

Critical Reading

Sébastien Tixeuil
UPMC

Sept. 2015

WHY AND HOW TO READ RESEARCH PAPERS?

2

Why ?

- You were asked to
- Literature survey / new field / problem
- Be up-to-date on current research in your field
- Allows you to replicate / extend the results
- Provides you with useful data
- Gives you « pre-digested » thoughts
- To decide whether (and where) to publish
- Teaches you how to write
- Review for a conference or a class
- ...

3

Why read research papers?

- **Read for a conference or a class**
- **Keep current in your own field**
- **Get “up to speed” in a new field**
 - Learn about a sub-field (e.g., wireless)
 - Learn about another discipline that may offer solutions to a problem

From: N. Feamster, A. Gray, “Reading and Reviewing Papers”
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

4

Types of research papers

- **Conference papers**
 - Most recent, hot off the press information
- **Technical reports**
 - Expand on the information in a conf. paper
- **Journal papers** (a.k.a. articles)
 - Expand and combine results from several conf. papers
- **Book chapters**
 - Expand a conf. or journal paper
- **Workshop papers**
 - Very hot/new topics, work in progress, preliminary results and ideas

5

Peer-Reviewed ?

- **Conference and journal papers are almost always « peer-reviewed »**
 - Examined by other computer scientists (3-5)
 - Public, single-blind, or double-blind reviews
- **Technical reports are typically not peer reviewed**
 - Still excellent sources of detailed information
- **Posters / Workshop papers are peer-reviewed**
 - Criterium usually differs from conference and journal papers

6

Too many papers...

- Example: Networking Conferences
 - SIGCOMM: ~ 40 papers
 - SIGMETRICS: ~ 20 papers
 - IMC: ~ 40 papers
 - CoNext: ~ 30 papers
 - ICNP: ~ 30 papers
 - Infocom: ~ 100 papers
 - Journals, workshops, ...
- Per year: More than 2,000 pages to read
- Impossible to read it all...
 - doesn't even count cross-disciplinary reading

7

Types of papers

- **Theoretical**
 - Describe / prove a theory / algorithm
- **Engineering**
 - Describe an implementation of an algorithm, or part or all of a computer system or application
- **Empirical**
 - Describe an experiment designed to test some hypothesis
- **Survey**
 - Review recent results in a field of research

8

Step I: Deciding *what* to read (1/3)

- **Purpose:** Learn about “hot topics” of current research in an area. (searching for problems, etc.)
- **Approach:** Scan papers in latest conference proceedings

From: N. Feamster, A. Gray, “Reading and Reviewing Papers”
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule ⁹

Step I: Deciding *what* to read (2/3)

- **Purpose:** Get up to speed on sub-field
- **Approach:** Transitive closure of related work of papers in a top conference

From: N. Feamster, A. Gray, “Reading and Reviewing Papers”
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule ¹⁰

Step I: Deciding *what* to read (3/3)

- **Purpose:** Learn about an area that is further afield
- **Approach:** Ask expert colleagues

From: N. Feamster, A. Gray, “Reading and Reviewing Papers”
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule ¹¹

Step 2: Deciding *How* to Read

- Always “top down”
 - First: Abstract, introduction, conclusion
 - Rest of paper if necessary
 - If you want to do follow-up research
 - If you want to better understand the methods/conclusions
- Next steps depend on specific purpose
 - News reading
 - Deep diving
 - Literature survey

From: N. Feamster, A. Gray, “Reading and Reviewing Papers”
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule ¹²

The Three-Pass Method (Keshav)

- **Pass I**
 - General idea
- **Pass II**
 - Basic content, but not details
- **Pass III**
 - In depth understanding

Reading a paper: Pass I

- **Bird's eye view: 5-10 minutes**
 - *Title, abstract, introduction*
 - *Section and subsection headings*
 - *Conclusions*
 - *Glance over references*

After Pass I

- *You should be able to answer:*
 - **Category**
 - What type of paper is it ?
 - **Context**
 - What other papers is it related to ?
 - **Correctness**
 - Are the assumptions valid ?
 - **Contributions**
 - What are the main contributions ?
 - **Clarity**
 - Is the paper well written ?

Reading a paper: Pass II

- **Read carefully, but ignore details** (e.g. proofs, complicated formulas, etc.)
- ~1 hour
 - Identify key points, or make comments in the margins
 - Figures, diagrams, illustrations, graphs
 - Properly labeled? error bars? etc.
 - Mark relevant unread references

After Pass II

- **Sometimes, you still don't understand the paper**
 - subject is *new* to you, *unfamiliar* terminology and *acronyms*
 - *proof* or *experimental technique* that you don't understand
 - *poorly written, unsubstantiated* assertions and numerous *forward references*
 - it's late at night and you are *tired*

From: S. Keshav, "How to Read a Paper" 17
ACM SIGCOMM Computer Communication Review, July 2007

After Pass II

- *Three choices:*
 - **set the paper aside**
 - hope that you don't need to understand the paper to be successful in your career
 - **return to the paper later**
 - perhaps after reading background material
 - **go on to Pass III**

From: S. Keshav, "How to Read a Paper" 18
ACM SIGCOMM Computer Communication Review, July 2007

Reading a paper: Pass III

- **Fully understand the paper**
 - 1(experienced)-5(newcomer) hours
 - Great attention to *detail*
 - Virtually *re-implement* the paper
 - Using same assumptions, *recreate* the work
 - Identify and *challenge* every assumption in every statement
 - Jot down ideas for *future works*

From: S. Keshav, "How to Read a Paper" 19
ACM SIGCOMM Computer Communication Review, July 2007

After Pass III

- You should be able to **reconstruct** the entire structure of the paper from memory
- You should be able to **identify** its *strong* and *weak* points
- You should be able to pinpoint **implicit assumptions, missing citations** to relevant work, and potential **issues** with experimental or analytical techniques

From: S. Keshav, "How to Read a Paper" 20
ACM SIGCOMM Computer Communication Review, July 2007

Invariant comprehension questions

- What is the **problem**?
- What are the **contributions**?
- What are the **conclusions**?
- What is the **support** for the conclusions?

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²¹
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Invariant evaluation questions

- What is the « **intellectual nugget** » ?
 - Each paper should have a single key intellectual contribution
 - Remembering this key idea will also give your brain a way to « index » the paper
- What is the main **contribution**?
 - New *finding* ?
 - New *method* ?
 - New *perspective* ?

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²²
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Reading the News

- **Conference proceedings**
 - Goal: Grasp main idea of a collection of a large number of papers. Keep informed about problems and recent solutions
- Top-Down Method
 - Skim table of contents: Papers are clustered into "sessions" which typically identify the main areas
 - Consider authors
 - Prioritize by (1) area of interest (2) reputable authors

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²³
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Deep Diving

- **Goal:** seek to understand some problem area in greater depth
- Find the seminal paper in the field
- Read carefully, including evaluation

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²⁴
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Literature Surveys

- Create the seed
 - Recent paper from top conference
 - Survey paper, if one exists
 - Seminal paper, if it is different from the above
- Perform transitive closure of cited work
 - Read related work sections of above papers

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²⁵
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Literature Surveys in the Dark

- **Step I**
 - Use an **academic search engine**, and some well-chosen **keywords**, to find 3 to 5 recent papers in the area
 - e.g. *Google Scholar*, *ACM digital library*
 - Run **Pass I** on each paper, then read their *related works* sections
 - Summary of recent work, perhaps a recent survey paper

From: T. Elsayed, "How to Read a Research Paper" ²⁶
May 2013

Literature Surveys in the Dark

- **Step II**
 - Find **shared citations** and **repeated authors** in the bibliography
 - These are the *key papers* and *researchers* in the area
 - **Download** the key papers and set them aside
 - Go to the **websites** of the key researchers, and see what they published recently
 - Permits to identify *current top conferences* in that field

From: T. Elsayed, "How to Read a Research Paper" ²⁷
May 2013

Literature Surveys in the Dark

- **Step III**
 - Go to the website of these top conferences and browse their **recent proceedings**
 - identify *recent high quality related work*
 - These papers, plus the ones you set aside in Step II, constitute the **first version** of your survey
 - Run **Pass I and II** on those papers
 - If they cite a *key paper* you did not cite earlier, obtain and read it, **iterate as necessary**

From: T. Elsayed, "How to Read a Research Paper" ²⁸
May 2013

Keeping Notes

- One-sentence summaries are infinitely better than nothing at all
- Primitive approach: Single file of notes
- Better: Database with BibTeX
 - There are some existing tools for bibliography management
 - Will also help you more quickly construct related work sections for your papers

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ²⁹
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Reference management software

- Helps you organize your bibliography
- Several paid and free solutions
 - Mendeley
 - Bibloscape
 - Endnote
 - Bibdesk
 - Bookends
 - Pybibliographer
 - Etc.

30

From Reading to Research

- A major reason to read research papers is to obtain new research ideas
- How can we arrive at new research ideas by studying papers that describe "solved problems"?

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ³¹
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

Some Questions to Ask

- **Time travel:** Will the solution apply n years from now?
- **Context switch:** Does the solution or technique apply to other problem domains?
- **Unfinished business:** Does the paper describe future work or directions? Open problems?

From: N. Feamster, A. Gray, "Reading and Reviewing Papers" ³²
http://www.gtnoise.net/classes/cs7001/fall_2008/syllabus.html#Schedule

More Questions

- **Follow Up:** Can the claims in the paper be better supported using other methods? Or, perhaps refuted?