CMPT 983 Fall 2022

Martin Ester

How to Write a Good Machine Learning Paper

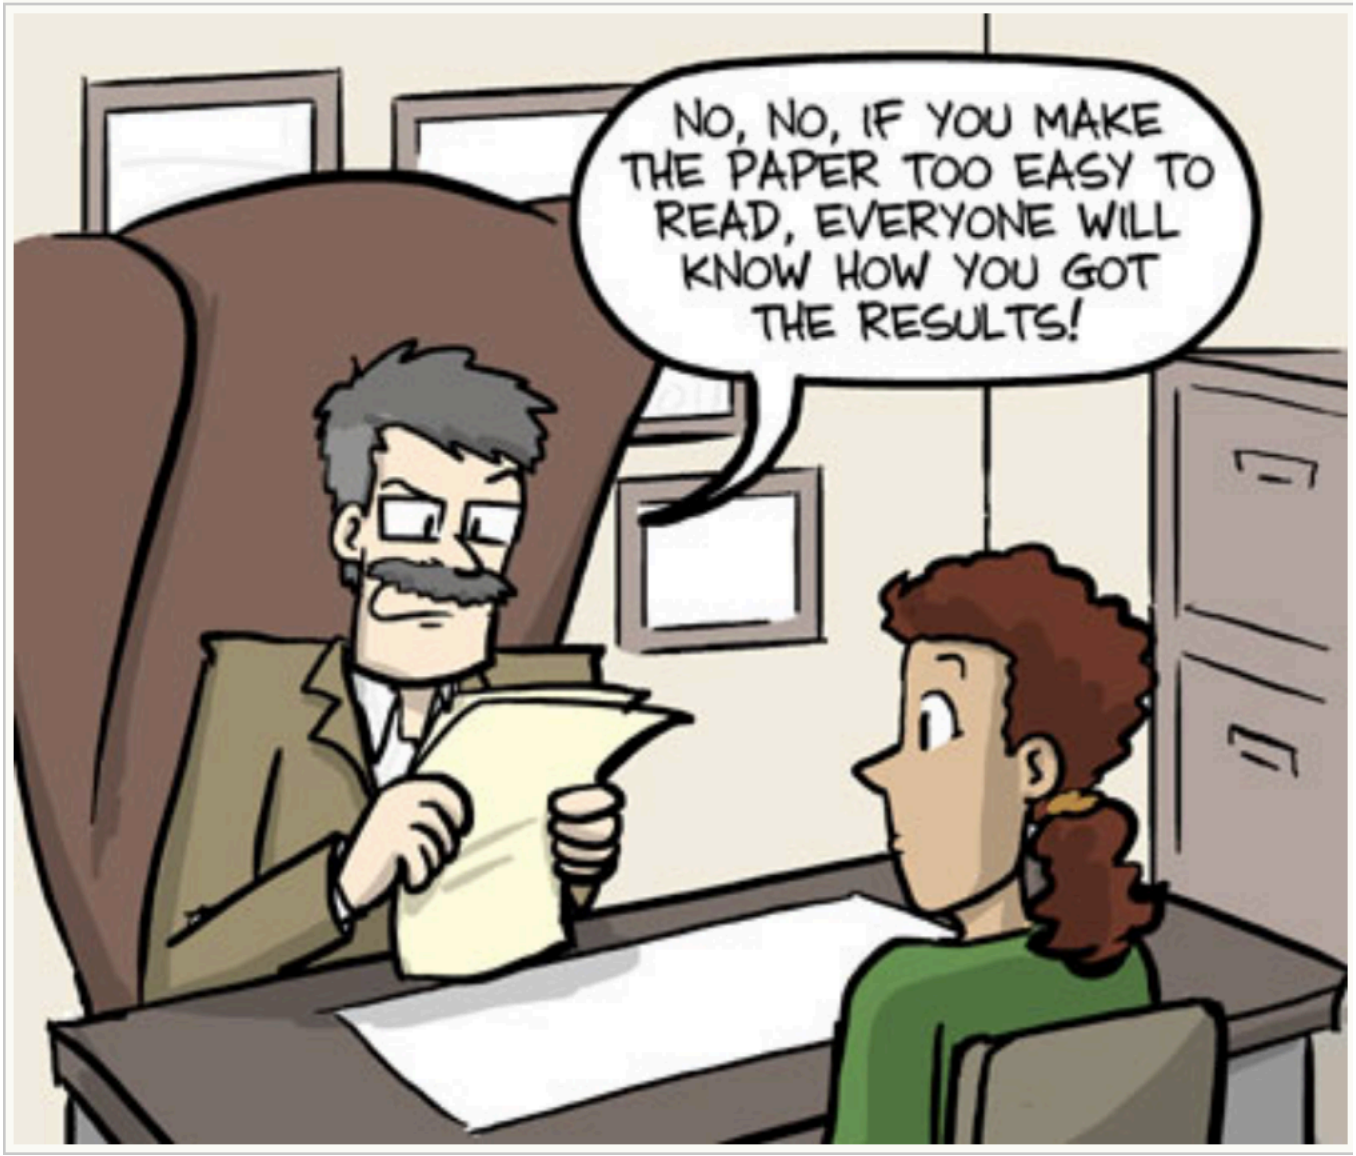
NO, NO, IF YOU MAKE THE PAPER TOO EASY TO READ, EVERYONE WILL KNOW HOW YOU GOT THE RESULTS!

Introduction

Bad papers

- Focus on the authors instead of on the readers.
  - fail to clarify the motivation for the work or
  - include unnecessary details.
- Attempt to impress the readers rather than inform them.
- As a result, they are interesting to or understandable by only a small set of highly specialized readers.

Good papers

- Focus on the readers.
- Interesting and useful to many readers, including newcomers to the field.

→ Examine your attitude!

→ Put yourself in the shoes of the reader!

Introduction

Good papers tell a good story

- Narrative
  - a spoken or written account of connected events
  - a representation of a particular situation or process in such a way as to reflect an overarching set of aims or values
- Metaphor of a hike

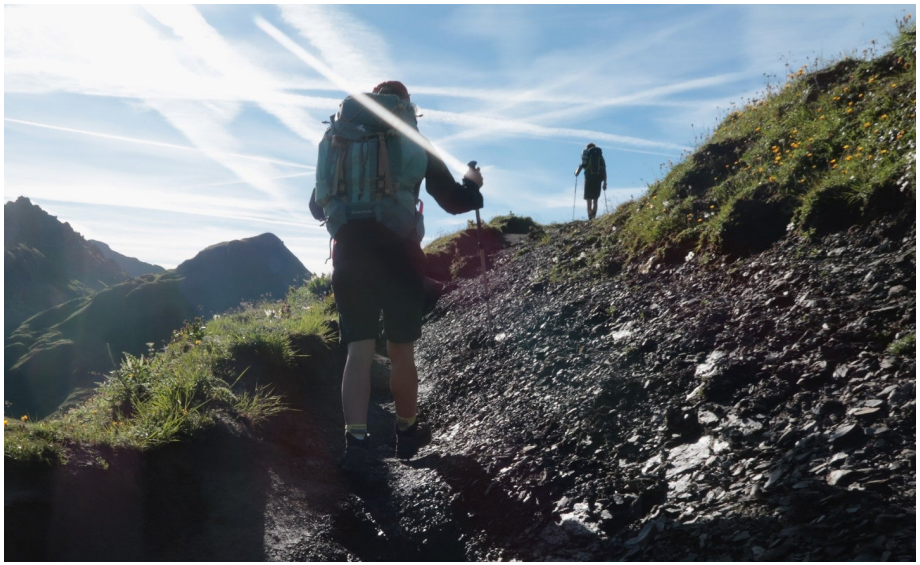

Structure of a Science Paper

- Introduction

  Clarifies the research question (motivation) for the work and prepares readers for the structure of the paper.
- Materials and Methods

  Provides sufficient detail for other scientists to reproduce the experiments presented in the paper.
- Results and Discussion

  Presents and discusses the research results.
- Conclusion

  Summarizes the outcome of the work by interpreting the findings at a higher level of abstraction and by relating these findings to the research question stated in the Introduction.

→ The findings are more important than the methods used!

Structure of a Computer Science Paper

- Title
  Catchy and indicative of your research contribution.

→ The longer a paper title, the lower its acceptance chance.

- Abstract
  A summary of the research problem, your claim, and the evidence.
- Introduction
  Background, motivation for your work, significance, an outline of the rest of the paper.
- Related work
  - critical review of the existing approaches,
  - discussion of relationship to your own work
    commonalities and differences

Structure of a Computer Science Paper

- Problem statement and algorithm design
  Describe the problem, the approach of the solution and the technical details of the solution.
- Evaluation
  Evidence to support the claim of your research contribution. Unless you can provide proofs for a theoretical paper, experimental results are always expected.
- Conclusion
  A summary of the research contribution, a discussion of its significance, and a mention of future work.
- References
  List and cite related work.

→ The method proposed is more important than the experimental results!

Introduction

- Makes the "claim": how does the paper advance the state-of-the-art in its discipline?

→ Tells a short version of the "story".

- The most important section of the paper.
- Most readers will only read this section.
- A reviewer may reject a paper based only on the Introduction.

→ May not read the rest of the paper (carefully)!

- The paper will be accepted only if the rest of the paper provides enough evidence to support the claim.

Introduction

Possible claims

- Your technique solves a new problem for the first time.
- Your technique performs better, in one or more of the following dimensions, than its rivals:

- Behaviour
  is more accurate, easier to understand, more similar to human outputs, easier to program, etc.
- Coverage
  is applicable to a wider range of domains.
- Efficiency
  is faster or uses less space.
- Usability
  Users find it easier to use than its rivals.

→ You should avoid claiming too many dimensions, but one or two with in-depth evidence.

Introduction

- Provide some background to orient those readers who are less familiar with your topic and to establish the importance of your work.
- State the need for your work, as an opposition between what the scientific community currently has and what it wants.
- Indicate what you have done in an effort to address the need (this is the claim / contribution to advance the state-of-the-art).
- Preview the body of the paper to mentally prepare readers for its structure.

Introduction

Checklist

- What exactly is the problem being solved?
- Why is the problem significant?
  - Has it not yet been solved?
  - Is it hard?
  - Does it have important applications?
- What are the state-of-the-art methods to solve the problem?
  What are their strengths?
  What are their limitations?
- What is the idea of your method?
- Why do you expect your method to outperform the state-of-the-art methods?
- How have you evaluated your method?
  And what were the main results?

Method

- Explain your solution first at an intuitive level, with examples.
- Use figures to illustrate the intuition.
- Do not simply put your algorithms in pseudo code.
- Then provide enough technical detail so that the interested reader can understand the solution: definitions, formulas, pseudo code, . . .
- Highlight your contributions. Show your novelty.
- Define terms before using them.
- Use meaningful and specific terms, not generic ones.
- Use the same term for the same thing consistently throughout the paper.
- Minimize the number of acronyms.
- Give your method a name and refer to it by that name.

Experiments

- Experimental design
  Goal of the experiments is to provide evidence supporting the claims of the paper.
  What research questions are the experiments going to answer?
- Aspects of the experimental design
  - Datasets
    Including real-life datasets
    Possibly also synthetic datasets
  - Baselines
    State-of-the-art methods for the same / similar problem
    The method that you are extending / improving
  - Metrics
    Quantitative measures of the performance
    Qualitative measures of the performance

Experiments

- Results
  Quantitative and qualitative results.
  Use figures and tables.
- Interpretation
  Summarize the main findings.
  Explain why your method outperforms the baselines.

→ Relate it to the motivation in your Introduction.
  Be fair to the baselines.
  Do not exaggerate the gains of your method.

→ No method consistently outperforms all state-of-the-art methods in all scenarios!

Experiments

- Enough details for your experimental design so that other researchers can verify your results.

→ The requirements on repeatability are becoming stricter!

- Why and how have you made the choices for your experimental design?
- Are the chosen baselines state-of-the-art?
- Are the experimental results consistent and conclusive?
- How did you tune the hyper-parameters of your method and of the baselines? How do these hyper-parameters impact the performance of your method?
- How do the experimental results correspond to the motivation of the paper?
- What have you found surprising and hard to explain?

NeurIPS Evaluation Criteria

Quality

- Is the paper technically sound?
- Are claims well-supported by theoretical analysis or experimental results?
- Are the authors careful (and honest) about evaluating both the strengths and weaknesses of the work?

Clarity

- Is the paper clearly written?
- Is it well-organized?
- Does it adequately inform the reader?
  A superbly written paper provides enough information for the expert reader to reproduce its results.

NeurIPS Evaluation Criteria

Originality

- Are the problems or approaches new?
- Is this a novel combination of familiar techniques?
- Is it clear how this work differs from previous contributions?
- Is related work adequately referenced?

Significance

- Are the results important?
- Are other people (practitioners or researchers) likely to use these ideas or build on them?
- Does the paper address a difficult problem in a better way than previous research?
- Does it advance the state of the art in a demonstrable way?
- Does it provide unique data, unique conclusions on existing data, or a unique theoretical or pragmatic approach?

NeurIPS Evaluation Criteria

Repeatability

- Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?
- Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?
- Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?
- Did you include the amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?

NeurIPS Evaluation Criteria

Potential negative impact

- Examples:
  - Better generative models used to generate Deepfakes for disinformation.
  - Deployment of ML methods that are biased against certain sub-populations.
- Consider different stakeholders that could be impacted.
- Consider possible harms that could arise when the technology is being
  - used as intended and functioning correctly,
  - being used as intended but gives incorrect results, and
  - being misused (intentionally or unintentionally).
- If there are negative societal impacts, you should also discuss any mitigation strategies.

My Advice

- Read other people's papers from the venues that you are targeting.

→ Gain valuable experience as paper reviewer!

- Choose a topic that has the potential for a lot of positive impact and little negative impact.
- Presentation is becoming more and more important. Spend a lot of time on good presentation!

→ Reviewers spend little time on your paper.

- Discuss related work carefully: commonalities and differences.

→ Improper discussion (or omission!) of related work is perhaps the most frequent argument for rejecting a paper.

- Start by writing the Introduction.
- Write and implement / evaluate in parallel.

→ The writing of the paper is the research!

- Get friendly reviews and revise based on their feedback.

References

English Communication for Scientists, Unit 2: Writing Scientific Papers, https://www.nature.com/scitable/ebooks/english-communication-for-scientists-14053993/writing-scientific-papers-14239285

Xindong Wu: How to Write and Publish Top-Quality Research Papers, http://www.cs.uvm.edu/~xwu/PPT/Writing16-Web.pdf