Report Writing and Presentation

He He

New York University

April 18, 2023

Logistics

Project presentation (10%): in class on May 2

- Format: 3 minute talk and 1 minute Q&A with the audience
- Recording: slides should be shared through Zoom
- Grading (10%):
    - Clarity (problem 2% + approach 3% + evaluation 2%)
    - Style (clean an legible slides) (2%)
    - Q&A (1%)

Logistics

Report (30%): due on May 5

- Each group only needs to submit one report on Gradescope
- Format:
  - Page limit: 4
  - File format: .pdf
  - Style: ACL format (see templates on website)
- Grading (30%):
  - Format (2%)
  - Clarity (6%)
  - Literature survey (4%)
  - Technical content (10%)
  - Evaluation (8%)

Write the paper first

Jason Eisner's advice:
`http://www.cs.jhu.edu/~jason/advice/write-the-paper-first.html`

Writing is not just translating thinking into words; it is part of the thinking.

- Writing is the best use of limited time
  - The audience only pays attention as far as they can understand
  - The idea (the writeup) is often more useful than the software

Write the paper first

Jason Eisner's advice:

`http://www.cs.jhu.edu/~jason/advice/write-the-paper-first.html`

Writing is not just translating thinking into words; it is part of the thinking.

- Writing is the best use of limited time
    - The audience only pays attention as far as they can understand
    - The idea (the writeup) is often more useful than the software
- Writing helps you plan what to do next
    - Writing things down forces you to come up with a coherent story
    - It helps you see the holes in your logic (if you are honest)

Write the paper first

Jason Eisner's advice:
`http://www.cs.jhu.edu/~jason/advice/write-the-paper-first.html`

Writing is not just translating thinking into words; it is part of the thinking.

- Writing is the best use of limited time
    - The audience only pays attention as far as they can understand
    - The idea (the writeup) is often more useful than the software
- Writing helps you plan what to do next
    - Writing things down forces you to come up with a coherent story
    - It helps you see the holes in your logic (if you are honest)
- Writing makes collaboration easier
    - Meetings are more effective
    - Someone can continue to work from where you left

Typical structure of an NLP paper/talk

1. **Title**: a precise summary of the paper's contribution
2. **Abstract**: high-level overview of the paper
3. **Introduction**: extended abstract
4. **Method**:
   4.1 Problem statement: formulation of the problem
   4.2 Approach: technical details of your method
5. **Experiments**:
   5.1 Setup: data, baselines, hyperparameters etc.
   5.2 Results: report all results (positive or negative)
   5.3 Analysis: how should one read the results
6. **Related work** (may also come after intro): prior approaches to the problem
7. **Conclusion**: summarize the highlights and point out future directions

Title

- Used as an index to retrieve or refer to the paper
- Be precise and specific
    - Too generic: “A New Model for Image Caption Generation”
    - Better: “Neural Image Caption Generation with Visual Attention”
    - Think about the keywords your audience might put in the search bar
- Don’t overclaim your contribution
- It’s nice to have a catchy phrase but don’t force it
    - Include it only if it’s relevant and delivers the key message
    - Good examples:
      “*Show, Attend and Tell*: Neural Image Caption Generation with Visual Attention”
      “*Know What You Don’t Know*: Unanswerable Questions for SQuAD”
    - Often a precise, “plain” title is good enough
- Avoid extra formatting

Abstract

Include all key messages: abstract has the most number of readers!

1. **The problem**:
   1.1 The general topic: “Deep Neural Networks (DNNs) are powerful models that have achieved excellent performance on difficult learning tasks.”
   1.2 The specific problem: “Although DNNs work well whenever large labeled training sets are available, they cannot be used to map sequences to sequences.”
2. **Your contribution**:
   2.1 One-sentence summary: “In this paper, we present a general end-to-end approach to sequence learning that makes minimal assumptions on the sequence structure.”
   2.2 Your method: “Our method uses a multilayered Long Short-Term Memory (LSTM) to map the input sequence to a vector...”
   2.3 Impact/results: “Our main result is that on an English to French translation task from the WMT-14 dataset...”

Introduction

Follow the same structure as the abstract but with more details

1. **The problem**: context, motivation, important related work
2. **Your approach**: intuition, high-level description
3. **Results**: summary of experimental results, takeaways

- Don’t spend too much space on general introduction and related work. The reader should know what you did by the second paragraph.
- Use examples and figures to illustrate your method.
- All claims should be supported by experiments/theories/analysis/citations.
- Use forward reference:
  “Our algorithm has linear time complexity at inference time (Section 2.1).”

Approach

- Many different ways to structure the section depending on the content
- Top down
  - A documentation of your final method: “Our model consists of three components: A, B, C.”.
  - Easier to write but hard to read: where does each component come from?
- Bottom up
  - Build up the method as if the readers are developing the method with you
    “A really simple way to solve the problme is to use A.”
    “However, it doesn’t consider X, so we add B.”
    “To further improve Y, we add C.”
  - By the time the reader finishes reading, they should think: “This is obvious! I could have come up with the idea”.
- Often a mix of the two

Approach

General suggestions:

- Always give high-level ideas and intuitions before go into technical details
- Put the “why” before the “how”
- Use running examples when illustrating complex concepts or procedures
- Notations and terminologies should be consistent throughout the paper
- Use equations only when it adds additional information.
    - Sometimes an idea can be described precisely in words.
    - Even when you do need an equation, decribing it in words first is helpful.

Experiments

- **Setup**: datasets, metrics, baselines
- **Implementation details**: preprocessing, hyperparameters of the model and the algorithm
- **Results**:
  - Full results are often shown in tables and charts.
  - In the main text, highlight important numbers and takeaways.
- **Analysis**: provides better understanding of the results, e.g.
  - Ablation study of the model: are all components equally important?
  - Error analysis: what are the limitations of the method?
  - If the results are negative: what are possible reasons?
    (Often there is a mismatch between the assumptions you made and the data.)

Related work

If some work directly motivates your work, it should go into the intro. E.g. you are

- extending a previous work,
- applying a prior method to your problem.

Why put related work at the end?

- The reader has not gained a good understanding of the problem and your approach to make judgments after the introduction.
- It breaks the flow from the introduction to the approach.
- It provides an overview of the area and transitions well to the conclusion.

Don’t just survey the area; describe how your work situates in the broader context.

Conclusion

- Highlight the key findings from your work.
- Here is the place where you can take some freedom to state your opinions
  - What are the limitations of your approach?
  - What are the important next steps?
  - What are promising future directions for the problem?

General suggestion

- Think about the key message you want to deliver and “repeating” it throughout the paper.
- Put yourself in the reader’s shoes.
- Get feedback from people who are less familiar with your work.

Talks vs papers

- The talk mirrors the paper so much of the advice still applies.
- But, the talk contains *much fewer details*.
  The paper is similar to the manual of a tool, while the talk is your explanation of how to use the tool.
- You have more freedom on how to communicate in the talk
  - Use visual aids: video/audio/animation, memes/stories etc.
  - Interact with the audience (ask questions, voting, demo)
- Purpose of the talk
  - Convey the key idea of your paper (so that they will go read it)
  - Entertaining (it’s a performance)

Giving the talk

General advice

- Number one mistake to avoid: go over the time limit
  (Generally, prepare *n* slides for an *n*-minute talk.)
- Know your opening sentences by heart (the rest will be easy)
- Use a lot of examples, especially given limited time
- Be brief on related work (they may be mentioned during Q&A)

For the project presentation

- You only have 3 minutes: practice is important!
- If more than one member will be presenting, practice the transition.
- Test audio and screen sharing to avoid technical issues.
- Suggested structure:
  problem/motivation, key ideas, evaluation, conclusion (1 slide each)

Opening

Main goal: engage the audience

- What is the problem? (Audience: what’s this about?)
- Why is it important? (Audience: why should I care?)
- What’s the challenge? (Audience: this is not trivial?)

They should be excited to hear your approach by now.

Be creative: use examples, data, quotes, anecdotes etc.

Key ideas

- *Use the bottom-up strategy*: solve the puzzle with the audience.
  It doesn't have to follow how you actually solved the problem, but there should be a story to connect all the pieces.
- *Think about what the audience may be wondering*: ask a question, then answer it.
  "Now, how do we optimize it given that the expectation is intractable? We use ..."
- *Leave the details to the paper*: what is the minimal set of things they need to understand?
- *Create references* to your key ideas/concepts.
  - Use icons, figures, running examples, toy problems, analogies etc.
  - This should be in their head when they retrieve your idea.

Results

- Don’t paste a wall of numbers. Use bar charts.
- Guide the audience on where they should pay attention.
- Put an explicit takeaway/tagline for each result you show.
- Think about it as a sequence of question and answers.
  “How well does the approach work on benchmark A?”
  “What about a different domain?”
  “Is component B really necessary?”
  “How fast is the inference?”

Conclusion

- Acknowledgement to collaborators.
- Show the takeaway messages you want the audience to leave with.
- Show the highlights/summary of the talk to start Q&A.

Style

Basics

- Font: sans-serif; sizes: 24+.
- Number slides (for easy reference during Q&A).

Layout

- Avoid clutter, wordiness (and full sentences).
  The lecture slides are bad examples (because it’s used both for the talk and for reference).
- Use keywords, diagrams, figures: the audience should pay attention to *you*, not the slides.
- If the content is complex, use animation to reveal the full slide.

Visual guidance

- Color coding should be consistent throughout the presentation.
- Use arrows and boxes to help the audience read figures and tables.
- Highlight (bold, background color) key messages.

Q&A

- Common types of questions
  - Clarification: “Can you explain X again?”
  - Interpration of the results: “Why is there a dip in the learning curve?”
  - Comparison: “How is this different from X?”
  - Extensions and new perspectives of the work (this is the type of questions we like)
- It’s okay to pause before you answer the question.
- Provide your thoughts even if you have no idea what’s the answer.