or variants like F7xx) is used, informally speaking, to denote an **Four or Five or, um, Six-thousand level Full-Fledged elective**. It does **not** mean a 4000-level course (otherwise, we would have written "4" instead of "F".) Specifically and technically, it stands for any course that's at the 4000-6000 level, at least 3 credits, not an independent study, and not required by the CS major. For example, Fxxx includes 4700, 5620, and 6110, but not 4820 (required by the major), 4999 (independent study), or 4121 (fewer than 3 credits). As another example, CS F4xx includes CS 4450 and 5430, but not CS 4410 (required by the major), 4411 (too few credits) or 4700 (doesn't match the pattern!)*

Vector name: RENAISSANCE (BASIS)

Encourages breadth of study across computer science, including education in fundamentals beyond the core that support multiple pursuits.

Explanation of requirements: Four upper-level CS electives among which at least two areas must be represented, and at least one of scientific computing and theory is represented among the areas (note that an additional course in each of these areas was required of all CS majors before Spring 2009).

Required courses: Four total (4 full and 0 half)

- Four elective courses drawn from the following list: CS Fxxx, CS/ENGRD 3220, CS 3810 (even if required by the major; this exception handles transitioning from the major requirements prior to Spring 2009), or cross-lists of any of these courses, where
 - at least two different hundredths digits are represented, and
 - at least one of the "hundredths digits" is either 2 or 8 (representing scientific computing and theory, respectively)

Recall by the definition of the "Fxxx" notation that the following courses (and others) do not count: CS 4820 (CS core course), CS 4999 (independent study), CS 4090 (teaching for credit), CS 4812 (if taken prior to AY 2010/11 when only 2 credits), to pick a few examples.

Vector name: ARTIFICIAL INTELLIGENCE

Focuses on understanding and creating intelligent systems.

Explanation of requirements: the core AI course and associated practicum, machine learning, an AI elective, and either connections to human intelligence or further grounding in techniques. Can be tailored to focus on machine learning. See also the human-language technologies vector.

Required courses: Five total (4 full and 1 half)

- CS 4700 (472), Foundations of Artificial Intelligence
- CS 4701 (473), Practicum in Artificial Intelligence
- An additional course numbered CS F78x or CS 4758 (Robot Learning); this includes, due to cross-lists, BTRY 4790/6790 (Probabilistic Graphical Models)
- An additional course numbered CS F7xx or CS/INFO 4300 (Information Retrieval) or CS F67x (Computer Vision) or CS 5846/ECON 4760/ECON 6760 (Decision Theory); this includes, due to cross-lists, COGST 4740/LING 4474, BTRY 4790, CEE 5290/ORIE 5340, INFO 6300, BTRY 6790
- An additional course chosen from the following list: COGST, PSYCH or LING Fxxx, CS/ENGRD 3220, CS F2xx (or the cross-lists MATH 4250 or MATH 4260). (Recall that independent studies such as LING 4493 or courses below 3 credits are disallowed according to the Fxxx naming convention.)

Recommended: scientific computing (e.g., CS/ENGRD 3220, CS 4210 or 4220), game theory or related economics topics.

- Recommended for focus on machine learning: graphical models, scientific computing/optimization, statistics and/or data mining

Vector name: COMPUTATIONAL SCIENCE AND ENGINEERING

Focuses on the numerical algorithms that support computer modeling and simulation to guide experimental and design work in data-intensive scientific computing.

Explanation of requirements: Effective numerical computing in science and engineering requires an understanding of the fundamental algorithms that have been devised to solve the continuous problems of applied mathematics. Requirements include a year-long sequence in numerical analysis that covers data fitting and approximation, nonlinear systems and optimization, ordinary and partial differential equations, and the essential methods of numerical linear algebra. An applied mathematics course that covers differential equations and/or optimization together with a programming versatility requirement completes the vector.

Required courses: Five minimum (3 full and 2 or more half). Note that a prerequisite for this vector is MATH 2930 (Differential Equations for Engineers) or Math 2220 (Multivariable Calculus).

- Two (different) courses numbered F2xx or MATH 4250 or MATH 4260, with the following exceptions: CS/ENGRD 3220 (Scientific Computation) may be counted with a grade of B+ or

better; and one cannot count both CS 4220/MATH 4260 (Numerical Analysis II: Linear and Nonlinear Problems) and CS 6210 (Matrix Computations), due to overlapping content.

Note that one can take CS 4210/MATH 4250 (Numerical Analysis I: Numerical Analysis and Differential Equations) and CS 4220/MATH 4260 in either order.

- One of the following applied mathematics courses: OR 3330 (Optimization I), TAM 3100 (Applied Mathematics I), Math 4200 (Differential Equations and Dynamical Systems), Math 4240 (Wavelets and Fourier Series), Math 4280 (Introduction to Partial Differential Equations), AEP 3210 (Mathematical Physics I), CEE 3310 (Fluid Mechanics), CEE 3710 (Structural Modeling and Behavior), MAE 3230 (Introductory Fluid Mechanics).
- A minimum of four-credit hours of programming experience obtained through any combination of CS 2022 (Introduction to C), CS 2024 (C++ Programming), CS 2042 (Unix Tools), CS 2044 (Advanced Unix programming and Tools), CIS 4205 (Effective Use of High-Performance Computing), CIS 4206 (Introduction to Scripting in Python and Perl).

Recommended: CS3300 (data-driven web applications), CS 4320 (introduction to database systems), computer graphics courses with significant mathematical content numbered (typically numbered CS F6xx or corresponding cross-lists), CS F78x (machine learning courses).

Vector name: DATA-INTENSIVE COMPUTING

Focuses on the foundations of managing, processing, and analyzing large datasets

Explanation of requirements: One course on data management and the associated practicum, one course on large data-driven systems, one course on machine learning/data mining, one course on information retrieval.

Required courses: Five total (4 full and 1 half)

- CS 4320 (Database systems)
- CS 4321 (Practicum in Database Systems)
- CS/INFO 5300 (The Architecture of Large-Scale Information Systems)
- A course numbered CS F78x, CS 4758 (Robot Learning) or ORIE 4740, where this list also includes due to cross-lists BTRY 4790/6790 (probabilistic graphical models)
- CS/INFO 4300 (Information Retrieval) OR CS/INFO 4302 (Web Information Systems) or CS 6740/INFO 6300 (Advanced language technologies)

Vector name: GRAPHICS

Focuses on computing with visual images

Explanation of requirements: The introductory graphics course and practicum, scientific computing, and two electives (in various combinations) in graphics and related fields

Required courses: Five total (4 full and 1 half)

- CS 4620 (formerly 465), introduction to computer graphics (also ARCH 3704, formerly 374)

- CS 4621 (formerly 466), computer graphics practicum
- CS/ENGRD 3220 or CS F2xx (includes, because of cross-listing, MATH 4250 and MATH 4260)
- A course drawn from the following list: CS 5625 (or CS 5620 prior to fall 2011), CS 5643, CS 6620, CS 6630, CS 6650. Note that the 5000-level courses have a comparable level of difficulty with 4000-level CS courses.
- An additional course drawn from the following list: CS F6xx (graphics and vision courses) or CIS 3000 or CIS 4002 (game design courses)

Recommended: courses in visual art or photography; additional 4000-level mathematics courses.

Vector name: HUMAN-LANGUAGE TECHNOLOGIES

Focuses on creating natural-language processing and information retrieval systems that can analyze and generate information in unstructured, natural-language format (e.g., English). Applications include search engines, automatic translations between languages, and human-computer interaction, among others.

Explanation of requirements: One course on the fundamentals of computational linguistics, natural language processing and/or information retrieval; and an additional course from those areas or in core AI or in programming languages; a course on machine learning (the fundamentals of AI course may serve as a substitute); the practicum in artificial intelligence; an upper-level linguistics course

Required courses: Five total (4 full and 1 half)

- One course from the following list: CS F74x, or CS 3740/COGST 4240/COGST 4240, CS/INFO 4300, and the following courses because they are cross-lists of the aforementioned: LING 4474/COGST 4740, INFO 6300.
- An additional course either drawn from the above list or numbered CS F110 (programming languages and logics) or LING 4485 (topics in computational linguistics).
- A course numbered CS F78x (machine learning courses) or CS 4700 (Foundations of Artificial Intelligence) or CS 4758 (Robot Learning)
- CS 4701 (Practicum in AI)
- An additional course numbered LING Fxxx (includes CS 3470/COGST 4240 and COGST/PSYCH 4280 as cross-lists, but no course can be used to fulfill both this and any of the above requirements)

[Recall that independent studies such as LING 4493 or courses below 3 credits are disallowed according to the Fxxx naming convention.]

Recommended: CS 2042 (UNIX Tools) or 2044 (Advanced UNIX Programming and Tools), statistics, scientific computing (courses number x2xx), a foreign language.

Vector name: NETWORK SCIENCE

The social, technological, and natural worlds are connected; the study of networks sheds light on these connections. This vector provides preparation in the network- and data-analysis tools to understand and develop predictive models of physical, biological, and social network phenomena.

Explanation of requirements: Two core courses in the analysis of networks; a course in machine learning; and a course on application areas

Required courses: Four total (4 full and 0 half)

- Two different courses numbered CS x85x (complex networks/information access) or (because they are CS x85x cross-lists): ECON 2040/INFO 2040/SOC 2090, INFO 6850. To be clear: CS 2850 counts towards this requirement.
- One course numbered CS F78x (machine learning courses)
- One course chosen from the following list: ORIE x350 (Game Theory), ECON 3670 (Game Theoretic Methods), ECON 3680 (Game Theory), SOC 3040 (Social Networks and Social Processes), SOC 4250/5270 (Artificial Societies), CS F84x (game theory courses)

Vector name: PROGRAMMING LANGUAGES

Spans logics and language semantics, language design, compilation and optimization.

Explanation of requirements: Requirements include a course on programming languages; a course on compilers and the associated practicum; an elective in either applied logic, computability and complexity, or advanced programming languages; and a (short) course on a particular language.

Required courses: Five total (3 full and 2 half)

- CS F110 (programming-language and logics courses, formerly 411/611)
- CS 4120 (introduction to compilers, formerly 412)
- CS 4121 (practicum in compilers, formerly 413)
- An additional course chosen from the following list: CS/Math 4860 (applied logic, formerly 486) or CS F810 (computability and complexity, formerly 481/682) or CS 6110 (611) if not already used to satisfy the requirements above.
- One short course, CS 202x, on an additional programming language. For example, 2022 (C), 2024 (C++), and 2026 (C#) qualify.

Vector name: SECURITY & TRUSTWORTHY SYSTEMS

Addresses fundamental problems in ensuring the security and reliability of our global critical computing infrastructure.

English description of requirements: Required courses include the operating systems practicum, mathematical tools for cryptography or verification, system security, and the second course in the systems core sequence.

Required courses: Four total (3 full and 1 half)

- CS 4411 (Practicum in Operating Systems)
- CS 4830 (Introduction to Cryptography, formerly 487) or CS/MATH 4860 (Applied Logic) or MATH 3360 (Applicable Algebra)
- CS 5430 (System Security)
- CS 5412 (Cloud Computing) or 5414 (Distributed Computing Systems) or CS 5410 (no longer offered)

Note: undergraduates should not have difficulty taking the CS54xx courses provided that they had no difficulty in CS4410/CS4411.

Recommended: INFO 5150 (Culture, Law and Politics of the Internet) or other relevant law or policy courses.

Vector name: SOFTWARE ENGINEERING / CODE WARRIOR

Provides students with the fundamentals of software engineering and extensive implementation experience in a variety of application areas.

Explanation of requirements: a masters-level course on software engineering, two programming-intensive practica in different areas, and an additional course with a heavy implementation component. Experience in C++ is strongly recommended (note that many CS majors are exposed to Linux through the core course CS 3410).

Required courses: Four total (2 full and 2 half, but at least one additional full course corresponding to one of the half-course practica is strongly recommended (and in some cases required))

- CS 5150 (Software Engineering)
- Two different courses numbered 4xx1 (practicums, such as 4121 (compilers), 4321 (databases), 4411 (operating systems), 4621 (computer graphics), 4701 (artificial intelligence)).
Note: It is strongly recommended (and in some cases required) that one also take the corresponding "full" course; for example, if one takes 4621 (the graphics practicum), it is recommended that one also take 4620 (introduction to computer graphics)
- One of CS F45x (computer networks), 5300 (architecture of large-scale information systems), CS 5412 (Cloud Computing), 5414 (Distributed Computing Systems), CS 5410 (no longer offered), or the following implementation-heavy computer graphics courses: CS 5620, 5643, 6620, 6630, 6650.

Recommended: Exposure to C++, either through 2024 (C++ programming) or practica taught in C++; CS 2022 (introduction to C) if 2024 not taken, 2026 (introduction to C#), 2042 (UNIX Tools), 2044 (advanced unix programming and tools); CS/INFO 2300 (intermediate design and programming for the web), CS/INFO 3300 (data-driven web applications).

Vector name: SYSTEMS

Focuses on design and implementation of the fundamental software systems that comprise our computing infrastructure.

Explanation of requirements: the core operating-systems practicum, and three systems electives (areas include networking, architecture, and so on). The launch-point for this vector is CS4410 (operating systems, a core CS course) and its practicum, CS4411. Many students follow on with CS4450/5450 (networks, not offered every year), CS5430 (security) and/or CS5410 (intermediate computer systems: cloud and edge computing). Note: undergraduates should not have difficulty taking these CS54xx courses provided that they had no difficulty in CS4410/CS4411.

Required courses: Four total (3 full, 1 half)

Recall: "Fxxx"=4000-6000 level, 3+ credits, non-independent-study, non-required course; for example, CS F4xx includes CS 4450, CS 5430, CS 5412, and CS 5414 (and CS 5410 - no longer offered), but not CS4410 (since 4410 is required of all CS majors) or CS 4411 (below 3 credits); similarly, CS 4321 is not included under F32x, and CS 4121 is not included under CS F12x.

- CS 4411 (CS415 OS practicum; normally taken with CS4410 (CS414), operating systems)
- 3 additional courses chosen from the following list: CS F4xx, CS F12x (compilers courses), CS F32x (database courses), CS/INFO 4302 (web information systems, formerly 431), CS/INFO 5300 (the architecture of large-scale information systems), ECE 4450 (computer networks and telecommunications), and the following ECE courses because they are CS F4xx cross-lists: ECE 4750(=CS 4420, computer architecture), ECE 5720(=CS5420, old 516, parallel computer architecture)

Recommended: Although any group of courses from the acceptable list will do, many systems students find it especially useful to get a solid background in databases.

Vector name: THEORY

Focuses on the study of efficient computation, models of computational processes, and their limits.

Explanation of requirements: a course in computability/complexity, two theory electives and an upper-level technical mathematics course

Required courses: Four total (4 full, 0 half)

- CS 3810 (offered up through Fall 2010) or CS 4810 (new course on computability and complexity, expected fall 2011 and thereafter). CS 3810 counts even for students for whom it is required for the major due to having taken CS 2800 (formerly 280) fall 2008 or previously.
- Two additional courses numbered CS F8xx or ORIE 6330 (Graph Theory and Network Flow) or ORIE 6335 (Scheduling Theory). (Recall that CS 4820 and CS 4812 (Quantum Information Processing) do not satisfy this requirement, since the "F" notation indicates a 4000-6000 non-CS-major-required 3+ credit course.)
- An additional 3+ credit course numbered MATH THxx, where the thousands digit T is between 3 and 6 inclusive and where the hundredths digit is between 1 and 9 inclusive, or MATH 4010 (Honors seminar: Topics in Modern Mathematics) or, due to cross-list with MATH 4860, CS 4860 (Applied Logic).

Vectors for CS Majors – Fall 2010

Recommended: CS 3220, 4210, or 4220 (Scientific computation, numerical analysis); MATH 4320 or 4340 (Abstract algebra); MATH 4810 (Logic).