

Doing Research in Natural Language Processing

Session 2: Scientific Writing: Audience

Frank Keller

22 September 2023

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

Introduction

Audience

Purpose

Occasion

Reading: Alley (2018), Chapter 1.

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

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

Introduction


Writing Matters: Example from Alley (2018)

On January 28, 1986, the Space Shuttle *Challenger* took off from Cape Canaveral, Florida. On board were seven astronauts, including teacher Christa McAuliffe, the first civilian in space. Millions of school children across the US watched.



Writing Matters

The Shuttle exploded 73 seconds after takeoff, killing all seven astronauts on board.



A subsequent investigation found that the solid rocket boosters were the source of the explosion. Two O-rings, seals to prevent the fuel from escaping, had failed.

Long before the fatal launch, engineers had reservations about the design of the O-rings on the boosters. NASA management requested they seek opinions from O-ring experts.

So NASA engineers visited two manufacturers of O-rings. They found that both manufacturers had serious concerns about the design of the O-rings.

The engineers wrote up a report about their visits. It contained strong warnings about the design of the O-rings and was entitled:

Subject: Visit to Precision Rubber Products Corporation and Parker Seal Company

But no one responded to the report. NASA's paper trail ends here.

Over to You

Exercise 1

The title of the reports on the Shuttle O-rings was:

Subject: Visit to Precision Rubber Products Corporation and Parker Seal Company

Can you see anything wrong with this title?

Exercise 1

The title of the reports on the Shuttle O-rings was:

Subject: Visit to Precision Rubber Products Corporation and Parker Seal Company

Can you see anything wrong with this title?

This is a weak title:

- It does not tell you what the report is about.
- The authors clearly haven't thought about their audience.

To write successfully, you need to understand your audience.

Exercise 1

To come up with a better title, let's think about the audience of this report:

- Who is the audience?
- Why is the audience reading?
- What does the audience know?

Based on this, what title would you suggest?

Audience

Audience, Purpose, Occasion

Before you write a scientific document, analyze:

- audience
- purpose
- occasion

These will greatly influence how you will write the document.

We will look at each aspect in turn.

Who is the audience?

- conference paper: audience has very similar background to yourself; experts in your area
- journal article: audience is typically broader, and depending on the journal may include generalists
- grant proposal: mixture of experts (reviewers) and generalists (panel members)
- podcast: general audience with an interest in scientific issues

The broader or more mixed the audience is, the harder the document will be to write.

Why is the audience reading?

- Once you know who your audience will be, ask what they want to get out of the document.
- Make sure this information is there, is detailed enough, and is structured so as to be easy to find and digest.
- For a grant proposal, look at the review form to see what the reviewers will look for; for a journal paper, look at other papers published in the same journal, etc.

Audience: Knowledge

What does the audience know? Thinking about this will tell you:

- how to arrange the content
- which terms to define
- what background to include

This tells you how to structure your document. Particularly hard if you have a mixed audience!

Think about your **primary** audience; maybe put content for your **secondary** audience in an appendix.

Over to You

Exercise 2

But it's not rocket science, it's NLP! So how about the following title:

Attention Is All You Need

This is the title of Vaswani et al. (2017), one of the most famous papers in the NLP literature. It currently has 89,126 citations on Google Scholar.

But is this a good title, one that targets the audience correctly?

Get together in pairs discuss this, if necessary re-write the title. Have a look at the abstract, too.

You have 5 minutes!

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Lukasz Kaiser*
Google Brain
lukaszkaizer@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.0 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature.

Purpose

Purpose: Inform

Most scientific writing has two specific purposes: *to inform and to persuade*. The level of persuasion varies: instructions require very little, a grant proposal requires a lot of persuasion.

Alley's analogy: a scientific document provides path that leads the reader up the mountain of your scientific expertise. If your purpose is merely to inform:

- you need to provide a path up the mountain
- it can be gentle (simple content) or steep (complex content)
- but you need to make sure readers can follow, break down the information, provide “vistas of understanding”

Purpose: Inform



Image: Alley (2018)

Purpose: Persuade

To inform, you need to answer *what, where, when, how*. To persuade, you also need to answer *why*. You need to build *credibility* with the audience:

- expend extra words to persuade; it's not about being maximally efficient
- you may not take most direct path up the mountain; it's more like navigating a boulder field
- the writing style changes: from lists (informative) to longer paragraphs (persuasive)

Persuasive writing explains why this is the right topic, research question, method, and technique.

Occasion

The occasion for which you're writing the document determines its:

- form
- formality
- politics (rarely relevant for PhD students)

Occasion: Form

Form refers to style and grammar, but also length and format of the document. We will discuss this later in the course.

Alley provides advice on grammar, punctuation, and usage in *Appendices A–C* of his book. Useful for both native and non-native speakers!

Important not to be prescriptive. Alley gives *advice*, not based on right/wrong, but on unsettles/distracts readers.

Note differences of British and American spelling. Important not which one you use, but to be consistent.

The *length of your document* has an obvious effect on how you write: a conference paper has half a page of literature review, a PhD thesis has a whole chapter.

The format of your document is often fixed (style file of conference or journal, thesis template). But Alley's *Appendix D* has general recommendations for formatting scientific documents (“for situations in which no graphic designer is available”).

Occasion: Formality

A certain level of formality is expected in scientific writing. Examples:

Too informal	Accepted
<i>a lot</i>	<i>much or many</i>
<i>get</i>	<i>obtain</i>
contractions (<i>don't</i>)	written out (<i>do not</i>)
<i>And ...</i>	<i>Also, ...</i>
<i>But ...</i>	<i>However, ...</i>

Normally, don't address the reader with *you* (exception: instructions).

Alley, Michael. 2018. *The Craft of Scientific Writing*. Springer, New York, NY, 4 edition.

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.