Teaching Technical Writing Using the Engineering Method

A Handbook for Groups

Norman Ramsey
Tufts University

August 2016

Preface

Many students at American universities have trouble with technical writing. To help students over this trouble, we often use inefficient methods. These methods, like other aspects of graduate education, sometimes remind me of medieval apprenticeships. For example, I have known smart, capable teachers who were reduced to "teaching a student to write" by taking that student's first paper and rewriting it from start to finish. This booklet describes better methods.

Instead of spending most of my "student time" working on writing—teaching the same material to seniors, graduate students, and postdocs—I now prefer to teach writing to a group that meets weekly. Teaching a group uses my time more effectively, and working in a group shows students that they are not alone in their difficulties. The problems they have are problems that everyone has, and they see these problems even in published papers. But we do not emphasize problems; instead we emphasize useful principles and practices that students can learn to apply to their own manuscripts.

- I emphasize principles that can be applied successfully by a beginning writer. Especially for students in science and engineering, a principle is easily applicable when there is a simple, experimental way to decide if the written words obey the principle. (For example, I do not try to teach "omit needless words," because I know of no simple way to decide if a word is needless.) In this approach, I have been greatly influenced by Joseph Williams (1995).

- I emphasize practices that have been shown, again by experiment, to lead to productive writing. For example, I explain the difference between "binge writing" and "brief, daily sessions." In this approach, I have been greatly influenced by Robert Boice (2000).

What both approaches have in common is that even a beginning student can apply a simple test to see whether he or she is applying a given principle or following a given practice. This focus on testable ideas seems to work especially well for engineering students.

This booklet explains why you and your students might want to have a writing group. More important, it explains what we do in enough detail that I hope you can replicate the experience.

Norman Ramsey
Medford, Mass.

Contents

Present to inform, not to impress; if you inform, you will impress.
— Fred Brooks

1 Why bother?

Faculty have more than enough to do. So why spend time teaching writing? Because your students have to write, and nobody else will teach them.

Is there a problem? A solution? Writing can be a trouble spot for both students and faculty.

- Many students dread writing; many are unhappy with the results but don't know what to do. Others may not be aware of a problem, but they may find it difficult to get their work published or write their dissertations.
- Faculty's efforts to help students write better can consume lots of time and energy.

It is possible both to make the process of writing more pleasant and to improve the product. And if you want your students to write better but you don't want a new full-time job, there is good news:

- Teaching writing to a group, not to individuals, is especially effective.
- It is relatively easy to help students make significant improvements.
- You needn't be an expert in teaching writing.

How can you appeal to scientists or engineers, who may not be interested in writing? My writing groups are geared to *technical* writers.

1. I emphasize an empirical approach using the *engineering method*. My students and I analyze writing using empirical principles (what Koen (1985) calls *heuristics*), which can be inculcated by exercises. Students can use the principles to change their writing and then evaluate the results right away: an author can see whether readers grasp the important ideas and whether they enjoy a paper. Readers' reactions help students decide for themselves whether applying the principles (and participating in the group) improves their writing.
2. I leverage students' knowledge about how research is conducted, students' training in careful reasoning, and students' appreciation of scientific or mathematical elegance. I emphasize parallels between research and writing: first thoughts are often off the mark; there are many dead ends; revision is essential to progress; elegance derives from—and often conceals—hard work.
3. In group, we continually return to the connection between writing and thinking. Although we start by looking at words in a row, we inevitably find ourselves discussing deep technical material. Better writing becomes better research.[1]

[1] Instructors benefit too. My contribution to the register-allocation paper I wrote with Mike Smith and Glenn Holloway emerged directly from a group's analysis of an earlier paper.

Doesn't study of writing take time away from research? Because students spend so much time and energy in the lab, they can come to view writing as a separate, secondary activity—a simple translation of their research into prose, which happens once the research is done. This misconception is terribly constraining. I try to teach students that

1. The purpose of technical writing is not to describe events that have happened in the lab but to cause events to happen in readers' minds.

2. Writing is thinking, and the best scientists view writing not as something separate from research but as a critical stage in that research.

3. Writing is a tool for discovery—to learn what we really know about a problem, we can write about it.

New views of writing help students use their writing to focus, refine, enhance—and thus ultimately speed up—their research.

Why a group? A writing group doesn't just make it possible to teach many students at once—it actually provides a better environment for students. Students do not focus relentlessly on their own writing. They learn that others share their difficulties. They see lots of writing at their own level, but they also get to scrutinize professional writing. In this way, they can evaluate the utility of the principles they are learning and decide for themselves what works.

A group also helps because it's easier to learn new principles by practicing on other people's writing. Such practice gives students a better chance of applying the same principles to their own writing.

A group helps by *reading* students' work and responding to it. For an author, it is invaluable to see where readers do not understand; where they miss the point; and where they feel distracted, bored, or confused. One's advisor's reaction might be idiosyncratic, but when a group of one's peers reacts the same way, it is easier to accept that the flaw might lie in the text.

The group has another, less tangible benefit: it can help students feel part of a community of technical writers; it can help make technical writing a subject for discussion inside your department; in short, it can help make writing less isolating and lonely.

She, who had, by the age of three hundred,
mastered all the known techniques developed
in tens of thousands of years of painting,
still believed her technique was weak,
and not only always strove to improve it,
but believed, the results of her own hand
to the contrary, that nothing except barest
technique was of any importance whatever.

— Steven Brust, *The Phoenix Guards*

2 What I teach and with what ends in mind

Just as the purpose of writing is to cause events to happen in readers' minds, the purpose of teaching is to cause events to happen in students' minds. The events I wish to cause can be grouped under three particular goals:

1. My first goal is to change each student's *technique*—I want students to be able to apply the editing principles described below. Most students, if they apply just a few principles consistently, improve their writing dramatically.

2. My second goal is to change each student's *belief*—I want students to believe that paying attention to writing is good for them.

3. My third goal is to change each student's *attitude*—I hope that at least some students will come to enjoy writing.[2]

My long-term goal is that by the time he or she is ready to start a dissertation, each student will be a confident, clear, fluent writer.

Despite these lofty goals, if you observed my classroom, you would see that I focus almost exclusively on technique—and low-level technique at that. I do this because mastering technique will improve students' writing, seeing the improvement will show them the value of technique, and with luck, success will teach them to enjoy it. And eventually, perhaps, good technique will lead to good judgment. In the rest of this section, I describe the technique I have found useful and the outcomes you can expect when teaching it.

2.1 Summary of technique

The engineering method applied to writing

In American universities, writing instruction seems to be grounded in the humanities. Such instruction often leaves science students bewildered, or even thinking that writing instruction is a waste of time. I dare not pit my scholarly authority against that of other writing instructors, so I don't try; instead, I tell students to use the engineering method: to develop a portfolio of useful technique, try things out and observe the results.

I draw an analogy between a piece of writing and a controlled laboratory experiment: there is an independent variable (a text) and a dependent variable

[2] I've been able to write relatively easily since I entered graduate school, but for many years I hated writing—though I enjoyed *having written*. I didn't start to enjoy *writing* until maybe five years after my PhD. I'd like my own students to start sooner.

(a reader's reaction). After suggesting a particular kind of change in a text, I ask students to evaluate whether the corresponding change in the mind of the reader is an improvement. If so, this is experimental evidence in favor of the change—a persuasive argument for changing one's writing.

Writing instruction by hypothesis testing is a perfect fit for students in engineering and the sciences, because these students want claims that can be tested objectively. After all, the whole culture of science and engineering demands objective evaluation of claims.

Principles and practices

The technique I teach is summarized by the editing principles and working practices in Table 1 on page 15. An *editing principle* describes some desirable property of a text; I try hard to phrase each principle so that a student can easily decide whether a particular text respects it. Each student can then evaluate the hypothesis that a text which respects the principle produces a better effect in the reader's mind.

The practices focus not on properties of texts but on the means by which texts are produced. A *working practice* recommends a behavior that, when followed by the writer as he or she works, can lead to more fluent and successful production. Unfortunately, a student cannot easily evaluate which practices make a difference, but we do read about and discuss some of Robert Boice's (2000) controlled experiments, which I have found convincing.

2.2 Expected outcomes

Groups vary, but there are certain things you can expect to happen:

- Almost all students will improve. Some of the exercises in Section 4.2 are easy enough that students will see results immediately.
- Many students won't improve right away, or even for quite a while. There may be significant delays between the times a student understands a principle, recognizes related flaws in others' work, sees how to correct those flaws by applying the principle, and can apply the principle in his or her own work.
- In a 75-minute meeting, you will be able to discuss a surprisingly small amount of text. Don't be disappointed; a narrow focus (even just a few sentences!) often yields the deepest insights.
- Even when the focus is very narrow, discussions will often be great fun.
- Your first meetings will probably focus on mechanics. But after a month or two, you can expect discussions of mechanics to lead to discussions of ideas. As a group gains experience, it moves more often (and more quickly) towards ideas.
- Most of the time, you will find that readers prefer texts which respect the principles in Section 4.1. But you will probably encounter one notable exception: there are texts that readers like, but that completely disregard Williams's advice about subjects and verbs (Principle 3). Such texts tend to be technical *description*, in which it is hard to identify any real actions taking place.

- Your students will be able to learn a lot about writing in a few hours per week; students should be able to leave each meeting feeling they understand something significant about their own and others' writing.

- In one semester, you will probably find it difficult to teach any particular student more than just 2 or 3 principles or practices—and even that much will be hard work. But if each student really internalizes 2 or 3 useful principles and consistently applies them to his or her own work, you (and your colleagues) will be impressed at how much better their writing gets.

2.3 Detailed goals and learning objectives

As of July 2014, I am about to offer, for the second time, a year-long, for-credit course entitled *The Engineering Method of Technical Writing.* As part of the course design, I have developed some more specific goals and objectives. At the highest level,

- I want students to be able to apply proven techniques of writing and revision to their own texts, in real situations outside the classroom.

- I want students to become so comfortable and fluent that writing becomes a satisfying, normal part of their process of research or development.

- I want students to be able to assess their own mastery and to identify what steps they will need to take, after the course is over, to continue to improve.

In service of these goals, I've identified a number of measurable learning objectives.

TO BE FILLED IN LATER.

3 Mechanics and pragmatics

Leading the writing group may be my most improvisational teaching. I cannot tell you how to improvise successfully, but in this section I can give a few suggestions about how to organize your group and how to structure and run each meeting. The first meeting is always special, but after that, plan for two different kinds of meetings: ones focused on principles and ones focused on practices. I describe both kinds, and I also give some suggestions for a syllabus.

Before reading the rest of this section, you may want to glance at Table 1 on page 15.

3.1 Expectations: preparation and meeting

I have found it ideal to meet once a week for around an hour and a half each time. And preparation is imperative; when students begin, they may need an hour or two for exercises and another hour or two for supplemental reading. But by the end of a year's study, almost every student can prepare in under an hour.

Attending the meetings is important; unlike other kinds of knowledge, which can be acquired independently from a text or by listening to a taped lecture, writing can really only be learned by doing and by getting feedback. A student who doesn't show up learns nothing, and it's unfair for such a student, who seldom comments on others' work, to benefit from others' reading of his or her own work.

Meetings about editing principles Most meetings will focus on a principle of clear writing and on its application to editing a text. A meeting focused on a principle will probably be organized around an exercise that students will have completed before the meeting. Each exercise requires analysis, and sometimes editing, of a sample text written by a student author or by a professional. A typical exercise focuses on just a few paragraphs, or at most a section. The exercises might remind you of problem sets, except that because they often call upon students to evaluate events that happen in a reader's mind, not everyone will agree on the answers.

My main goal is that students be able to complete relevant exercises successfully. If they can do the exercises well (and if they can *remember* them in their own writing), they know enough to apply the principles to the drafting and revisions of their own texts.

Meetings about working practices Some meetings will focus on a working practice observed in successful writers and on the group's own experience with the practice. A meeting focused on a practice may be organized as a discussion of some reading. The discussion may resemble discussions in graduate seminars. The most effective discussions analyze how the ideas in a reading might apply to students' own writing practices.

There is a better way to organize a meeting about practices, but it requires preparation in advance: ask your students to keep a notebook that records data about how they tried to use a practice and how it affected their productivity. Analysis of the data will reveal not only what works, but even more important, how to make it work.

3.2 The first meeting

For the first few meetings of a new group, I like to analyze professional papers. Such papers provide an opportunity for tremendous frankness—students can complain bitterly without hurting anyone's feelings. Such papers also provide a good platform for testing our principles scientifically: Do the writers we like respect our principles? With what effects in the minds of the readers?

To engage the students in this enterprise, I invite them to bring examples of published papers they like and dislike:

> At our first meeting, I invite you to bring in a sample paper from the professional literature. If in the course of your classwork or your research, you have encountered a paper that you especially liked (or disliked), please bring it tomorrow, and during the term, we'll discuss some of them. I confess that it is often more satisfying to discuss bad papers, because it can be so easy to improve them.

Before the first meeting, I lend each student a copy of Williams,[3] and I ask the students to read the first chapter. With this background in place, here is what I do to launch a new writing group:

- I provide some introductions and explanations: what the group will be about, why we are doing this, and how we will use the principles, practices, and exercises.
- I give out copies of the student's edition of this booklet.
- I talk a bit about Principle 0 (correct English). Most of my students have no difficulty writing correct English, but many of them have had teachers who gave them an unreasonably narrow idea of what is correct. I rarely spend class time on questions of usage, but at the first meeting I come armed with a stack of authorities, and I encourage students to do some browsing and to broaden their ideas of what is acceptable.

  My favorite authorities are Fowler (1968) and Garner (2003). For questions related to punctuation and typesetting, I like *The Chicago Manual of Style* (Chicago 1993). Finally, many computer scientists seem to like Dupré (1998). Although I myself dislike the book, I recommend it anyway.
- I ask each student to talk a bit about papers he or she has brought to the meeting—at minimum, I collect the papers that are especially liked and especially disliked.

After that, my main goal is to talk with students about their impressions of Williams, Chapter 1, and to be sure that they understand the idea of using the experimental method to evaluate a principle:

1. Learn how to decide if a text complies with the principle.
2. Compare a compliant text with a non-compliant text.
3. Decide which, as a reader, you like better.

Then, if you find a principle you like, you can go on to advanced work:

4. Transform a non-compliant text into a compliant text.

[3] My students' work in writing group is above and beyond their usual courseloads. To avoid provoking resentment, I don't require them to buy anything. A few smart ones always wind up buying their own copies.

3.3 How to organize a meeting about a principle

Before the meeting

1. Choose a text and an exercise. Not every text works well for every exercise; if something goes wrong, it is usually because you are trying to do a relatively advanced exercise using a text in which the basics are not right. (You can take "advanced" to mean "later in Williams.") You therefore should make your choice in one of two ways:

   - Have an exercise in mind and choose a text to fit. A good reason to do this is that you are trying to proceed logically through some sort of syllabus.
   - Have a text in mind and choose an exercise to fit. A good reason to do this is that you are trying to give a student feedback on his or her work.

   Beginning students may feel less threatened if they discuss a paper from the literature. My students, who are experienced and who are accustomed to discussing writing, are much more motivated to study and discuss a text written by a member of the group. I therefore try to help students master *Subjects, verbs, and objects* (Principle 3) and *Information flow* (Principle 5) as soon as possible. Once they get those principles down, their own texts are usually suitable for a variety of exercises.

2. Prepare the text for distribution. For most exercises, it is useful to label each such paragraph with a letter and to number each sentence within those paragraphs (there's an example in Section 5.4). During discussion, participants can then easily refer to "sentence 2 of paragraph A," for example. If you are doing a long exercise such as Exercise H, it's sufficient just to label each paragraph with a letter.

3. Send the text and exercise to the students, with instructions that they spend about 30 minutes on it. Distribute the labeled paper at least 48 hours in advance—one cannot have a productive discussion of a paper that students have only skimmed.

During the meeting I structure the meeting around the exercise.

- I typically work my way around the room, asking each student about one sentence or one paragraph, depending on the scope of the exercise. After each student's answer, I ask for other students' opinions. I often write all suggestions on the board.

- I listen carefully to help students avoid any pitfalls that I associate with the exercise. (Section 4.2 discusses potential pitfalls.) I continually ask students to distinguish a text's objective, measurable properties from their feelings about that text, then to correlate them.

- I carefully listen to and shape the students' comments about the text. When the text is written by a student author, handling comments is especially important; an analysis appears below.

- If the group feels that revisions are called for, I will sometimes ask for suggestions. Depending on context, I may stop the process after general suggestions, or we may write new sentences.

A meeting focused on an exercise feels something like a problem session that one might conduct for a class in math or physics.

3.4 Commenting on a text

A key part of our approach is to ensure that group members' comments support empirical evaluation of a text. Here are some examples of useful comments:

- "I believe that the most important idea in the paper is the idea of using a finite automaton to model the infinite space of possible signatures."
- "At the end of paragraph A, I was happy, but but the time I got to sentence 3 of paragraph B, where it says that a machine register has a weight that is equal to the number of resources it consumes, I felt that I no longer understood what was going on."
- "I don't understand the distinction between an 'argument' and a 'parameter'."

Each of these examples provides evidence of important events happening in the reader's mind, and two of them tie those events to specific words or locations in the text.

An example of a less useful comment is "the third section is not well written." The comment is about the text, not about what is in the reader's mind, and it is not focused on any specific part of the text.

One of the leader's most important roles is to help develop vague comments into something useful. Students need a vocabulary with which to discuss elements of their papers and to analyze texts that are poorly received. They need ideas for creating alternatives. It's often helpful for the leader to ask readers to put themselves in the author's mind: What did the author think was most important? What effect was the author hoping to achieve in the reader's mind?

When the author is a member of the group, comments must be handled especially carefully. When a student's paper is being discussed, we may ask him or her not to speak. In part, we do this to be sure that any comments are made in response not to the author but to the words on the page. In part, we do it to mimic the process of professional reviewing; after all, when an author sends a paper out for review, the author can't enter into a discussion to explain what he or she really meant. And in part, we do it to give the author space to focus his or her full attention on what the readers are saying, without being distracted by thinking about how to formulate a response. (It also helps that the readers can focus on the text without being distracted by the author.) After the discussion, we always invite the author to respond. Be sure to leave time for this response—the participants can learn a lot from hearing what the author intended.

When a text is being discussed, we encourage the author to take notes. It is surprisingly difficult to remember what is said; no matter how sympathetic the group and how mature the author, hearing one's work discussed is such an intense experience that it can be hard to hear and remember everything as it is being said. We have even used a microphone to record discussions, especially when discussing a draft that is nearly ready to be submitted.

3.5 How to organize a meeting about a practice

In a meeting focused on practices, my main goals are that students become aware of their own writing practices, and that each student may identify a new practice that he or she might wish to try—or an old practice that might be best abandoned. Because the external supporting readings are not as crisp, and because my own material is less well developed, these meetings are less structured and more difficult to conduct.

Before the meeting

1. Choose a reading. I usually choose a chapter from Boice (2000). I have also sometimes used Becker (1986) or Beck (2003). Your university writing center may also have some suggestions.

   The chapters from Boice are short but dense; in a 75-minute meeting, don't try to cover more than one.

2. Distribute the reading.

3. Prepare a few questions to guide discussion during the meeting. (A sample of the kinds of questions I have used can be found in Section 5.6.)

During the meeting I generally arrange the meeting as a discussion, and I conduct the discussion using the same teaching methods that I use in my advanced graduate seminars. For these kinds of classes, Bruffee (1999) recommends small-group discussion of prepared questions, followed by a plenary meeting of the entire class. Although I do not share Bruffee's obsession with the social construction of knowledge, I have found his teaching methods effective.

3.6 How to plan a syllabus

A short syllabus If you and your students want just to try out a writing group and see if it does anything for you, you might be able to do something interesting in two meetings.

- The first meeting, which is described at length on page 9, should focus on the empirical method.
- For the second meeting, I recommend Exercise A. In my experience, this exercise is the most likely to show students how discussions of writing can lead to discussions of research ideas.

A long syllabus Fixing a long, rigid syllabus in advance isn't productive. If students haven't mastered basic principles, there's no point in teaching advanced ones. And it's well known that under pressure, writers typically forget the new things they've learned, making it useful to go back to basics. So instead if a syllabus, the most I can give you are some ideas about what to do when.

- If you meet once a week for an academic year (about thirty meetings), that's enough to get students exposed to all of the principles and practices on page 15, and to get them real experience with a significant fraction (maybe two-thirds of the principles and one-third of the practices). To expect mastery would be unrealistic; if your students master two or three principles and maybe a practice, be very happy.

- During the course of a semester (or year), plan for your students to follow three arcs: from small scale to large scale, from technique to ideas, and from discomfort to comfort.

- The principles on page 15 are organized from small scale to large scale, but this is not the right order in which to learn to apply them. Instead, follow the order in Williams. In particular, plan explicit exercises on *Subjects, verbs, and objects* (Principle 3, Exercise A and possibly also B), *Information flow* (Principle 5, Exercises C and D), *Coherence* (Principle 7, Exercises E and F), and *Abstract* (Principle 11, Exercise L), in that order. Plan at least two or three meetings for *Subjects, verbs, and objects* and *Information flow*, one or two meetings for *Coherence*, and two meetings for *Abstract*.

- *Structure of a section* (Exercise H) is an exercise in search of a principle, but it is nevertheless worth doing at least twice. This exercise can really help a student organize a chapter in a thesis or a difficult section in a paper. Postpone the exercise until your students have a good chance of writing reasonably coherent paragraphs with well-chosen subjects and verbs.

- *Parallel structure* (Principle 10, Exercise I) is a bit tricky. If you're lucky, you'll get a student's text in which you can plan to discuss parallel structure. Watch for an opportunity.

- *Consistent names* (Principle 1, Exercise B) and *Singular* (Principle 2, Exercise J) are best covered opportunistically on the fly, as they arise in some text.

- I can't give you much guidance about practices—which practices are helpful depends on the needs and even the personalities of your students. All the practices are valuable, but the one that is really essential, because it runs so counter to established myths and practices in our field, is *Write in brief daily sessions* (Practice 2). The experimental data are all there in Boice's research.

  Many students have trouble getting started; for them you may want to discuss *Prewrite* (Practice 4) and *Write a Shitty First Draft* (Practice 6). Our undergraduate students are often very worried about whether their initial work is good enough (thank you, Harvard); for them, we discuss *Focus on the process, not the product* (Practice 3). Beck's (2003) article on perfectionism can be interesting reading.

  In my own writing, my most valuable weapons are brief daily sessions (Practice 2) and index cards (Practice 5).

I tried to integrate research into the ways readers read with my experience working with professional writing in a variety of fields, in order to create a system of principles that would simultaneously diagnose the quality of writing and, if necessary, suggest ways to improve it.

— Joseph Williams, *Style*

4 Principles, practices, exercises, and guidelines

This section sets forth what I actually teach: editing principles, working practices, and guidelines for successful writers, together with exercises that help students master the principles. The principles and practices are listed in Table 1 on page 15. I introduce them with short explanations of the thinking behind them, but neither the principles nor the practices are self-explanatory. Explanations can be found in Williams (1995), in Boice (2000), and also in the exercises.

4.1 Principles and practices for technical writers

Editing principles An editing principle is useful only if a beginning writer can test to see if a text obeys it. Here are some examples of principles that are difficult to test for (all real advice from real writers):

- Omit needless words.
- Pay attention to the rhythm of the paragraph.
- Group ideas into sentences in the most logical way.

Here are some principles that are easier to test for:

- The agents and actions that you want to appear most important in the mind of your reader should be used as the subjects and verbs of your sentences.
- The old information in a sentence should appear at the beginning, and the new information should appear at the end.
- Don't use different words to mean the same thing, especially for technical terms. For example, don't use both "stack frame" and "activation record."
- In technical text especially, prefer singular to plural. For example, in the sentence "lexical analyzers translate regular expressions into nondeterministic finite automata," how will you know if a single lexical analyzer translates one expression or many? Singular is clearer.
- To clarify the meaning of mathematical or terminological definitions in which no action is taking place, illustrate the definitions with plentiful examples.

By using testable principles, we stay within the educational culture of science and engineering: for each principle, students can test the hypothesis that applying the principle makes writing clearer.

The principles in Table 1 are organized more or less by scale; in general, earlier principles apply to smaller parts of a manuscript. I have starred principles that I consider especially valuable.

Editing principles

0. *Correctness.* Write correct English, but know that you have more latitude than your high-school English teachers may have given you.

⋆1. *Consistent names.* Refer to each significant character (algorithm, concept, language) using the same word everywhere. Give a significant new character a proper name.

⋆2. *Singular.* To distinguish one-to-one relationships from n-to-m relationships, refer to each item in the singular, not the plural.

⋆3. *Subjects, verbs, and objects.* Put your important characters in subjects, and join each subject to a verb that expresses a significant action. Understand what object, if any, is acted upon.

4. *Definitions.* Mathematical definitions lack significant actions, so clarify them using examples.

⋆5. *Information flow.* In each sentence, move your reader from familiar information to new information.

6. *Emphasis.* For material you want to carry weight or be remembered, use the end of a sentence.

⋆7. *Coherence.* In a coherent passage, choose subjects that refer to a consistent set of related concepts.

⋆8. *Purpose.* Give each paragraph a purpose, which is the effect you want in the mind of your reader.

9. *Paragraph = Issue + Discussion.* Begin a paragraph with one to three sentences, which end by emphasizing the *issue*. Finish by using the issue in in coherent subjects or *thematic strings*.

⋆10. *Parallel structure.* Order your text so your reader can easily see how related concepts are different and how they are similar.

11. *Abstract.* In an abstract, don't enumerate a list of topics covered; instead, present the essential information found in your paper.

Working practices

⋆1. *Pause mindfully, frequently.* Mind your body, thoughts, and feelings—and the stage of your work.

⋆2. *Write in brief daily sessions.* Ignore the common myth that successful writing requires large, uninterrupted blocks of time—instead, practice writing in brief, daily sessions.

3. *Focus on the process, not the product.* Don't worry about the size or quality of your *output*; instead, reward yourself for the consistency and regularity of your *input*.

4. *Prewrite.* Before you write, think, talk out loud, and jot down notes, diagrams, outlines, and so on.

5. *Use index cards.* Use them to plan a draft or to organize or reorganize a large unit like a section or chapter.

⋆6. *Write a Shitty First Draft*™. Value a first draft not because it's great but because it's there.

7. *Don't worry about page limits.* Write the paper you want, then cut it down to size.

8. *Cut.* Plan a revision session in which your only goal is to cut.

Table 1: Principles and practices of successful writers

Working practices It's surprising how many books on writing talk only about the words on the page and not about what the writer is actually doing—how the writer behaves. These books are missing many important questions: Where do you write? When? How often? For how long? With what goals? How do you know when to stop? How do you think about writing? The answers to these questions affect not only your attitudes but also the amount and quality of the text you produce. Some of the research is astonishing (Boice 2000).

I call useful behaviors *practices*. Good practices change students' behavior, which in turn can change their attitudes, which in turn can change behavior, and so on in a virtuous cycle. Writing practices are personal, and a teacher ought not to prescribe given practices but rather should help students discover which ones are best for them. I have found each of the practices helpful, but my two favorites are to write in brief daily sessions (Practice 2) and to prewrite and revise using index cards (Practice 5). Use the group to help each student discover which practices work best for him or her.

Your working practices should vary depending on where you are with your manuscript. Every manuscript, and every part of a manuscript, starts out as thoughts about what to write and ends up as a finished, edited product. But the evolution of a manuscript is not continuous; it passes through four distinct *stages* or *phases*, which are shown at the top of Figure 2 on the following page:

- *Active waiting* is a structured way of planning to write. It produces ideas.
- *Prewriting* puts your writing plans into a tangible form. It produces "stuff," like outlines, scribbles, pictures, notes, or index cards.
- *Drafting* is the process of producing a first draft. It produces *private* text; only the author sees the text
- *Editing* is everything that happens after you have a draft. It produces *sharable* text; edited text may be shared with collaborators, supervisors, and eventually, reviewers.

Although we focus a great deal of attention on the editing stage and the application of editing principles, each stage calls for some unique practices. And as shown at the bottom of Figure 2, each stage benefits from the practices of mindfulness and of brief, daily sessions.

The practices in Table 1 are organized somewhat by stage. Practices 1, 2, and 3 apply to work at any stage. Practices 4 and 5 apply to work in the prewriting stage, And Practices 6 and 7 apply to work in the most difficult stage: the production of the first draft.

4.2 Exercises

The exercises below comprise most of what we do when I teach writing. Almost every exercise is designed to teach one of the principles in Table 1; an important exception is Exercise H, which although valuable, does not come with an articulated principle. Not all exercises are equally good; among the best are Exercises A, C, and H, which you can profitably do more than once.

The exercises are listed in an order in which it may be useful to do them. Most have been tested thoroughly; those not so tested are clearly identified.

<table>
<tr><td>1 Active waiting</td><td>2 Prewriting</td><td>3 Drafting</td><td>4 Editing</td></tr>
<tr><td colspan="4">Brief, daily sessions</td></tr>
<tr><td colspan="4">Mindfulness</td></tr>
</table>

Figure 2: Four stages of a manuscript, with foundational practices

Why include direct objects in Exercise A?

Williams talks only about the agents (characters) and the actions they take, not about the objects acted upon. It may be that using agents and actions as subjects and main verbs is enough to help you construct good sentences. But when you are planning a larger work—say to describe a complex experiment or a system in which hardware, software, and people all play a role—getting a handle on the direct objects is invaluable. It's common, for example, to find a character that is the agent of one action but the direct object of another action (the IDE controls the compiler but the user controls the IDE). It's also common to find relationships in which there may be agency in both directions (the n-gram statistics direct the classifier and the classifier updates the n-gram statistics). Capturing the relationships in table (or even in a directed, labeled graph) is a first step toward explaining them clearly.

Exercise A: Who does what to whom This exercise is based on Chapter 2 of Williams's *Style: Toward Clarity and Grace*. The big lesson from Chapter 2 is this: if you have certain ideas in your head about what agents and actions are most important, you will communicate those ideas most clearly if you make those agents and actions the subjects and verbs in your sentences. To help you learn how to apply this principle, here is an exercise in three parts.

The first part is about the ideas that form in your head as you read, not about the words on the page. Take the text, and as you read each paragraph, identify

- The important characters in the story
- The actions taken by those characters
- When applicable, the direct objects of those actions

Use *your own words* to identify the characters and their actions, not necessarily the words in the text. I recommend using a clean sheet of paper to create a table listing agents (who), actions (does what), and objects acted upon (to whom).

The second part is to go through the text again. Make a distinctive mark on the main subject and verb of each sentence. (I like to underline the subject and double-underline the verb.) Whenever the main verb is a form of a *light verb* like “make,” “do,” “have,” “bring,” “put,” “take,” or especially “be,” make a distinctive mark in the margin.

The third part is to compare. How consistent are the important characters and actions with the subjects and verbs used in the text? If you felt good about the text and enjoyed reading it, did you find that the characters and actions were consistent with the subjects and verbs? If you didn’t enjoy the text, did you find that the characters and actions were inconsistent with the subjects and verbs? Did you find lots of light verbs?

For the instructor: *My experience with this exercise varies widely, but at its best, this is by far the most interesting of the exercises. Discussion of agents and actions can produce deep insight into research. The exercise can lead to significant revisions in the text, including revisions that cross paragraphs.*

The big pitfall in this exercise is to use the exact words found in the text as one’s description of agents or actions. Coming up with descriptions in one’s own words often requires real intellectual effort, but the effort can be repaid by a better understanding of the material.

Other notes:

- *A question that sometimes comes up here is whether two different words are used to mean the same thing. Less often, a single word may be used with two different meanings. I usually suggest rewording to avoid such problems (Exercise B).*

- *As part of this exercise, when some agents and actions are on the board, it can be quite helpful to get the group to work out what entities are truly distinct and how they relate to one another.*

- *The list of light verbs comes from Pinker (2014, page 105). The most pernicious of these is “be.”*

Exercise B: Diction Sometimes it can be hard to work with agents and actions because the agents or actions are difficult to identify by name. Writing about research in computing can be especially difficult because there are so many new things for which there are no established names. To help with these problems, here are the three parts of Principle 1 (*Consistent names*):[4]

- *Give it a name.* Some writers try to dodge the issue by carefully avoiding naming things. Saying “our language,” “the prototype system,” or “the algorithm” doesn’t do the job. Do your reader a favor and give your language, your system, or your algorithm a name.

- When you are talking about one idea, *always use the same word* or phrase. For example, don’t call your idea “data dispersal” in one place and “revealing secrets” in another.

- When you are talking about different ideas, *never use the same word.* For example, don’t use the word “system” to talk about a model, an algorithm, and a software artifact.

[4] All the examples are from papers we have discussed.

The exercise is to scrutinize a manuscript and identify places where names are misused or an important thing is unnamed. Reduce the number of names as needed, and choose effective names for each concept, agent, action, and object.

For the instructor: *Beginners, especially undergraduates, consistently have trouble with diction. And even an experienced writer can forget to name the new baby.*

I've never devoted a meeting to this exercise; instead, I handle the principle when it comes up in a student's paper. But depending on the text, a class that intends to do Who does what to whom *(Exercise A) may well wind up doing this exercise instead.*

Exercise C: Old and new information This exercise is based on material from Chapter 3 of Williams's *Style: Toward Clarity and Grace*. The idea is to make text easier to read by considering the flow of information from one sentence to the next sentence within a paragraph. Williams argues that information flows best when old information is at the beginning and new information is at the end.

To learn how to apply this principle, here is an exercise in two parts.

- The first part is to go through the text and mark the old and new information in each sentence. If there is more than one piece of new information, mark the *most important* new information. I usually mark with a dotted underline for old information and a solid underline for new information, but you should mark using a system that works for you.

  Remember that although we look at each sentence in isolation, good information flow *connects* adjacent or nearby sentences. Often, the old information at the beginning of a sentence is the new information delivered at the end of the preceding sentence. And the new information at the end of a sentence becomes available to be used as old information in the succeeding sentence or sentences. To see a perfect example of how old and new information can be used to connect sentences so they flow smoothly, look at the sentences about the black hole on page 47 of Williams.

  Finally, the opening sentence of a paragraph is special. The old information in the context is not as strong. It's most likely that you adjust the structure of the first sentence so as to end with the information that you want to be old when you write the second sentence. You can learn much more about how to open paragraphs in Chapter 5 of Williams.

- The second part is to identify one or two sentences that you would like to revise based on information flow, and to suggest a revision for each.

As students are learning, they may struggle with the ideas of "old" and "new" information. Here are a couple of other ways of thinking about flow:

From *familiar* to *unfamiliar*
From *input* to *output*

It's also possible to structure a sentence, particularly the very first sentence, to flow from general information to specific information—but this trick is really an application of Principle 6, which recommends to put the most important information at the end of a sentence.

For the instructor: *When looking at old and new information, my classes often see two kinds of problem sentences:*

- *The sentence leads with the important new information and finishes with old information. This style is not good for technical writing, but it matches how most people converse—just listen, or look at the dialog in a novel.*[5] *So this is a case where the advice, "write the way you talk," is not entirely helpful. A key difference between conversation and technical writing is that sentences spoken in conversation typically don't pile up into groups of four, six, eight, or more. But in writing, paragraphs of that size are routine. And if you write a long sequence of sentences, each of which leads with new information, readers are going to feel that your writing is choppy. The treatment for this symptom is just what Williams prescribes: lead with the old information and finish with the new information.*
- *The sentence may or may not have the old and new information in good positions, but no amount of rearrangement is going to make the sentence easy to understand, because it contains too much new information.*[6] *The treatment for this symptom is to break up the sentence into multiple independent clauses, each with its own subject and verb. Those clauses might then be expressed as a compound sentence, as multiple sentences, or any combination. Once they have the idea of independent clauses, students can figure this out.*

For the instructor: *One more note about this exercise: I'm always surprised at how much trouble students have with it. The difficulty seems to be in identifying what information is old and what information is new. If your students also have trouble, you may want to try the next exercise instead.*

Exercise D: Important information Students sometimes have difficulty identifying old and new information in Exercise C. A useful simplification is simply to mark the *most important* (new) information in each sentence. Discussion can proceed based on how many sentences place the most important information at the end—or how sentences might be revised by moving the most important information to the end.

Exercise E: Coherent subjects This exercise is also based on material from Chapter 3 of Williams's *Style: Toward Clarity and Grace*. The idea is to make a paragraph feel more coherent by considering the string of topics within the paragraph (Williams 1995, page 56).

To learn how to apply this principle, here is an exercise in two parts.

- Go through a paragraph and underline the first five or six words of each sentence.

[5] Look at a novel, not a play. Playwrights, unlike novelists, seem to be aware of the impact that information has when it is placed at the end of a sentence. Or maybe they are just less interested in sounding "natural."

[6] I've always been annoyed by teachers, editors, and colleagues who advise writers to avoid long sentences. It doesn't matter how long a sentence is; what matters is how hard the sentence is to understand. Regardless of its length, if a sentence delivers three, four, or more items of new information, then unless those items are closely related, readers may find the sentence hard to understand.

- Study the topics for coherence. First, eliminate every outlier that simply refers to information from the end of the preceding sentence. Do the remaining topics seem to form a coherent sequence? If not, please *suggest* a sequence of topics you like better, then *revise* each sentence as needed to move the suggested topic to the beginning of the sentence.

When doing this exercise, especially to a well-written paper, be aware that for the opening sentence in each paragraph, special rules apply. An opening sentence might begin with an uninteresting subject, or even with metadiscourse. The goal is to put more emphasis on the new information that appears at the *end* of the opening sentence. This technique is shown at some length in Chapter 4 of Williams (1995). If you haven't yet read Chapter 4, peek at the section "Some Syntactic Devices" on pages 71 and 72. That section shows a few tricks that writers use to weaken the opening of a sentence in order to "add weight to the end of [the] sentence." These constructions, provided they are used sparingly, can "stress those ideas that you intend to develop in the following sentences."

To handle the special rules that apply to opening sentences, you can extend the exercise in this way:

- When the first sentence of a paragraph doesn't begin with a few words that seem to be a topic, look instead to the end of that sentence. If the end of the sentence introduces new ideas that are then developed in the rest of the paragraph, put a box around those new ideas.

- When you study topics for coherence, include the boxed ideas on your list of topics.

For the instructor: *I've had mixed results with this exercise. In the professional literature, we have read some truly impressive paragraphs, which stuck to one or two coherent subjects without seeming boring or repetitive. On the other hand, when we try to do the trick ourselves, it can require complete disassembly and reassembly of several paragraphs, which in turn requires real insight into the author's intentions. When it works, it's great, but I haven't yet figured out when it can be done without destroying the paragraph. In other words, I have had a hard time figuring out in advance whether a text will work well with this exercise.*

Exercise F: Quick start Underline the first seven or eight words of each sentence. If the underlined portion does not contain an agent as subject and an action as verb, that sentence is a candidate for revision.

For the instructor: *Williams suggests this exercise, but I have not actually used it.*

Exercise G: Issue and discussion This exercise is based on Chapter 5 of Williams's *Style: Toward Clarity and Grace*. The idea is that a paragraph is divided into two parts. The paragraph opens with a kind of introduction or overture, which Williams calls the *issue*. The last sentence of the issue should end, in the position of emphasis, with the topic of the paragraph. That topic can then engender *thematic strings* that recur throughout the rest of the paragraph, which Williams calls the *discussion*. Thematic strings are especially effective when they appear as coherent subjects, but a thematic string may appear anywhere in a sentence. The clear emphasis at the end of the issue, coupled with the recurring thematic strings in the discussion, give your reader a strong sense of coherence and what the paragraph is about.

The exercise proceeds as follows:

1. Choose a paragraph and as best you can, draw a dividing line between issue and discussion. If you cannot find such a dividing line, write new sentences for the beginning of the paragraph, and draw the dividing line after your new sentences.
2. Underline the phrase that appears in the position of emphasis at the end of the issue. In your own words, restate the subject, object, concept, idea, or theme embodied in that phrase.
3. Throughout the discussion, underline each sequence of words that you think is thematically connected to the emphasized phrase at the end of the issue.
4. Decide if you think the paragraph respects the principle. Decide how coherent you feel the paragraph is—that is, how much does the paragraph feel to be about one thing. Are the two correlated.

In the advanced version, revise:

5. Choose another paragraph. Looking at what the paragraph is about, choose a set of thematic strings that (a) are all related and (b) can reasonably appear in the paragraph.
6. Rewrite the issue, or write a new issue, to end in a phrase chosen from among those thematic strings.
7. Rewrite the discussion to use the thematic strings.
8. Compare the original with the revision, and decide if the original is improved. Ask your readers!

Exercise H: Structure of a section This exercise can help an author with the structure of a section in a conference or journal paper or with the structure of a chapter in a thesis. The preparation is simple: each member of the writing group reads each paragraph of the section or chapter and answers two questions.

1. What is the purpose of this paragraph?
2. How well does it fulfill its purpose?

The first question is more important than the second. The hard part is distinguishing the *purpose* of a paragraph from the *content* of that paragraph. Roughly speaking, content is what a paragraph is about, while purpose usually has to do with causing an event to happen in a reader's mind. During the meeting, you will find you must continually keep the group focused on purpose.

The fun comes in the meeting. *The author is not allowed to say anything.* Instead, the text has to speak for itself. The writing-group moderator will help the group form its collective impressions of the paragraphs. We've had good experiences with this exercise.

- It can be eye-opening for the author to learn how others read the text.
- It can let the author know how successful the section is in general.
- It can identify several kinds of structural problems:
  - Paragraphs that try to serve two or three purposes at once
  - Paragraphs the purpose of which is not obvious
  - Introductory material at the end of the section
  - Paragraphs serving the same purpose that are widely separated in the text
  - Redundant paragraphs

  There's generally no need to try to identify such problems in advance; these identifications emerge naturally from the discussion.

We've used this exercise successfully with sections of 10–20 paragraphs.

After doing this exercise, it may be helpful to use a deck of index cards to reorganize the section (Practice 5).

For the instructor: *This exercise can be very helpful to the student who has a draft that may be submitted within a few weeks. As noted above, you'll have to lean hard on people to force them to distinguish* contents *from* purpose*.*

A paper suggested with this exercise in mind may also be suited to Parallel structure *(Exercise I). Prefer Exercise I; you can do* Structure of a section *with any paper, and opportunities to teach parallel structure should not be overlooked.*

Exercise I: Parallel structure Much of scientific writing is about making comparisons. When two or more complex things are compared, a reader can follow more easily if the comparison uses parallel structure. I encourage you to establish parallel structure as a two-part process:

- You design a *template* for describing each thing to be compared. The fixed parts of the template are the same in each description; the variables vary.
- For each thing you describe, instantiate the template by substituting text for the variables.

These two parts can be iterated in either order. Your job is to identify an effective template, but both designing the template and judging its effectiveness are easier if you look at instances.

The exercise proceeds as follows:

1. Begin with a piece of scientific writing in which two or more alternatives are compared.

2. Break each alternative down into its atomic elements. (The choice of what elements are considered atomic is up to the author.)

3. For each atomic element in each alternative, consider how it relates to elements in other alternatives. There are three possibilities.

   - Elements are parallel because they are the same in multiple alternatives. For example, both XEmacs and GNU Emacs are free software released under the Gnu Public License.
   - Elements are parallel because although they are not the same, they are directly comparable; e.g., you have to choose among them. For example, XEmacs draws the screen by using the Xt toolkit, whereas GNU Emacs uses the X protocol directly and can work without a toolkit.
   - Elements in one alternative have no parallel in another alternative. For example, GNU Emacs is controlled by Richard Stallman, who is a well-known, controversial figure. No person of similar characteristics is associated with XEmacs.

4. Finally, and most difficult, design a template that can be used to describe each of the alternatives, such that parallel elements appear as variables in the template.

Identifying parallel and non-parallel elements can be done outside of group, but choosing a good parallel structure is best done during the group meeting.

Here is an example of parallel structure from a paper on visualizing time-varying data. I use a three-sentence template; the variables are in italics and angle brackets:

1. The ⟨*Nth*⟩ category of time-varying data is ⟨*mumble*⟩.

2. ⟨*Mumble*⟩ behavior ⟨*does something characteristic*⟩.

3. The ⟨*mumble*⟩ category includes ⟨*example*⟩.

And here are the instances:

- The first category of time-varying data is *regular*. Regular behavior grows, persists, and declines in several (distinct) stages. The regular category includes many natural phenomena and their simulations, such as the earthquake.

- The second category of time-varying data is *periodic*. Periodic behavior shows patterns that recur at fixed intervals. The periodic category includes climate data, which normally follow a daily, monthly, or yearly pattern.

- Finally, the third category of time-varying data is *turbulent*. Turbulent data feature ubiquitous, spontaneous fluctuations distributed over a wide range of spatial and temporal scales. The turbulent category includes computational simulations of fluid dynamics.

The original text read as follows:

> The first category of time-varying data is regular, which usually involves a certain phenomenon that grows, persists, and declines in several (distinct) stages. The rate of change at each stage could vary dramatically in space and time. Many natural phenomena and their simulations, such as the earthquake, fall into this category. The second category of time-varying data is periodic. For this type of data with recurring patterns, special attentions are paid to space-time abnormal events. For example, climate data normally follow a daily, monthly, or yearly pattern. Occasionally, however, the data may also fluctuate out of expectation, creating an abnormality that requires attention or investigation. Finally, the third category of time-varying data is turbulent. A number of computational fluid dynamics (CFD) simulation data are turbulent, featuring the ubiquitous presence of spontaneous fluctuations distributed over a wide range of spatial and temporal scales

If you want to go deeper into sentence structure, Williams's Chapters 8 and 9 show many examples of how large sentences can be structured, including using parallel structure. But be warned that I found these chapters more advanced and more difficult to learn from than the earlier chapters.

For the instructor: *Parallel structure is a very powerful technique, and in technical writing it comes up often. When applied to a few pages of a student's manuscript, the exercise can be extremely effective. When your students have drafts, watch for an opportunity to teach parallel structure. You may also try to help them use parallel structure to discover whether their comparisons are fully thought out.*

Exercise J: Singularity A common fault in computer-science writing is to use plural everywhere. In technical text especially, prefer singular to plural. For example, in the sentence "lexical analyzers translate regular expressions into nondeterministic finite automata," how will you know if a single lexical analyzer translates one expression or many? Singular is clearer.

The exercise is to tackle several paragraphs and eliminate as many plurals as possible (without changing the meaning of the text).

For the instructor: *I've never found it worthwhile to spend a whole meeting on this exercise. Instead, when the opportunity presents itself, I hammer away at the principle.*

Exercise K: Cutting Many professional papers are limited to a fixed number of pages. To produce a paper within the limit, it is sometimes necessary to write a longer paper and then cut. Because it is so difficult to cut one's own work, we suggest practicing cutting on someone else's work. We plan to experiment with cutting a section to $\frac{3}{4}$ or even $\frac{1}{2}$ of its original length.

Start this exercise with a section in which each paragraph has been labeled with its purpose, as in Exercise H. Use this information to decide how many jobs the section does within the paper as a whole. Based on this decision, cut in one of two ways:

- If the section does multiple jobs, perhaps one or more of those jobs can be eliminated. In this case, identify the paragraphs doing the work, and cut those paragraphs.
- Perhaps the section does only one job, or perhaps each of the jobs it does is essential. In this case, assign a *relative value* of each paragraph; an easy

measure of value would be the ABC scale. Now cut the C paragraphs, followed by as many of the B paragraphs as needed to reach your length goals.

Make these kinds of cuts, repeating if necessary, until the text is at or just under the target length. Now re-examine and rewrite the section to be sure that it is still coherent, that transitions make sense, and so on. If this rewriting pushes you over the length limit, go back and cut again.

When cutting a technical paper, it is tempting to keep all the "real content" and to remove motivation and examples. Resist this temptation.

For the instructor: *It is very difficult to cut one's own work. Unless a member of the group has a paper that needs cutting to meet a length limit, it may be better to practice these techniques on a published paper. It should not be hard to find a journal paper that could be improved by vigorous cutting.*

I haven't actually used this exercise.

Exercise L: Writing the abstract Writing a scientific abstract is a specialized art. To practice this art, follow the advice given by Landes (1966): make sure the abstract includes the *essential information* presented in the paper (Depending on whether your students have significant experience reading and analyzing technical papers, you may wish to begin instead with Exercise M.)

To prepare for the exercise, the group leader should take a technical paper and remove the abstract. Since this is one of the few exercises for which group members will have to read an entire paper, it helps if the paper is well written, easy, and of interest to most members of the group.

The exercise has two parts:

- To prepare for writing group, read the paper and mark those points that you think constitute the "essential information" that should be presented in the abstract. Highlight, make a list, or do whatever works for you.
- In group, attempt to prepare abstracts at two of the more common lengths: 200 words and 50 words. As time permits you may also try 300 or 100 words. If you are motivated to write an abstract ahead of time, by all means do so.

If you are pressed for time and cannot read the whole paper, you may do almost as well by abstracting what you find in the introduction and conclusion.

Our experience is that an interesting paper usually requires two sessions: In the first session, we agree on what constitutes the essential information in the paper. In the second session, we write abstracts. Because the actual writing requires that we choose suitable subjects and verbs and manage the flow of information well, it helps to do this exercise after your students can apply these techniques successfully.

For the instructor: *Abstract writing is a niche skill but an important one. If you have the luxury of many meetings, you may well wish to teach it. If you wait until late in the year, you'll be able to see if your students can successfully apply other principles (like information flow) to the writing of the abstract itself.*

Exercise M: Reading the abstract Finding the essential information in a paper requires real intellectual work, and if your students don't already have practice doing this kind of work, it may be that you won't have time to teach it. A reasonable substitute for Exercise L is to read some abstracts instead.

Choose a handful of abstracts. For easy reference, number each sentence of each abstract. The exercise is as follows:

- For each numbered sentence, say whether the sentence actually *presents* information or whether it merely *promises* information.
- Find the abstract that the group likes best, and the one that the group likes least. Is there any relationship between what the group likes and an abstract's ratio of promises to presentation?

4.3 Supplementary guidelines

We have developed some useful guidelines that we do not yet know how to turn into crisp principles or exercises.

Guideline 1: Explaining a technical concept Science and engineering often involve explaining new concepts. To help decide if the explanation of a new concept is adequate, here are some questions:

- Have I enumerated all the properties of the thing?
- Have I said whether the thing is completely characterized by those properties?
- Does each property have a name?
- If mathematical, does each property have a symbol?
- Have I said what *kind* of value each property is? (Integer, real number, string, symbolic expression, list, tree, graph, etc etc)
- Have I explained relationships that hold among the properties? Who or what guarantees that these relationships hold?
- If I have definitions in which no actions are taking place, have I illustrated each defined thing with examples?

Guideline 2: Checklist for technical exposition Here are some ideas, questions, and techniques we have found helpful when planning, organizing, and assessing a paper as a whole. Some of these ideas apply only to computer science.

- Have you identified the target audience?
- Have you told your reader what you expect? For example, should he or she just understand high-level ideas, or is it important to get all the details? What should a reader take away? For example, should your reader be led to draw a conclusion? Acquire a new skill?
- Do you have examples? They are helpful, and they should

  - Be plentiful
  - Use parallel structure
  - Be connected to each other when possible

  An ideal, when possible, is to use a single running example that appears in each section of the manuscript. (It may be supplemented by additional examples.)

- Is every general, abstract declaration illustrated by an example? For example, is a declaration such as "A constructor is used at compile time to build an abstraction" illustrated by an example such as "for example, a compiler might use the constructor `gbind` to build an environment that binds `main` to a procedure."?

Computer scientists often create artifacts that are too complex to be easily described. Here are some notes for presenting complicated technical abstractions:

- You may well have a nest of interrelated concepts for which there is no obvious order of presentation. To come up with an order, you may have to tell lies, i.e., make simplifications for pedagogical purposes. Such simplifications should be announced. For example, you could claim for pedagogical purposes that a variable stands for a number, not a location.

  Another technique is to mention a concept without defining it. For example, you might say "Let's assume that l is a location on the stack, without going into the details, which are in Section 12."

  Checklist: Is every concept mentioned before it is used? Are most concepts defined before use?

- Types help. Do you give the type of every operation?

- Do you explain the name of each variable? Do you explain what each Greek letter may stand for? For example, do you explain that ρ stands for an environment?

- When presenting abstract data types, we are aware of two styles. Hoare's style talks about the abstraction represented by a type and explains the concrete operations by their effects on the abstraction. For example, Hoare might explain an environment by using the abstraction of a set of bindings, and he might explain lookup by finding a binding with a given left-hand side.

  Algebraic or equational style (owing much to Goguen and Guttag) gives equations that relate concrete operations on the type. Equations can usually be turned into a term-rewriting system that can specify results returned by observers. For example, algebraic style might rewrite a lookup operation into the value looked up (by substituting equals for equals at every step).

  Checklist: Do you know what style you are using? Are your definitions and examples all consistent with that style? Do you wish to use both styles? If so, have you explained the redundancy to your reader?

5 Sample materials from my group

This section presents materials I have used with classes or groups. Each subsection is devoted to a text, a principle, or a practice, together with materials or ideas I have used with my students. I've gathered these materials over a span of many years. Early on, I rarely saved more than a text, some suggestions, or some questions for discussion. But once I realized the value of fully worked examples, I made sure to gather not only a text but also my analysis, my revision, and an account of my classroom experience. These more comprehensive materials appear early in the section; the sketchy stuff follows.

5.1 Exercise (Who does what to whom): Sri Lankan inscriptions

Figure 3 shows a single paragraph, marked C, on which my class did Exercise A, analyzing the agents, the actions, and the objects acted upon, then revising the paragraph. My solution follows.

Table of agents, actions, objects

Agent	*Action*	*Object*	*From*
archaeologist?	identifies	letters	C1
faults in stone	distort/obscure	physical shapes (of letters)	C2
stonecutter	influences	physical shapes (of letters)	C3
tool	influences	physical shapes (of letters)	C3
researchers	want	science???	C5
letters	have	ideal shapes	
researchers' software	infers	ideal shapes	C10

Reflections and analysis

I found the key characters in this paragraph very hard to identify. The letters and their shapes are clearly key characters. Archaeologists, stonecutters, tools, researchers, and software all show up as characters.

My table mentions only sentences C1 to C3, C5, and C10. I don't mention the other sentences because I couldn't find anything new that they contributed. And I'm not sure quite what the authors mean by "more scientific." Something about using human senses is unscientific?

The important characters in the story are seldom reflected in the subjects of the sentences. (They do appear as "letter" in C3, as "alphabet" in C7, and as "shape" in C9, but it's not enough.) The important actions don't appear as main verbs. Instead, the paragraph shows the problem of *nominalization* that is mentioned by Williams: of the ten main verbs in this paragraph, five are forms of "is".

My revision cuts the paragraph down to six sentences, and as the central characters in the story, it focuses on the letters and their shapes.

Revision

Archaeologists who wish to read stone inscriptions must identify what letters have been inscribed. Each letter has an ideal shape, but real letters differ from the ɔtions

19 SEP 2014

Atsushi Minato[1] and Satoru Ozawa[1]

Ibaraki University, Hitachi 316-8511, Japan
ɡy, Colombo 7, Sri Lanka
ı@yahoo.com

Sri Lanka is fortunate because of the diversity and richness of historical sources, especially literary sources like *Mahavamsa* and *Deepawamsa* [6]. In addition, there are the foreign sources in several languages which are useful for studying history in Sri Lanka[6]. All these literal sources were written in later periods. Inscriptions sources are of inestimable value because they are contemporaneous unlike most of the literary sources. If they have escaped from the ravages of man and nature, they remain in their original form. The early inscriptions are free of poetic embellishments and merely record events without didactic or pedagogic objective, unlike the *Mahavamsa* which was written for a particular purpose. Therefore, the historical value of inscriptions is very much enhanced[7].

When we get a newly unearthed inscription, we must first identify the letters inscribed on the stone surface. The stone surface is usually contaminated by various kinds of noises such as scratches, cracks, voids, etc. And also, the same letter sometimes takes different shapes depending on the skill of inscriber and the tool of inscribing. Up to recent days, the work of identification of letters has been carried out by human sense. We like to introduce more scientific method for the letter identification. In order to have a scientific method, the first step is to produce alphabet fonts of ancient scripts. The alphabet of early Brahmi script has already been created by archaeologist by studying the common features of letters found in inscriptions[5,8]. Their method is based on human sense and hand writing. Therefore, the shape of the alphabet can be slightly different depending on the creator. The aim of this research is to create more precise alphabet fonts of early Brahmi scripts in Sri Lanka by analyzing a lot of inscriptions by computer without depending too much on human sense.

II. Evolution of Letters in Sri Lanka

Figure 3: Paragraph C, on Sri Lankan inscriptions

ideal; their shapes are influenced by the skill of the inscriber and by the nature of the inscribing tool. And letters' shapes are usually distorted by scratches, cracks, voids, or other defects in the stone. At present, inscribed letters are identified by human archaeologists. But ideally, letters could be identified using computer software. In this paper, we present a first step: an algorithm that analyzes images of inscriptions and computes an ideal shape for every letter in the Brahmi alphabet.

Experience in the classroom

This text was too difficult for beginning students. It took over an hour to fight through the words the authors had written in order to discover the key ideas of letters, ideal shapes, and physical shapes. I can recommend this text only for experienced students.

5.2 Exercise (Information Flow): Data visualization

A student brought us the paper *Importance-Driven Time-Varying Data Visualization*, saying that she found it hard to read because it "jumped from idea to idea." This symptom, according to Williams, can be treated by improving information flow. I also found that the parallel structure was incomplete, and that by completing it and improving it, we could also mitigate a reader's feeling of being yanked from idea to idea. Figure 4 shows the text we analyzed; Figure 5 shows the same text with old and new information marked.

Reflections and analysis

I found many parts of the text that respected Williams's ideas about information flow, but I also found some significant "reversals," in which new information came first and was followed by old information. In my revision I repaired these defects. I also corrected two other deficiencies:

- Where the original text uses inconsistent diction, sometimes saying "behavior" and sometimes "phenomenon," I made the diction consistent, using "behavior" throughout.
- Where the original text compares the three categories of time-varying data, I tried to establish parallel structure. Here's the structure I picked:
    1. The Nth category of time-varying data is *mumble*.
    2. *Mumble* behavior *does something characteristic*.
    3. The *mumble* category includes *example*.

So as to avoid packing too many writing principles into one example, I've not improved the text as much as I wished to. (I especially wished to improve the choice of nouns.)

Revision and commentary

Here it is, sentence by sentence:

A1. Almost every major scientific discipline has time-dependent simulations or time-varying data.

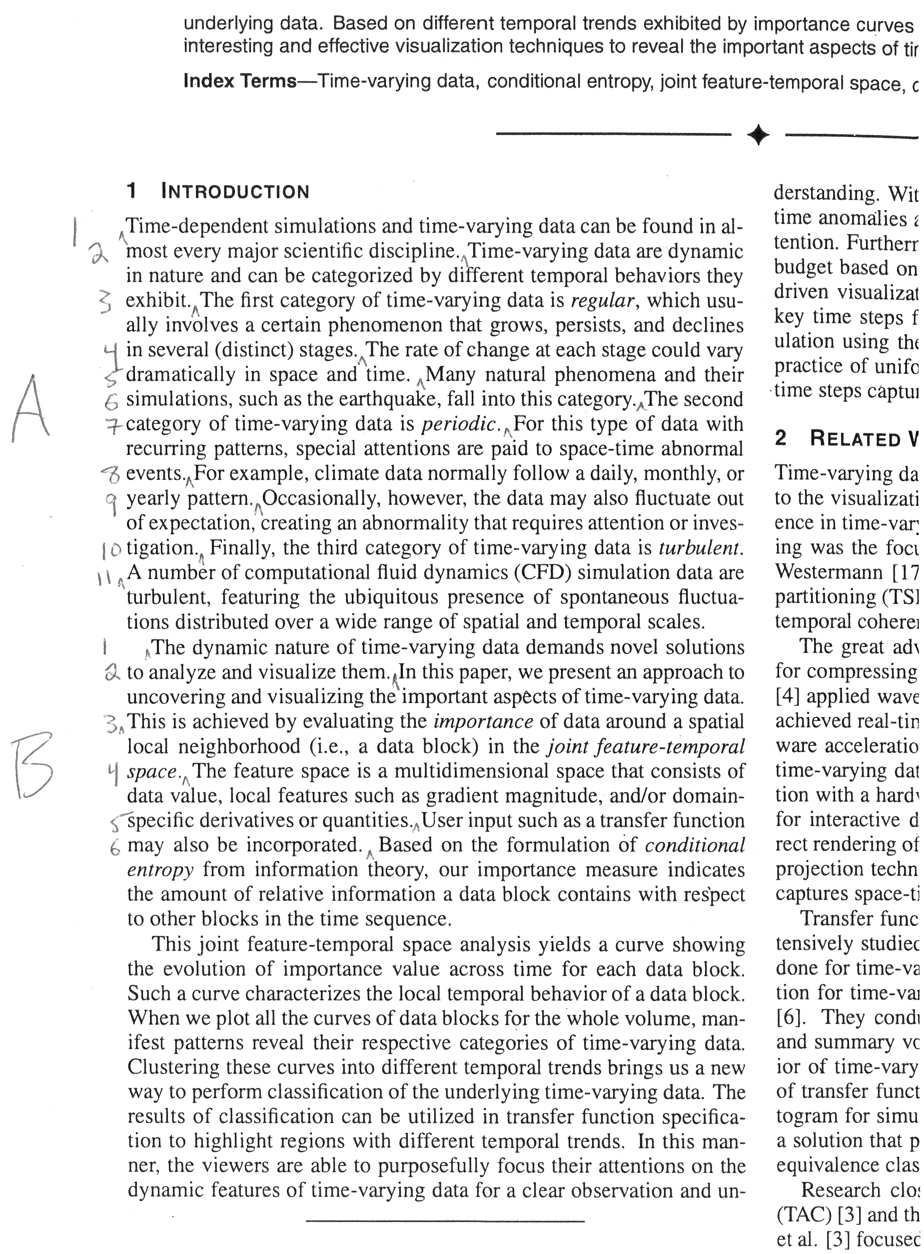

underlying data. Based on different temporal trends exhibited by importance curves
interesting and effective visualization techniques to reveal the important aspects of tir

Index Terms—Time-varying data, conditional entropy, joint feature-temporal space, c

1 Introduction

Time-dependent simulations and time-varying data can be found in almost every major scientific discipline. Time-varying data are dynamic in nature and can be categorized by different temporal behaviors they exhibit. The first category of time-varying data is *regular*, which usually involves a certain phenomenon that grows, persists, and declines in several (distinct) stages. The rate of change at each stage could vary dramatically in space and time. Many natural phenomena and their simulations, such as the earthquake, fall into this category. The second category of time-varying data is *periodic*. For this type of data with recurring patterns, special attentions are paid to space-time abnormal events. For example, climate data normally follow a daily, monthly, or yearly pattern. Occasionally, however, the data may also fluctuate out of expectation, creating an abnormality that requires attention or investigation. Finally, the third category of time-varying data is *turbulent*. A number of computational fluid dynamics (CFD) simulation data are turbulent, featuring the ubiquitous presence of spontaneous fluctuations distributed over a wide range of spatial and temporal scales.

The dynamic nature of time-varying data demands novel solutions to analyze and visualize them. In this paper, we present an approach to uncovering and visualizing the important aspects of time-varying data. This is achieved by evaluating the *importance* of data around a spatial local neighborhood (i.e., a data block) in the *joint feature-temporal space*. The feature space is a multidimensional space that consists of data value, local features such as gradient magnitude, and/or domain-specific derivatives or quantities. User input such as a transfer function may also be incorporated. Based on the formulation of *conditional entropy* from information theory, our importance measure indicates the amount of relative information a data block contains with respect to other blocks in the time sequence.

This joint feature-temporal space analysis yields a curve showing the evolution of importance value across time for each data block. Such a curve characterizes the local temporal behavior of a data block. When we plot all the curves of data blocks for the whole volume, manifest patterns reveal their respective categories of time-varying data. Clustering these curves into different temporal trends brings us a new way to perform classification of the underlying time-varying data. The results of classification can be utilized in transfer function specification to highlight regions with different temporal trends. In this manner, the viewers are able to purposefully focus their attentions on the dynamic features of time-varying data for a clear observation and un-

derstanding. Wit
time anomalies a
tention. Furtherr
budget based on
driven visualizat
key time steps f
ulation using the
practice of unifo
time steps captur

2 Related V

Time-varying da
to the visualizati
ence in time-var
ing was the focu
Westermann [17
partitioning (TS]
temporal coherei

The great adv
for compressing
[4] applied wave
achieved real-tin
ware acceleratio
time-varying dat
tion with a hard
for interactive d
rect rendering of
projection techn
captures space-ti

Transfer func
tensively studiec
done for time-va
tion for time-va
[6]. They cond
and summary vc
ior of time-vary
of transfer funct
togram for simu
a solution that p
equivalence clas

Research clo
(TAC) [3] and th
et al. [3] focusec

Figure 4: Paragraphs on time-varying data

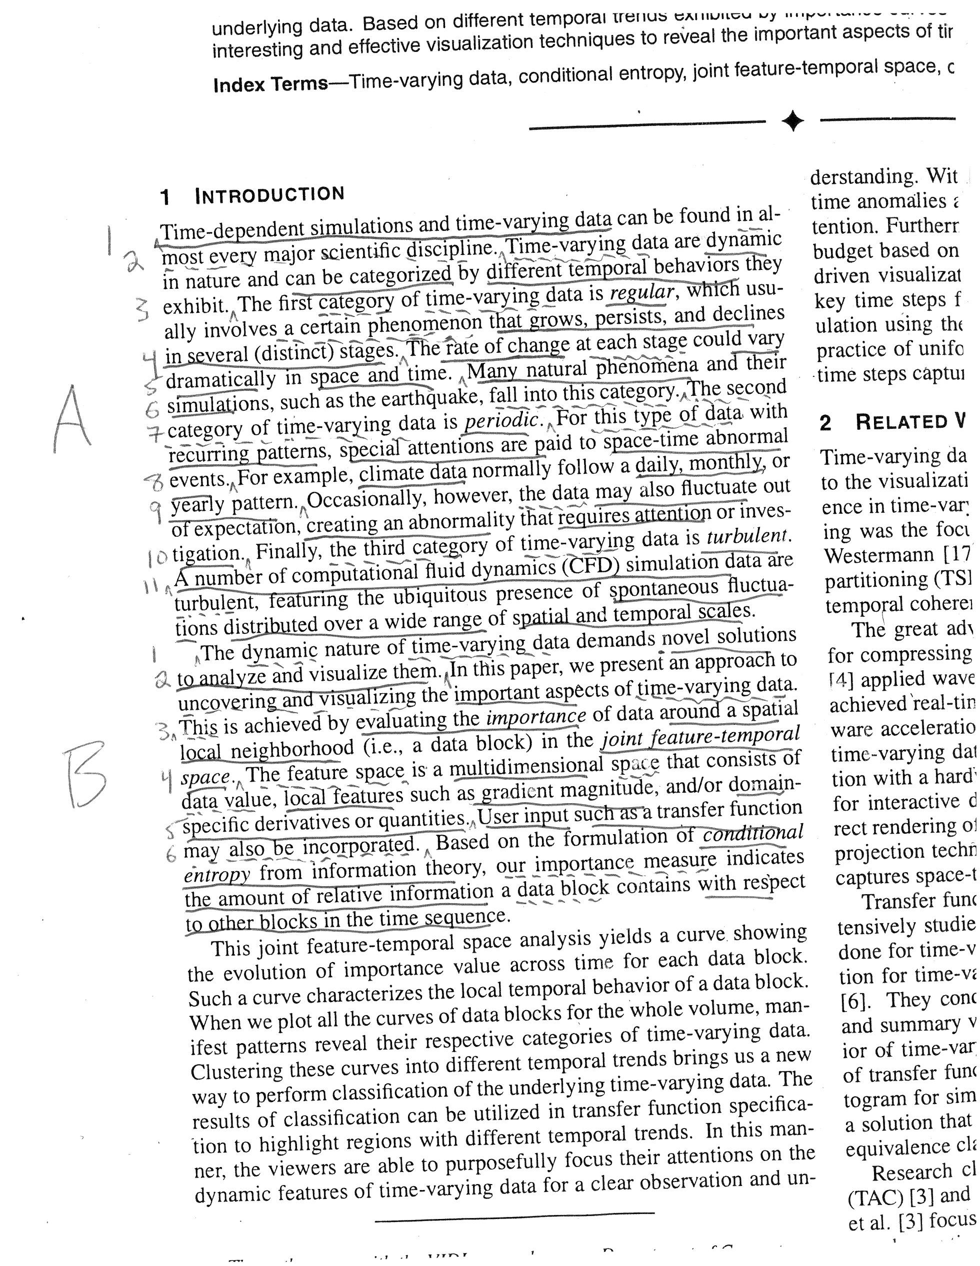

underlying data. Based on different temporal trends exhibited by importance curves ...
interesting and effective visualization techniques to reveal the important aspects of tir

Index Terms—Time-varying data, conditional entropy, joint feature-temporal space, c

1 INTRODUCTION

Time-dependent simulations and time-varying data can be found in almost every major scientific discipline. Time-varying data are dynamic in nature and can be categorized by different temporal behaviors they exhibit. The first category of time-varying data is *regular*, which usually involves a certain phenomenon that grows, persists, and declines in several (distinct) stages. The rate of change at each stage could vary dramatically in space and time. Many natural phenomena and their simulations, such as the earthquake, fall into this category. The second category of time-varying data is *periodic*. For this type of data with recurring patterns, special attentions are paid to space-time abnormal events. For example, climate data normally follow a daily, monthly, or yearly pattern. Occasionally, however, the data may also fluctuate out of expectation, creating an abnormality that requires attention or investigation. Finally, the third category of time-varying data is *turbulent*. A number of computational fluid dynamics (CFD) simulation data are turbulent, featuring the ubiquitous presence of spontaneous fluctuations distributed over a wide range of spatial and temporal scales.

The dynamic nature of time-varying data demands novel solutions to analyze and visualize them. In this paper, we present an approach to uncovering and visualizing the important aspects of time-varying data. This is achieved by evaluating the *importance* of data around a spatial local neighborhood (i.e., a data block) in the *joint feature-temporal space*. The feature space is a multidimensional space that consists of data value, local features such as gradient magnitude, and/or domain-specific derivatives or quantities. User input such as a transfer function may also be incorporated. Based on the formulation of *conditional entropy* from information theory, our importance measure indicates the amount of relative information a data block contains with respect to other blocks in the time sequence.

This joint feature-temporal space analysis yields a curve showing the evolution of importance value across time for each data block. Such a curve characterizes the local temporal behavior of a data block. When we plot all the curves of data blocks for the whole volume, manifest patterns reveal their respective categories of time-varying data. Clustering these curves into different temporal trends brings us a new way to perform classification of the underlying time-varying data. The results of classification can be utilized in transfer function specification to highlight regions with different temporal trends. In this manner, the viewers are able to purposefully focus their attentions on the dynamic features of time-varying data for a clear observation and un-

derstanding. Wit
time anomalies a
tention. Furtherr
budget based on
driven visualizat
key time steps f
ulation using the
practice of unifo
time steps captu

2 RELATED W

Time-varying da
to the visualizati
ence in time-vary
ing was the focu
Westermann [17
partitioning (TS
temporal coheren
The great adv
for compressing
[4] applied wave
achieved real-tim
ware acceleratio
time-varying dat
tion with a hard
for interactive d
rect rendering o
projection techn
captures space-t
Transfer func
tensively studie
done for time-v
tion for time-va
[6]. They conc
and summary v
ior of time-var
of transfer func
togram for sim
a solution that
equivalence cl
Research cl
(TAC) [3] and
et al. [3] focus

Figure 5: Paragraphs on time-varying data, marked with old and new information

Comment: I'd prefer better verbs, e.g., "almost every major scientific discipline simulates time-varying behaviors or gathers time-varying data."

A2. Time-varying data are dynamic in nature and can be categorized by different temporal behaviors they exhibit.

Comment: I didn't care for this sentence. But the information flow is OK.

A3. The first category of time-varying data is *regular*. Regular behavior grows, persists, and declines in several (distinct) stages.

A4. At each stage, the rate of change could vary dramatically in space and time.

Comment: I find this sentence still far from clear. What I believe the authors mean is this:

> Within a single stage, behavior only grows, persists, or declines, and it does so at a relatively constant rate. But the rates of change in different stages can differ from each other dramatically.

Still more work is needed to capture the ideas about space and time.

A5. The regular category includes many natural phenomena and their simulations, such as the earthquake.

A6. The second category of time-varying data is *periodic*. Periodic behavior shows patterns that recur at fixed intervals.

Comment: To improve parallel structure, I've made up the second sentence from what I understood of the text. And I've dropped sentence A7, which I believe is redundant with A9.

A9. But sometimes the data may break the pattern, creating an abnormality that requires attention or investigation.

A8. The periodic category includes climate data, which normally follow a daily, monthly, or yearly pattern.

A10. Finally, the third category of time-varying data is *turbulent*.

A11. Turbulent data feature ubiquitous, spontaneous fluctuations distributed over a wide range of spatial and temporal scales. The turbulent category includes computational simulations of fluid dynamics (CFD).

B1. The dynamic nature of time-varying data demands novel solutions to analyze and visualize them.

Comment: The information flow is not bad, but I hate this sentence. I prefer "The dynamic nature of time-varying data demands novel analyses and visualizations."

B2 & B3. We uncover and visualize the important aspects of time-varying data by evaluating the *importance* of data around a spatial local neighborhood (i.e., a data block) in the *joint feature-temporal space.*

Comment: I couldn't stand the metadiscourse or the verb "present." But even though the sentence contains a lot of new information, I felt I could understand it fairly easily.

B4. The feature space is a multidimensional space that consists of data values, local features such as gradient magnitude, and/or domain-specific derivatives or quantities.

Comments: Although I find the information itself quite vague, I think the information flow *is OK. And I was compelled to fix the disagreement in number.*

B5. The feature space may also incorporate user input such as a transfer function.

B6. Our importance measure indicates the amount of relative information a data block contains with respect to other blocks in the time sequence. The measure is based on the formulation of conditional entropy from information theory.

Comment: In this last sentence, notice how the flow of the paragraph has broken down completely. The new information in B3 is the joint feature-temporal space, and that's the subject of sentences B4 and B5. And when we reach B6, it has disappeared completely! Here's a minor improvement:

> The feature space enables us to assign an importance to each data block. The assignment indicates the amount of relative information each block contains with respect to other blocks that are nearby in feature space. Relative information is measured using the formulation of conditional entropy from information theory.

And here's a version I like even better. My final version eliminates the redundancy of saying both "relative information" and "with respect to," and it says "how much information" as opposed to "the amount of information."

> The feature space enables us to assign an importance to each data block. The assignment indicates how much information each block contains relative to other blocks that are nearby in feature space. Relative information is measured using the formulation of conditional entropy from information theory.

Revision without commentary, as whole paragraphs

Almost every major scientific discipline has time-dependent simulations or time-varying data. Time-varying data are dynamic in nature and can be categorized by different temporal behaviors they exhibit. The first category of time-varying data is *regular*. Regular behavior grows, persists, and declines in several (distinct) stages. At each stage, the rate of change could vary dramatically in space and time. The regular category includes many natural phenomena and their simulations, such as the earthquake. The second category of time-varying data is *periodic*. Periodic behavior shows patterns that recur at fixed intervals. But sometimes the data may break the pattern, creating an abnormality that requires attention or investigation. The periodic category includes climate data, which normally follow a daily, monthly, or yearly pattern. Finally, the third category of time-varying data is *turbulent*. Turbulent data feature ubiquitous, spontaneous fluctuations distributed over a wide range of spatial and temporal scales. The turbulent category includes computational simulations of fluid dynamics (CFD).

The dynamic nature of time-varying data demands novel analyses and visualizations. We uncover and visualize the important aspects of time-varying data by evaluating the *importance* of data around a spatial local neighborhood (i.e., a data block) in the *joint feature-temporal space.* The feature space is a multidimensional space that consists of data values, local features such as gradient magnitude, and/or domain-specific derivatives or quantities. The feature space may also incorporate user input such as a transfer function. The feature space enables us to assign an importance to each data block. The assignment indicates how much information each block contains relative to other blocks that are nearby in feature space. Relative information is measured using the formulation of conditional entropy from information theory.

Experience in the classroom

In analyzing and revising this text, my students were tackling this exercise for the second time. Most of them were able to work effectively with this text, understanding where information flow was OK and where it was not OK. Not all students, however, were confident enough to propose revisions based on their analyses.

One sign that this text worked well was that several readers suggested revisions that I liked better than my own revisions, and that I could also see were faithful to the original text. One such revision was *Almost every major scientific discipline exhibits time-varying data.* As another example, one reader provided a much-needed signpost by following sentence A2 with a new sentence *There are three categories of time-varying data*. As my class promptly identified, this sentence presents the information in the wrong order—the new information is at the beginning—and by correcting the information flow, we arrived at *Time-varying data can be divided into three categories.*

Most of my students spotted the significant difficulties with sentences B5 and B6. Sentence B5 was so unclear that different readers disagreed about its meaning: some readers felt that user input was incorporated in "feature space," and others felt that it was incorporated in some kind of calculation of "importance." And all readers felt the floor drop out from under them when the "feature space," after being introduced in B3 and elaborated in B4, was completely abandoned in B5 and B6.

Even after all our revision, my students and I still found these paragraphs difficult to understand. A good next step would be to improve the parallel structure in the revision. For example, with periodic data, we're told what people look for. What should they look for in regular data and turbulent data?

Finally, everyone wished the authors had chosen nouns and verbs with greater care. The authors use similar but not identical nouns repeatedly in sentences B3 to B6, and we couldn't tell if the authors were just sloppy, or if they varied the nouns to express some sort of meaning. We were left feeling unhappy and like we didn't understand what was being said.

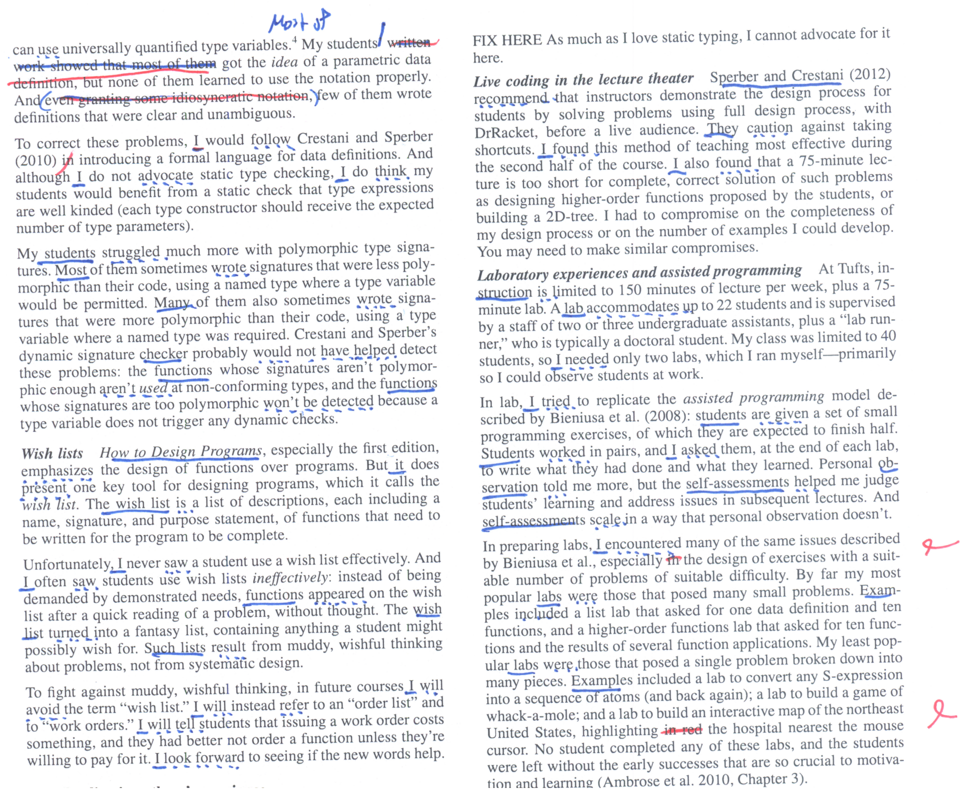

can use universally quantified type variables.[4] My students' written work showed that most of them got the *idea* of a parametric data definition, but none of them learned to use the notation properly. And even granting some idiosyncratic notation, few of them wrote definitions that were clear and unambiguous.

To correct these problems, I would follow Crestani and Sperber (2010) in introducing a formal language for data definitions. And although I do not advocate static type checking, I do think my students would benefit from a static check that type expressions are well kinded (each type constructor should receive the expected number of type parameters).

My students struggled much more with polymorphic type signatures. Most of them sometimes wrote signatures that were less polymorphic than their code, using a named type where a type variable would be permitted. Many of them also sometimes wrote signatures that were more polymorphic than their code, using a type variable where a named type was required. Crestani and Sperber's dynamic signature checker probably would not have helped detect these problems: the functions whose signatures aren't polymorphic enough aren't *used* at non-conforming types, and the functions whose signatures are too polymorphic won't be detected because a type variable does not trigger any dynamic checks.

Wish lists *How to Design Programs*, especially the first edition, emphasizes the design of functions over programs. But it does present one key tool for designing programs, which it calls the *wish list*. The wish list is a list of descriptions, each including a name, signature, and purpose statement, of functions that need to be written for the program to be complete.

Unfortunately, I never saw a student use a wish list effectively. And I often saw students use wish lists *ineffectively*: instead of being demanded by demonstrated needs, functions appeared on the wish list after a quick reading of a problem, without thought. The wish list turned into a fantasy list, containing anything a student might possibly wish for. Such lists result from muddy, wishful thinking about problems, not from systematic design.

To fight against muddy, wishful thinking, in future courses I will avoid the term "wish list." I will instead refer to an "order list" and to "work orders." I will tell students that issuing a work order costs something, and they had better not order a function unless they're willing to pay for it. I look forward to seeing if the new words help.

FIX HERE As much as I love static typing, I cannot advocate for it here.

Live coding in the lecture theater Sperber and Crestani (2012) recommend that instructors demonstrate the design process for students by solving problems using full design process, with DrRacket, before a live audience. They caution against taking shortcuts. I found this method of teaching most effective during the second half of the course. I also found that a 75-minute lecture is too short for complete, correct solution of such problems as designing higher-order functions proposed by the students, or building a 2D-tree. I had to compromise on the completeness of my design process or on the number of examples I could develop. You may need to make similar compromises.

Laboratory experiences and assisted programming At Tufts, instruction is limited to 150 minutes of lecture per week, plus a 75-minute lab. A lab accommodates up to 22 students and is supervised by a staff of two or three undergraduate assistants, plus a "lab runner," who is typically a doctoral student. My class was limited to 40 students, so I needed only two labs, which I ran myself—primarily so I could observe students at work.

In lab, I tried to replicate the *assisted programming* model described by Bieniusa et al. (2008): students are given a set of small programming exercises, of which they are expected to finish half. Students worked in pairs, and I asked them, at the end of each lab, to write what they had done and what they learned. Personal observation told me more, but the self-assessments helped me judge students' learning and address issues in subsequent lectures. And self-assessments scale in a way that personal observation doesn't.

In preparing labs, I encountered many of the same issues described by Bieniusa et al., especially in the design of exercises with a suitable number of problems of suitable difficulty. By far my most popular labs were those that posed many small problems. Examples included a list lab that asked for one data definition and ten functions, and a higher-order functions lab that asked for ten functions and the results of several function applications. My least popular labs were those that posed a single problem broken down into many pieces. Examples included a lab to convert any S-expression into a sequence of atoms (and back again); a lab to build a game of whack-a-mole; and a lab to build an interactive map of the northeast United States, highlighting in red the hospital nearest the mouse cursor. No student completed any of these labs, and the students were left without the early successes that are so crucial to motivation and learning (Ambrose et al. 2010, Chapter 3).

Figure 6: A text with subjects and verbs labelled

5.3 A text labeled for finding coherent subjects

Figure 6 shows an excerpt from a late draft of one of my own papers, *On Teaching* How to Design Programs*: Observations from a Newcomer*. In blue ink, I have labeled the subject and verb of each sentence. (There are some other revisions, mostly in red.) I was looking for opportunities to improve the coherence of the subjects, as described in Exercise E. Many paragraphs show relatively coherent subjects; some do not. If you wished, you could compare this text with the final version of the paper as it appeared in ICFP 2014.

5.4 A text for applying principles

Here are two paragraphs from a beginning student's draft, marked up for use in an exercise. Such a text might be suitable for Exercises A, C, D, E, or F. I believe I chose this particular text for Exercise A.

Paragraph A

1 One solution to this problem and related problems associated with
the impact of large trades is the use of so-called "upstairs markets"
in which broker/dealers shop around large orders to others without
2 revealing the details of the transaction to the public. Market makers
conduct large transactions by mutual agreement, often at a discount
to the prevailing market price, without having to use the more obvious
3 "downstairs markets." Both upstairs and downstairs markets suffer
from the problem of "front-running," in which someone who knows
about an upcoming order is able to execute a trade that takes advan-
4 tage of that knowledge. Sometimes, this is illegal, when the front-
running is based on inside information; in other cases, it is merely
parasitic, when one trader recognizes the habits of another trader or
sees an incoming large order and is able to act quickly to place a
5 trade before it hits the marketplace. Quite recently, important mar-
ket businesses have been investigated and charged with violations of
SEC regulations due to illegal front-running practices.

Paragraph B

1 In upstairs markets, trading large blocks of stock can take more time
and is more prone to negotiation; fewer participants are able to eval-
2 uate the transactions. This ultimately may lead to less competitive
3 prices for those buying or selling. Our solution to this problem is the
use of a homomorphic encryption scheme, similar to that described
by Paillier (1999), to preserve the secrecy of the details of investors'
limit orders while proving that the market is operating according to
4 its published rules. For each order, we encrypt the price and the
5 quantity and place them in an open marketplace. When two orders
"meet," i.e., a buy and sell order have compatible price (buyer's bid $\geq$
seller's ask), the market notifies the traders and the orders are removed
6 from the marketplace. An appropriate clearinghouse completes the
transaction.

5.5 Papers from the literature

You can pick almost any paper and learn something useful. But if you want a recommendation, here are two:

- Xavier Leroy's 1994 POPL paper "Manifest types, modules, and separate compilation" beautifully illustrates many of the principles.
- I hesitate to say this in public, but if you want good ideas badly described, try David Parnas's classic paper "On the Criteria for Decomposing Systems into Modules."

5.6 Discussion questions for writing practices

This section contains some discussion questions I have used in meetings focused on practices of successful writers.

Mindful ways of writing Here are questions I asked of students who had read Boice (2000), Introduction to Part II: *Writing in Mindful Ways*, and Chapter 9: *Wait*.

1. On pages 103–105, Boice enumerates eight reasons why people have difficulty with writing. Which of these are yours? What others are missed? What, at present, is the greatest impediment to your writing?
2. On pages 111–112, Boice enumerates six practices of successful writers. He then restates these practices as bullet points. Do you apply any of these practices in your own work? If there are practices you do *not* apply, which one do you find most appealing?

   (This question may be difficult because the practices are not so easy to understand.)
3. What role do your emotions play in your ability to write?
4. I find the advice in Chapter 9 (*Wait*) to be among the most difficult to understand of all of Boice's book. Can you restate it in your own words? It may help to state the advice negatively: that is, what is it that Boice suggests you *stop* doing?

The questions above are more than enough for a 75-minute session. But here are some more:

5. Do you have a "writing site?" If so, what do you do when you first arrive at your writing site?

   Do you have distinctive tools or materials? Do these help? Do you think they would help?
6. Does your writing suffer from "prolonged procrastination followed by rushed beginnings under deadlines?"
7. How might you apply "precommitment" and visualization to your own writing?

Active waiting In a later class, I asked students also to read Boice (2000), Chapter 9 (*Wait*), and then to practice 5 to 10 minutes of active waiting every day for twelve days, while taking daily notes in a lab notebook. I then followed up with a class discussion based on this handout:

> Our productivity and job satisfaction are profoundly affected by the feelings and thoughts that come to us as we prepare to write.

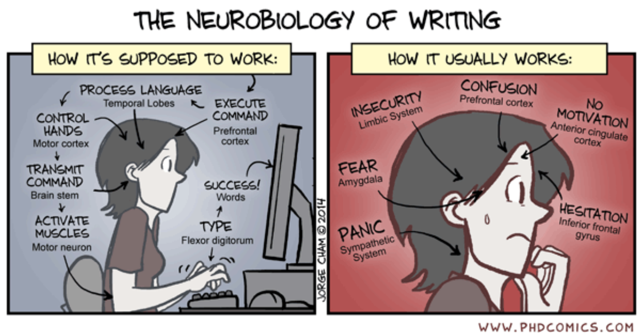

Figure 7: Feelings of student writers

> Please consult your lab notebooks and your experience and answer these questions:
>
> 1. In your time spent actively waiting, what feelings came to you?
> 2. What feelings do you imagine Bob Boice would hope would come to you?
> 3. In your time spent actively waiting, what thoughts came to you?
> 4. What thoughts do you imagine Bob Boice would hope would come to you?
> 5. On a path to becoming a fluent and productive writer, where does active waiting fit in?

My students' experience was very diverse. Many of them were anxious about the exercise and uncertain about the purpose. Some students reported being calm. Some students reported being both anxious and calm—presumably at different times.

During discussion, students reported both that 5 to 10 minutes was far too short a time to accomplish anything, and also that they were far too busy to make time for active waiting every day. Wiser students prevailed, and they eventually concluded that active waiting might indeed be useful, but that the issue is you have to *plan* for it. Simple, but not easy.

Some students figured out that Bob Boice would probably hope that they would get a few ideas about their writing. And one lucky student actually got some new ideas for some experiments to try in his research!

Most students were unclear about how active waiting might fit into a writer's practice, but a few guessed correctly that it can be used effectively before tackling the first draft of a new section or paragraph.

Most interestingly, during discussion *no* students referred to any practices or experiences of mindfulness: no awareness of the breath, no awareness of tension in the shoulders or face, no awareness of uneasy feelings in the gut. And therefore, no awareness of how these things might change through a few minutes of active waiting.

References

Martha Beck. 2003 (July). Ready... Aim... Oh, well.... *Oprah* magazine.

> This short article suggests ways of overcoming (or working around) one's perfectionism. It may be useful for writers who have trouble producing.

Howard S. Becker. 1986. One right way. Chapter 3 of *Writing for Social Scientists: How to Start and Finish Your Thesis, Book, or Article*, Chicago Guides to Writing, Editing, and Publishing. University Of Chicago Press.

> This chapter debunks a few myths and contains helpful advice about writing practices.

Robert Boice. 2000. *Advice for New Faculty Members.* Allyn & Bacon.

> This book is full of priceless data and practical suggestions about what sorts of behaviors characterize writers who are successful, productive, and take pleasure in their work. Unfortunately, I find much of the material poorly presented or difficult to understand. It nevertheless repays careful reading. (Boice's older title, *Professors as Writers*, is well reviewed, but I have not read it.)

Joan Bolker. 1998. *Writing your dissertation in fifteen minutes a day: A guide to starting, revising, and finishing your doctoral thesis.* Macmillan.

> Bolker explains many of the same writing behaviors and habits that Boice (2000) does, but in a form specialized for doctoral students. As you might guess from the title, Bolker, like Boice, advocates for brief, daily sessions. If you find Boice difficult to read—full of jargon and written from a strange social-science point of view—then you might want to try Bolker. If you can translate her prescriptions from the doctoral setting to your own setting, or if you happen to be writing your doctoral dissertation right now, you will probably value her recommendations.

Kenneth A Bruffee. 1999. *Collaborative Learning: Higher Education, Interdependence, and the Authority of Knowledge.* Second edition. Johns Hopkins University Press.

> This book is somewhat controversial for its views of how people learn and what higher education really means, but I have found it invaluable for its suggestions about how to teach discussion classes.

Chicago Editorial Staff. 1993. *The Chicago Manual of Style: The Essential Guide for Writers, Editors, and Publishers.* University of Chicago Press, fourteenth edition.

> I find this edition more readable than the fifteenth. The fifteenth edition does have a lot more information about citing electronic sources.

Lyn Dupré. 1998. *BUGS in Writing, a Guide to Debugging Your Prose*. Addison Wesley Professional, revised edition.

> Though I can't fathom why, many computer scientists love this book. I much prefer Fowler (1968) or Garner (2003).

H. W. Fowler. 1968. *A Dictionary of Modern English Usage*. Oxford University Press, second edition. Revised by Sir Ernest Gowers.

> Fowler's masterpiece is the original on which many later guides are modeled. Avoid the evil, permissive third edition.

Bryan A. Garner. 2003. *Garner's Modern American Usage*. Oxford University Press, second edition.

> In my view, Garner has the best claim to be a legitimate successor to Fowler.

Billy Vaugh Koen. 1985. *Definition of the Engineering Method*. American Society for Engineering Education, Washington, DC.

> Koen has written extensively on "the engineering method," which he identifies with the use of *heuristics*. All the principles and practices I teach are heuristics.

Kenneth K. Landes. 1966 (September). A scrutiny of the abstract. *Bulletin of the American Association of Petroleum Geologists*, 50(9):1992–1993. Accompanied by The Royal Society's "Guide for Preparation and Publication of Abstracts.".

> Writing an abstract is an important but specialized skill. This one-page article explains what good and bad abstracts look like. Don't miss the author's lament about his science.

Steven Pinker. 2014. *The Sense of Style: The Thinking Person's Guide to Writing in the 21st Century*. Viking.

> Once you've got what you can from Williams (1995), this is a good next book. Don't be put off by the first chapter, which is a bit strange. From Chapter 2 on, Pinker recapitulates quite a bit of what you'll find from Williams, but you'll also find new things. About you and your readers, you'll get more about how you might be trying to affect your reader, and how the "curse of knowledge" makes us say things in ways that we understand, not necessarily in ways that readers will understand. You'll get a nice chapter on the value of grammar and syntax, which is really interesting because the notions of grammar have been updated to account for what's known today about linguistics. Good, accessible writing about grammar is always hard to find. Finally, you'll get yet another dictionary of usage—130 pages' worth. This one seems to have the right attitude about what "correctness" really means, so it may be better than Garner's (2003), but I haven't evaluated it to make sure.

Joseph M. Williams. 1981. The phenomenology of error. *College Composition and Communication*, pages 152–168.

This brilliant paper supports my contention in Principle 0 that you don't have to pay much attention to the rules your high school English teachers taught you. If you've been browbeaten about rules, this paper will give you the ammunition you need to fight back—especially if you make it all the way to the end.

Joseph M. Williams. 1995. *Style: Toward Clarity and Grace*. University of Chicago Press.

Repeatly described as a "master teacher" of writing, Williams has written widely on the topic. Of his works, this volume is among the most accessible and reasonably priced. And it's better than some of the later books written in collaboration, especially the posthumous ones.

Acknowledgments

For ideas, suggestions, and contributions to an early draft of this document, I'm very grateful to Kim Cooper. Among other tricks, Kim taught me to ask students for papers they had read and disliked. I also thank the Harvard Writing Project for supporting this work with a Gordon Gray Faculty Grant for Writing Pedagogy, and Dean Venkatesh Narayanamurti for additional financial support.

Marva Barnett introduced me to Bob Boice, and Christian Lindig introduced me to Joseph Williams. Without them, I wouldn't have anything to teach.

I learned Principle 2 (*Singular*), which I use often, from an anonymous referee.

In their Course Design Institute of January 2014, the Tufts Center for the Enhancement of Learning and Teaching gave me a number of intellectual tools which I have used to clarify my goals and to sharpen the exercises.

Finally, I am extremely grateful to the many students who made it possible for me to go from a bunch of vague ideas to a teaching method that works—and whose writing convinced my colleagues that I wasn't wasting everyone's time. I can't name them all, but I want to thank by name each of the research students who agreed that, as a condition of working with me, they would attend writing group once a week: João Dias, Paul Govereau, David Hiniker, Andrei Lapets, Clayton Myers, Reuben Olinsky, Daniel Peng, and Kevin Redwine. Not one of them complained.

Afterword

I wrote the first version of this booklet in 2006. In 2011 I changed a few things, of which the important was a title that might have frightened the horses in the street. In 2014, as part of improving a course to be taught at Tufts, I made a few more changes, mostly to clarify goals and objectives, but also to add Principle 4. During that class, I developed the summary in Figure 2 on page 17, which I added to this booklet in 2016.

Even the original booklet was a long time coming. I would love to explain the material in more depth and (of course) craft the writing to a higher degree of clarity. But the best is the enemy of the good, and I chose to get the thing out the door. Although there are some significant gaps—I would especially like to write an explanation of each of the practices in Table 1—I believe that the booklet will be useful for its principles, practices, and exercises. If you find it so (or not), please drop me a line.