Introduction to Neural Networks
U. Minn. Psy 5038
Daniel Kersten

Outline

Last time

Non-linear, multi-layer, feedforward networks and error back-propagation.

Today

Science writing

Mid-term exam

Science writing

General points

■ Where I'm at

I have spent a lot of time critiquing my own writing and that of others. But I'm still learning. It is one thing to know what to do and another to actually do it. It is also one thing to do something well, and quite a different thing to explain how you do it well. I don't claim to always write well or to know how to write the best. But I do write better than I used to. I've learned from my advisors, journal reviewers, and from my own students. I continue to learn from self-evaluation. So here I'll summarize my ideas and advice.

Two favorite sources:

> Gopen & Swan: Provides an excellent discussion of general principles. Not a cookbook.
>
> A Pocket Style Manual 2nd Edition. Diana Hacker: A good cookbook.

I will talk about science writing, rather than other kinds of writing, for three reasons: 1) This is a science course and you'll get practice doing it; 2) I'm your teacher, and it is the only kind of writing that I know; 3) At a general level, the kind of writing doesn't matter too much, because the main principles generalize to any accurate communication of information. So even if after this course you never write about scientific results again, you will have learned something about clear communication that will pay off later.

■ Why bother?

Why bother learning to write well? After all, you might argue, a good scientific result should stand on its own. There are a few cases where just a hint of a major discovery is so motivating that exactly how the information is conveyed is relatively unimportant compared to the result. (The structure of DNA, Einstein's theory of special relativity, the proof of Fermat's last theorem.) But for most of us who write up scientific results, good writing is essential to success. Here are a few reasons.

Most scientific papers contribute only small increments to the overall body of knowledge. A published paper may get read by only 2 or 3 people. On the one hand, this figure is depressingly low. On the other hand, this is an opportunity to gain 2 or 3 advocates. An incremental contribution can be amplified or minimized by the qualtiy of the writing.

Good writing is altruistic and polite--it respects the reader.

Good writing is also selfish--someone else who writes better can get the credit. I know of papers in my field whose (important) results never gained widespread recognition, and were superceded by a better articulated paper. The author of the second paper is the one that gets remembered in citations and textbooks. Granted it isn't just the clarity of exposition, but often the depth and breadth of the second paper. However, this depth and breadth can emerge in part from an effort to clearly communicate thoughts and intuitions that are otherwise to vague to be useful. Thus the best reason to practice writing well is that it will help you think more clearly. And if as a scientist you think more clearly, you may end up discovering more. As Gopen and Swan put it: "*Improving the quality of writing actually improves the quality of thought*". There are many times that I've thought that I've thoroughly understood something, only to be stumped when having to explain it.

As an aside, the same goes for programming. It is one thing to have a good idea, and quite another to design computer program that implements it. There are many apparently "good" theories in cognitive science that stumble at the implementational stage. That is why I teach this course using computers. Programming forces you to make your assumptions and theories sufficiently explicit to be tested. Similarly fleshing out your ideas with sufficient clarity to be understood by another forces you to be explicit about your assumptions and the logic of your arguments.

■ Determine your audience

Good writing is good communication. Thus an essential first step is to determine your audience. You could be writing for the general public, kids, undergraduates, graduate students, scientific laypersons, scientific professionals, scientific colleague in your general field (e.g. cognitive science), scientific colleague in your speciality (e.g. vision), or even sub-speciality (e.g. motion perception).

Even after deciding the type of audience you will be writing for, it isn't always easy to figure out where your audience is at. It may take some research and discipline to understand your audience.

Because of specialization, there are big differences amoung members of the category of professions that know and use science, for example:

Social scientists, physical scientists, biological scientists, scientific professionals, engineers, grant reviewers, scientists in industry, business users of science.

You may want to write to reach two audiences. There are good and bad ways of doing this too, and we'll get to that below.

As Gopen and Swan put it: "*If the reader is to grasp what the writer means, the writer must understand what the reader needs*"

Tip: Write as if to someone you personally know who represents your audience. If you have a chance, have that person read a draft. Encourage the person to be frank about what is not understood.

And to help you out, in your final paper for this course your audience is not me--the audience is your peers!

■ Good writing is hard work for most of us

Good writing is hard because it is the result of clear thinking, and clear thinking takes work. Good writing is also hard because it requires thinking about expectancies and logical flow at multiple levels of abstraction. Let's understand what is meant by expectancy, logical flow, and levels of abstraction.

The Expectancy Principle: Provide on a familiar context, and add small, even slightly surprising, instructive deviations

A good joke is a caricature of what you want to accomplish when communicating and writing. You set up the audience with a story line or context that produces a set of probable outcomes, based on past knowledge and experience. Then you supply the new information. If the new information (or punch line) is too close to what the reader expects, the story (or joke) is boring. If the new information is too much of a logical leap, the audience gets lost (the joke falls flat, because the listener couldn't "get it").

Th expectancy principle applies to: humor, creativity in art, beauty, music, learning, fiction, non-fiction, and science writing.

As Gibbons (1743-1793) wrote: "The powers of instruction are seldom of much efficacy except in those happy dispositions for which they are almost superfluous"

The trick is to find the fine line defining "almost superfluous".

Avoid spurious information. But spurious is relative. If details obstruct understanding for most of the audience, but are important to a minority, put the details elsewhere (e.g. in an appendix). Determining your audience helps determine what goes where.

Opening sets context & motivation and general expectation. Closing (stress position) brings fulfillment to expectation.

When you write up your final projects, I want you to follow a particular form, not because "that's just the way it is", but rather because the form makes explicit the role of expectation. Expectation is why there is often a standard form in science papers: Abstract, Intro, Methods, Results, Discussion.

Violating the order can disregard the tacit structural knowledge of our audience. When is that good?

Many scientific journals have changed their format to put the Methods at the end. Why?

(To be nerdy, the expectancy principle is also analagous to predictive coding by the nervous system...determine what fits the current model of the input and pass on the information that is new)

"Logical flow" & expectancy

"Logical flow"

Any piece of writing should carry the reader along without extraneous mental intrusions into the flow of the description. This is true of both non-fiction and fiction writing (the maxim in fiction to avoid intrusions that interrupt the "suspension of disbelief").

For science writing, the way to do this is to follow the principle of logical flow.

What is "logical flow"?

From Gopen & Swan: "Put in the topic position the old information that links backward; put in the stress position the new information you want the reader to emphasize".

The topic position is typically at the beginning to set context. The stress position is usually at the end to provide the new information. Old to new.

Work towards logical flow within units and between units (a unit can be a sentence, paragraph, or section). Use the lead subject word to link with previous sentence or paragraph. Or you could set up a list of expectations at the beginning of a unit, and then follow through in sequence with the fullfilments. At a coarser scale of abstraction, a paragraph or section plays the role of a unit. E.g. use the opening concept or paragraph to link backwards. Use the closing paragraph to emphasize/summarize the new information.

Tip: Watch your "referents"!! Especially at the beginning of a new paragraph.

Here are two examples where the referents differ in their degree of ambiguity:

"Put the old information in the topic position. The topic position is typically at the beginning. This is where the reader usually expects old information."

"Put the old information in the topic position. The topic position is typically at the beginning. This will help you to improve your writing."

What did "this" refer to? Don't assume the referent is as obvious to your reader as it is to you.

Logical flow at multiple levels

Analogous to "multiresolution": Structural similarities: Within a sentence.Within a paragraph.Within a section.Within the paper.

An outliner can help to make the logical flow at multiple levels explicit. E.g. this subsection should have a logical flow, and if you close the cells, the section should also have a logical flow. You can use parts of an outliner like a scaffolding that gets removed once the structure of the text has been built. Sometimes I use many more outliner levels while writing than I need for the final draft. It helps me organize my thoughts at multiple levels. But once I've figured out the structure, the text stands on its own without an overdose of subsubsections, subsections, and sections. Too many of these can also disrupt the flow.

Paper structure

■ Abstract

Mini-version of the whole paper. Still a unit of discourse with logical order: set context & motivation, raise question, give answer. Minimize details from the methods, unless the method is a primary contribution.

■ Introduction

Motivation, motivation, motivation. Articulate the question for the target audience.

"Funnel" principle

Start intro by motivating with the big picture, and gradually focus in on the scientific question/hypothesis being tested. The metaphor is that your paper should provide a channel into which readers with diverse backgrounds can be brought into an appreciation and understanding of the focused scientific question of the paper. Adjust the "width" of the funnel to the diversity of the audience. Usually the length of the "funnel" is fixed, determined by scientific journal convention.

Here is another way of thinking about the "funnel principle". The average number of readers per scientific journal article is small. Further, the number that actually make it all the way to the end is even smaller. The introduction can serve to motivate and reach an opening audience that might be different than your closing audience.

By the end of the introduction, the funnel is narrow--it is where you, the expert on the material, are at.

Example

"Soon after I started physiological research, I was lucky enough to make an interesting, but as it turned out unoriginal discovery (Barlow, 1950). I was repeating some experiments that Hartline (1938, 1940) had done on the frog's retina, with the idea that the very large receptive fields he had discovered might not be simple spatial integrators of light, as he had suggested, but might have some form of pattern selectivity. The experiment was to measure the threshold for eliciting impluses from a retinal ganglion cell as a function of the area of the stimulus spot. If spatial integration occurred, and the sensitivity over the receptive field was uniform, ..."

From: Barlow, H.B. Perception: what quantiative laws govern the acquisition of knowledge from the senses?

The Punch-line of the introduction

Context at the beginning. The "punch-line" at the end. The "punch-line" is the crucial, exciting scientific question, i.e. end the introduction section with a clear statement of the hypothesis to be tested. By the end of the introduction, your readers should be waiting with bated breath to find out both how and what your answer is. Don't give away the whole story. But no hard and fast rules. Sometimes a titillating preview of the result can be a good idea--the equivalent of a movie trailer. Again, consider the diversity of the audience.

■ Methods

Explain how you answer the question posed in the introduction. Some of this explanation may go in the Introduction, but the details go in the Methods.

■ Results

Context is the data and analysis. The stress is their interpretation in the light of the question(s) posed at the end of the Introduction. Most good results sections have good figures, and the writing centers around explaining the figures. Work on making sure that the scientific conclusions can be drawn naturally and clearly from the figures and the figure captions, and the Results section writing will come naturally.

■ Discussion

Describe the broader implications of your results. The place for humble speculation.

■ General principles: Economy, symmetry and elegance

Principle of (translational) symmetry. If points 1,2,3 are raised and highlighted in the intro, those points should be followed up in the experiment, results,and discussion. (special case of using reader expectation)

Don't change the order the points are addressed.

When introducing a key concept, give it a name and use only that name when referencing it later

If jargon is necessary,pick one word per concept and stick with it. Don't use multiple jargon words that mean the same thing.

Sentence and paragraph structure

■ Rhetorical principles from Gopen & Swan:

Subjects should be followed as soon as possible by their verbs

Information to be emphasized should be placed at "points of syntactic closure". E.g. stress positions in a sentence are typically at the end.

i.e. "save the best for last"

Place the person or thing whose "story" a sentence is telling at the beginning of the sentence, in the topic position.

Place context or "old information" at the beginning for linkage backward, and contextualization forward.

Unit of discourse (e.g. sentence, but regardless of size) should serve a single function or make a single point (applies to more than sentences).

(Tip. You may have several points you want to make in a paper. But it is rare that all points have equal importance. Determine the priorities. There should always be a "take home message".)

Readers expect the action of a sentence to be articulated by the verb. Use informative verbs (as in information theory. "is, "are presumed to be", "has" are low information words.)

Provide context before asking reader to consider anything new ("motivation motivation motivation")

Try to ensure that the relative emphases of the substance coincide with the relative expectations for emphasis raised by the structure.

Avoid the problem **The Foundations** sang about: "Why do you build me up (build me up) Buttercup ,baby Just to let me down (let me down) and mess me around"

■ Economy, "redundancy reduction"

I often review scientific journal submissions that are too wordy. One of the best ways of avoiding unnecessary verbiage is to follow the principle of logical flow at multiple levels. There are too many sentences (or paragraphs or subsections) that just seem to be elaborating on what went before, or that are tangential. In other words, the article fails to follow the principles of logical flow. One could advocate a kind of minimax rule: "maximize information transfer with the minimum number of units", but I don't think advice is that useful. But I do think it is the result of working towards logical flow.

But, don't over-do-it! I.e. avoid telegraphic writing.

Great quick overview here:

A Pocket Style Manual 2nd Edition. Diana Hacker.

Practical advice on the role of figures and captions

A reader should be able to understand the basic take-home-message of a paper from its abstract, figures, and figure captions.

A paper can often be substantially shortened and clarified with a good figure. As we've all heard, "a picture is worth a thousand words", and beautiful examples can be found in Tufte's classic books (see References).

Don't inundate the reader with lots of figures. Figures are often your best chance to win an advocate.

Examples of good writing:

Scientific audience:

I've suggested Hopfield's paper as an example in neural networks. Other favorite writers of mine are: Horace Barlow, Gordon Legge, Anya Hurlbert

Scientific layperson: George Gamow, Freeman Dyson, Steve Pinker, Matt Ridley

Examples of bad writing:

Gopen and Swan provide some examples of bad writing, and it is well worth the read.

I should be practicing what I preach, and write a well-written notebook about writing. So far it is mainly a lecture outline. For me, I allow much more work to write a paper (40+ hours) than to prepare a lecture. This notebook will evolve, and then hopefully some day I can look back at earlier drafts and see how well the notebook follows the advice it contains.

In the meantime, feel free to point out inconsistencies!

Next week's topics

■ Neural representation: examples from vision

■ Neural representation & coding

References

Hacker, Diane. A Pocket Style Manual 2nd Edition. (Third edition from Amazon)

A good cookbook. Short, handy, good examples, with the most important and useful stuff at the beginning.

Gopen, G.D., & Swan, J.A. (1990). The Science of Scientific Writing. *American Scientist, 78*, 550-558. (See American Scientist online, html)

Provides an excellent discussion of general principles. Not a cookbook.

Tufte, E.R. (1983). The visual display of quantitative information. (p. 197). Cheshire, Conn. (Box 430, Cheshire 06410): Graphics Press.

Tufte, E.R. (1990). Envisioning information. (p. 126). Cheshire, Conn.: Graphics Press.

Tufte, E.R. (1997). Visual explanations : images and quantities, evidence and narrative. (p. 156). Cheshire, Conn.: Graphics Press.