Scope

Currently, this document only offers recommendations for navigating the Joint NC/ML track. Anyone wishing to help add guidance for students navigating the PNC or the NC/Statistics tracks is eagerly welcomed to contribute (see last page).

Overview

The document was compiled with input from PNC students from different years. It aggregates our collective "tribal knowledge". Some of us struggled with coursework while others did not. The diversity of our experiences was a result of our backgrounds, research workloads as well as personal strengths and weaknesses. After many conversations that sounded like "**what would I do differently if I knew in 1st year what I know now?**", it became clear that these lessons ought to be summarized into a plan that many of us wish we had when we began.

On average, following the plan below should **minimize coursework related stress and maximize research productivity/grades**. The plan recognizes that there are **no prizes for finishing coursework faster** and prioritizes **timely completion of research milestones and meeting minimum grade requirements** that are both critical for the Joint NC/ML Program. If the plan below works out for you, that is great. If it is harder than you thought that is not a sign that something is wrong with you. Talk to students ahead of/around you. You may find you are not alone and discover tips and tricks that can ease your path.

Lastly, as you progress, please share your experiences and help keep this guide relevant as program requirements/classes evolve. This will help students who follow after you. See the end of this document for contact details.

Joint NC/ML coursework plan

Table 1: (C) denotes MLD core courses that are compulsory and (E) denotes MLD elective courses. Typical elective courses are mentioned below, but there may be additional options that are not listed here. All other courses are CNBC requirements. Grade distributions for classes are from past years and may vary based on instructor.

Semester # (Overall Difficulty)	**Course (Difficulty)**	**Course Logistics & Details**
1 (MODERATE)	(C) 36-705 Intermediate Statistics (Moderate)	HIGHLY RECOMMENDED IN 1ST SEMESTER A PRE-REQ. FOR MULTIPLE DOWNSTREAM CLASSES. **Workload:** Moderate/Heavy –9-11 HWs (20%): Do these well if you wish to do well in exams. –3 exams (25%,25%,30%): Everything is solvable, but not always within 50 minutes. –Doing well is critical. Practice all HWs beforehand. **Comment:** A theory/technically intensive foundational class. A worthwhile investment of time. Only class of its kind in the coursework (heavy reliance on exam performance). Siva is an excellent teacher, that said, you'll get out what you put into the class.
	36-765 Statistical Models of the Brain (Easy)	**Workload:** Easy The exact details may vary, but expect something like –Weekly Readings: 55% Readings –4 HWs: 20% of the grade –1 Project: 20-25% Exploring a simple computational neuroscience problem is sufficient. **Commment:** Geared to help non-computational students learn computational thinking and expose non-neuroscience students to neuroscientific problems.
2 (HARD)	(C) 10-716 Advanced Machine Learning: Theory & Methods (*Hard*)	**Workload:** Heavy –50% 4 HWs –25% 1 Midterm Exam (may be open book) –25% Project (2 person teams) - see 10-715 comments about projects **Commment:** Risk Minimization, Non Parametric statistics form much of the content.
	03-763 Adv. Sys. Neuroscience (Easy)	**Workload:** Light. Grading scheme may vary from year to year, here is an example – 20% 1 Journal Article Presentation in class – 20% Class participation/attendance (grad section) – 15% Weekly discussion question sheets – 15% 1 News and Views article – 30% Final exam (must pass) for undergrad section (memorization needed).

	Math 3370: Mathematical Neuroscience (Easy)	**Workload:** Moderate –6 HWs –1 project **Commment:** This is a dynamical systems modeling class you'll simulate neuron circuit models, bursting dynamics, oscillations, how different channel types affect firing etc. Not hard, but time consuming to do all the work. Familiarize yourself with XPPAUT.
	First Milestone Prep.	Devote time to your first milestone project. Avoid too big a project.
Summer	***First Milestone***	*Complete First Milestone*
3 (HARD)	(C) 10-715 Advanced Introduction to ML (Hard)	**Workload:** Heavy Exact format varies, but expect something like the following –HW: In 2018 (Prof. Balcan): 4-6 big HWs (∼50%), with theory & coding questions. In 2019 (Prof. Shah): Addition of mini-HWs (9-10) taking 10-15 minutes of work, each focused on assessing understanding of lectures. – 1 or 2 exams (20-30%): these are hard & time runs out. – 1 Team Project (∼25%): High cost if you get into a complex project. Keep it simple. **Comment:** Students without a formal introduction to ML may find the course difficult but also value its broad coverage of topics. Those familiar with basic ML may get broader exposure to newer techniques, but also find familiar content repeated.
	85-765 Cognitive Neuroscience (Moderate)	**Workload:** Moderate 3 exams only (33% each): High time cost for exam prep. Lots of content to memorize. Precision and detail in your answers on exams is CRITICAL to do well. **Comment(s):** This class is a stark contrast to Cell. & Sys. Neuroscience in difficulty. MD/PhDs report finding this easy. Those with limited memorization skill find it hard. For those lacking neuroscience background, this can be both very difficult & valuable. Be religious about lecture. Yes 10:30 AM may be painful, but it'll be worse if you don't go because Prof Olson makes many finer points not be spelled out in notes, which show up on exams.
4 (EASY)	(E) 10-707 Deep Learning (Moderate) OR (E) 10-708 Probabilistic Graphical Models (Moderate)	**Workload:** Moderate (10-707) OR Moderate (10-708) **10-707(With Russ):** 3 HWs (60% of grade) + 1 Project (40% of grade) HWs are 40% theory + 60% code (experiments won't finish without vectorized code) – HW 1: Write Neural Network from scratch, implement variants (dropout, regularization, sparsity). Once finished run many experiments and explain results. – HW 2: Write CNN from scratch. Same routine as above. – HW 3: (May Vary): Recurrent NN, Language models, Beam Search, Line Search. – Project: Advice from 10-715 applies. **10-707(With Andrej):** HWs (45%) + 2 midterms (15% each) + final (25%) Theory and Coding HW are separated and difficult (even TA's may struggle). **OR** **10-708:** 4 HWs (50%) + 1 Project(40%) + Scribe a Lecture (10%) **Comment:** All things graphical models. Recent editions of the class include a focus on other areas of ML and how they relate to graphical models.
	Second Milestone Prep.	Devote time to your second milestone project to get draft manuscript in place.
Summer	***Second Milestone***	*Complete Second Milestone*
5 (MODERATE)	(E) 10-725 Convex Optimization (Hard)	**Workload:** Heavy 5-6 HWs(50-60%)+5-6 Quizzes(20%)+1-2 Exams(20%)+1 Project(deprecated 2019+) – HWs include hard theoretical/proof problems followed by related programming problems. Optimization Algorithms like GD, SGD, Newtons Method, Primal Dual Interior Point, Coordinate Descent, ADMM, Frank Wolfe, Projected Gradient Descent are to be written from scratch. Convex programming demands more clarity and precision than coding for other classes. HWs offer great value for time invested. – 5-6 Quizzes (20%): MCQs (20 questions each) accompanying HWs. – 1 or 2 Mini-Tests (20%): Like quizzes, but longer. Don't be deceived. People make lots of small mistakes on MCQs due to lack of precision/clarity. 1 Project: If projects come back, advice for 10-715,707,708 applies. In 2019, project was replaced with 2 open ended questionsadded to last 2 HWs. Students graded each other in a peer-review system. **Comments:** This class is technical/theory and time intensive.However, it is worthwhile An extra credit option (lecture scribing) is best done early in the semester
	03-762 Adv. Cell. Neuroscience (Easy)	**Workload:** Light: Grading scheme may vary from year to year, here is an example – 35% 1 Journal Article Presentation in class – 20% Class participation/attendance (grad section) – 10% Weekly discussion question sheets – 30% 1 News and Views article – 5% Final exam for undergrad section (memorization needed). – Final Exam is 5%, but is a MUST pass, otherwise you'll fail the class.

Dos and Dont's, best practices and other thoughts

This is a catch all section for insights/opinions that don't fit into the course template.

- Make a habit of talking to others in the program, particularly those on the track you are on. You may find you are not all that alone in terms of problems you face. This may seem trite, but it works and makes a big difference.
- If you have a lot of free time, do more research, don't add classes.
- If you are going through the Joint Program in NC/ML, 10-701 should be avoided. The material overlaps significantly with 10-715, though 10-715 is a bit harder. MLD does not consider the two equivalent, so you will need to take 10-715 regardless of whether you took 10-701 or not. In case you feel like you lack ML background, talk to academic advisors first. 10-701 is not the gentle introduction you might be looking for. If you really are unfamiliar with ML/math, then it may also be important to think carefully about whether the ML track is right for you.
- Don't take heavy quantitative classes together (705 & 715, 715 & 725, 705 & 725 or 708 & 707, 716 & 707, 716 & 708). You may survive the elevated stress, but your grades may drop/research progress may slow which will matter in Joint Program Admissions. Even if these negative outcomes don't come about, there is no advantage to be had by doing this.
- Don't assume classes will teach you what you need for research. Classes tend to cover a broad range of topics (breadth) and often don't delve too deep into each.

We need your help to keep this document useful

Program and coursework requirements change over time. The students who first wrote this document did not navigate exactly what you will (or have). We need your help in order to keep this resource up to date, as well as to make it useful by expanding it so it covers all PNC tracks. If you have input/suggestions that might be beneficial to others, please send an email to the current maintainer. Your input will be treated as confidential and will be aggregated into the document.

Current Maintainer: Arish Alreja (aalreja@andrew.cmu.edu)