

*CS 2001: Research Topics*

# **Writing up Your Research**

Prof. Adriana Kovashka  
University of Pittsburgh  
October 6, 2021

# Plan for this lecture

- Structure of a paper
- Getting started
- Timeline
- Writing style
  - Consistency, aesthetics, emphasis, inclusivity
- Making paper fit the limit
- Common mistakes
- Activities: Examples, accepted/rejected versions, literature reviews

# Homework (Reminder)

- Write a literature review on a topic of your choice
- Identify a claim you want to make: e.g. “X is a good approach”
- In class, we will pair up and give each other feedback on the literature review
- Aim for a 3-5 page review
- Due on Oct. 11 before class! – be ready to share with a partner (also upload to Canvas)

# Mechanics: Overleaf (Demo)

# Reminder: Structure of a Paper

- Introduction
  - Motivation
  - Gaps in science, deficiencies in prior work
  - Brief overview of key ideas
- Related work
  - Group by topic, state how proposed work different
- Approach
  - Explain clearly each step
  - Distinguish standard aspects from your contribution
- Experiments
  - Describe setup, data, metrics, etc.
  - Compare to papers from the recent literature
  - Verify contributions of each method component through ablations

# Getting started

- Download conference/journal template
- Set up Overleaf (project directory)
- Write extended outline
  - What will each paragraph say? – summarize with one bullet
  - (Show to advisor, iterate based on feedback)
- Then fill in details
  - Show to advisor, iterate based on feedback

# Timeline

- You don't have to wait for all results to be finished before you start writing
  - Write parts of intro, including motivation and claims, as early as you have a clear idea—helps you identify the work you still need to prove claims
  - Write related work—you might find more relevant baselines, methods to compare to
  - Write parts of approach—good way to make sure approach is sound, and your advisor is on the same page
  - Write plan for experiments—to help identify what experiments still need to be conducted
- Iterating on writing with advisor can take 2-4 weeks for a first-time submission, so leave plenty of time!

# Writing Style

- Your advisor may be used to writing papers in a particular way
- A certain style may be common in your research community (and not others)
- Talk to you advisor about their writing style
- Talk to them about their pet peeves
  - “minor annoyance that an individual finds particularly irritating” (Wikipedia)
  - These might be about consistency, aesthetics, etc.



# Consistency

- Strive to describe and run things in a way that is consistent across methods etc.
  - Don't run one method with one set of settings and another with other settings, unless you have a reason
- Strive to describe and format text consistently
  - Don't use Capitalization in one section Title but not Another
  - Don't use bold for emphasis in **one section** and underline *for another*
  - Use the same string for the same conference, not "CVPR" and "32<sup>nd</sup> CVPR" and "cvpr"

# Aesthetics

- Make sure tables and figures are legible
  - Font large enough
  - Colors in lines in figures different enough
- Make sure it's easy to navigate through paper
  - Provide section headers, paragraph headers, where appropriate

# Inclusivity

- Use neutral pronouns – “they” rather than “he”/”she”
- Avoid potentially hurtful terms, e.g. “master and slave”
- Consider including examples that a broader audience will find intuitive: do not assume understanding of any particular culture
- Use colors everyone will read easily; avoid red/green and blue/purple

# Emphasis

- Writing is a craft—think of yourself as a woodcarver—every detail matters
- Lose sentences (in intro) that are not crucial for the point you're trying to make—you're just losing the reader's attention and an opportunity to emphasize your contribution
- Be careful about *how* you say what you—what is the accent?

# Making Paper Fit the Page Limit

- Your draft is 9-12 pages but limit is 8 pages—no problem!
- Remove papers from related work that aren't all that related, or trim your descriptions of them
- Read through whole paper, and rephrase sentences as needed to make your narrative more concise, *punchier*
- Revise paragraphs ending in a few words (trailing lines)
- Remove unnecessary words and characters
  - Examples: “first X and second X” -> “first and second X”, “that”
  - Change tense to lose characters, replace subsection headers with paragraph headers, etc. — it's easy, but it takes practice!
- Last resort: use negative vspace `\vspace{-2mm}`
  - May be disallowed by conference
- You will take many passes through your paper, and each time, you'll find ways you can say it more concisely

# Note on Equations

- Explain all variables
- Explain purpose, idea in English
- Refer to equations by name rather than just ID later in the text

# More from readings: Jones

- Don't wait: write
  - Idea -> write -> research
- Identify your (one, sharp) key idea
  - “infect mind of reader with your idea, like a virus”
  - “can you hear the ping?”
- Tell a story
- Nail your contributions to the mast
  - Contributions should be listed, and should be refutable
- Related work later
- Put readers first
  - Present the intuition, examples, then details
- Listen to your readers
- Use simple, direct language

# Common Mistakes in Writing

- Starting introduction with platitudes (“deep learning has greatly improved the state of the art...”)
- Too long related work, plus not saying what’s different about proposed work
- Method mixes standard techniques with novel contribution
- Not clear what motivation, key contribution are
- Results—unclear where gains in performance come from
- Results—focusing just on the negative outcomes—results won’t be perfect, but the question is whether there is evidence that your proposed idea has benefits



# Other Paper Submission Logistics

- Paper ID
- Title
- Author list and order
- Subject areas

# Writing Rebuttals: Audience

“We write rebuttals for two different audiences

1. The reviewers, who have read your paper (to varying degrees), but may have forgotten some of the details or didn't understand them in the first place.
2. The AC, who is likely even less familiar with your work, and a good guiding principle is to assume that all they will read is the set of reviews and the rebuttal.”

# Writing Rebuttals: Goals

- “1. For the reviewers: clarify doubts, answer questions, correct misunderstandings, push back on mischaracterizations, and make a good-faith effort to incorporate feedback and improve your work.
2. For the AC: convince them that you have made a good-faith effort, present a representative summary of the reviews, help them understand if the reviewer concerns were addressed, call out bad-faith reviewing, and ultimately, *help them make a decision*.
- [...] Would a neutral third-party be able to tell if the reviewer concerns were addressed *purely based on your rebuttal* (without reading the paper or the reviews again)?”

# Writing Rebuttals: Example Process

1. Itemize reviewer comments
2. Brain dump possible responses
3. Write a draft rebuttal
4. Review and revise

# Writing Rebuttals: Example Process

Reviewer comment	Author 1	Author N
1. There existing several related papers discussion of the using human attention map in image captioning and visual question answering. For example, (1) Liu et al. Attention correctness in neural image captioning. (2) Qiao et al. Exploring human-like attention supervision in visual question answering. Please illustrate the differences with these papers.	The papers mentioned provide attention supervision over the attention layer. Our central argument for this and the next point will be that Grad-CAM is more faithful than attention. In order to show this I am planning on doing occlusion studies in the proposal space and compare that with the attention weights and the Grad-CAM proposal importance weights. Also with attention supervision only the layers before the attention layer can be updated, but with HINT all layer weights can be updated. Include lines from the paper. Att supervision doesn't work.	We should first very clearly say what you say in the first sentece of your response. And expand on that a bit if needed to make the point clearly. Does the paper differentiate our work from these work R2 cites or other such works? If so, we should clearly say in the response "As discussed in LXYZ-ABC..." You can then make the point about which layers can be updated. We can then additionally make the "central argument" point. But the direct response should be clear / not confused with the description of a new experiment and such.
2. It seems that the ground-truth attention map is used for the VQA task. For the captioning task, although no ground-truth attention map is used, the segmentation maps are used. As such compare with other methods, strong information about the image are incorporated, which should results in performance improvements.	Human attention or segmentation maps are used only during training and not during testing. While we agree that this is extra information used during training, we show why other approach fail to utilize this information to achieve improvements in performance during test time. Only a fraction of images in VQA have Human attention. Also if it is possible to such a good boost with just human attention, people would start collecting. Also HATs are important to know if models are making the right decision for the right reasons.	"we show why other approach fail to utilize this information to achieve improvements in performance during test time." You'll have to point to a specific experiment in the paper / lines in the paper / table in the paper and reproduce the curcial numbers here to support this claim. Then you can say this is only at training time, not at test time. (I think the reviewer already knows this. So starting with this response is not a strong start.)
3. For the alignment between human attention and network importance, it seems that there only exists one human attention map. However, for the network importance, specifically the VQA task, when the question is different, we should pay different attentions to different local regions. Therefore, how to make the corresponding alignments? The human attention is stable, while the network importance varies dynamically.	This is incorrect. HAT is question dependent. i.e. there exists different maps for different questions.	Yup. Say there are different human attention maps for different questions. So the human attention map is also question dependent.
The state of the art and beyond in the field is moving away from such human guided approaches. Localization is already being done in a wholly unsupervised fashion using embeddings for example. . Also the proposed approach is not scalable. Verification is done through human studies which is fine but again not scalable.	Highly opinionated with no citations. I don't agree that state of the art "and beyond" (whatever that means) is moving away from human guided approaches. Disagree that Localization is done in a wholly unsupervised way. Approaches for semi-supervised localizations exist, but they are still significantly worse than fully supervised approaches (xx% diffence in ILSVRC localization). Also disagree with the comment that verification is done only through human studies. In section 5 we quantitatively evaluate task performance , and in section 6 we quantitatively evaluate grounding, both of which show the effectiveness of HINT without requiring human workers. Our Human studies are required to show that our HINTed models are more trustworthy to humans than base models, which is needed not just for generalization but also necessary as more algorithmic decisions are made in the society	Also maybe make some point about how without human guidance sure models have good accuracies, but they can be heavily biased (and give examples from VQA maybe)?
It would be helpful to have an ablation like study in which you increase or decrease the level of HINT's and see what happens to get a deeper insight into what you are doing.	I can set this up. We can vary the amount of HATs used and examine how the performance varies	Yup

# Writing Rebuttals: Tips

1. **Start positive.** Provide a summary of the reviews, highlighting positive things that reviewers said about your work. Rebuttals focus mostly on responding to perceived negatives, don't let RACs forget about the strengths along the way.
2. **Order matters.** Start with the biggest concerns that you have good answers for and work your way to less clear-cut responses and minor points.
3. **Let reviewers speak for themselves, then respond directly.** Quote the core of the reviewer's question or concern concisely and completely. Then before saying anything else, respond to it directly. And then give details, describe context, or explain your position.
4. **Be conversational.** Notice the conversational nature of the example responses above. It makes it easier for RACs to follow, and the responses are less likely to be perceived as being combative.

# Writing Rebuttals: Tips

- 5. Respond to the intent of the questions.** Don't feel trapped to only discuss the quoted concern — also address the intent of the comment. For example, “Why didn't you evaluate on GLORP3?” may generally be calling your experiments into question. Answer, but then point out that you've already evaluated on X,Y, and Z which should be sufficient! Note that it is useful for other RACs to be reminded of your extensive experimental evaluation. A first glance at a reviewer comment suggesting otherwise could leave a false impression.
- 6. Don't be afraid of emphasis.** “Row 2 in Table 4 shows *exactly* that.” “We do NOT need a human-in-the-loop at test time.” Notice that many of the responses above are not just direct, but also have emphasis (in tone if not formatting of text).
- 7. Feel free to set the stage.** If it seems like all reviewers missed a central point, a concise, crisp recap of the main point could help.
- 8. Keep things self-contained.** Assume RACs don't remember much about your paper and that they likely won't read it again in detail. Re-introduce any acronyms, remind them of relevant details of an experimental setup. Notice that all the responses above likely make sense to you even though you may be unfamiliar with the papers (and in some cases the names and details are made up).

# Writing Rebuttals: Tips

9. **But get credit for details you already included.** That said, if something a reviewer asked for was already in the paper, say so. Give them line/Table/Figure numbers, and then restate it in the rebuttal. The references back to the main paper are to establish credibility with all RACs that the paper was not lacking important details. (They are not necessarily to have RACs go back and look at the paper.)
10. **Consolidate common concerns.** Save space by responding to multiple reviewers at once if they share related concerns.
11. **Color-code reviewers.** Notice above the trick to color-code reviewers. Make it as easy as possible for reviewers to spot responses that are relevant to them — even when things are merged or not in reviewer order.
12. **Stats speak louder than words.** Rather than argue with RACs, give them data/stats to back your claim up. These can be statistics/analysis of your data or results. Or the results of additional experiments you run to respond to their concern (if allowed by the venue). Every time you find yourself having a different opinion than the reviewer, ask if you can establish that with data. You can always provide the intuitive arguments after you've settled the issue with data.



# Writing Rebuttals: Tips

- 13. Don't promise, do.** Instead of saying “We will discuss Singh et al. in the paper.”, provide a discussion in the rebuttal. Instead of saying “We will explain what  $D_{RT}$  stands for in the paper”, explain what it stands for in the rebuttal. And then also add that you will add it to the paper. It makes it significantly easier for RACs to trust that you will make the promised changes.
- 14. Be receptive and reasonable.** Most RACs will appreciate it. Plus, it is just the better thing to do — these are your colleagues! :)
- 15. Be transparent.** Reviewers hinted at an additional experiment but the venue doesn't allow it? Say so. They asked about intuitions about a trend but you don't have any? Say that you've thought about it but don't have any good ideas, and will continue to investigate it. Don't have enough GPUs to run the experiment they asked for? Say so.
- 16. Shine a spotlight on reviewers acting in bad-faith.** In some circumstances, a reviewer may not be adhering to reviewing best practices or may not have taken the reviewing role seriously. It can be important to make sure the other RACs realize this and appropriately discount their review. Pointing out unreasonable or unsubstantiated comments and referencing other reviewers that disagree can help. This can also include confidential comments to the AC (where applicable).

# Writing Rebuttals: Tips

- 17. Acknowledge reviewer efforts.** On the other hand, if a reviewer goes above and beyond to be constructive, thank them for it. Typo list? Thank you. Pointers to relevant work? Thank you. Detailed musings about future work? Thank you. Add at least a short blurb acknowledging these things!
- 18. Don't forget the humans on the other end.** Keep in mind that this is not just a scientific but a sociopolitical interaction with other humans :) So decide whether you'd like to be argumentative and risk your reviewer taking a strong stand against the paper, or if you'd like to work towards a common ground. Finding points where you do agree with the reviewer and acknowledging them can help with the latter.

Look at readings

# Activity 1: Focus on Writing

- Pick your favorite iconic paper
- How do the authors tell their *story* in the intro?
- How do the authors explain novelty, good results, bad results?
- What do you like about how they are written?
- What do you not like, can be improved?
- Is something surprising about how they are written?

# Activity 2: Accepted/Rejected Versions

- Skim (20-30 min total) the accepted and rejected versions of one paper
- Compare/contrast – what was similar, what was different?
- Why was one version accepted, the other rejected?
- Share with the class (15 min total)

# Activity 3: Peer Feedback on Literature Reviews