Notes on Reading Papers

1 Surface Level: What Does the Paper Say?

As you're first reading a paper, you should keep the following questions in mind:

- What problem is this paper trying to solve?
- Why is this problem important / what is the motivation for solving this problem?
- How does this paper propose solving it?
- How does this paper evaluate its solution, and what are the criteria for success?
- What are the contributions of this paper?

Any scientific paper should answer these questions, and answers to these questions speak to the essential parts of what the paper is. Well-written papers will directly tell you the answers to these questions, often with summaries in the introduction, with phrases like:

- Problem: "...is a problem...", "A common problem in ...is ...", etc.
- Motivation: "[A particularly large group of people] encounter this problem ...", etc.
- Solution: "We propose to ...", "In this paper, we ...", etc.
- Evaluation: "We evaluate our solution by ...", "Our evaluation shows ...", "We find that [our solution] outperforms [existing solutions] ...", etc.

- Contributions: "Our observation(s) is/are ...", "Our insight is ...", "Our contributions are ...", etc.

If you cannot answer these questions, then you do not understand the most important aspects of the paper you're reading.

2 Beyond Face Value: Critical Thinking

Once you understand the core aspects of the paper, you can start to think critically about it and dissect the work. For each one of the basic questions in the previous section, you can ask further questions:

- Is this actually a problem?
- Is this actually an important, well-motivated problem?
- Does the proposed solution actually solve the problem?
- Is the evaluation designed to measure anything meaningful?
- If the evaluation does measure something meaningful, is that measurement actually indicative of a solution's success or failure?
- Are the contributions novel?

Most papers are lacking in at least one of these categories, sometimes many. These questions can be hard to answer, and are usually dependent on background knowledge. For example:

- Sometimes we have "solutions in search of problems", where the problem the work claims to solve is artificial. Wider knowledge of background information helps inform whether or not something is a real problem; if this is your first time hearing of the problem, this is suspicious.
- Many problems affect few people, and are only mildly inconvenient.
- Some papers entirely fail to evaluate their work, or provide only anecdotal evidence that the solution is successful.

- Some papers will collect metrics which do not correspond to anything meaningful, or do not indicate anything about the solution itself. Knowledge of alternative evaluation strategies is helpful here.

- Sometimes work has been done elsewhere, perhaps in a very different context. Generally the only way to know this is familiarity with background information.

Just because papers are lacking in one or more areas, it does not immediately invalidate the work, nor does it mean the authors are lazy. Publications generally happen when there is some worthwhile bump in knowledge which others can build on in an incremental fashion. Individual bumps are imperfect, sometimes severely so. Collectively, these bumps form a solid whole.

Research often happens around these sorts of critical thinking questions. For example:

- If you think a paper's solution only partially solves a problem, a more complete solution may be a novel result.

- If you think a paper performs the wrong evaluation, an experiment performing what you think is the right evaluation may get surprising, publishable results.

- If you think a series of related papers consistently perform the wrong evaluation, this could be a symptom of a wider problem. Explicitly exposing this problem and deriving a better evaluation can be a publishable result.