PRISM

Lecture 2 - Learning about Prior Work

Roger Grosse and Ishtiaque Ahmed

University of Toronto, Winter 2021

"A month in the laboratory can often save an hour in the library."

— Frank Westheimer

Today

- Today: how to learn about prior work in a research area
  - Taking courses and reading textbooks
  - Reading papers and watching talks
  - How to do a lit search
- Examples of how to read a computer science paper
- Q&A with grad student mentors

Where to Learn Things

Where to Learn Things

Learning by Taking Classes

- **Pro:**
  - interact with experts in the area
  - subject matter arranged to be read linearly
  - force yourself to learn it well by doing exercises and studying for exams
  - you've already been practicing this for years
- **Con:**
  - large time commitment, if you need one specific thing
  - course might not be offered when you need it
  - doesn't cover the cutting edge of the research topic you're interested in (unless you're really lucky)
  - course may be aimed at students with very different background and motivations

Where to Learn Things

Learning by Reading Textbooks

- **Pro:**
  - random access
  - more comprehensive
  - more likely to find a book specialized to what you need
- **Con:**
  - not meant to be read linearly
  - shortest path to reach an idea not immediately obvious
  - requires a lot of motivation and self-discipline
  - no teaching staff to turn to if you're stuck
  - covers ideas that have been around > 10 years, not current research

Where to Learn Things

Learning by Reading Papers

- **Pro:**
  - probably the only place you can find the very cutting edge research
  - get more insight into where the ideas come from
  - learn about open questions
  - be able to evaluate the evidence for yourself
- **Con:**
  - a lot of background knowledge assumed
  - maybe nobody's figured out how to explain it clearly yet
  - not all papers are correct (even if they've passed peer review!)
  - authors may oversell the method and de-emphasize the drawbacks to get past the reviewers

Where to Learn Things

Learning by Attending Conferences

- **Pro:**
  - hear short talks on the papers
  - interact with the authors at poster sessions
  - meet other researchers interested in similar topics
- **Con:**
  - expensive
  - probably not an option when you're just starting out
- **Note:** many fields (e.g. machine learning) post videos of conference sessions publicly
  - Research talks are aimed at experts in the subfield
  - Keynote talks and tutorials are a great way to get started on a topic — given by top researchers, more polished, aimed at a broader audience

Textbook Knowledge

Textbook Knowledge

- Courses and textbooks are a great way to learn the foundational concepts needed for your research
- Academic papers usually assume a level of background equivalent to a graduate course in the subfield (e.g. CSC413/2516 for deep learning papers)
- Differences from what you're used to
  - "just-in-time" rather than "just-in-case" learning
  - might need to learn really specific things from a wider range of fields (e.g. physics, control theory, etc.)
  - having the context of a concrete research problem helps a lot!

Textbook Knowledge

- Choosing the right textbook is really important
  - see what's assigned on course web pages (especially MIT OCW)
  - check out multiple books from the library and pick the one you like the best
- Decide how much time you're willing to invest
  - one class $\approx$ 100 hours, which you can't take on lightly
  - one strategy
    1. spend a few hours casually reading textbook chapters, Wikipedia articles, etc. on topic X
    2. if you're still getting stuck because you don't understand X, then go back and learn it more carefully (do exercises, seek out additional resources, etc.)

Textbook Knowledge

The Princeton Companion to Mathematics is an amazing reference for pure math topics.

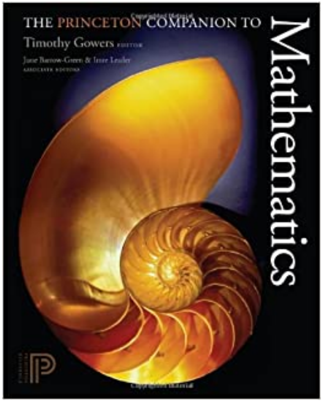

Textbook Knowledge

- Learning on your own is challenging because it requires *metacognitive* skills
  - judging if you understand something
  - diagnosing why you're having trouble with something
    - e.g. maybe you have too high a cognitive load because you didn't learn linear algebra carefully enough
  - considering different cognitive strategies to solve a problem
  - judging if your answer is correct
  - knowing how memory works ("will I still remember this a month from now?")
- "Talent" is just the word we mistakenly apply to skills that are learnable, but no one's figured out how to systematically teach
  - no simple recipe for self-directed learning — you have to figure out how your mind works!

Textbook Knowledge

- *A Mind for Numbers: How to Excel at Math and Science (Even If You Flunked Algebra)*, by Barbara Oakley
    - insights from cognitive science on how to learn more effectively
    - one student told me this book turned his life around
- Also an associated Coursera course, “Learning to Learn”

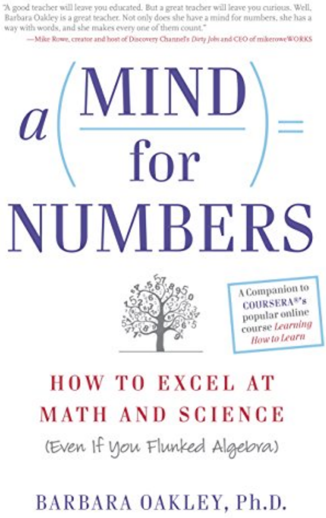

Reading Papers

Reading Papers

Types of academic papers:

- Conference papers
  - e.g. International Conference on Machine Learning (ICML)
  - the main type of peer-reviewed venue in CS
  - less important in most other fields
  - page limited (e.g. 8 pages), details often relegated to the Appendix
  - rapid turnaround
- Journal papers
- Preprints (arXiv)
- Monographs, review papers

Reading Papers

Types of academic papers:

- Conference papers
- Journal papers
  - e.g. Journal of Machine Learning Research (JMLR)
  - the main type of peer-reviewed venue in most other fields
  - in CS, often a more complete and polished exposition of work previously published in a conference ($\geq 30\%$ new material)
  - Often a multi-year publication delay
- Preprints (arXiv)
- Monographs, review papers

Reading Papers

Types of academic papers:

- Conference papers
- Journal papers
- Preprints (arXiv)
  - not peer reviewed (but hopefully to be published in a conference/journal soon)
  - the most cutting-edge
  - encouraged in some fields (e.g. machine learning), viewed with suspicion in others (e.g. natural langauge processing)
- Monographs, review papers

Reading Papers

Types of academic papers:

- Conference papers
- Journal papers
- Preprints (arXiv)
- Monographs, review papers
  - e.g. *Foundations and Trends in Machine Learning*
  - longer (e.g. 30–100 page) works summarizing a large body of literature
  - describes previously published results, not novel research
  - more specialized, research-oriented than textbooks

Reading Papers

Typical structure of a conference/journal paper in machine learning (will vary in other fields):

1. **Introduction:** motivates the problem, summarizes the main contributions
2. **Background:** describes prior work that's needed to understand the paper
3. **Methods:** presents the main technical contribution of the paper
4. **Related Work:** situates the work in the context of what's already known (different from background, in that not a prerequisite for understanding the paper)
5. **Experiments:** testing the method empirically
6. **Discussion/Conclusions:** summarizes the main contributions again (to a more informed reader), sometimes outlines open questions

Note: some of this is for the "benefit" of the reviewer

Reading Papers

- Choosing the right papers to read is really tough and important
  - papers vary a lot in quality
  - even many excellent papers aren't very accessible
- Find the top conferences and journals by seeing where professors at top institutions publish
- Highly cited papers (according to Google Scholar) are often good places to start (but this is a noisy signal!)
- Check the reading lists for graduate topics courses (these papers are selected for impact and readability)

Reading Papers

Multiple levels of reading:

1. quick skim to decide if you want to read it (5 min)
    - Read the abstract & intro, look at the figures
    - Skim the main body to see what sorts of techniques they use
    - What do they claim to have done?
2. casual read to get the main idea (30 min–1 hour)
    - Read through the paper linearly, skip things that don't make sense
    - What are the main ideas?
    - What evidence do they provide? Is it believable?
3. deep dive (many hours)
    - Read it carefully, follow the references, work through the math
    - Maybe implement a simple version
    - This is for papers you might build on

Reading Papers

HOW TO READ A RESEARCH PAPER

Among the questions that you should ask yourself when reading a research paper are the following.

1. What is the research paradigm that the author is using? Example paradigms are psychological experiments, formalization and theorem proving, and artifact design and construction. If the paper is part of a well established field, you should describe the field and its current state.

2. What is the problem area with which the paper is concerned? For example, "Automatic Generation of Compilers from Denotational Semantic Descriptions of the Source Code" would describe a research paper on compilation.

3. What is the author's thesis? That is, what is he/she trying to convince you of?

Source: `https://courses.cs.washington.edu/courses/cse590k/09sp/how_to_read.txt`

Reading Papers

4. Summarize the author's argument. That is, how does the author go about trying to convince you of the thesis?

5. Does the author describe other work in the field? If so, how does the research described in the paper differ from the other work?

6. Does the paper succeed? Are you convinced of the thesis by the time that you have finished reading the paper?

7. Does the author indicate how the work should be followed up on? Does the paper generate new ideas.

8. Some papers implicitly or explicitly provide a new way of doing things or of thinking about problems. If your paper does so, describe the approach.

Source: `https://courses.cs.washington.edu/courses/cse590k/09sp/how_to_read.txt`

Literature Searches

Literature Searches

- How to find relevant prior work on a topic?
- In my own field (deep learning), thousands of papers are published every year
- **General strategy:**
  - Start with papers you already know about (e.g. recommended by your supervisor), or do a Google search for the obvious keywords
  - Follow the citation graph: papers cited by these, and papers that cite them (using Google Scholar's "cited by" feature)
  - Look up the authors' web pages to see if they've published anything else relevant
  - Find review papers if possible
  - Skim a lot of papers, identify a smaller number to read carefully
  - Repeat
- Highly cited papers are helpful to start with since you'll see them referenced more often
- **Rule of thumb:** was your search thorough enough to find your original batch of papers if you didn't already know about them?

Literature Searches

Some questions to ask yourself:

- Why are people interested in the topic?
- What were the influential ideas and big advances?
- How would you categorize the approaches? (usually many ways to slice and dice)
- What are the strengths and weaknesses of various approaches?
- How do researchers in the area validate their ideas? (Proofs? Experiments?)
- What is the "state-of-the-art"?
- What are the open problems?
- What other research problems are closely related to this one?

Literature Searches

- Note that this is not a *prerequisite* for contributing to a research project
  - As an undergrad, you will probably rely on your supervisor or grad student mentor for day-to-day tasks
  - They will give you bite-sized tasks that you can do without a high-level picture of the research area
- Instead, do this *while* you're doing your research, so you can gradually become able to contribute to the project at a higher level
- Good idea to spend half your "research time" learning about the topic
- When you interview for an MSc/PhD program, you will want to impress the interviewers with your breadth and depth of understanding of the problem you worked on.

Example paper:
http://noodle.cs.huji.ac.il/~yweiss/iccv01.pdf