

Doing Research in Natural Language Processing

Session 4: Scientific Writing: Structure

Frank Keller

29 September 2023

School of Informatics
University of Edinburgh
keller@inf.ed.ac.uk

Title

Introduction

Middle

Conclusion

Reading: Alley (2018), Chapter 7.

Please also look at Alley's web site, which has a lot of videos and additional materials:

<https://www.craftofscientificwriting.com/>

Organizing Content

Scientific documents are almost always structured as:

- Title
- Introduction
- Middle
- Conclusions

The structure of the middle can vary. In experimental papers, it's typically:

- Method
- Results
- Discussion

We will discuss abstract, appendices, figures and tables separately.

Title

We've already talked about titles. Quick recap:

- “short and sweet” is great advice for the title of novels, but it doesn't work well for titles of scientific papers
- the title should enable the audience to decide whether to read the paper or not
- in a bibliography or a web search readers will *only* see the title
- use the title to specify the scope of the document: identify the field of work and separate it from others in the field

Title: Examples

The Bluest Eye.

Sula.

Song of Solomon.

Tar Baby.

Beloved.

Jazz.

Paradise.

Love.

A Mercy.

Home.

God Help the Child.

Title: Examples

The Bluest Eye.

Sula.

Song of Solomon.

Tar Baby.

Beloved.

Jazz.

Paradise.

Love.

A Mercy.

Home.

God Help the Child.

Titles of novels of Tony Morrison (Nobel Prize in Literature, 1993).

Title: Examples

Titles don't have to be phrases, they can be sentences:

Colonoscopic screening for colorectal cancer improves quality of life measures

This title is a one-sentence summary of the result of the study. In some fields, this is the conventional way to write titles.

And remember our example from session 1 (Vaswani et al., 2017):

Attention Is All You Need

This title is memorable and stands out, but doesn't say much about the content of the paper. How about:

Attention Is All You Need: The Transformer as an Efficient Sequence Processing Architecture without Recurrence

A semicolon can often be used to structure a title.

Title: Examples from ACL 2021

How Did This Get Funded?! Automatically Identifying Quirky Scientific Achievements

Aspect-Category-Opinion-Sentiment Quadruple Extraction with Implicit Aspects and Opinions

Deep Differential Amplifier for Extractive Summarization

COSY: COunterfactual SYntax for Cross-Lingual Understanding

Cascaded Head-colliding Attention

Is Sparse Attention more Interpretable?

Relative Importance in Sentence Processing

To POS Tag or Not to POS Tag: The Impact of POS Tags on Morphological Learning in Low-Resource Settings

Introduction

Introduction

The introduction prepares the reader for the main content of the document, answering the following questions:

- What exactly is the work?
- Why is the work important?
- What is needed to understand the work?
- How will the work be presented?

Not all question may be present, and not always in this order.

What exactly is the work?

- Describe the scope and limitations of the work.
- Provide more detail than the abstract.
- State any underlying theoretical or methodological assumptions.

Why is the work important?

- Give the audience a reason to starting reading, and continue reading, the document.
- If the document is a proposal: why should this be funded?
- The importance of the work can derive from its applications, but also from pure curiosity.
- In that case, you need to instill this curiosity in the reader!
- Bear in mind who your *audience* is: experts, general readers, funders, managers.
- This determines how you need to justify the importance of your work.

Why is the work important?

In size, density, and composition, Ganymede and Callisto (Jupiter's two largest moons) are near twins: rock-loaded snowballs. These moons are about 5000 km in diameter and contain 75 percent water by volume. The one observable difference between them is their albedo: Callisto is dark all over, while Ganymede has dark patches separated by broad light streaks. This paper discusses how these two similar moons evolved so differently.

Here curiosity is the main motivation for the work.

What is needed to understand the work?

Present the background required to understand the main part of your document:

- review literature and related work
- show that your work is novel, unique
- identify gaps in the literature (respectfully!)
- tailor this to your audience and what they know
- boost your credibility as an author

Can be a separate section or chapter, or just a part of the intro.

In a thesis, you can have both an upfront background chapter and a background section for each content chapter.

What is needed to understand the work?

A table is sometimes a good way to present the related literature:

Dataset	Task	#L	#V	Obj	Imgs	Sen	Des	Cln	ML	Resource	Example Labels
Ikizler (Ikizler et al., 2008)	AC	6	6	0	467	N	N	Y	N	–	running, walking
Sports Dataset (Gupta et al., 2009)	AC	6	6	4	300	N	N	Y	N	–	tennis serve, cricket bowling
Willow (Delaitre et al., 2010)	AC	7	6	5	986	N	N	Y	Y	–	riding bike, photographing
PPMI (Yao and Fei-Fei, 2010)	AC	24	2	12	4.8k	N	N	Y	N	–	play guitar, hold violin
Stanford 40 Actions (Yao et al., 2011)	AC	40	33	31	9.5k	N	N	Y	N	–	cut vegetables, ride horse
PASCAL 2012 (Everingham et al., 2015)	AC	11	9	6	4.5k	N	N	Y	Y	–	riding bike, riding horse
89 Actions (Le et al., 2013)	AC	89	36	19	2k	N	N	Y	N	–	ride bike, fix bike
MPII Human Pose (Andriluka et al., 2014)	AC	410	–	66	40.5k	N	N	Y	N	–	riding car, hair styling
TUHOI (Le et al., 2014)	HOI	2974	–	189	10.8k	N	N	Y	Y	–	sit on chair, play with dog
COCO-a (Ronchi and Perona, 2015)	HOI	–	140	80	10k	N	Y	Y	Y	VerbNet	walk bike, hold bike
Google Images (Ramanathan et al., 2015)	AC	2880	–	–	102k	N	N	N	N	–	riding horse, riding camel
HICO (Chao et al., 2015)	HOI	600	111	80	47k	Y	N	Y	Y	WordNet	ride#v#1 bike; hold#v#2 bike
VCOCO-SRL (Gupta and Malik, 2015)	VSRL	–	26	48	10k	N	Y	Y	Y	–	verb: hit; instr: bat; obj: ball
imSitu (Yatskar et al., 2016)	VSRL	–	504	11k	126k	Y	N	Y	N	FrameNet WordNet	verb: ride; agent: girl#n#2 vehicle: bike#n#1; place: road#n#2
VerSe (Gella et al., 2016)	VSD	163	90	–	3.5k	Y	Y	Y	N	OntoNotes	ride.v.01, play.v.02
Visual Genome (Krishna et al., 2016)	VRD	42.3k	–	33.8k	108k	N	N	Y	Y	–	man playing frisbee

Over to You

Exercise 1

Let's again look at the introduction of Vaswani et al. (2017) (next page):

- What exactly is the work?
- Why is the work important?
- What is needed to understand the work?
- How will the work be presented?

Which of these questions are present in this excerpt?

Exercise 1

Recurrent neural networks, long short-term memory [13] and gated recurrent [7] neural networks in particular, have been firmly established as state of the art approaches in sequence modeling and transduction problems such as language modeling and machine translation [35, 2, 5]. Numerous efforts have since continued to push the boundaries of recurrent language models and encoder-decoder architectures [38, 24, 15].

Recurrent models typically factor computation along the symbol positions of the input and output sequences. Aligning the positions to steps in computation time, they generate a sequence of hidden states h_t , as a function of the previous hidden state h_{t-1} and the input for position t . This inherently sequential nature precludes parallelization within training examples, which becomes critical at longer sequence lengths, as memory constraints limit batching across examples.

In this work we propose the Transformer, a model architecture eschewing recurrence and instead relying entirely on an attention mechanism to draw global dependencies between input and output. The Transformer allows for significantly more parallelization and can reach a new state of the art in translation quality after being trained for as little as twelve hours on eight P100 GPUs.

Middle

Middle

The middle of a document presents the work in a logical and persuasive fashion. To achieve this, choose a *strategy* for presenting the material:

- chronological
- spatial
- classification and division
- cause-effect
- comparison-contrast

The strategy you choose must be suitable for the audience. Use it to group the work in sections, to make it more digestible.

Descriptive Headings

Section headings should be descriptive, make the strategy of the document obvious, serve as a roadmap for the reader.

Sections provide *whitespace*, which allows readers:

- pause and reflect on what they've read
- jump to the information that interests them, skip those parts that don't

Don't over-section the text, consider if a paragraph break is better than introducing a subsection.

Stylistic considerations for section headings:

- avoid cryptic one-word titles also for sections
- use parallelism (all subsections are titled using noun phrases, participles, etc)
- Don't use dangling subsections (1, 2, 2.1, 3, 3.1, 3.2, 4)

Look at the **table of content** to see if your sectioning works.

Put `\tableofcontent` at the beginning of your paper (even if the final version won't have a ToC).

Descriptive Headings

Weak Headings

Introduction

Debris Recovered

Cataloguing

Interpretation

Results

 Placement

 Bomb Makeup

Work to be Done

 Interpretation

Strong Headings

Introduction

Completed Work

 Recovering Debris

 Cataloguing Debris

 Interpreting the Debris

Preliminary Results of Work

 Placement of Bomb

 Construction of Bomb

Future Work

Depth is the level of detail in your document. Affected by:

- **occasion:** determines length, e.g., conference paper vs. journal article
- **audience:** satisfy the readers interest, anticipate their questions, don't raise questions the document doesn't answer
- **purpose:** if you want to persuade, you will need to discuss advantages and disadvantages, rebut objects, etc.

Adapt paragraph and section length depending on the depth of your document.

Avoid very long or very short paragraphs (exception: instructions).

Over to You

Exercise 2

Here's the sectioning of the Vaswani et al. (2017) paper:

1. Introduction
2. Background
3. Model Architecture
 - 3.1 Encoder and Decoder Stacks
 - 3.2 Attention
 - 3.2.1 Scaled Dot-Product Attention
 - 3.2.2 Multi-Head Attention
 - 3.2.3 Applications of Attention in our Model
 - 3.3 Position-wise Feed-Forward Networks
 - 3.4 Embeddings and Softmax
 - 3.5 Positional Encoding
4. Why Self-Attention

Exercise 2

5. Training

5.1 Training Data and Batching

5.2 Hardware and Schedule

5.3 Optimizer

5.4 Regularization

6. Results

6.1 Machine Translation

6.2 Model Variation

6.3 English Constituency Parsing

7. Conclusion

Which *strategy* do the authors use to structure the middle? Do they follow Alley's advice on *sectioning*? Are the headings *descriptive*? Which level of *depth* do they use?

Conclusion

Conclusion

Summarizes the middle and provides a future perspective:

- provide an analysis of the most important results
- analyze the results overall, not individually (do that in the middle, in a Discussion section)
- do not present new evidence or new results here
- provide a future perspective:
 - recommendations that derive from your work
 - future direction of you work
 - re-iterate the scope and limitations of your work (already in the intro)
- ACL conferences now require a separate section on limitations

Conclusion ties together loose ends, provides closure.

References

- Alley, Michael. 2018. *The Craft of Scientific Writing*. Springer, New York, NY, 4 edition.
- Gella, Spandana and Frank Keller. 2017. An analysis of action recognition datasets for language and vision tasks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, Volume 2: Short Papers*. Vancouver, pages 64–71.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*. Curran Associates, Red Hook, NY, pages 5998–6008.