

How to Write a CS Paper

Voicu Popescu

Overview

- A lecture on how to write CS research papers
- A systematic approach—a recipe, a formula, an algorithm

Motivation

- Writing a paper is difficult
 - Complex topic
 - New results
- Paper writing rarely taught explicitly in graduate school
 - Learned by reading papers
 - Learned through painful trial and error

Misconceptions about paper writing

- *“Writing a paper takes a couple of hours”*
 - No. It takes an experienced writer a week w/ sleep and 36h w/o sleep to write a paper.
- *“Writing a paper takes literary talent”*
 - No. Keep poetry and metaphors out of the paper.
- *“Writing a paper is a mysterious, amorphous process”*
 - No. There is a method for writing papers.
- *“English proofreading services can fix a poorly written paper”*
 - No. English proofreading fixes language problems, not exposition problems.

When to start writing

- Option 1: once you have proof of concept
 - Pro: plenty of time available for writing
 - Con: not all results available, writing has to anticipate results, writing cannot accurately emphasize strengths demonstrated in results
 - Recommended for conference submissions, and for novice writers
 - Might require a second writing pass (i.e. a major revision) to fine tune paper to final results

When to start writing

- Option 2: once all results are obtained
 - Pro: writing reflects results with high fidelity, including in abstract and in introduction
 - Con: little time available for writing, due to imminent (conference) deadline
 - Recommended for conference submissions for experienced writers, and for journal submissions (no hard deadline)
 - *Warning: can lead to submission delays*

Formatting

- Use template provided by targeted venue
 - Word
 - LaTeX
- Format from the beginning
 - Accurate estimate of paper length
 - Avoids formatting nightmares close to the deadline

Tell a story

- A well written paper tells a story
- The story has to
 - flow from the “introduction” section all the way to the “conclusions and future work” section
 - be easy to read
 - be exciting
 - clearly state contributions
 - not overstate contributions
 - provide sufficient detail for reproducibility
 - not follow the work timeline proportionally

Tell a story

- The story has to
 - reiterate important points (title, abstract, introduction, method, and conclusions) without being repetitive
 - be consistent, no contradictions
 - contain no ambiguities; no “would”, “could”, “should”, “might”; everything described outside the future work section should have been actually implemented; no speculations

Figures

- Whenever something is hard to describe, use a figure (i.e. diagram, image, graph)
- Have enough figures, with detailed captions
 - Someone looking only at figures should get the main idea of the paper
- Figures should be of very high quality
 - Use professional software, e.g. Visio
 - Be prepared to invest time (multiple hours, revisions)
 - Start with canvas of final size
 - 8pt font in the final paper layout (no scaling)

Philosophy

- Your method is assumed to be bad until you prove that it is good
- Your paper is assumed to be rejected until you prove it has to be accepted
- It is not enough to not provide good reasons for the paper to be rejected
- You have to provide good reasons for the paper to be accepted

Paper components

- Title
- Authors list
- Abstract
- Keywords
- Introduction
- Prior work
- Method overview
- Method details 1
- Method details 2
- ...
- Results and discussion
- Conclusions and future work
- Acknowledgments
- References
- Appendices
- Video

Title

- Important
 - First thing a reader sees
 - Together with abstract and keywords used to decide reviewers
- Desired qualities
 - Informative
 - Accurate
 - Not too long
 - Catchy, easy to remember, impressive
- Formatting
 - Capitalize every word except for prepositions
 - “Reflected-Scene Impostors for Realistic Reflections at Interactive Rates”

Title architecture

- Most frequently
 - Nickname: New-Thing for What
 - “The WarpEngine: An Architecture for the Post-Polygonal Age”
 - “GEARS: A General and Efficient Algorithm for Rendering Shadows”
 - New-Thing for What
 - “Simplification of Node Position Data for Interactive Visualization of Dynamic Datasets”
 - “Reflected-Scene Impostors for Realistic Reflections at Interactive Rates”
 - What by (using) New-Thing
 - “CAD Visualization by Outsourcing”

Title architecture

- New-Thing
 - A new paradigm; radically new approach to solving a problem or set of problems
 - “Forward Rasterization”
 - “Camera Model Design”
- What
 - A breakthrough: finally a solution to a long standing problem
 - “Efficient Large-Scale Acquisition of Building Interiors”

Authors list

- Typically sorted on contribution
 - Rarely done alphabetically (in our field)
- First author should
 - Understand all the work reported in paper
 - Be able to present the paper
 - Know how every aspect of the method works
- Collaborators to include
 - Anyone who has contributed a significant idea
 - This leaves out those whose contribution is exclusively in the implementation, in making figures, or in collecting data (they go in acknowledgment section)

Abstract

- The longer type of abstract
 - Two paragraphs
 - First paragraph
 - Problem
 - Problem importance
 - Why problem is difficult
 - Limitations of state of the art
 - Second paragraph
 - Brief description of method contributed by paper
 - Method scope (i.e. input for which it works, assumptions)
 - Brief description of method evaluation
 - Results highlights

Abstract

- The shorter type of abstract
 - Just the second paragraph of the longer type
 - Brief description of method contributed by paper
 - Method scope (i.e. input for which it works, assumptions)
 - Brief description of method evaluation
 - Results highlights

Abstract

- Length of abstract is usually regulated
- Abstracts are expected to be dense
 - Start from something twice as long and condense
 - Tip: you could write the introduction first and then condense that into an abstract

Keywords

- Used to determine reviewers
- Used for readers to find your paper in future
- Some conferences / organizations (e.g. ACM) provide list to choose from
 - Choose carefully
 - Add your own if at all possible
- Sort based on generality
 - Usually ascending order

Paper components

- Title
- Authors list
- Abstract
- Keywords
- **Introduction**
- Prior work
- Method overview
- Method details 1
- Method details 2
- ...
- Results and discussion
- Conclusions and future work
- Acknowledgments
- References
- Appendices
- Video

Introduction

- The **most important** part of the paper
 - Often the only part of the paper a reader/reviewer will read closely from beginning to end
 - Many reviewers decide on acceptance by the end of the introduction and use the other sections as a source of evidence for their decision
 - Be prepared to spend a long time writing it (one day) and revising the introduction (throughout the writing process)

Introduction formula

- Five plus two paragraphs
- Together with title, teaser figure, author list, keywords, abstract should cover at most the first two pages of paper.
- Paragraph 1
 - Problem
 - Problem importance

Introduction formula

- Paragraph 2
 - Why is problem hard?
 - Summary of prior work approaches and of their shortcomings
 - OK to have references
 - I prefer not to have references
 - Ask reader/reviewer to extend their trust until prior work section where all prior work claims are backed up with references
 - This allows reader/reviewer to focus on story

Introduction formula

- Paragraph 3
 - Details on shortcomings of prior art that take similar approach as taken by present paper
 - What are the problems that need to be solved, for the approach to succeed?
 - This should lead to insight that created method described in current paper. Clearly understanding the problem, in detail, leads to inspiration, to good idea.

Introduction formula

- Paragraph 4
 - Introduce method presented by paper
 - Start with “insight”, “inspiration”, “key observation”
 - No implementation details, just high level ideas and concepts used

Introduction formula

- Paragraph 5
 - Summary of examples where method was tested
 - Summary of results
 - If you have an accompanying video, mention it explicitly—otherwise reviewers might miss the video!

Introduction formula

- Paragraph 6 (optional)
 - List of contributions
 - At least two, at most three, bullets recommended
 - Simplifies reviewer's job finding the contributions (they are asked by the review form to list contributions)
 - Well written paragraphs 4 and 5 could make this paragraph unnecessary
 - Reviewers could be annoyed by the list of contributions
 - contributions of a well written strong paper are self-evident
 - explicit list of contributions can be interpreted as an attempt to manipulate reviewers

Introduction formula

- Paragraph 7 (optional)
 - Paper organization (list section titles and what each section does)
 - More useful when there are multiple “method details” section (i.e. longer papers)
 - Usually omitted for shorter papers

Prior work

- One of the most boring sections to a reader
 - Typically very poorly written
- Prior work section should be
 - Well organized
 - Comprehensive
 - Relevant to paper at hand
 - Fair

Prior work

- Convince reviewers that are expert in the area that you too are an expert in the area
- Help reviewers outside the area catch up on the state of the art
- Nothing worse than a poorly written prior work section
 - No knowledge of prior work
 - No understanding of prior work
 - No good delimitation of the contributions of the current paper

Annotated bibliography

- You write a little bit of the prior work section every time you read a paper
 - Collect an annotated bibliography
 - For every paper you read
 - Collect the citation
 - Write a summary paragraph
 - Write a strengths paragraph
 - Write a weaknesses/limitations paragraph
 - The annotated bibliography will be an invaluable help when writing prior work sections, your thesis, etc.
- Start from recent major conferences and venues
- Take one step back (i.e. look at their references)
- Take several steps back for the most relevant work

Prior work

- Organize prior work section on approaches
 - Define each approach
 - Cite early, recent, and best known paper for each approach
 - For each paper cited write a sentence
 - On what it does
 - Another one on what it excels at
 - And another one on its shortcomings
- End approach discussion with summary of strengths and weaknesses
 - If your paper takes different approach, contrast approaches
 - If your paper takes same approach, contrast your method with other methods in the approach
 - Devote more space to the approach to which your method belongs

Prior work

- Do not reuse prior work from other papers
 - Prior work section should be designed and detailed for the present paper
- Prior work section should be about one page
 - You never lose points for too many references
 - You can lose points if references are not enough
 - However, the total length of the paper has to be commensurate to contribution
 - Prior work can be condensed
 - Do not use a reference as a noun
 - “[2] describes a method”, “same approach as in [2]” are incorrect

Paper components

- Title
- Authors list
- Abstract
- Keywords
- Introduction
- Prior work
- **Method overview**
- Method details 1
- Method details 2
- ...
- Results and discussion
- Conclusions and future work
- Acknowledgments
- References
- Appendices
- Video

Overview

- Gives a high-level view of your entire method
- Use a diagram
 - Blocks for the various stages of your method
 - Arrows indicating the data flow
 - Label arrows with the type of data
- Use a pseudocode description of the main steps of your algorithm
- Each stage or step is later described in a section
 - Refer to the future section

Overview

- Gives reviewers essential help
 - Reviewers volunteer their time
 - You are responsible for making their job as easy as possible
 - Do not expect reviewers to spend hours and hours trying to make sense of your poorly written paper
 - Reviewers will simply say in the review: “I tried but I could not understand the paper, and I am an expert in the area; what chances does a regular reader have?”

Method details k

- These sections are the easiest ones to write
 - It's your work, it's what you did, you know it all too well
 - You love what you did, and you can't wait to tell people about it
- Level of detail
 - Sufficient for a skilled graduate student to reproduce your work
 - Not overly verbose—concise and to the point
 - No innovation should be left unexplained
 - No simple implementation details should be provided

Method details k

- Use references when you use an existing tool
 - Make sure you explain what the algorithm/tool does
 - OK to summarize (in one sentence) how the tool does it to make paper self contained
- Use figures
- Use present tense

Method details k

- Remember, do not use “can, could, should, would”
 - Nothing worse than giving the reviewer an uneasy feeling that some of the work described is only proposed and that it was not actually done
- Do not overuse “very”, “highly”, they end up weakening what is claimed
 - E.g. “very accurate” is less accurate than “accurate”

Method details k

- Double-blind review
 - You cannot disclose your identity
 - OK to reference your prior work
 - Use third person
 - “they did this and that” not “we did this and that”
 - Do not include 10 references to your work
 - It will amount to a blatant disclosure of your identity

Paper components

- Title
- Authors list
- Abstract
- Keywords
- Introduction
- Prior work
- Method overview
- Method details 1
- Method details 2
- ...
- Results and discussion
- Conclusions and future work
- Acknowledgments
- References
- Appendices
- Video

Results and discussion

- You talked the talk, now you walk the walk
- Everything you promised has to be substantiated by results
 - High quality should be supported by high quality
 - Interactive rates should be supported by interactive rates
 - Overcoming shortcomings of prior art should be supported by a favorable comparison to prior art
 - Any discrepancy substantially weakens the paper

Results

- First paragraphs
 - Describe applications and scenes where you tested your method
 - Describe machines on which you collected timing information
- Subsection 1: quality
- Subsection 2: performance
- Subsection 3: comparison to prior art
- Subsection 4: limitations

Results and discussion: quality

- Provide evidence as to how well your method works
- If your method resorts to approximation, resort to truth

Results and discussion: performance

- Measure performance accurately
 - Relevant data sets
- Measure performance thoroughly
 - Identify parameters affecting performance and measure performance for various values
 - Discuss numbers obtained; discuss best and worst cases
 - When appropriate derive asymptotic cost of your method
- Show performance with graphs and tables

Results and discussion: performance

- Give some information on implementation
 - High level, do not give boring details
 - Get into details only if you did something very clever that brought a lot of performance gain
- Remember
 - Paper does not cover linearly the work you put in
 - Things that took months to implement might not even be mentioned

Results and discussion: comparison to prior art

- Try to find implementations of most prominent prior art methods
 - It saves you having to implement them
 - It brings more credibility to the comparison
 - Ask authors if they are willing to share their code
- Show quality and performance differences
 - Conduct a thorough analysis
 - Do not avoid cases where your method doesn't do so well
 - Performance analysis for same quality
 - Quality analysis for same performance
- Discuss the comparison
 - Explain the differences
 - Explain the tradeoffs—e.g. more speed, less quality

Results and discussion: limitations

- Reviewers have to list the limitations of your method
- A strong paper is expected to self-report its limitations
- Fundamental limitations, which you might inherit from the general “approach” taken, and say so
- Limitations specific to your method, explain what you gain for those limitations, i.e. the tradeoff
- Be unapologetic—your method works for some types of input, and it’s OK that for some it does not
- Explain how some limitations might be removed through future work

Conclusions and future work

- Closing arguments in defense of your paper
 - Closing statement. The last time you talk to reviewers
 - Remind them how good your paper is
- State one more time very succinctly what the method does
 - Emphasize the strengths
 - Emphasize the difference to prior art
- Summarize the comparison to prior art one more time

Conclusions and future work

- Sketch directions for future work
 - Short term fixes and extensions were already mentioned in the limitations subsection
 - Do not make it sound like “paper is incomplete, but accept the paper please, and we promise we will do all these things”
 - Think big and think far into the future
 - Big improvements
 - Applications of method to new contexts

Acknowledgments

- Withheld for double-blind reviews
- Acknowledge all who helped, in decreasing order of contribution
- Acknowledge your group
- Acknowledge your sponsors

References

- Format well
- Do not include references not used in paper
- Include all references used in paper
- Sort according to instructions (appearance, alphabetically)

Appendices

- Put in an appendix text that is not essential to the exposition
 - Proofs
 - Additional results tables
 - Comments from users
 - Questionnaire used in user study
- Do not put in an appendix anything that you want to make sure a reviewer reads

Paper components

- Title
- Authors list
- Abstract
- Keywords
- Introduction
- Prior work
- Method overview
- Method details 1
- Method details 2
- ...
- Results and discussion
- Conclusions and future work
- Acknowledgments
- References
- Appendices
- **Video**

Video

- Typical but not unique to graphics papers
- A lot of additional work
- It can take as long as writing the paper
- Video and paper need to be consistent
 - Emphasis
 - Method description
 - Result illustration
- Title, introduction, and results of paper on one hand and video on the other hand are strongly interdependent

Video

- Length
 - At most five minutes
 - Some conferences have limits, usually 5min
 - Reviewers lose patience
 - 5min are enough to make your point

Short video

- Video components
 - Best results

Medium video

- Video components
 - Split-screen two-way comparison between method and prior art
 - Or, split-screen two-way comparison between method and truth
 - Additional examples of method

Long video

- Video components
 - Limitations of prior art
 - Preview of best results
 - Illustration of proposed method
 - Split-screen two-way comparison between method and prior art
 - Split-screen two-way comparison between method and truth
 - Or Split-screen three-way comparison between prior art, method, and truth
 - Additional examples
 - Conclusion

Video

- It's not an action movie!
 - Camera should move very slowly, and even slower in the case of split screens
 - The sequences should be as long as possible
 - Go back and forth several times to make important points
 - Put a red box around an important detail you want to make sure the viewer sees
- For real-time methods include a real-time sequence
 - Side by side comparisons should be done from stills for perfect synch

Video

- Audio voice over is essential
 - Video is difficult to understand without audio
 - Use audio to guide the viewer's attention to the most important qualities of your method
 - Audio has to be well synchronized to video
 - Mentioning a concept should slightly precede the visual illustration of the concept
 - Audio script should be well aligned with paper introduction, results, and conclusions

Thank you

- Good luck with paper writing
- If this lecture was helpful, acknowledge me in your paper!