The Importance of Good Writing in Research (and CEO-level memos)

Daniele Quercia

1 Introduction

Writing isn't some pointless fluff you do at the end of a research project - it's half the battle. Actually, it's more than half. Bad writing will bury even the best ideas. Here's how it works:

- **Good ideas + garbage writing** = Rejection. Straight in the bin.

- **Mediocre ideas + solid writing** = Acceptance. Welcome to the club.

- **Brilliant ideas + great writing** = Not just acceptance, but a lot of citations. People actually read it!

Do yourself a favor and spend at least 50% of your research time making sure your paper doesn't read like a toddler's diary.

2 Team Activities to Improve Writing

At our team, we've got a couple of tricks up our sleeve to help you polish our work:

Reading Groups. Two reasons we do this: to steal good ideas and to figure out how good papers are actually written. Here is the golden rule:

1. Read well-written papers and **copy their style**. That's right, copy. Don't reinvent the wheel.

2. Ignore most papers because, let's be honest, they're not great.

Paper Clinics. Your colleagues are there to roast your paper before the reviewers do. Here is the best practice: *The better your draft, the better the feedback. If your first draft looks like nonsense, expect comments like, "What is this?"*

3 Writing Your Paper: A Step-by-Step Guide

Before You Start Writing. Read Chapter 3 of *The Pyramid Principle*. .

Fixing Dodgy Sentences. If your sentences make no sense:

1. Read The Science of Scientific Writing;

2. Check Chapter 10. of *The Pyramid Principle*. If you can't visualize what you're writing, neither can your reader. Draw it first, then write.

Structuring Your Paper. The typical structure include:

1. **Abstract & Introduction** There's one structure that actually works. Follow it. It's basically what reviewers use to judge your paper:

SITUATION: Problem X is very important because . . .
COMPLICATION: In tackling problem X, related work failed in doing Y
PROPOSAL: To partly tackle Y, we make N contributions [list of contributions]

Golden rule: Before writing any abstract (and intro), you need to know what X, Y, and N are. More than one X or Y? You don't know what you are writing about. Too vague? Nobody cares.

2. **Related Work**

   - Don't review everything under the sun. Focus only on Problem Y.
   - Keep it to one page. If it drags on, you're overcompensating.
   - End with this line: *"To sum up, previous work has failed to address Y."* Boom.
   - Can't pinpoint a clear Y? **Rewrite your Abstract/Intro/Related Work.**

3. **Methods.** Here you explain how you've solved the problem.

4. **Evaluation.** Here you explain how you've tested that your solution actually works. This could be also called "Results" section. This section typically shows the results of your experiments and is generally structured as:

   - **Goal.** The goal of my system is to . . .
     (example from Section 4 in [1] "The goal of our algorithm is to predict trust ratings on portable devices.)
   - **Metrics.** To ascertain whether my proposal meets that goal, we answer the following questions (each question is answered by a metric): . . .
     (example in [1]: "To ascertain the effectiveness of our algorithm at meeting this goal, our evaluation ought to answer three questions:
     (a) (Predictive Accuracy) How accurate is our algorithm in predicting trust ratings?
     (b) (Prediction Robustness) What is the impact of uncooperative users upon the algorithm's accuracy?
     (c) (Overheads) What time, storage, and communication overheads does our algorithm impose on a mobile phone?
     . . . ")
   - **Evaluation Setup.** To measure my metrics, we gathered data and went through the following steps . . .
     (example from Section 4 in [1]: "Using this data set (described next), we evaluate whether our algorithm is effective in predicting real trust ratings (Section 4.1). Then, to evaluate how robust our algorithm is, we emulate how users may rationally turn to be uncooperative (Section 4.2). Finally, we implement our algorithm to assess whether it is usable on a mobile phone (Section 4.3).")
   - **Evaluation Results.** We find that our metrics "go up and down" because . . . [explain why].

5. **Discussion.** Here you discuss how your results are: (1) in-line with previous work; and (2) differ (expand) on previous work. Also, you can list the limitations of your work.

Final Draft Check. Before "embarrassing" yourself, run through *How to Read an Engineering Research Paper*: PDF link.

How Long Should This Thing Be?

- CS papers: about 12 pages. More? You better have a good reason.

- A survey paper? Fine, stretch it out a bit.

Additional Resources. Here are two amazing resources:

- *How to Write a Great Research Paper*: Microsoft link.

- **Video Lecture of the previous slides**: Watch on YouTube.

Final Thought: If your writing is unclear, people will assume your research is bad. Fix it.

References

[1] Daniele Quercia, Stephen Hailes, and Licia Capra. Lightweight Distributed Trust Propagation. In *ICDM 2007*.