Writing Better Papers

Paul H. Dear

November 2011

This is very short summary of the key points of a workshop that I run for students and post-doctoral scientists in the life sciences. The course teaches scientific writing, with particular emphasis on getting papers published in good journals. The full course lasts two or three days, and includes a variety of exercises.

1. Know the enemy!

1.1 – The Editor

The gatekeeper to the journal is the Editor. It is his responsibility to set the tone and scope of the journal, and to ensure that only the best papers are considered. Depending on the journal, up to 80% of manuscripts may not get past the Editor.

He is faced with a pile of manuscripts each month, and he is only human. He will first try to eliminate as many of them as possible, so that he can focus on what remains. (If you've ever advertised a job, think about what you do when faced with a pile of CVs – you start by looking for reasons to cull them.) Most authors assume that the Editor is there to welcome papers, but this is a grave misunderstanding: *the Editor wants to reject your paper, and it is your responsibility to convince him not to.*

The most common reason for Editorial rejection is that the paper does not appeal to a sufficiently large part of his readership. When writing and submitting a paper, bear in mind that it must be of immediate relevance to a good proportion of the journal's readers, but it must also be of at least *some* interest to *all* readers of that journal. This is why getting published in Nature or Science (whose readers range from geologists to virologists) is so difficult. We'll come back to this point a bit later.

1.2 – The Reviewer

If your paper gets past the Editor, it will be sent out to one or more reviewers. Whereas the Editor wants to reject your manuscript, the reviewers may start out with a favourable, neutral or hostile attitude. Also, whereas the Editor has a vested interest in giving conscientious and unbiased evaluation of all manuscripts (and may be paid to do so), reviewers are seldom paid and may have their own agendas; they may also be too busy, or just plain dumb.

Every author will, sooner or later, suffer at the hands of a referee who "completely misunderstood the whole ******* point!" As an author, this is *your* fault. If something can be misunderstood, it will be misunderstood, and it's your responsibility to ensure that the reviewer cannot misunderstand what you have written.

Whereas the editor is looking for the overall content and relevance of your paper, the reviewer will be scrutinizing points of detail and scientific quality. Nevertheless, his (or her) initial impression will be formed early, and it's difficult to change that first impression. Again, this is a point we'll return to later.

2. Who writes the paper, and when?

There is no "I" in "team", but there's an "I" in "science" and a couple in "writing". Good papers are never written by a committee. Even though the work may be a collaborative effort, one person needs to take responsibility for putting the manuscript together.

If that person is you, don't spend endless days asking people whether they like this wording or that. Start with a group discussion, then write the paper in the way you think best, and let them tear it to pieces afterwards. Also, do not assume that every last author wants to read, re-read and comment on every version of the manuscript; this takes forever. Give all the authors an opportunity to have their say, but not after every full stop and comma. When you need input from the other authors, ask them as precisely as possible for what you require – "I need two paragraphs on the ELISA assays and a figure for the hormone time-course".

As to when you should start writing the paper, my advice *should* be to start outlining the paper when you're about two thirds of the way through the research – this will help to highlight any experiments that need to be done to support an argument. However, nobody (not even me) *ever* does this – papers always get written at the end, and then there's a scramble to repeat an experiment and get a decent gel picture.

3. Finding the story (and a bit about ethics)

A paper is a story, but research isn't. Your research follows a tortuous path – the initial premise of the research may be unfounded; key experiments fail or the money runs out; the final experiment suggests that your first experiment should have been different.

In writing a paper, it's your responsibility to pull this mess into shape and present a coherent narrative that flows logically. You can't change the facts, but within reason you can (and should) present them in a way that makes sense.

There's a grey area here, between arranging the facts in a way that is easy to follow on the one hand; and distorting the truth on the other. It's a fine line to tread but, broadly, the reader is not interested in (and needn't know) all of the incidental blind alleys and mis-steps you took, unless they impact on the science itself. This holds true for all aspects of paper-writing: you need to make your work presentable and attractive to as many readers as possible, but you have to stop short of spin and distortion.

One other point about story telling. It is often the case that work which took 90% of the time and 90% of the budget only deserves to take up 1% of the paper. If you have sequenced and analysed a genome, for instance, the sequencing may have been by far the largest (and most expensive) part of the work. But, by and large, the sequencing can be summarized by a couple of sentences in Materials and Methods – the reader does not need to share your pain.

Always stop and ask yourself – is this in the paper because it's important, or because we spent a lot of time and money on it?

4. Appealing to a broader readership

The commonest fault in papers, and the commonest reason for rejection by the editor, is that they are too tightly focussed from the outset. You may have spent fifteen years working on a protease in a virus that infects cave-bats in Argentina, but there are probably only three other people in the world who are interested in that protease, of that virus, of those cave-bats.

So, step back. Why does the virus matter? Maybe it is a member of a class of viruses with much wider impact on other species, including humans. Maybe that protease is so peculiar to that class of viruses that it is a potential target for antiviral drugs. Maybe if the bats all die, the mosquito population rises and more people die of malaria. Maybe... well, it's up to you to find the "maybes". But you *must* put your work into the broader context.

An excellent idea is to listen to yourself next time you're explaining your work to a friend or relative who knows nothing about it. You will find yourself explaining *why* it's interesting, *why* it's relevant, *why* it matters. You need to put some of that into your paper.

5. The key parts of the paper

It is unlikely that anybody, ever, will read your paper from beginning to end – think about how you read papers. Certainly the editor will have formed an opinion after reading a dozen sentences or so, as will the reviewer. Most readers of the journal – if it makes it to print – will read a few sentences here and there and then move on, unless you can grab their attention.

For this reason, it is vital that you make your key points concisely, clearly and at *several* key places in the paper. The first place, obviously, is the abstract, and we'll come back to that much later. Other key places are at the beginning of the introduction (where you outline the problem and why its interesting); at the end of the introduction (where you briefly summarize your approach to the problem); at the beginning of the discussion (where you remind people of the problem); and finally at the end of the discussion (or the conclusion; where you summarize your findings and point out where things should go from here).

These key parts of the manuscript must contain the distillation of your paper. A busy editor will skim through these parts and, if you haven't grabbed his attention by then, you're sunk. Get those parts right, though, and the editor will be on your side and may forgive the odd problem in the bulk of the paper.

6. First draft – generic format

You may have already decided on your target journal, and know its format and length restrictions. Nevertheless, it's usually best to write the first draft in a "generic" format, and to ignore length restrictions or limits on the number of figures. Start by writing the paper the way you want it; you can worry about shoehorning it into the journal's requirements later. If you worry too much about format and length at the outset, it's difficult to find the right story for the paper.

A lot of first-time authors find it difficult to decide where to put things in the manuscript – does this belong in methods or results? As a rough guide:

- The Abstract is a short distillation of the paper (we'll come back to this later).
- The Introduction sets the broad context of the work and emphasizes its relevance to as wide an audience as possible; it then explains the specific problem being addressed and *it's* importance; it concisely summarizes what has already been done in the field; it explains the approach you've taken to answer the question; and it very briefly indicates the major findings ("Here, we show that....", "We find that...").
- The Materials and Methods contain all of the experimental details, including reagents, buffer compositions, and the protocols.
- The Results section is tricky. You cannot give *only* experimental results, or it's impossible to follow the whys and wherefores. You need to briefly explain why you did an experiment ("To see if this protease was active during infection, we..."), what you did ("...harvested cells 7hr after infection and prepared cell extracts as in M&M..."), what you found ("The levels of protease were significantly higher....see fig. 3"), what this means ("This suggests that the inflammatory response...") and how this leads to the next experiment. In other words, there needs to be enough discussion of each result in turn, to make a flowing narrative. The reader should be able to follow the entire results section by itself, referring to Materials and Methods only if they are curious or want to replicate the experiments themselves.
- The Discussion should begin with a gentle introductory sentence, reminding the reader of the problem (remember – an Editor may flick straight to the Discussion when he first sees the paper). Then, take the reader through the results, discussing and explaining their significance. Do *not* re-state the results in detail – just remind the reader of the key findings. Don't say "In the infected cells, the levels of cytokines were 56.3% higher after seven hours, and 87.3% higher after 12 hours";

instead, remind the reader of the trends and explain what they mean ("In infected cells, cytokines levels rose steadily, suggesting that..."). As you get toward the end of the Discussion, broaden it to return back to the wider problem, the relevance of your findings to the big picture, and perhaps where the research can go next. You can be a bit more expansive and speculative here, within reason. It's helpful to "signpost" the end of the Discussion with something like "In summary..." or "In conclusion..." – this helps an editor or reader to quickly find out what you've done.

- The Conclusion is kept as a separate section in some journals. If so, then your "In summary..." section belongs here instead of as the last paragraph of the Discussion.

7. Guiding the reader through the narrative

You may have spent a long time working out a flowing and logical narrative for your paper, but this beautiful structure may not be apparent to the reader. Therefore, it is very important to use "linking phrases" to explain how you get from one step to the next, and reminders of points you made a while ago. For instance: "The elevated cytokine levels suggested an inflammatory response, so we...", or "Since we had already shown that these cells did not proliferate, we...". These phrases can be made very compact, but make the paper much easier to follow.

8. Defensive writing

If your paper has weaknesses, it is unwise to leave them to the reviewer to point out. Instead, acknowledge the limitations of your work and explain why your conclusions are nevertheless valid.

For instance, if your sample size is very small, don't wait for a reviewer to tell you that your conclusions aren't warranted. Instead, say something like "Although this is only a small representative sample, the elevated cytokine levels seen in all cell types strongly suggests...".

Beware of saying something like "This proves that..." unless you can justify it; instead, say something like "This strongly suggests that...". If there is an alternative explanation for your results, don't ignore it: raise the possibility before the reviewer does, and then explain why you think your explanation is more likely.

Having said that, don't water down your findings unnecessarily.

9. Keeping the reader on board

You are, hopefully, an expert in the subject you're writing about. Your readers, on the other hand, will have a wide range of expertises. This raises the problem of how to give less informed readers the information they need, without boring or patronizing those who know the subject inside out. The trick is to slip information in, in passing. Consider this sentence:

> "GCL1 cells were treated with n-butyl tertalamide to induce expression of *KPL1."*

That sentence is impenetrable to anyone who doesn't know the subject in detail. You could get around this like so:

> "GCL1 cells (which are derived from primary fibroblasts) were treated with n-butyl tertalamide (which is a DNA damaging agent) to induce expression of *KPL1* (which is a protein involved in DNA repair)."

but this will irritate people who are familiar with the subject (or who think they are). The better option is to slip the information in more subtly:

> "The primary fibroblast cell line GCL1 was treated with the DNA-damaging agent n-butyl tertalamide to induce expression of the DNA repair protein *KPL1."*

In this way, readers who know the subject don't feel they're being lectured, and readers who don't know can pick up the information as they go along.

10. Writing style

Getting down to the nitty-gritty of writing style would fill an entire book. A few quick pointers:

- Write in the *active* voice in general; most journals prefer it nowadays, and it is clearer and more concise.
- Avoid contractions ("don't", "doesn't", "it's" etc) and colloquialisms.
- Your writing should be formal, but not contrived. Try reading a sentence aloud: if it sounds painfully unnatural when spoken, re-write it.
- Remember, if something can be misunderstood, it will be misunderstood, and it will be your fault.
- Buy (and use) a copy of "*The Elements of Style*" by Strunk and White; it is an excellent and concise guide to good writing.

11. Figures, and why they matter

Figures are of great importance. The editor (your first point of contact with the journal) wants his journal to look professional – sloppy or ugly figures will prejudice him against your paper. The reviewers need to be able to get the information they need quickly and easily, as of course do your readers if you make it to publication.

Figures have to be technically perfect. Now that everyone has access to graphics software and graphing programs, there is no excuse for sloppy figures – they just make your paper look amateurish. You should expect to spend many hours tweaking the figures.

The figures should show what is necessary, clearly and elegantly. Avoid unnecessary ornamentation or the "comedy special effects" courtesy of Excel. Use colour only where necessary, but bear in mind that if part of one figure uses colour, then additional colour on the same figure doesn't usually cost more.

Wherever possible, use "vector graphics" rather than "bitmap graphics". In vector graphics (such as Adobe Illustrator), lines, shapes and letters are represented by coordinates and can be scaled up or down as necessary. In bitmap graphics, everything is represented by pixels; the image may look grainy when enlarged, and is difficult to edit.

Compose multi-section figures carefully: keep all graphs in the same format, all labels in the same font, and (where possible) all sections the same size. Keep the various sections neatly aligned, as if you were hanging pictures on a wall.

12. The all-important abstract

I've left the Abstract until near the end of this guide for two reasons. First, it is of paramount importance; and second, its content can only really be finalized when you have worked out the rest of the paper.

The abstract should be a perfect "thumbnail" of your entire paper. You have perhaps a couple of hundred words in which to state the problem and its relevance, explain you approach, report your findings and emphasize their significance. On no account clutter up the abstract with minutiae.

Every single word in the abstract should be chosen with care. Once again, it's a good idea to listen to yourself explaining your work to non-expert colleagues or friends; the things you tell them are likely to be the backbone of your abstract.

Finally, remember that people finding your article online (for instance, in PubMed) may only have access to the abstract.

(A few journals have very strict rules regarding the precise format of the Abstract, in which case you have less flexibility.)

13. The Title

In general, avoid over-long and over-detailed titles. You can instead indicate the field of study (*"The inflammatory response to bornavirus infection in fibroblasts."*), or the essence of the findings (*"Bornavirus infection triggers a cytokine cascade in fibroblasts"*).

Unless it is an extremely specialized journal, avoid titles like *"Infection of GCL1 immortalized fibroblasts with bornavirus strain FS4 causes elevated expression of IL-1, IL-6 and TNF-alpha within seven hours."*

14. Authorships and author order

Deciding who should be authors on a paper, and in what order, is a subject that causes great angst. Policies vary from lab to lab but, as a rough guide, authorship should be given only to people who have had a significant hands-on or conceptual role in the work. "Significant hands-on" would perhaps be more than a day's work. The "conceptual role" typically includes the group head (who,

presumably, shapes the research themes in the lab); it may extend to the departmental head, but often it does not.

People who have contributed reagents or loaned equipment should generally be acknowledged rather than given authorships, but there are exceptions. For instance, in some clinical studies, the collection and archiving of clinical specimens used in the research involves considerable work and may warrant an authorship.

Your paper may make use of data from other researchers (for example, it may be a bioinformatics paper based on sequence data from elsewhere). If the data has already been published (in print or in a public database), it's not usually appropriate to give an authorship to the researchers who produced it. There are exceptions, however. For example, some genome centres make all of their raw sequence data freely available, but request that others do not publish extensive genome-wide analyses of the data without consulting them first.

Work on the manuscript alone doesn't warrant an authorship (but should be acknowledged instead), unless of course the paper is a review article, or the person has contributed greatly to the interpretation and theoretical framework.

Policies also vary widely concerning author order. The most common practice in the life sciences is to put the "hands on" people at the front (with the first author being the one who did the bulk of the benchwork), and the "ideas" people at the back (with the last author being the one who conceived the research). It's really only the first four and last four author positions that matter.

Some journals allow "statements of equal contribution" or "statements of joint senior authorship". These can help to soothe bruised egos, though it's doubtful if most readers notice them. The position of "corresponding author" carries a little more weight.

15. Standing back, and getting a realistic critique

Before you start the painful process of preparing to submit your paper – stop! Take a break from the paper for a few days, and then come back to it. It is very easy to get wrapped up in the details of the writing, and to miss the fact that the story has evolved since you began. Step back, and ask yourself if the overall structure of the paper is right. It may be heartbreaking to realize that the whole thing needs to be organized in a completely different way, but in fact this is fairly easy to do now that you've got the material down on paper.

Next, and of utmost importance, get your manuscript critiqued by someone else who has not been involved in the work or in the writing. Ideally, this should be done by at least two people: someone from a different field; and someone who is familiar with the field. They will be able to give you the "editor's perspective" and "reviewer's perspective", respectively.

Do not simply give them the paper and ask for their opinion: people (especially colleagues) are inherently polite, and you will probably only get minor comments that don't address the real problems. Instead, give them the paper and say "I want you to give me five reasons for rejecting this paper." In that way, they will feel free to find and report the weaknesses in the paper. If someone tells you that the paper is great and ready to send off, get another opinion.

16. Reformatting for the journal

If you've written the first draft in a generic format, reformatting for a specific journal is generally a mechanical process.

The main problem is in reducing the word-count of your paper to fit the journal's limits: it's quite usual to find that your paper is too long by a factor of three! There are various ways to cut words:

- Cut out sections. Often, a section is written early because it seems important but, in the end, it turns out not to be so. Whole paragraphs can often be omitted, or replaced with a single sentence. You may bruise some egos (for instance, if work which took someone six months is reduced to a sentence), but you are writing for the readers, the editor and reviewers, not for the authors.
- Move material. Lots of detailed results can often be moved into "Supplementary Material". This has the advantage of making the important points stand out more clearly in the main text. Figure legends can also, in many journals, contain a lot of information taken from the main text.
- Shorten sentences. The overwhelming majority of written sentences are capable of being reduced in length very considerably. Or, to put it in only 1/3rd as many words, "*Most sentences can be shortened*."

17. The covering letter

The covering letter is your way of introducing yourself and your work to the editor, and needs considerable care. Typically, you start with something like "*We are pleased to submit our manuscript entitled '...' to the Journal of ...*" : if you're not pleased, why should the editor be?

Briefly summarize the problem, your approach and your findings (in outline only), and explain why you feel the work is novel and important. It also doesn't hurt to flatter the editor (or his journal and his readers) a little, "*We feel that publication in the Journal of ... would enable our work to reach the most appropriate readership.*"

18. Responding to referee's comments

If your paper is one of the 20-30% that make it past the editor, and is amongst the 50% of *those* that don't get rejected after the reviews, then you are doing pretty well. However, you will almost certainly have a list of points raised by the reviewers that need to be addressed to the editor's satisfaction. (Editors tend

not to send papers back to reviewers after they've been amended, unless the points in question are complex or contentious.)

You must respond to each reviewer's point in turn, either by making the necessary changes or (in a very few cases) by explaining why they are not needed.

It is helpful to the editor if you quote each of the reviewer's points in turn, and then give your response. Be courteous to the reviewers (through gritted teeth, if needs be), and never dismiss their points lightly. If the reviewer has misunderstood, you need to clarify the text as well as explaining the reviewer's mistake to the editor. A typical response letter might go like this:

> "Dear Sir,
>
> We are very grateful to the referees for their careful reading of the paper, and are delighted that they have broadly recommended publication. We have addressed each of their points as follows:
>
> Referee 1
> "*Not enough detail in figure 2*" - we have amended and clarified the figure.
>
> "*Some of these conclusions were reached in 1978 by J.P. Smith*" *- we are grateful to the referee for reminding us of this. We have added text in the Introduction to correct this oversight, and added the necessary references.
>
> [*so now you know who the referee was!]
>
> "*They have failed to do the necessary control for the haemoglobin data presented in figure 3*" - The control experiment was described in para. 7 of Results, but we have modified the text to make this clearer.
>
> Referee 2
> "*The interpretation of the data in figure 4 is doubtful. The data could also suggest that this gene is upregulated in T-cells as a result of a the immune response.*" We had considered this possibility, but on balance the weight of evidence suggests otherwise (see figure 5). Nevertheless, we have added a sentence in the discussion both to raise and to argue against this interpretation.
>
> [etc.]
>
> In conclusion, we feel that the manuscript has been greatly improved now that we have addressed all of the referees' points, and we look forward to your response.
>
> Yours sincerely "

19. Final words

Writing papers is hard work, and writing papers that get published is even harder. Always try to look at your paper from the perspective of the editor, the reviewers, and of course the readers.

Good luck!
Paul H. Dear