



Writing Modern Geoscience

Randolph T. Williams  *1

¹Department of Geoscience, University of Wisconsin-Madison, Madison, WI, USA

Abstract Scientific writing is a cornerstone component of modern geoscience practice. Despite this, the subject rarely benefits from any explicit instruction during graduate training. This editorial attempts to address that gap by providing a simplified approach to writing modern geoscience papers. Implementation of this approach requires aspiring scientific writers to first consider why we as scientists write papers and what our goals should be for our audience. From there, specific and detailed advice is provided on how to best structure the idea flow and presentation within each of the primary components of a peer-reviewed manuscript (e.g. introduction, methods, results, discussion, conclusions, and abstracts). Examples of annotated, illustrative text are provided throughout. Considerable attention is paid to the notion of a scientific paper as an argument (as presented primarily in its discussion section) in favor of some proposition, and the key aspects of what makes an argument compelling. The goal is to provide new or struggling scientific writers with an explicit mechanism to construct impactful peer-reviewed manuscripts more easily.

Technical Editor:
Mohamed Gouiza

Reviewers:
Maëlis Arnould
Robin Lacassin
Craig Magee

Published:
11 January 2024

1 Prelude

Developing proficiency in scientific writing is arguably the most difficult aspect of scientific practice. Most graduate programs offer little explicit instruction on the subject, which is surprising, given that the potential for solid writing is one of the more important qualifications for a successful scientist to have. Much of our formal education as it pertains to writing is provided in our undergraduate years, and it is arguably focused primarily on sentence structure and grammar. If nothing else, sentence structure and grammar are things that most graduate students I encounter do at least passably from the outset (recognizing of course that English is a non-native language for many, which imposes additional difficulties on the writing process)¹. A larger issue with scientific writing I encounter more regularly relates to issues of writing structure and idea flow, under- or over-describing methods and/or results, a wayward sense purpose (i.e. what is the actual point of this paper anyways), and ineffective or spurious arguments in favor of preferred conclusions. I am also led to believe that many people, particularly those early in their careers, are keenly aware of (and indeed, fixated on) their limitations as scientific writers, and that this is a source of significant stress. This was the case for me as well. What follows here is my attempt to simplify

the process of writing modern geoscience papers, in the hopes that these ideas may facilitate more efficient construction of scientific manuscripts that are clear and compelling.

I suspect that at least some of the stress of scientific writing comes from the fact that we know it is done primarily by scientists. Scientists are supposed to be smart, and writing is something that smart people do well. The preceding statement is, of course, nonsense. Writing is a skill like anything else and it needs to be practiced to be learned. It is a subject that, in many ways, cannot be mastered. There is always room for improvement, and in my experience the spectrum of proficiency in scientific writing is quite broad regardless of career stage. At this moment, any particular person may be a weaker or stronger writer when compared to their peers, but this says little about what is possible (or likely) in the future. In short, it is best to cast aside any notions of what one *should* be able to do in terms of scientific writing. It is much more effective to focus instead on what one *can* do in terms of scientific writing, and I can say with confidence that virtually anyone who understands the topic they are writing about can write about it reasonably well².

It is important at this point to appreciate that there is a difference between an exceptional paper and a good paper, that most papers are “average”

*✉ twilliams@wisc.edu

¹For those potentially struggling with issues of style on a more granular level (i.e. sentences and words as opposed to paragraphs and arguments), I highly recommend William Strunk and E.B. White's classic text "*The Elements of Style*" (*Strunk and White*, 2000), which is now free as part of the public domain

²Sometimes a resolute inability to write about a topic in a clear and concise way is an indication that one does not yet fully understand it. This is a lesson that can be difficult to accept, but it is very useful. I have returned to it many times throughout my career.

by definition, and that average does not necessarily mean “bad”. Here I will argue that the most reasonable goal for an aspiring scientific writer should not be to write an exceptional paper - it should be to write a good paper. If one accepts this framework, then some exceptional papers are likely to happen spontaneously along the journey, but these will be noteworthy (and hopefully impactful) exceptions within a large body of otherwise solid work. There is also a serious practical advantage to this approach in that writing “good” papers is not very difficult. Understanding why depends on recognizing the purpose of publishing papers in the first place.

The purpose of a scientific paper is to convey information. It really is that simple. Ultimately, it is possible that your readers decided to engage with your paper because they are interested and curious, but more likely it is because they suspect that it has something that they need. Some piece of information, or some framework of ideas that allows them to better contextualize their own work or interpret their data. Reading scientific papers, at least outside of peer review, is therefore an inherently self-centered activity. As an example, try to recall the number of papers you read in the last year for no purpose other than idle curiosity. The purpose of most papers today is therefore not to inspire your reader, or induce them to change their opinions, or convince them that you are particularly insightful. Rather, these are things that may well happen on their own if you simply show your reader some data/analyses/interpretations that are useful to them in a package that is easy to understand. This, in my opinion, is the primary difference between an exceptional paper and a good paper. The former *does* manage to inspire its readers and change their conception of some aspect of geoscience. The latter just easily and accurately conveys information that the reader needs, perhaps with a bit of convincing as to why they need it. It is entirely possible of course that a good paper may also inspire or act as a catalyst for change locally, but the potential to do so is likely more a function of the person reading it than the quality of the writing itself.

With the aim of helping those in need to write better papers more efficiently, I will provide here a few simple and explicit mechanisms for moving from a blank document to a coherent draft of a manuscript in relatively short order³. A variety of examples will be provided from the published literature, many of which come from papers on which I am an author. These selections are not included for the thought that they are inherently better than any others, but rather that they came more easily to mind when I was searching for illustrative text. Much of what is presented here are the results of what I learned

³For those looking for a deeper dive into many of the subjects discussed in this editorial, I recommend “*Writing Science*” (Schimel, 2012) by Joshua Schimel and “*Writing About Archeology*” (Connah, 2010) by Graham Connah (which contains a variety of useful writing information outside of its specific disciplinary focus on archeology).

as a graduate student, postdoctoral fellow, and as an independent, soft-money research scientist in academia. At all stages of this process I have had the privilege to work with some exceptional writers from whom I learned a great deal (see acknowledgements). Although this guide is described throughout in terms of constructing peer-reviewed papers, many of the concepts are just as readily exported to proposals, peer-review reports, etc. Many of the examples provided will be from the fields of structural geology and tectonics, although the techniques should have wide applicability in in other areas of earth science and beyond.

1.1 A Word on Manuscript Structure and Leveraging Figures

The vast majority of papers published within the geosciences follow a roughly consistent structure for first-order sections: Introduction, Methods, Results, and Discussion (sometimes referred to as the “IMRaD” format). Most journals will also require a preceding abstract and many authors opt for a final Conclusions section as well. The reason for this consistency is simple - it is very effective at conveying information in a way that maximizes the potential for reproducibility and minimizes the intermingling of results and interpretations. Even when a paper omits some of these exact terms the general flow is much the same. The reader needs to know why the work was done before they will care about how it was done. They need to know how it was done before they can assess if the results are valid. They need to know what the results are and if they are valid before they can be made to understand what those results mean within the broader picture of the field of research. In truth, there is rarely a need to move away from this structure. Recall that the goal is to produce a good paper and the requirements of a good paper are quite simple. Peer-reviewed papers are not eligible for Pulitzers or other popular literature prizes. Novelty in paper structure can be effective if done well, but they are not required to produce a top-notch manuscript. This editorial will therefore be premised on the use of the IMRaD structure. Strategies for effectively constructing abstracts, conclusions, and titles will be provided at the end of the editorial, and indeed it is generally best to construct those components of your manuscript after completing at least a rough draft of the main body of text.

Although this editorial is focused on *writing* of manuscripts for peer-reviewed publication, no amount of text, no matter how elegantly and concisely constructed, can replace the need for effective illustrations in the form of figures. In fact, one common and effective strategy to structuring the flow of information and writing is to build the presentation of a manuscript around figures. At a minimum, rough drafts of key figures illustrating important aspects of our methodology; fundamental data and/or results; and synthesizing results into interpretations via schematic or conceptual diagrams

should be constructed before writing begins. There are (at least) two reasons for this approach. The first is that there is a large visual component in our ability to digest and understand information, particularly when it is spatially or numerically complex. The second is that by building the presentation of a manuscript around our figures, the figures and the text can share the load of conveying information. This enables us to write in a more clear and concise manner because we can effectively replace (what would otherwise be) large blocks of descriptive text with a citation to an illustrative figure. This is not to say that one should not provide any description of the results or interpretations within the text. Rather, one should attempt to build synergy between the text and associated figures for maximum clarity. Although detailed strategies for constructing clear and compelling figures are beyond the scope of this editorial, a variety of resources on the subject are available⁴

2 Writing Introductions

The introduction is, quite possibly, the most important section of a paper. A good one primes the reader to receive whatever information/results/conclusions you are hoping to convey. In an increasingly attention starved community, however, it also needs to strike a balance between setting up your work for proper engagement/appreciation and boring or overloading the reader with superfluous details or bait-and-switch topics. It is also a somewhat difficult section for those earlier in their careers, in large part because a good introduction moves seamlessly and naturally from the broad field of research in which your work is making a contribution to the specific way that contribution was made, which in practice often entails answering somewhat niche and specific questions. Where is one to begin? The answer of course, is “at the beginning”, although ironically it is very often useful to write the introduction last after you have a working draft of the other important sections. We will return to this last point later in this section.

2.1 Introduction Structure

Thankfully, a very basic formula for writing compelling introductions quickly and easily is available. I would argue that it is on some level quite popular already, as its result is demonstrated nearly ubiquitously in the peer-reviewed geoscience literature (I encourage you to have a look after reading this section), although it is rarely (if ever) explicitly dissected as a matter of practice. I will refer to it as the 2-3-4 formula. At the grossest scale, your introduction should be composed of either two, three, or four paragraphs (which is “best”

⁴For a good starting point on the subject of figure construction, see “How to Write and Illustrate a Scientific Paper” by Björn Gustavii (Gustavii, 2017)

depends on where the manuscript will be submitted, as discussed below). The three-paragraph format will likely be the most effective approach for most papers. Taking the three-paragraph format as the most general case, its structure is quite simple:

- Paragraph 1: Set the stage for your work by providing a concise overview of the general field of research. If possible, forecast a problem or unknown within that field of research
- Paragraph 2: Define and expand on a problem or unknown in the field of research you introduced previously that your current research can potentially help to address. If possible, forecast a potential solution in a general way.
- Paragraph 3: Describe the way in which your work attempts to provide that solution, a brief synopsis of the degree to which that effort was successful (i.e. the results), and tie it back to the general field of research you introduced in the first paragraph.

We can consider the two- and four-paragraph structures as special case deviations from the three-paragraph structure. For example, the two-paragraph structure is generally most useful in short format publications focused on “low-hanging fruit”. Stated alternatively, it is useful in the happy case that the importance of your research is going to be so obvious to the target audience that you can convey in a single paragraph the general area of research and the problem that needs to be addressed before moving on to how your work addresses it (i.e. you only really need paragraphs 1 and 3 from the three-paragraph example above). The four-paragraph structure, conversely, can be useful in comparatively long-format venues covering a broad range of ideas/results/data, as is often the case in review or synthesis type papers (i.e. you need all the paragraphs shown in the three-paragraph structure above, with an additional paragraph between the general overview and expanding upon the problem or unknown to better contextualize your work).

2.2 Introduction Idea Flow

To effectively implement the 2-3-4 formula, I recommend framing each paragraph in terms of a “topic sentence” and a “forecast sentence”. Many have likely encountered the idea of the topic sentence in undergraduate writing courses. In short, the topic sentence expresses the main point of the paragraph. Everything that follows in the paragraph should be resolvable back to the topic sentence. The forecast sentence, in contrast, sets up the ideas in the next paragraph. Perhaps more mechanistically, we could consider the forecast sentence of one paragraph as setting up the topic sentence of the next paragraph. Think of it this way - *well constructed topic and forecast sentences in your introduction should flow naturally to the extent that an expert in your field could get a good*

sense of the goal of your research, why it is important, and what was learned from those sentences alone. The sentences between the topic and forecast sentences in each paragraph are intended to guide a reader with less experience along a slightly longer path to the same endpoint understanding.

Consider the following three-paragraph structure example from Phillips and Williams (2021) where the topic sentence is highlighted in blue and the forecast sentence is highlighted in green (citations and figure calls omitted).

Cataclasis is a fundamental mechanical mechanism shaping fault-rock evolution and is driven by continual fracturing and particle size reduction during sliding along a fault interface. This process is recorded in part by the abundance of fragmented rocks in and around fault zones, and controls fault-rock porosity, permeability, frictional resistance to sliding, and the formation of new surface area (which consumes some portion of the earthquake energy budget). *Thus, mechanistic models of cataclastic deformation are essential for understanding the evolution of fault rocks in the brittle crust.*

Although several mechanistic models for cataclastic deformation have been proposed, "constrained comminution" has gained the most wide-spread acceptance. Constrained comminution describes fracturing of nearest-neighbor grains of similar size through stress bridges during cataclasis. Sammis et al. (1987) argued that this process should lead to power-law distributions of particle size, a signature which might be identified in the rock record. In log-scaled plots of cumulative particle-size frequency, power-law distributions manifest as straight lines. These log-linear relationships imply self-similarity within a data set, from which the fractal dimension (D) can be estimated by the slope of the distribution. Many previous studies have recognized this fact and attempted linear fits to measured particle size distributions to demonstrate the presence of power-law scaling (i.e. inferring the action of constrained comminution). This approach, however, suffers from limitations. First, previous studies have relied on either visual (qualitative) linear fits of particle size distribution data, or attempted linear regression analysis, which is not robust when performed on log-transformed data. Second, and more critically, other distributions (e.g. log-normal, exponential, stretched-exponential) may also appear linear in log space. Linearity in these plots alone is therefore an insufficient metric for inferring power-law scaling.

Thus, the presence of power-law scaling in previously examined fault rocks has not been definitively demonstrated, and it remains unclear if constrained comminution is a viable descriptive model for cataclastic deformation.

In this paper, we test for the presence of power-law scaling in naturally and experimentally deformed fault rocks, thereby assessing the validity of the constrained comminution model. We employ a well-established statistical approach, which determines a best-fit power-law function for a particular data set, tests the data set against a suite of synthetic data sets with identical power-law parameters, and then compares the best-fit power law to other possible distributions. When applying this statistical approach to a suite of 61 samples we find that power-law distributions are a poor fit to the data in all cases. *We propose that fault-rock particle-size distributions are better fit by log-normal distributions and discuss the implications of this distribution for interpreting the micromechanical mechanisms of cataclasis.*

Now read just the topic and forecasting sentences contiguously without the intervening sentences. They go something like this (paraphrasing):

Cataclasis is a central component of brittle deformation in the crust, but more explicit models describing the process of cataclasis would help us better understand the micromechanics of fault slip. The "constrained comminution model is one such possibility, but it is not clear if the data from actual fault rocks support that model. We tested to see if real fault rocks support the constrained comminution model using statistical analysis, and found that they do not.

In effect, this is a very simple, one paragraph synopsis of the entire paper using less formal language. Someone with experience in researching brittle faults and fault rocks would likely have no problem following this and developing an immediate (albeit basic) understanding of the goal of the paper and its results. The popular term to describe something like this is an "elevator pitch". Your introduction is then an expanded elevator pitch that entices the reader to understand why your work is important and continue forward. The intervening sentences in the actual paper are intended to provide more specifics for the expert reader and a more general basis for understanding for the novice. At a more granular level, each sentence in the elevator pitch also accurately summarizes its corresponding

paragraph in the text.

This example leads to the most concrete advice I can provide on constructing effective introductions - build them forward from just a few simple sentences/ideas, generally only one for each paragraph using the 2-3-4 formula described above. Start very basic and then work toward more complex ideas. When you have constructed a simple elevator pitch that would resonate with an expert in your field, fill in the ideas to provide better guidance for more novice readers.

It is possible that some may view the ideas above as overly simplistic or too prescribed. To that I can recommend a simple activity: pull up some of your favorite papers, read their introductions, and see if you can reconstruct their elevator pitch. What is the point of each paragraph? How do the ideas lead into the next paragraph, and how does it all come together into a coherent picture? I will wager that for any paper you personally feel has a well written introduction, it will be quite easy to deconstruct it into its component elevator pitch and core ideas. If you find indeed that most can be deconstructed in this way, then it follows that the forward process of starting from the core ideas you are hoping to get across must also work.

Now to circle back to some earlier ideas in this guide, we can ask ourselves if *Phillips and Williams (2021)* provide a brilliant, ground-breaking, beautiful introduction that will inspire countless readers for generations to come? Or did they simply provide a "good" introduction that serves the reader's need for information and context? I suspect most will agree it is the latter. *Phillips and Williams (2021)*'s introduction simply sets the stage and context of investigation, describes a problem of clear community relevance that needs addressing, and states clearly how that problem was addressed, all in three very simple paragraphs using (largely) plain language that can be read in less than 2 minutes. That being said, it is much easier to take this approach when you do not have to forecast an entire paper that does not exist yet. There will be times in your career when the purpose and direction of your work is so clear in your mind that you can write the paper all the way through from introduction to conclusions. More often though, writing papers is as much a mental exercise for our understanding as it is a report for our peers. We learn more about our topic and how best to present it through the process of writing. In this way, it is generally easiest to write an introduction after you have an otherwise complete manuscript draft that you are comfortable with, *such that the content and requirements of the manuscript draft can be reflected naturally in the introductory material.*

3 Methods and Results

There is a little known fact that can make writing compelling methods and results in our papers

considerably easier. It lies in recognizing that in reality, the results/data we generate are dependent on the methods we employ. In writing, however, the methods section is more readily informed by the results we want to report. The reason for this is simple: cogency between the methods and results sections is achieved in part by presenting our methods in broadly the same order that we present the data they produce. The results, however, generally need to be built in a specific logical order. For example, for your reader to be able to accurately contextualize your microstructural results, they need to first understand the more fundamental field observations of the outcrop the samples were collected from. To contextualize your *in-situ* geochemical or geochronological results, the reader likely needs to understand key features of the sample microstructure and where in it the analyses were conducted. As such, it is often more effective to have a complete draft of the results section before working on the methods.

3.1 Methods

This is arguably the most straightforward section of a manuscript to produce (at least once some draft results are in place). The methods section in the main text of a manuscript should state simply what methods were used to produce the manuscript's results, including only as much detail as is necessary to allow the validity of those results to be evaluated by the reader and no more. More detailed methodological descriptions (which are important to allow others to reproduce your approach), should be provided in the Supplemental Material. The differentiation between these two categories of methodology is discussed more below. As stated above, the methods should be presented in broadly the same order that their corresponding products will be described in the results. Within that, it is also generally best to move from more observational and/or qualitative types of methods (e.g. outcrop or map-scale observations, general microstructural features) to more analytical or quantitative types of methods (e.g. kinematic analysis, geochronology, and statistical analysis). This is intended in part to respect the natural flow of results described above, where it is generally our qualitative observations that lead us to select samples and approaches for more detailed and/or quantitative scrutiny. Moving from observational/qualitative to analytical/quantitative approaches also allows us to build complexity slowly, as the latter typically have more specific parameters that need to be provided. This is why statistical analyses, for example, are generally described last, as these are inherently done on data that you collected previously.

This organizational approach can also be done on a more granular level if the manuscript is describing larger, more involved research efforts with a large number of different methods. In that case, subsections may be used to organize the information,

but all of the above rules will still be applicable. More specifically, if your methods section needs to encompass multiple subsections, then move from those that are more observational/qualitative to those that are more analytical/quantitative both across and within subsections.

It is also important to note that all methods are subject to limitations, both in general and as applied to specific instances. These often propagate into issues of the quantity and quality of data that those methods produce. It is of course important to be forthright about these issues, but there are several ways that can be approached when writing a manuscript. If there are any issues that are inherent to a particular methodology, it is reasonable to address them explicitly in the methods sections. This can be done in narrative form within the overall methodological description or using explicit subsections such as "Method Limitations". Descriptions of how sources of uncertainty or data quality impact our ability to make inferences about Earth history and/or processes, however, should be contained within the manuscript's discussion section (described in detail below).

The current practice observed in most peer-reviewed papers today is moving away from providing highly detailed descriptions of research methodology in the main body of a manuscript. Instead, more general descriptions are provided in the main manuscript whilst more detailed and granular descriptions are relegated to the Supplemental Materials (which can be referred to as needed). This is particularly the case with analyses that are now relatively commonplace throughout the geosciences. For example, the reader generally does not need to know what kind of scanning-electron microscope you used or its exact beam conditions when you collected an image. Similarly, if the target audience of your paper on sediment provenance analysis is primarily sedimentologists, then they likely do not need to know the exact reference material you used in your detrital zircon dating or how many dozens of times you rechecked its composition during mass spectrometry. Those details can and should be contained in the Supplemental Materials so that particularly curious readers and/or experts can check your work or attempt to reproduce your results. In the examples above, however, it is recommended to omit them from the main manuscript. There are at least two exceptions to this rule: 1) you utilized a commonplace analysis but it required atypical parameters to function appropriately or which may affect the interpretation of results (i.e. using a very low accelerating voltage to analyze some relatively fragile phase on an electron microprobe); and 2) your manuscript is specifically presenting a new or modified method for community consideration in which case more details are required for validation.

One common mistake I see in many manuscripts is

an overly complicated and descriptive presentation of methods. This is understandable. The nature of life in the sciences is to live very intensively within a microcosm of small details. Producing good and reliable data, particularly from the natural world, is very difficult and this skill is most typically found in people who concern themselves greatly with the details. It is also true that human nature is to play to our strengths in public exhibitions, and many who are earlier in their careers are generally a great deal more comfortable describing how and why they did something rather than what the results mean or why those results are important. In short, although it is admirable when a person spends so many hours categorizing and cataloging their samples before preparing thin sections, it is unlikely that their readers will be overly impressed with (or even interested in) those details. Keep your methods simple and restricted in detail to only those required for the reader to understand the results.

3.2 Results

As mentioned in the previous section, a well constructed manuscript will exhibit a synergy between the presentation of its methods and results. The results, however, are the driver of that relationship. The reason for this is that relaying observations to other human beings fundamentally requires context. As mentioned previously, your reader cannot possibly grasp the validity or importance of your in-situ U-Pb ages if they know nothing about the microstructure within which those ages were obtained, etc. As was the case with the methods, there are two distinct levels of hierarchy in presenting results. The first is that we need to build forward with context, providing information in a logical order such that the building blocks are there for each new data type (e.g. describe the microstructural observations of your sample before the in situ U-Pb analyses). The second pertains to cases where distinct results or datasets are coequal in terms of presenting a logical succession of data. Taking again the example of needing to contextualize in-situ U-Pb results, microscale variations in mineralogy observed via optical microscopy may be of equal importance to various deformation fabrics determined via electron backscatter diffraction (EBSD). In that case, it is generally best to present more qualitative datasets or observations first (e.g. the optical mineralogy), before moving on to the more quantitative results (e.g. the EBSD data), much as described in the previous section.

Achieving a more coherent granularity in the presentation of research results works much the same way as was described in the previous section on research methods. More specifically, relatively complex datasets may be best approached by dividing them into subsections for easier understanding, but how we approach the ordering of information is identical (i.e. present results that

are needed to contextualize other observations first, and then work from more qualitative and/or simple to quantitative and/or complex). One way in which taking a more granular approach to research results differs from the presentation of methods, however, is that the results section may require you to present effectively the same data types and/or observations in different subsections. An example of this might be a case where you need to compare observations or measurements between distinct field sites and/or associated samples, or perhaps between different experimental treatments. In that case, the best way to achieve greater clarity when presenting your results is to use a consistent order and style. If you begin with foliation/lineation data when presenting the results from site A, for example, then begin by presenting the results from site B in the same way. This is one area of your manuscript where monotony and/or redundancy are acceptable, and in many cases even desired. See for example, the first sentences from two successive paragraphs in the results section *Williams et al. (2017a)*, describing outcrop observations at two key locations along a fault (figure calls removed for brevity):

In the North fracture zone, the exposed, well-cemented hanging-wall damage zone extends 2–8 m from the fault core. It is exposed for ~200 m along strike, immediately south of the fault's intersection with the Rio Salado. A well-developed foliation similar to that exposed in other sections of the fault is visible in the hanging-wall damage zone...

The South fracture zone also has a well-cemented hanging-wall damage zone that extends 1–10 m from the fault core where exposed. Its exposure is continuous for 400 m along strike, beginning ~500 m north of where the fault intersects Arroyo Chante. The fault curves abruptly to the east near a soft-linked relay between two segments at the southernmost exposure of the South fracture zone. A well-developed foliation similar to that displayed elsewhere in the fault is visible in the hanging-wall damage zone...

These consistent, repeating presentation structures allow the reader to more readily compare and contrast important differences or similarities in your data, and more importantly the context in which they occur. The same strategy should be employed for reporting and comparing the characteristics of specific samples or experimental treatments.

Finally, one of the cardinal rules of manuscript construction is that the results section is to clearly and objectively state only the results. Interpretations based on those results should be strictly relegated to the discussion section. There are many reasons for this separation. The first is simply that it is good scientific practice. The aggregate balance of our results should inform our interpretations, and we should do our best to avoid biasing how we

view or present our data based on our preferred interpretation for what those data say about the physical world. Avoiding this tendency also allows the reader a chance to form their own opinions prior to hearing those of the author. The second reason pertains to the difference between results and interpretations as “products” of your work. The results describe fundamental data or observations. In principle, these should accurately characterize the nature of the physical world. The best data and/or observations are therefore such that they would not vary greatly between different observers when employing identical methodologies. Interpretations, on the other hand, are simply our opinions on what exactly our data imply for how the physical world operates. These can, and very often do, vary greatly between different observers when presented with an identical dataset. The reason we attempt to publish our results is because we believe that our data and interpretations, the products of our work, have value. We need to respect, however, that each may be more or less valuable to any particular reader. Even if we fail to convince the readers that our interpretations are correct, our data may still make a valuable contribution to future work. Therefore, we should present our data in a manner that is not conflated with our interpretations.

There is, however, an isolated case when an interpretation can be presented along with the results, specifically when it is required to avoid the presentation of spurious data to the reader. Consider the following example from the results section of *Williams et al. (2019)*, where the blue highlighted sentence introduces what is effectively an interpretation (citations and figure calls omitted):

Of 12 distinct vein samples targeted for analysis, seven yielded age pairs that were sufficiently well resolved to allow calcite growth rate estimation. Of the five age pairs that did not yield usable constraints, three exhibit large analytical uncertainties that prevent accurate estimation of the time difference between t_1 and t_2 . The remaining two pairs were excluded based on anomalously large differences between t_1 and t_2 , which reinspection of these veins revealed to be the result of a punctuated (noncontinuous) growth history recorded by subtle, euhedral-growth boundaries overlooked during initial sample screening.

4 Discussions

Although all of the sections of a manuscript are important, the discussion is really the core of the science. The purpose of the discussion is to contextualize your results, synthesize them with the broader literature, and make inferences about geologic processes or history. In this way, it is often

the case that the point of a discussion section is to build up to the idea that our data say something useful about the world. Perhaps our data and analyses provide positive evidence in support of some aspect of geologic history proceeding, or some physical process functioning, in a particular way. Alternatively, the main implication of our data may be somewhat negative, casting doubt on or disproving a previously held conceptual model, hypothesis, or history. It is worth noting in that case that it is often useful to provide some alternatives that may be more consistent with the data, but this is not strictly required. Merely pointing out the limitations or errors of previous ideas is useful in bringing them to light, after which the community is incentivized to investigate further and chart a path forward. That being said, our audience of fellow scientists are inherently skeptical, and convincing them that we have learned something new with any real degree of certainty is difficult. It is for this reason that our discussion section is, more than anything else, an argument. We need to make it as strong as possible.

So to construct a good discussion section, we may first need to ask ourselves what exactly are the characteristics of a good argument? There are many answers to this question and I highly recommend exploring some of them elsewhere. For simplicity sake here, I will focus on just three aspects of a good argument that I think are of particular importance to writing a compelling discussion:

- A good argument is **self critical**. It honestly portrays limitations in data quality and quantity in addition to providing a good faith consideration of reasonable alternative interpretations.
- A good argument has **broad reach**. It effectively leverages the available literature to provide supporting evidence and/or counter examples as necessary. It is faithful to the literature as it stands today, not as we would wish it to be.
- A good argument is **measured**. It uses careful and appropriate language to link observations and interpretations in addition to making sure that the significance of its conclusions are adequately balanced by the quality and quantity of the evidence.

It is necessary to first discuss these ideas and how they relate to writing an effective discussion. Afterwards, I will provide some more practical, concrete advice for how to actually structure a discussion section and techniques for producing a defensible draft a bit more easily

4.1 Constructing a Self Critical Discussion

There are two primary aspects of self criticism that I will address here. The first involves the need to provide an honest assessment of data quality/validity/limitations/sources of uncertainty.

Most data sets have at least some of these issues, and leaving your reader to infer them on their own is arguably dishonest, but certainly erodes trust in your scholarship when they invariably discover the limitations you chose not (or forgot) to explicitly address. Depending on the nature of a dataset, however, at least some of this information may reasonably be provided in the results section (as discussed above). That being said, presentation of information related to data quality and/or quantity in the results is generally provided without any additional commentary as to how these limit our ability to make good inferences about Earth processes, which is more appropriately addressed in the discussion.

How we approach the consideration of data quality, sources of uncertainty, etc in our discussion section depends in part on the complexity of our dataset and associated interpretations and conclusions. Regardless, you need to anticipate your readers' concerns but also be honest with yourself regarding the limitations of your work. Are your data biased in some way (common where field measurements are collected across lithologies that are variable in their tendency to outcrop)? Are there any potential artifacts in your geochronology that could produce errors that are not accounted for by the analytical uncertainty (common when your sample contains multiple phases of differing age that all include the parent and/or daughter isotope targeted by the analysis)? This is the place to explicitly address those issues. The goal is to present your reader with some questions about the validity of your work (some of which they may not have even thought of yet) and then alleviate those concerns by stating how you will act to mitigate them or why they are of little importance. Perhaps your outcrop measurements are subject to some bias, but you still managed to collect a large/robust data set for the underrepresented portion of the system such that statistical analyses can speak to the likelihood of error. Perhaps your geochronology sample does contain contaminant phases, but you have independent data to suggest that they are volumetrically minor and/or have very low concentrations of the parent/daughter isotopes being targeted relative to the phase of interest.

For relatively straight forward data sets, it can often be a good idea to simply address the issue in total at the beginning of the discussion. Explicit subsections on "Sources of Uncertainty" or "Data Limitations" or similar are direct and helpful. Consider for example this short excerpt from the first discussion subsection (called "Sources of Uncertainty") in a paper by *Williams and Kirkpatrick (2022)*, where a potential source of error is raised (blue highlighted text) and dismissed as unlikely to be of major importance using plain language and basic reasoning (green highlighted text):

Our analysis is purely statistical, and relies

on previously published descriptions and interpretations of inclusion-trail spacings. As such, we cannot attest to the reliability or accuracy of those previous estimates, which may vary between each of the data sources. We emphasize, however, that measurement of inclusion-trail spacing in veins is a straightforward and well-established approach, and only requires that a reasonably accurate count of individual inclusion trails is obtainable over a known linear distance. Thus, we assume that the primary data input in our analysis is accurate.

The approach of using one overarching subsection to describe sources of uncertainty becomes quite confusing, however, when applied to more variable and/or complex datasets. It forecasts too much about the upcoming discussion that the reader is unlikely to remember as they proceed forward. In that case, discussion of data quality and/or sources of uncertainty need to be addressed “on the fly” as data are linked with interpretations and your overall argument builds in each discussion subsection. The general idea given above still applies though, and it is good to address issues of data limitations and/or sources of uncertainty first before moving on to making interpretations based on those data. Applying this to the “on the fly” approach means that you address these issues at the beginning (or at some otherwise early point) of each discussion subsection where data will be invoked to make an interpretation.

The second key aspect of constructing a self-critical argument, and thus a good discussion, is to conduct a good faith consideration and presentation of reasonable alternative interpretations, even if just to eventually dismiss them as unlikely. There is virtually always another way one could interpret a particular dataset. Some of these will be subtle variations imposed on a more general interpretation. Others may differ dramatically. It is important to note that we are referring to these as “alternative” interpretations for a reason, namely that you likely have an interpretation that you prefer and wish to focus on. It is human nature to do so, and it does not need to be a problem from a scientific perspective. This is exactly why we consider alternatives, as succinctly and beautifully described by Stephen Jay Gould: *“The best form of objectivity lies in explicitly identifying preferences so that their influence can be recognized and countermanded”*. That being said, we can only reasonably defend a preferred interpretation in our papers if it is uniquely consistent with the evidence, but to do that we have to first introduce competing interpretations and explain to the reader why those are not consistent (or at least are less consistent) with the available evidence. Consider this discussion provided by *Roberts and Tikoff (2021)*, where the authors examine

the possibility that fabrics in a magmatic dome reflect overprinting during regional shortening as opposed to primary dome formation:

The presence of a strong preferred orientation of foliation within the Mt Edgar dome raises the question: Do foliations and lineations represent a northeast-southwest shortening overprint that is unrelated to dome formation? Previous structural studies have noted the presence of thrust faults in the surrounding greenstone belts that record northeast-southwest shortening. Several aspects of our data indicate that the field foliations and AMS ellipsoids predominantly reflect the structures associated with Mt Edgar dome formation. First, the northwest striking foliation is best developed in the Cleland Supersuite, and is more scattered or non-existent in other parts of the dome. This variation is not consistent with a regional, penetrative foliation overprint. Second, the Cleland unit is itself elongate parallel to the strike of foliation, and lineations plunge radially inward toward a central location. These two observations suggest that the preferred alignment within the Cleland unit resulted primarily from internal flow and associated syn-doming strain rather than a planar, externally imposed deformation. Third, foliations and lineations are parallel to Tambora boundaries in the southwest and western regions of the dome. The fact that foliations and lineations deviate from the dome-wide orientation at these contacts is not consistent with a late-stage overprint.

The blue highlighted text introduces a simple but reasonable alternative interpretation. The green highlighted text considers that alternative interpretation and its implications in light of the available data. The yellow highlighted text then dismisses the alternative interpretation as unlikely.

Finally, when addressing alternative interpretations in your discussion, it is important to always engage with the most robust possible form of that alternative interpretation. Few things can erode confidence in an argument faster than a dismissive and derogatory description of alternative interpretations or those that advocate for them. Do not include leading language intended to cast doubt on the alternatives from the outset, or beg the question regarding their validity. Do not attribute a particular interpretation to be the product of its advocates’ backgrounds, personalities, or biases. These things amount to something akin to ad hominem attacks, and they have no place

in a reasoned argument. It is always better to be intellectually charitable to alternative interpretations and those who advocate for them. It is also important, however, to not omit critical aspects of an alternative interpretation, even if you do not like or agree with them. Rather, your presentation of alternative interpretations should be as though you are (at least briefly) acting as a proponent for that position. Present it faithfully and fully as its advocates might, then address it on its merits and in light of the relevant data/observations. This is an argumentative tool sometimes referred to as a “steel person” argument (i.e. the antithesis of a “straw person” argument).

4.2 Constructing a Broad Reaching Discussion

In an ideal world, your data could stand alone as evidence in support of your interpretations/conclusions. We should always do our best to collect data in a way that maximizes this potential. It is rarely the case, however, and even if it was the case, drawing in at least some additional supporting evidence from other sources would still make for a stronger argument. This is primarily what is meant by an argument having “broad reach”. It also illustrates one way in which reading actually makes us better scientific writers. We need to maintain a mental catalog of work that has been done in our field of research (and sometimes beyond) so that we can place our results within the larger scheme of geologic understanding and context. This will most frequently be done by drawing on the literature to provide support for some of our more specific interpretations, as is shown in this excerpt from the discussion of Williams *et al.* (2021):

*Illite/smectite, however, is not detected in protolith granodiorites, suggesting an authigenic origin during fluid-rock interaction. This interpretation is supported by apparent increases in I/S concentration with decreasing grain size in gouge samples (Figure 6; cf. Boles *et al.*, 2015; Schleicher *et al.*, 2010), and suggests growth of sub- μm clay grains during alteration of less stable phases. It is also consistent with increases in Mg and Fe content with fault-rock development (Harder, 1972), which cannot be explained by mechanical mixing with mafic enclaves alone (Figure 7e).*

This simple paragraph offers several interpretations and justifies them with two different sources of evidence: the primary data presented in the paper (as given in the figure calls) and other peer-reviewed papers that report similar observations or conclusions (as given in the citations).

It is vital, however, that we not “cherry pick” supporting data from the literature. This point relates as much to general scientific practice and integrity as it does to writing. A useful maxim is to remember that we need to respect the peer-reviewed literature

as it stands today, and not attempt to depict it in a way that we wish it to be. We all have a responsibility to engage fully with other sources that are pertinent to our interpretations and conclusions. This responsibility is not alleviated by the fact that we may disagree with some of those sources or their advocates. If that is the case, then acknowledge the potential relevance of the source in question and state in plain language (with supporting evidence) why you disagree with its conclusions. In other cases, the nature of our data and the available literature may be such that there simply is no single, unambiguous interpretation, as is shown in this more complex example from the discussion in Williams *et al.* (2017b) (modified slightly to be one continuous paragraph):

The fact that such euhedral terminations and equant crystals are the exception, and continuous fluid and solid inclusion trails separating generations of elongate calcite crystals are the norm, we infer that these veins most likely formed through a crack-seal process [cf. Ramsay, 1980; Fisher and Brantley, 1992, 2014; Oliver and Bons, 2001; Hilgers and Urai, 2005; Holland and Urai, 2010; Bons *et al.*, 2012]. However, Hilgers and Urai (2005) cautioned that not all fracture-parallel inclusion trails record crack-seal, particularly where trails are discontinuous. Even continuous fluid inclusion trails have been shown to form over length scales on the order of a centimeter during analog experiments in the absence of crack-seal processes (e.g., Means and Li, 2001), and could potentially result from variations in fluid chemistry during disequilibrium growth in nature as well (Wiltschko and Morse, 2001; Hilgers and Urai, 2005).

In this case, the blue highlighted text makes a specific interpretation on the basis of the paper’s primary observations (this is the authors’ “preferred” interpretation, as discussed previously), while the green highlighted text reveals some additional support for that interpretation from the peer-reviewed literature. The yellow highlighted text, however, brings in an alternative interpretation with its own sources of support in the peer-reviewed literature. It is likely clear at this point that the notions of self critical and broad reaching arguments/discussions are not entirely separable. This intertwining of the notions of self criticism and broad reach is also true of the notion of constructing a measured argument, as will be discussed in the next section. It also shows that despite some of the more simple examples above, where we introduce an alternative interpretation and then (hopefully) dismiss it on the basis of available evidence, this

is not always possible. Sometimes we just have to proceed forward with some aspect of the argument being left open to multiple interpretations. Ambiguity is sometimes just a fact of life (and science).

4.3 Constructing a Measured Discussion

Writing a measured discussion revolves around two main ideas: accurate use of language as it pertains to certainty and balancing the weight of our conclusions with the available evidence. Of these, the former is considerably more straightforward to achieve. In short, it is of critical importance to make judicious use of language as it pertains to the certainty of our interpretations and the degree to which they flow from our observations. *A good discussion never states an interpretation as a fact, because all of our interpretations are subject to change in the future when/if new data become available.* As such, we need to exercise a bit of humility and think very carefully about the degree to which we feel confident in our interpretations, and convey that confidence clearly to the reader. One potential way of doing this is to consider the (un)certainly categorizations of “Permissive”, “Suggestive”, and “Compelling” and modifying some of the definitions provided by *Bateman et al.* (2022) accordingly:

- Permissive indicates a particular property or inference cannot be ruled out, but it is not the only available solution.
- Suggestive indicates that some positive evidence is available for a particular property or inference, but that the evidence also allows for the possibility of other properties or inferences.
- Compelling indicates that a significant amount of positive evidence is in support of a particular property or inference, and that this evidence would be difficult to reconcile with other properties or inferences.

In practice, writing with respect to these ideas often revolves around careful consideration of the language we use to link our observations to our interpretations. Consider the examples provided in Table 1 where the language used to link observations and interpretations with varying degrees of certainty is given in bold italics.

The second aspect of constructing a measured argument is to make sure that the significance of our interpretations/conclusions is balanced by the quality and quantity of the available evidence. This process is arguably more tenuous to describe than appropriate language use as it pertains to interpretation certainty, but its spirit is encapsulated by Carl Sagan, who said that *extraordinary claims require extraordinary evidence.* If some aspect of your argument is broadly consistent with what most already believe on the basis of previous work or established theory, then relatively little evidence is generally required to convince your readers it is true. The more it

deviates from previous work and established theory or knowledge, however, the more you will need to rely on a large amount of data of exceptional quality/reliability. The reasons for this are both sociological and scientific. It is generally difficult to change a person’s opinion on an issue once they have come down on a particular side of it, regardless of its nature. Moreover, when making an argument to our fellow scientists, many (but sadly not all) of their opinions as they relate to the natural world will be based on previous research/data/observations.

So how do we know if our argument is well measured? One way is by simply considering the degree to which it deviates from established knowledge or thought in the field of research. Suppose you are studying a small magmatic system that is widely understood through multiple, previous, and independent geochronological research efforts as having been intruded in the Miocene. Suppose then that a single sample in a suite of those you date with U-Pb geochronology yields an estimated age of 1.2 billion years. If you interpret that the old age reflects some sort of infidelity in U or Pb content in the sample through time, leading to an erroneously old age, you are unlikely to have any major problems during peer review (at least with that particular point). This is something most would consider to be a parsimonious argument - it is arguably the most simple interpretation available given all of the available evidence. If you decide, however, to interpret that single, anonymously old age as being evidence that the magmatic system in question instead reflects punctuated formation, beginning in the Mesoproterozoic and culminating over a billion years later in the Miocene, and premise your entire manuscript around this point, you will likely have a substantial problem in review. This is not to say that you should not propose incredible new ideas or interpretations through your work. Rather, that the veracity with which you propose those ideas or interpretations should be consistent with the quality and quantity of evidence that is available. For a real world example of how achieving this balance can change the course of earth science, see *Atwater* (1970).

We can also consider the “scope” of the interpretations relative to the data used to support them. In general, it is frowned upon to use isolated, local or regional scale observations to discount regional or global scale interpretations, conceptual models, or phenomena. The key word here is “isolated”. Finding one area of active faulting that deviates from Gutenberg-Richter behavior in its magnitude-frequency scaling does not mean that the general premise of Gutenberg-Richter is incorrect. Finding one outcrop for which paleomagnetic analyses record a movement vector that is inconsistent with the prevailing tectonic reconstruction does not necessarily discount that reconstruction. It is entirely possible that what we are looking at in these cases is bad data, or

Table 1 – Examples of judicious use of language as related to the certainty of interpretations.

Permissive	<ul style="list-style-type: none"> • “It is possible that the metamorphism is retrograde in origin.” • “An interpretation that the observed deformation is the result of the same collisional event that led to regional folding is feasible” • “We consider it plausible that the change in depositional style records the gradual progradation of a delta lobe.”
Suggestive	<ul style="list-style-type: none"> • “Our EBSD data suggest that shear was accommodated predominately by diffusion creep.” • “It is probable that the remarkably low coefficient of friction observed in some of our samples is associated with the presence of smectite clays.” • “We conclude that a causal relationship between ground-rupturing earthquakes and increased channel runoff is likely.”
Compelling	<ul style="list-style-type: none"> • “Our geochronological results show that the fault was active until at least the mid-Cretaceous.” • “Recrystallization is indicated by the presence of abundant subgrains decorating quartz porphyroclasts.” • “A maximum clockwise rotation of 23° since the late-Miocene is demonstrated by our paleomagnetic and geochronological analyses.”

even something of an aberration. A special case of a counter example within an interpretation or conceptual model that is, at least generally, correct. None of this is to say you cannot use your dataset to cast doubt on prevailing interpretations or conceptual models. You certainly can, and often, you should. But there is a difference between casting doubt on those ideas, and inviting the community to reconsider and reinvestigate some positions, and unilaterally declaring them to be false based on what are ultimately isolated observations. Similarly, you can also use your dataset to propose grand new ideas about Earth history and processes, but it is rarely reasonable to unilaterally declare them to be correct based on your individual observations alone. In this way, the idea that we should balance the significance of our interpretations with an appropriate quantity and quality of evidence ties back to the language we use to present those interpretations. A little appreciated fact about scientific writing is that *speculation is generally considered permissible insofar as you state explicitly to the reader that you are speculating.*

Many will recognize at this point that the examples above represent somewhat extreme examples of arguments that are definitely not measured. In that way, they say something about what one definitely should not do, but relatively little about what one should do. More colloquially, many of the examples above could be referred to as “overreaching”. Most of us are quite capable of recognizing some of the more egregious examples with little formal training, and this gives us at least some guidance in how we avoid overreaching when constructing our own arguments. Recognizing more subtle cases, however, takes time and experience. Even then, disagreements about

what constitutes sufficient evidence for a particular claim can and will arise between reasonable people and capable scientists. It is also the case that spending months or years collecting and interpreting data and writing up the results of those efforts often puts us a bit too “close” to the argument, and limits our ability to recognize overreaches in our own science and writing. This is actually part of the reason why peer review exists. In this way, it is useful to remember that peer review can also happen on your terms by reaching out to colleagues and/or members of your community to see if they are willing to provide an informal review of your manuscript. Receiving feedback after presentations at conferences can be similarly useful, and in general if your audience there is unconvinced by your overall arguments and conclusions there is little reason to believe things will be different when subjecting a written version to peer review.

4.4 At Long Last, Practical Advice for Writing a Discussion

The practicalities of writing a good discussion section are built on the recognition that our readers are humans. Their attention spans, now more than ever, have an upper limit. Your argument cannot gain traction in your readers’ minds if you violate that limit. To this end, your discussion section, and in fact your overall paper, usually needs to be limited to conveying one or two key conclusions or “take home” messages that you want your readers to internalize. In shorter format venues such as *Nature* or *Geology* it is more likely that you can only really convey a single, central conclusion. Think of the entire purpose of your discussion (and in fact, your entire manuscript) as building an argument toward these conclusions.

In this way, the first step in writing the discussion is to decide what those conclusions actually are. Take a minute and write down each of your primary conclusions in as simple a manner as possible. If you require more than one sentence for each, it is likely that they are too complicated. Try to simplify them.

One simple, but often overlooked tool that will allow you to quickly establish a framework for the discussion with your reader is to forecast it right at the start. This approach is particularly effective when the discussion is going to be long and/or complex, describing a variety of disparate data types and associated interpretations. Consider, for example, the first short paragraph of the discussion from *Titus et al.* (2011):

The geologic and geodetic observations described in this paper provide complementary information about deformation in the fault borderlands measured over different time scales. We next interpret each set of observations in the context of the other. The comparison allows us to suggest alternative approaches to future models of GPS data from central California. We also hypothesize that the progressive deformation of some geologic structures in central California is linked to the onset of fault creep along the San Andreas fault.

This approach is frequently not required in shorter, more straightforward discussions, but that does not mean it cannot be effective there as well. One way to think about it is this: if the journey is short and simple, you can probably start it immediately. If it is long and complex, it may be best to give the reader a sense of where things are heading before setting out. As an aside, I would argue that the example text from *Titus et al.* (2011) also forecasts a very simple encapsulation of the two main conclusions of the paper.

At the broadest scale, there are two distinct structures that a discussion section can adopt. The first is that of a single section of narrative argument that discusses all of the relevant uncertainties, interpretations, and conclusions of the manuscript. The second is a more explicitly constrained and hierarchical approach that divides the same information into a number of discrete subsections (or even subsubsections). Both have their place depending on the length and subject matter of the manuscript in question. The former is arguably less common today, and it is certainly more difficult to construct in an effective way, but it can be useful in short format publications focused on a single, relatively straightforward topic or dataset. The latter is generally easier to write and to read, largely because it breaks information and interpretations up into more manageable chunks, and the subsection headers themselves carry some of the weight of guiding the reader through the logic of various interpretations, as we will see below.

In either case, the elevator pitch approach used for writing introductions is equally effective at organizing a discussion. It differs in just two primary ways: 1) it likely needs to be a bit longer and/or more complicated; and 2) it is not possible to prescribe in advance the number of paragraphs/sections that will be required, or the order in which specific ideas should be presented. Toward this latter point, however, it is useful to remember that we can think of a discussion section as being to our interpretations what a results section is to our data - there is generally a logical need to present our interpretations in a particular order. For example, the reader needs to understand how your paleomagnetic and geochronologic data are integrated to interpret a rotation rate and axis location before they can reasonably evaluate your local tectonic model. They also need to understand that tectonic model before they can grasp the significance of your work for more regional- or global-scale tectonic processes. Thus, we need to build the discussion forward, step-by-step in a way that gives the reader the information that they need to follow your inferences through to the overall conclusions.

As an example of the ideas above, we can consider how the first-order points of our elevator pitch may be reflected in a discussion's subsection headers, and how we can use these judiciously to maximize clarity and impact. In the example just above, we might have a subsection header called "Inferences on the rate and axis of rotation", or similar, where we take the reader through the relevant logic, any associated assumptions, and (of course) relevant uncertainties. Consider also, however, another example from the discussion provided by *Scharer et al.* (2011) in a paper titled "A reevaluation of the Pallett Creek earthquake chronology based on new AMS [Accelerator Mass Spectrometry] radiocarbon dates, San Andreas fault, California", where their subsection headers are as follows:

1. *Comparison of AMS and Conventional Dates*
2. *Sedimentation Rates*
3. *Comparison of Earthquake Ages*
4. *Impact of Dating and Stratigraphic Assumptions on [Apparent Earthquake Periodicity]*

These subsection headers effectively convey much of the message of the entire discussion on their own. One of the primary reasons this structure is effective is that it builds the argument forward in a way that is easily digestible, moving logically from relatively simple ideas to those that are more complex. For example, the discussion begins with a simple presentation of *issues related to the data alone* (i.e. comparison of new AMS ages and conventional radiocarbon ages). It then resolves those data back to the relevant geology and the processes being studied, specifically in thinking about how differences in dating method affect the determination of sedimentation rates

adjacent to surface-rupturing faults. Note that subsections 1 and 2 are providing commentary on what could reasonably be described as types of observations and/or measurements and their validity. They provide a coherent grounding for the interpretations to come. The third subsection, for example, moves to the province of *inference or interpretation that draws on the earlier observations and/or measurements*, in this case by using them to define a chronology of recurrent earthquakes. The difference is perhaps subtle, but important. It is true, for example, that both AMS radiocarbon dating and its application to determining sedimentation rates between organic-bearing layers rely on a variety of assumptions and thus involve some degree of inference or interpretation. The ideas that any particular sample has an age and that sedimentation occurs at a finite rate during the formation of surficial deposits, however, are not reasonably a matter of interpretation or dispute. They are facts/definitive phenomena that we seek to measure or otherwise quantify. In comparison, the very idea that a particular sedimentary horizon or disturbance in sedimentary layering necessarily records a past earthquake whose timing can be constrained by dating the stratigraphy is most certainly an interpretation (and it has indeed been the matter of significant dispute, at least locally). Another reason the structure provided by *Scharer et al. (2011)* is effective is because it drives inexorably to a single point - how the presented data/interpretations/ideas impact our assessment of earthquake periodicity as inferred from the geologic record. It is good practice to try to end on the interpretations that form the primary conclusion(s) of your manuscript, or otherwise the main ideas that you want the reader to take away from your work. Remember that a good scientific mindset of course works forward, moving from basic observations through to interpretations and conclusions. A good writing mindset, however, often works in the reverse, starting where you want your reader to be and charting a path backwards from there to where they most likely are when they pick up your paper.

5 Abstracts and Conclusions

In many ways, these sections actually accomplish the same goal, albeit with different levels of detail. The goal is to succinctly describe the results of your research and how they were achieved (i.e. what methods and data did you use to arrive at your interpretations). The primary difference between the two sections is the level of detail provided, but many of the rules described in other sections above still apply.

5.1 Abstracts

The goal of the abstract is to effectively summarize your manuscript such that the reader can get a sense of whether looking deeper would be worth

their while⁵. The primary issue from a practical writing standpoint rests on your ability to produce an abstract that is brief, direct, and compelling. Many journals today will limit your abstract to 250 words or less, and you will often need to use those words to summarize several years of detailed research. To put that number in context, this paragraph will already encompass 100 words by the end of the sentence you are reading. It can be a surprisingly complicated task, but fortunately one that graduate and postdoc life provide an opportunity to practice on a regular basis. Moreover, the vast majority of abstracts constructed by our community follow a surprisingly prescribed format, such that familiarizing yourself with that format can ease the path forward greatly.

When writing an abstract, it is important to remember that it is unique in terms of its form and function. It does not (and realistically cannot) obey many of our normal “rules” regarding the need for a narrative text. Abstracts are summations of key details whose only purpose is to rapidly give the reader a sense of the research in question. In principle, there are four questions that an abstract must answer for the reader. What is the nature of the research that was done? What are the data or observations that were collected? What do those data or observations indicate? And why are those conclusions important or useful? One surprisingly simple and effective way to construct an abstract is simply address these questions directly, immediately, and in that order. Consider for example, this abstract from *Williams et al. (2021)*:

We document the mechanical and geochemical processes of fault rock development in the shallow San Andreas fault (Mojave segment), and quantify their importance in shaping the mineralogy, grain size, fabric, and frictional characteristics of gouge. Through a combination of field and laboratory analysis of an extensive suite of shallow (less than 150 m) drill cores, we show that fault rocks evolved from a granodiorite protolith via three main processes: distributed microfracturing/pulverization; cataclastic flow and incipient fabric development; and production of authigenic illite/smectite during fluid-rock interaction. The interdependence of these mechanical and geochemical processes results in a diverse suite of fault rocks, and causes significant changes in frictional strength. Spatial variations in the effects of these mechanisms, as manifested in fault-rock mineralogy and geochemistry, indicate marked variations in their relative contribution to fault-rock evolution. These data reveal a complex San Andreas fault with multiple principal slip zones and damaged and altered rock hosting numerous interconnected

⁵Additional (and quite entertaining) resources for more effective abstract construction can be found in *Landes (1951, 1966)* and *Lowman (1988)*

secondary slip surfaces. The resulting picture of the San Andreas Fault zone suggests a substantial departure from the simple structures envisioned for near-surface seismogenic faults in numerical models is required, and may inform future efforts to forecast peak ground accelerations during southern California earthquakes.

It is also permissible, however, to provide a small amount of direction for the reader in terms of setting up the general research area and its importance before diving into what was done to advance it. Consider, for example, this abstract from Williams *et al.* (2019):

Fracture cementation is an important control on the recovery of prefailure levels of permeability and strength in faults. The timescales of this process, however, are almost entirely unknown from direct analysis of the rock record. We report U-Th dating results that quantify rates of fracture cementation in syntectonic calcite veins from the Loma Blanca fault, New Mexico, USA. Measured rates vary from ~0.05 to 0.80 mm/ka and exhibit a power function correlation with minimum fracture apertures. We argue that this correlation is the result of crystal growth in a transport-limited system, where cementation rates were proportional to rates of postseismic fluid flow in individual fractures. We further argue that such transport-limited, flux-dependent cementation necessarily leads to a heterogeneous distribution of permeability and strength recovery as fluids migrate through fault-zone fracture networks. These heterogeneities may influence rupture propagation pathways and the continual development of fault-zone architecture/complexity.

You may have noticed a similarity between these two examples, specifically in that they both make a definitive, single-sentence statement about the nature of the research and why it was done using active-voice language (i.e. “We show...”, “We present...”, “New high-resolution Ar geochronology of the Snake River plane reveals...”). They just do so at different locations. The first dispenses with any need for describing a broader area of research in favor of diving straight into the details of the paper. The second provides two brief sentences to set the stage before taking much the same course. If you look closely, the abstract in the vast majority of peer-reviewed publications in geoscience include similarly, if not identically, structured summary statements. For example, here are the first three sentences from the abstract of three different papers in the current issue of *Geology* at the time of this writing (v. 51, n. 3):

Episodes of slow and fast plate subduction, slab rollback, and backarc opening are widely

documented; e.g., in the central Mediterranean region. Pervasive damage by fluids is emerging as a possible weakening mechanism that could lead to slab segmentation and breakoff. We show that low-velocity anomalies within the Ionian slab present in along-dip seismic tomography profiles are traces of past damaging events generated by water penetration into the oceanic lithosphere when it was at the trench. Carminati and Chiarabba (2023).

*Extensional tectonics are marked by shallow magma crystallization depths, whereas compressional tectonics are associated with deeper crystallization depths. This implies that variations in crystallization depths can be used to track changes in Earth’s dominant tectonic regimes through time. We therefore developed a new “pressure of crystallization proxy based on the variation of the $^{176}\text{Lu}/^{177}\text{Hf}$ ratio in zircon. Moreira *et al.* (2023).*

*Subglacial abrasion drives erosion for many glaciers, inundating forefields and proglacial marine environments with glaciogenic sediments. Theoretical treatments of this process suggest that bedrock abrasion rates scale linearly with the energy expended through rock-on-rock friction during slip, but this assumption lacks an empirical basis for general implementation. To test this approach, we simulated abrasion by sliding debris-laden ice over rock beds under subglacial conditions in a cryo-ring shear and a direct shear device. Hansen *et al.* (2023).*

In fact, 14 of the 18 papers in this volume of *Geology* include active-voice, summary statements of this general nature. In the vast majority of those, the summary statement occurs within the first three sentences of the abstract.

5.2 Conclusions

The goal of the conclusions is to bring the sum total of your manuscript together in a way that concisely recapitulates your arguments and the evidence that are used to support them. It is worth noting, however, that the more recent approach to scientific writing (at least in the geosciences) does not really save the central conclusions of the work for the conclusions section. Rather, the central conclusions are often presented (or at least forecasted) in key sections of the discussion, and simply reiterated (along with the primary supporting observations) in a more succinct and direct way in the conclusions section. In short, we hope that we have effectively made our argument and convinced our readers of our conclusions *before* they make it to the conclusions section. This is the place to remind them what exactly was shown and perhaps provide a few specific take

away messages on why it is important. In a way, we can actually view the conclusions sections as a type of extended abstract. The actual abstract is intended to give people who haven't read your paper an overview of its contents. The point of the conclusions section is to give people who have read your paper a summary of its contents, largely as a reminder of how your work was conducted and what exactly was learned from it. As such, there is generally some assumption that they have read at least part of your manuscript, although the very best conclusions are generally able to function on their own.

Of all the sections in your manuscript, the conclusions are arguably the easiest to organize. This is because (presumably) you have already laid out the relevant information and arguments that support your central conclusions in the preceding text. The simplest path toward constructing the conclusions section is then to follow that same order. In fact, it is important to remember when writing a conclusions section that the laws of self plagiarism apply to copying text between individual papers, not within an individual paper. The conclusions are a section where you can feel free to use text from other sections in your manuscript, modifying as necessary to fit the paragraph structure and flow of their new placement. If you were able to construct a solid sentence or two in the main body of your manuscript that you think will really resonate with the reader and emphasize some key point(s), you can use them again (verbatim if need be) in your conclusions.

6 Titles

Constructing a good title for your manuscript can be a deceptively complicated task. A good title grabs a potential reader's attention from a veritable sea of other articles in a Google Scholar search. It also makes the overall topic of the manuscript abundantly clear. When possible, a good title may also forecast the fundamental conclusions or findings of the work. If the topic has some level of associated controversy or intrigue in its field, a good title will leverage those factors. The primary difficulty of constructing a good title then, when combined with the factors above, is that it should ideally be rather short. It should never be long. What exactly constitutes "long" can (and does) differ between individual writers, but 20 words is likely a reasonable maximum value. Considerably shorter than that is preferred.

One way to accomplish many of the axioms listed above is to make the title as direct as possible. Consider then, this example of a title that does not follow this advice, focused on coseismic boiling of pore-fluids as a potential mechanism for mineralization and fracture sealing following earthquakes:

Exploring the role of coseismic boiling in fault and fracture sealing

Although this would not necessarily be a *bad* title,

it is entirely neutral in tone. It also does not give the reader any clues as to whether anything overly useful will be learned about the process of interest after reading. An alternative approach that remains quite popular is to, in effect, use the research question being addressed as the title of the paper (perhaps with a subclause teasing an answer or highlighting why the question is important or how it was addressed). Consider then an alternative title for the example above:

Can coseismic boiling seal faults? A perspective from mineral kinetics

Again, this would certainly not be a bad title for such a paper, and it would be an improvement on the previous, more neutral iteration. It effectively teases a presumably open question in a field of research, thereby adding an entertaining air of mystery, whilst also forecasting how that question will be answered. Assuming of course that the question is eventually answered by the paper, however, a much more direct way to construct the title might be the following:

Coseismic boiling cannot seal faults: Implications for the seismic cycle

This is the title of the actual paper (Williams, 2019). The reason this title is more effective is because it makes a direct, strong statement of what the paper is about and what was learned by the research, in addition to teasing why the reader should care about it. Remember, however, that titles should obey the same rules presented above as they pertain to the need for a measured argument. There is an unfortunate tendency to use much stronger language in titles and abstracts than is warranted by the quantity and/or quality of the presented evidence, even if the main body of text avoids that problem. A good title should be as firm and direct as possible, but it should not overpromise with respect to a manuscript's results or importance.

7 Postlude

My hope is that this editorial has illustrated one of the most important aspects of good scientific writing - it need not be complicated. Indeed, our community has in many cases converged (perhaps unknowingly) on a variety of writing structures that make effective communication of research results almost formulaic. There remains, however, ample room for personal style and alternative approaches to flourish. In this way, good scientific writing incorporates aspects akin to both rigid engineering and unconstrained artistry. Any individual manuscript may well lean more toward one of these aspects than the other whilst not sacrificing any overall quality. This is all to say that despite some of the more concrete advice given throughout this guide, development of one's personal scientific writing style still requires time and considerable experimentation. More modern word processing suites that incorporate versioning (such

that no text is ever truly deleted) are very effective at facilitating the latter.

In conclusion, I would like to offer one final piece of advice for aspiring scientific writers: always keep the needs of the reader foremost in your mind. Who is your audience? What sort of background knowledge are they likely to have? What is it you hope to convince them of and what information do you think they will need to understand it? Without a firm sense of the answers to these questions, and a sincere desire to have the reader learn something from your manuscript, it is unlikely that the employment of any of the specific advice above will be of great use. Good writing, scientific or otherwise, should always consider the human component.

Acknowledgements

I would like to recognize Laurel Goodwin, John Farver, Charlie Onasch, Peter Mozley, Christie Rowe, and Basil Tikoff for their writing advice early in my career. In a more general sense, I would also like to thank anyone who has been gracious enough to share some of their knowledge with me in the form of co-authoring papers and proposals over the years.

Copyright notice

© Author(s) 2024. This article is distributed under the [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are credited, and any changes made are indicated.

References

- Atwater, T. (1970), Implications of plate tectonics for the cenozoic tectonic evolution of western north america, *GSA Bulletin*, 81(12), 3513–3536, doi: 10.1130/0016-7606(1970)81[3513:IOPTFT]2.0.CO;2.
- Bateman, K. M., C. G. Wilson, R. T. Williams, B. Tikoff, and T. F. Shipley (2022), Explicit instruction of scientific uncertainty in an undergraduate geoscience Field-Based course, *Science education*, 31(6), 1541–1566, doi: 10.1007/s11191-022-00345-z.
- Carminati, E., and C. Chiarabba (2023), Slab damage and the pulsating retreat of the Ionian-Apennines subduction, *Geology*, 51(3), 227–232, doi: 10.1130/G50676.1.
- Connah, G. (2010), *Writing about Archaeology*, Cambridge University Press, doi: 10.1017/CBO9780511845383.
- Gustavii, B. (2017), *How to Write and Illustrate a Scientific Paper*, Cambridge University Press.
- Hansen, D. D., J. P. Brooks, L. K. Zoet, N. T. Stevens, L. Smith, C. E. Bate, and B. J. Jahnke (2023), A power-based abrasion law for use in landscape evolution models, *Geology*, 51(3), 273–277, doi: 10.1130/G50673.1.
- Landes, K. K. (1951), A scrutiny of the abstract, *Bulletin of the American Association of Petroleum Geologists*, 35(7), 1660.
- Landes, K. K. (1966), A scrutiny of the abstract, II: GEOLOGICAL NOTES,

- AAPG Bulletin*, 50(9), 1992–1992, doi: 10.1306/5D25B6AF-16C1-11D7-8645000102C1865D.
- Lowman, P. D. (1988), The abstract rescruinated, *Geology*, 16(12), 1063–1063, doi: 10.1130/0091-7613(1988)016<1063:TAR>2.3.CO;2.
- Moreira, H., A. Buzenchi, C. J. Hawkesworth, and B. Dhuime (2023), Plumbing the depths of magma crystallization using 176Lu/177Hf in zircon as a pressure proxy, *Geology*, 51(3), 233–237, doi: 10.1130/G50659.1.
- Phillips, N. J., and R. T. Williams (2021), To D or not to d? re-evaluating particle-size distributions in natural and experimental fault rocks, *Earth and planetary science letters*, 553, 116,635, doi: 10.1016/j.epsl.2020.116635.
- Roberts, N. M., and B. Tikoff (2021), Internal structure of the paleoarchean mt edgar dome, pilbara craton, western australia, *Precambrian research*, 358, 106,163, doi: 10.1016/j.precamres.2021.106163.
- Scharer, K. M., G. P. Biasi, and R. J. Weldon, II (2011), A reevaluation of the pallett creek earthquake chronology based on new AMS radiocarbon dates, san andreas fault, california, *Journal of geophysical research*, 116(B12), doi: 10.1029/2010jb008099.
- Schimmel, J. (2012), *Writing Science: How to Write Papers That Get Cited and Proposals That Get Funded*, Oxford University Press, USA.
- Strunk, W., and E. B. White (2000), *The elements of style*, Longman, New York.
- Titus, S. J., M. Dyson, C. DeMets, B. Tikoff, F. Rolandone, and R. Bürgmann (2011), Geologic versus geodetic deformation adjacent to the san andreas fault, central california, *GSA Bulletin*, 123(5-6), 794–820, doi: 10.1130/B30150.1.
- Williams, R. T. (2019), Coseismic boiling cannot seal faults: Implications for the seismic cycle, *Geology*, 47(5), 461–464, doi: 10.1130/G45936.1.
- Williams, R. T., and J. D. Kirkpatrick (2022), Are low-frequency earthquake moments area- or slip-limited? a rock record examination, *Geophysical research letters*, 49(2), doi: 10.1029/2021gl095759.
- Williams, R. T., L. B. Goodwin, and P. S. Mozley (2017a), Diagenetic controls on the evolution of fault-zone architecture and permeability structure: Implications for episodicity of fault-zone fluid transport in extensional basins, *GSA Bulletin*, 129(3-4), 464–478, doi: 10.1130/B31443.1.
- Williams, R. T., L. B. Goodwin, W. D. Sharp, and P. S. Mozley (2017b), Reading a 400,000-year record of earthquake frequency for an intraplate fault, *Proceedings of the National Academy of Sciences of the United States of America*, 114(19), 4893–4898, doi: 10.1073/pnas.1617945114.
- Williams, R. T., P. S. Mozley, W. D. Sharp, and L. B. Goodwin (2019), U-Th dating of syntectonic calcite veins reveals the dynamic nature of fracture cementation and healing in faults, *Geophysical research letters*, 46(22), 12,900–12,908, doi: 10.1029/2019gl085403.
- Williams, R. T., C. D. Rowe, K. Okamoto, H. M. Savage, and E. Eves (2021), How fault rocks form and evolve in the shallow san andreas fault, *Geochemistry, Geophysics, Geosystems*, 22(11), doi: 10.1029/2021gc010092.