

Writing Tips

The University of Toronto [U.of.T] expects its students to write well, and provides numerous resources to help them, many of which are identified on the “Writing at the University of Toronto” website:

<http://www.writing.utoronto.ca/>

I recommend that students familiarize themselves with this site, which presents a wealth of material, describes the highly regarded U.of.T writing centres, and lists a variety of writing courses available free to all students. As faculty members, we are not in a position to *teach* students how to write well, but it is our responsibility to ensure *that they do so*. Also, an ability to write well is an incredible asset in almost any conceivable career. It is students’ responsibility to take advantage of all of the resources provided to them, to ensure that they reach this goal.

The present document should be viewed as a personal—and rather idiosyncratic—complement to those general resources. It does not attempt to replace any of the general aids provided on the website. Rather, it merely enumerates a number of tips and rules of thumbs relevant to the most common rhetorical and stylistic issues that I have encountered in students papers over the years. You should read through this once, before writing any papers—and then refer back to it, as necessary, as a reference. In grading papers, I may refer to points in this document by label—e.g., “B.4” if you have mistakenly used ‘which’ in a context requiring ‘that.’

I • Style

- A. **Brevity:** Be clear and concise. Say what you mean, sans drama or frills. Write plain prose. Don’t worry, especially at first, if the sentence structure seems boring. In professional writing, the content should carry the excitement, not the form.

It is said that Nadia Boulanger, the legendary Parisian composition teacher, was in the habit of throwing open her students’ compositions, pointing to a note selected at random, and demanding: “Why does that note need to be here?” This is an excellent discipline to apply to one’s writing. Every word should play an essential role. It is amazing how quickly loose or wordy phrasing can lose the reader’s allegiance. Even a single verbose phrase, in an otherwise well-formed sentence, will let the reader’s attention drift, or—worse—lead them to put the piece down.

- B. **Get to the point:** Start directly on topic. You don’t need a general introduction. (Academic and professional readers are easily put off by unnecessary introductions and generalizations.) Especially in an academic context, but in many professional contexts as well, you should write for someone who knows the literature, and assume all necessary general background.

For example, suppose an early draft of your paper were to start out as follows:

“For centuries people have attempted to construct technological artifacts that can speak and understand human language. The situation is no different with computers, and for dec-

ades numerous projects have been undertaken to develop computational systems that understand natural language. These projects have had to deal with daunting challenges. It turns out that human language is bewilderingly complex. It is often understood in terms of three aspects or dimensions: *syntax*, *semantics*, and *pragmatics*. It is relatively clear why semantics and pragmatics are hard to build into a machine, but even syntax has proved difficult to automate. A central difficulty is pronoun resolution. Two kinds of example show that pronoun resolution can depend on arbitrary real-world knowledge. ... <discussion of the examples>”

In a paper of anything up to several thousand words, *the first seven sentences of this paragraph should all be deleted!* All that should be retained is the final sentence, which with the paper should start:

“Two kinds of example show that pronoun resolution can depend on arbitrary real-world knowledge...”

- C. **Foreshadowing:** In a short paper, don’t summarize in the first paragraph (i.e., don’t say things like “In this paper I will first do α , then say β , and conclude by arguing γ ”). Explicit foreshadowing—traditional in scientific papers, though less so in the humanities—becomes appropriate only in longer papers or books. In an essay of only a few thousand words, *every word is precious*. Save them for the substantive content, for your own thinking and intellectual contribution.
- D. **100 stitches per inch:** In my experience, first drafts (including my own!) are often of about the right length overall—but nevertheless much too wordy, locally. The reason this is not a contradiction is that too few points are made, but those that are made are loosely framed, and too disconnected. The goal is almost exactly the opposite: to have the argument itself be very fine-structured, but to have the prose that conveys the argument do so in a very compact way.
- As a way to think about this, imagine your argument as leading its reader on a journey or path—say, “100 feet” in length in some conceptual space—by providing a series of stepping-stones. Suppose your paper is to be 2,000 words long. The wrong thing to do is to make ten points, each one ten feet from the last, spending 200 words on each. Ten feet is too large a step for anyone to take on their¹ own. Instead, you should provide much more closely-spaced stepping stones (i.e., more detail about the structure of the argument), by saying exactly what follows from what, what depends on what, what assumptions you are relying on, etc. If you do this well, you will end up with, say, 100 points—but with each one separated from the previous one by only (as it were) only a foot or so. If you did that while still employing 200 words per step, the paper would be 20,000 words—massively too long! But with careful writing, you can usually express each step or point in, say, 20 words, instead of 200. The net result would then be a paper of 2,000 words—the same length as the original.
- While both versions would be 2,000 words in total, in other words, the initial one made 10 points using 200 words for each, while the final version would make 100 points using 20 words each. The conceptually fine-textured version will be hugely better. Not only will it be more interesting to read; it will also have a far better chance of being conceptually sound—since you are more likely, while writing it, to notice any conceptual slips, non sequiturs, and outright mistakes.
- E. **First person:** Avoid first-person *content*: “*I have been thinking about ...*”, “*I learned from reading Ethelred’s book on consciousness that ...*”, etc. That is not to say that all first-person *pronouns* must be excised. The latter stricture is out of date, and can lead to prose that, to a contemporary

¹See B.2, below.

reader, comes across as pretentious and awkward. For example, it might be entirely appropriate to note an assumption by saying: “Human consciousness, I believe, requires no more of a substrate than a physically instantiated brain.” But note that the *topic* of that sentence is human consciousness, not your experience. That’s a good diagnostic test: judicious uses of the word ‘I’ are acceptable, but *never make yourself, or your personal experience, be the subject matter*.

- F. **Retrospective clarity:** A rhetorical style in which some people specialize is something I have come to call *retrospective clarity*: prose that is perfectly clear—does an excellent job of stating and summarizing the point—to someone who already understands the point, but that manages to convey nothing whatsoever to someone who does not already understand! The trouble in such cases is that one needs to understand the content, in order to understand the argument, and the argument in order to understand the language. What one wants is the opposite. Make your writing clearly comprehensible to someone who does not know what you are talking about, so that, by reading and understanding it, they can come to learn and comprehend the subject matter and argument—and through that, possibly even come to agree with your overall point.

In sum: the reader should understand the point by understanding the prose—not the other way around.

- G. **Struggle:** By and large, in academic prose,² it is good to convey a sense of the struggle or wrestling that you have done in coming up with your position. This is a tricky: a simple story about what you thought en route is not what anyone is looking for. Nor does the reader want to learn that you had to struggle to get up to speed on the literature, or to understand other people’s arguments. That is: the struggle conveyed should not be a record of your own personal travail. Rather, the idea is to convey the intellectual struggle that underlies the conclusions for which you are arguing. First-rate academic papers often bear elegant traces of that kind of conceptual sweat and blood.

II • Grammar and Usage

- A. **On usage in general:** By far the most important lesson about grammatical, stylistic, and rhetorical quandaries is the following: *there isn’t always a solution*. Some people think that English is—or should be—as well defined as logic or a programming language, and that there must be “a correct way” to handle all situations. But that is simply not so. Often—far more often than most new writers suspect—the proper response to a question about what to do in a tricky circumstance is: “Don’t go there! Rewrite!” For example, neither of the following is any good:

- a. (✗) Neither John nor his children is coming tonight.
- b. (✗) Neither John nor his children are coming tonight.

Rewrite! Something like the following would serve equally well:

- c. (✓) I don’t think we will be seeing either John or his children tonight.

- B. **On gender-neutral personal pronouns:** What does one say: *The person raised his hand? raised his/her hand? When (s)he comes, give her/him a tip?* The topic of gender-neutral personal pronouns is nothing if not contentious. Contrary to what is often taught, my recommendation is to use ‘they’ as a grammatically plural but semantically singular personal pronoun—e.g., as in “everyone should pick up their paper after class.” Here is how I described the issue in a book I once wrote:³

“Throughout, I will use ‘they’ and ‘them’ as syntactically plural but semantically singular third-

²This may not be so true in other professional contexts.

³*On the Origin of Objects*, MIT Press, 1996, fn. 3, p. 29.

person personal pronouns of unmarked sex. It seems the best alternative. ‘He/she’ and ‘his or her’ are irretrievably awkward, and it seems unlikely that any entirely new or invented word will ever be satisfactory. And there is precedent. In the second person case, we are entirely accustomed to using the syntactically plural ‘you’ to convey semantically singular as well as semantically plural meanings. In formal writing, the syntactically plural ‘we’ is even sometimes used as a stylized and somewhat detached form of singular first-person reference (though in this book I will use the singular ‘I’). So using ‘they’ fits into a general pattern of employing the plural form when pointed, individual reference is not justified. It appears moreover that exactly this pattern is naturally evolving in informal speech.

None of this is to deny that ‘the painter picked up their brush’ sounds awkward and informal— to say nothing of ‘the barber shaved himself.’ The awkwardness may pass, however, and anyway informality is better than artifice.”

This usage is also the recommendation of Geoffrey Nunberg, a long-time usage editor of the *American Heritage Dictionary*, and perhaps America’s preeminent contemporary authority on language use.

- C. **On anaphoric reference:** Make sure that all anaphoric referents (e.g., uses of ‘he’, ‘she’, ‘they’ and ‘it’, bare uses of ‘this’, etc.) are perfectly clear, semantically. Never carry anaphora across a section break. Someone can put a paper down, overnight, and pick it up the next day and begin reading a new section, expecting it to be clear. They don’t want to start out by reading a sentence that says: “The reason this is important is...”
- D. **On ‘which’ vs. ‘that’:** The following is a useful rule of thumb: use ‘which’ for non-restrictive relative clauses, ‘that’ for restrictive ones. *Restrictive* means that the subsequent qualifying phrase is necessary in order to identify the referent; *non-restrictive* means that it is not necessary, but provides additional information about that already-identified item. Thus consider the following two instructions:
- a. Go in and knock on the first door on the right that is red.
 - b. Go in and knock on the first door on the right, which is red.

If the first door on the right is blue, and the second is red, then (a) is fine, and means that you should knock on the second door. In the same circumstances, however, (b) is simply false.

Some notes:

2. The relative-clause-introducing ‘that’ always converts to ‘which’ if it is preceded by another ‘that’. Thus (ii) is correct, not (i):
 - a. (✗) He always destroys that that he loves.
 - b. (✓) He always destroys that which he loves.
3. You can’t always tell from a sentence itself whether the clause is restrictive or not. Thus consider:
 - a. Can you hand me the book that is on the coffee table by your knee?
 - b. Can you hand me the book, which is on the coffee table by your knee?

The second (ii) would be correct if the previous discourse or surrounding context had fixed the reference of ‘the book,’ and the fact that it is on the coffee table is additional information, to help the person find it. The first would be correct if the identity of the book being referred to is being established by that locational fact.

- E. **On split infinitives:** I can do no better than to quote Geoff Nunberg’s usage note in the *American*

Heritage Dictionary:

“The split infinitive has been present in English ever since the 14th century, but it was not until the 19th century that grammarians first labeled and condemned the usage. In the 20th century many linguists and writers have rallied to its defense. H.W. Fowler chided that class of people “who would as soon be caught putting their knives in their mouths as splitting an infinitive,” but whose aversion springs only “from tame acceptance of the misinterpreted opinion of others.” No plausible rationale has ever been advanced for the rule, though it may arise from a hazy notion that because the Latin infinitive is a single word, the equivalent English construction must be treated as if it were indivisible.

“Still, many people who dislike the construction avoid it without difficulty. The sense of the sentence *To better understand the miners' plight, he went to live in their district* is just as easily expressed by *To understand the miners' plight better, he went to live in their district*. In some cases avoidance of the split infinitive may result in a stylistic improvement. The sentence *We are seeking a plan to gradually, systematically, and economically relieve the burden* becomes clearer if the adverbs are placed at the end: *We are seeking a plan to relieve the burden on our employees gradually, systematically, and economically*. (In an earlier survey the example having the split infinitive was accepted by only 23 percent of the Usage Panel.) But in other cases the effort to avoid a split infinitive may have unfortunate consequences. In *The tenant coalition is planning to aggressively seek cooperative ownership of the apartments the city acquired*, any attempt to reposition the adverb *aggressively* would create an ambiguity. In *We intend to use every political favor we are owed to soundly defeat this bill and its riders*, any other position will create an unnatural rhythm. In *We expect our output to more than double in a year*, the phrase *more than* is intrinsic to the sense of the infinitive phrase, though the split infinitive could be avoided by use of another phrase, such as *to increase by more than 100 percent*. In this example the split infinitive is accepted by 87 percent of the Usage Panel.

“Excessive zeal in avoiding the split infinitive may result in an unnecessarily awkward placement of adverbs in constructions involving the auxiliary verbs *be* and *have*. When we read sentences like *I want this clearly to be understood*, we may suspect that the placement of *clearly* is the result of an effort to avoid the construction *to be clearly understood*, under the misapprehension that the latter involves a split infinitive. By the same token, there are no grounds for objecting to the position of the adverb in the sentence *He is committed to laboriously assembling all of the facts of the case*. What is “split” here is not an infinitive but a prepositional phrase.”

- F. **On strong verbs:** Especially during the last fifty years, people have a tendency to use weak verbs, followed by a nominal or prepositional phrase. Thus: *they ate dinner*, rather than *they dined*; *they made a start* as opposed to *they began*; etc. In writing (see A.1, above), it is strongly recommended that you use strong verbs. For example: “she distinguished three kinds of beech,” as opposed to “she made a distinction among three kinds of beech” (three words saved!). Similarly: “they apologized”, not “they made an apology” (two more!).
- G. **On adverbs and adjectives:** It is good writing practice to “lose” (i.e., delete) all adverbs and adjectives. Needless to say, this is sometimes impossible. Interestingly, however, it is possible far more often than one might think. And it is beneficial! The resulting prose is typically more vivid and compelling, not less.

As a strategy to achieve this, consider using a stronger main word (noun or verb) that builds in the qualifying meaning you have “outsourced” to the modifier. So: ‘crimson’ instead of ‘deep red’; ‘thunderously’ instead of ‘very very loudly’—and then perhaps lose the ‘thunderously’, too, as in “he thundered” in place of ‘he spoke thunderously’ (itself in place of ‘he spoke very very loudly’). Then, once the intensity of the words has been increased, and their number correspondingly reduced, see whether the sentence doesn’t perfectly well convene the meaning with no qualifiers at all. Keep the qualifiers if you absolutely must; but never hang on to them (§E.3).

III • Spelling and Punctuation

- A. **On spell checking:** Check all spelling! Spelling mistakes make it seem as if you didn’t take your paper seriously—so why should its readers?

Note: all word processors contain spelling checkers, which are invaluable. But do not rely on them alone. There are too many kinds of spelling mistakes they cannot catch (because they neither parse nor understand). Errors of the sort that get past automatic spelling checkers are much more common in student papers today than they were 20 years ago—including all three in the following single sentence:

- a. (✗) “Its true that their eating desert.”

- B. **On it’s vs. its:** ‘It’s’ is a conjunction, short for ‘it is’. ‘Its’ is a gender-neutral possessive pronoun. Thus: “It’s blowing its top.”

- C. **On the use of italics:** As everyone knows, spoken language allows one to emphasize (put more stress on) certain constituents of a sentence, thereby affecting the meaning. Thus consider how different are the connotations of the following one sentence, when spoken with four different stress patterns (the first is to be read with approximately even intonation):

- a. Pat didn’t really want ice cream.
- b. *Pat* didn’t really want ice cream.
- c. Pat didn’t *really* want ice cream.
- d. Pat didn’t really want *ice cream*.

The meaning is obviously different: (b) suggests that whereas Pat didn’t want ice cream, someone else did; (c), that Pat said or indicated that she wanted ice cream, but wasn’t being truthful; and (d), that Pat wanted something, but whatever it was, it wasn’t ice cream.

When writing, it is natural to press italics into service as a proxy for oral stress, to indicate that certain parts of a sentence should be emphasized, in order to convey different meanings. But compelling as this may seem at the time of writing (since most people “think” language complete with intonation), it is not a good idea. Papers written with a lot of italics are not much fun to read; they feel as if the author has taken you by the collar and is talking too loudly, hauling you in close and shouting in your face. A better strategy—the one that professional writers use—is to accept the fact that written language does not mark stress, and to compensate by using more complicated syntax to “front” or privilege different parts of the sentence, so as to achieve the same underlying semantic intent.

Thus the meanings conveyed by the italicized sentences above would be better written as follows, in italics-free variants (though note how pedantic it would normally be to speak these versions):⁴

⁴Do you know people who speak written language?

- b'. It wasn't Pat who wanted ice cream.
- c'. Pat's claim to want ice-cream was disingenuous.
- d'. Ice cream wasn't what Pat wanted.

Many people realize that written language is typically more syntactically complex than spoken language, but fail to recognize the reason why. Syntactic compensation for the absence of intonation is one of the primary reasons why this is so.

- D. **On punctuation and quote marks:** Purely logically, if the final word in a sentence is quoted, but the whole sentence is not quoted, it would seem as if the ending quote mark should precede the final period. Purely logically, that is, it would seem as if α would be the “correct” punctuation, as opposed to β :

- α . He called it “post-hoc rationalization”.
- β . He called it “post-hoc rationalization.”

On so-called monospaced or “fixed-width” fonts (fonts like Monaco and Courier, and those on old-fashioned typewriters, in which each letter is given the same amount of horizontal space on a line), the logical answer is also the correct answer. So if typing on a typewriter, follow α :

- α . (✓) He called it “post-hoc rationalization”.

According to professional typographers, however, the situation changes with so-called *proportionally-spaced* (or “variable-width”) fonts—that is, fonts that allocate different amounts of horizontal space to each character, depending on how wide it is (i.e., so that the character ‘i’ ends up using quite a bit less space than the character ‘w’). In such cases, which often make use of sophisticated kerning⁵ algorithms, typographers agree that the final period or comma should be “moved inside” the quote marks. Thus, in a proportionally-spaced font, β is correct, rather than α . This is the convention you will find enforced if you publish a book or article in a professional journal.

Why contravene logic? The typesetters’ answers are roughly these:

- Printed periods and commas in proportionally-spaced fonts are so small that, if they are “out on their own,” rather than nestled against a character, the eye tends to treat them as specks or bits of dirt. Naively, we often forget that a period is a contextual entity. What typographers emphasize is the very specific nature of that contextuality: to be (recognized as) a period is to be a small dot nestled to the right and towards the bottom of an alphanumeric character.
- When a period or comma is brought inside a quote mark, the subsequent character kerning means that, in the final output, the period and the quote mark will be almost on top of each other, vertically. As illustrated in the 300% enlargements in Figure 1 (above), whereas the monospaced version of β (β_m) puts the quote mark decidedly to the right of the period, in a proportional version (β_p) it ends up barely to the right at all. In fact it is almost directly above the period. It is evident that β_p lacks β_m 's manifest awkwardness.




Figure 1 — Monospaced vs. proportional printing

⁵Kerning is a process of adjusting the spacing between adjacent characters in a proportionally-spaced font so as to meet aesthetic guidelines and generally improve the look of the printing.

As regards the two versions of α , even though the period is closer to the ‘n’ in the proportional version (α_p) than in the monospaced version (α_m), it is still far enough away that it may trip up our perception. Even if more “logical,” α_p still doesn’t read as smoothly as β_p .

3. Human reading processes are not only automatic and subconscious, but proceed in parallel. Given (1) and (2) above, therefore, there isn’t really much psychological reality, especially in β_p , to the fact that the period is to the right of the quote mark; really, it is read almost simultaneously.

I don’t know what you make of these arguments. When I was young I was a strict rationalist, and thus preferred α over β . But as so often happens when one encounters real experts, working with professional typesetters has convinced me that they were right, and that I was both intransigent and naïve.

It is very rare that any of us any longer use monospaced fonts (except for programs and code). So as a general rule, β (i.e., β_p) is the way to go.

Two final comments:

4. It is only periods and commas that “move in inside the quote marks” in this way; semi-colons, colons, question marks, and exclamation points all stay outside.
5. In some circumstances, sense dictates that the foregoing rules should be over-ruled. Three obvious examples: (i) when spelling itself is the subject matter; (ii) when a single character is being quoted; (iii) in the case of URLs. Thus the left-hand version of each of the following is clearly to be preferred over the right-hand one:

<p style="text-align: center;">✓</p> <p>Your password is ‘cat23-45’.</p> <p>The second character is ‘m’.</p> <p>Go to “http://www.uoft.ca”.</p>	<p style="text-align: center;">✗</p> <p>Your password is ‘cat23-45.’</p> <p>The second character is ‘m.’</p> <p>Go to “http://www.uoft.ca.”</p>
---	---

- E. **On inter-sentential spacing:** It is a convention, when using monospaced fonts (e.g., of the sort used on old-fashioned typewriters—see C.4, above) to use two spaces, not just one, between sentences. This is because the inter-word and inter-character spacing on typewriters is in general so “loose” that using just a single space between sentences doesn’t sufficiently differentiate sentence breaks visually, making things hard to read. But in proportionally-spaced fonts, kerning makes inter-word and inter-character spacing much tighter—and so the extra space between sentences is no longer needed (and hence should not be used).

Again, since we almost always use proportionally-spaced fonts these days, use only a single space between sentences.

- F. **On ‘i.e.’ and ‘e.g.’:** The abbreviations ‘i.e.’ and ‘e.g.’ are both always followed by a comma (except when they occur at the end of a sentence, which is rare). Thus:
1. Ontogeny—i.e., the development of the embryo into an adult—recapitulates phylogeny.
 2. Some names—e.g., Pat, Randy, and Lesley—are used by both men and women.

G. **On commas**

1. Long nominal subject phrases are not followed by a comma, even if one would pause there if speaking the sentence. Thus the following is incorrect (the comma should be removed):
 - a. (✗) The argument that came up in class last week about whether brains are conscious, is illuminated by recent neuroscientific results.
2. If you start a sentence with “but”—not often a good idea, but occasionally OK (especially in in-

formal contexts)—it is not followed by a comma, in the way that ‘however’ is. Thus:

- a. (✓) However, we should leave by 2:00.
 - b. (✓) But we should split by 2:00.
3. Non-restrictive relative clauses are typically preceded by a comma, whereas restrictive ones should not be:
- a. I bought that book you told me about, which cost \$29.95.
 - b. Would you hand me the book that is sitting over there on the table?

H. **On placement of interruptions and appositives:** When using a commenting phrase—such as “in other words,” “however,” or an appositive—a useful rule of thumb, especially in polished writing and in long passages (not so much in informal comments), is that you should place it where its introduction will add *only a single comma to the overall sentence*, not two. Thus (b) is better than (a):

- a. He did not, however, know the name of the teacher, who was new.
- b. He did not know the name of the teacher, however, who was new.

since (b) adds only a single comma to the base sentence (c):

- c. He did not know the name of the teacher, who was new.

I. **On single vs. double quote marks:** A good basic rule of thumb is that one uses double quotes around words or phrases when part of the understanding of the expression involves understanding the quoted words or passage. This covers the case of quoting a passage from another writer, because the whole point of quoting them is to bring “what they said”—i.e., the content of their writing or utterance—into the present discourse. Similarly, when one uses quotes to distance oneself from the use of the word, but nevertheless wants to convey the meaning (sometimes called “shudder quotes”), one again uses double quote marks, as in:

- a. Recently, large corporations have taken to describing layoffs as “downsizing.”

When then does one use single quotes? Around uninterpreted words or items. For example:

- b. In the aphorism that “whether life is worth living depends on the liver,” the word ‘liver’ is meant to have a double meaning.

One good (but not infallible) way to tell whether you are quoting the meaning of the embedded phrase (double quote marks) or its pure syntactic or lexical form (single) is to see whether the expression within the quote marks must fit grammatically into the outer sentence—i.e., whether the sentence would remain grammatical if the quote marks were removed. If it must fit grammatically, then the embedded fragment is probably being used for its meaning, not its form, in which case one should use double quotes. Thus whereas (c) is well-formed, (d) is not—implying that the quotational context in this example *does* require grammatical conformity, implying in turn that double quote marks should be used:

- c. (✓) Nixon was heckled by some of Spiro Agnew’s famous “nattering nabobs of negativity.”
- d. (✗) Nixon was heckled by some of Spiro Agnew’s famous “are going down to.”

In the following example, however, the sentence is grammatical as it stands, but would not be if the quotes were removed, implying that it is the form of the words that are being quoted, not their meaning, implying in turn that single quotes are appropriate:

- e. (✓) The only words one could read on the torn paper fragment were ‘are going down to.’

Some exceptions to this rule include the following:

2. When double quote marks occur inside a passage that is itself being quoted, the inner double quote marks are converted to single quotes.
3. Long (multi-word) expressions are almost always double-quoted, whether interpretation is required or not. Thus one would write
 - a. The word ‘misspelling’ is itself misspelled.
 - b. The passage “it turns out that at least three varieties of gooseberry grow in Maine,” is somewhat British in using the singular ‘gooseberry’; Americans are more likely to use ‘gooseberries.’

J. On other quotational idiosyncrasies:

1. If you use italics to mention a word, you don’t need quotes as well. Thus these are fine:
 - a. (✓) He used the term ‘iatrogenic,’ meaning “caused by a doctor.”
 - b. (✓) He used the term *iatrogenic*, meaning *caused by a doctor*.
 but the following is not good:
 - c. (✗) He used the term ‘*iatrogenic*’, meaning “*caused by a doctor*.”
2. Note that various kinds of inter-item lexical rules cross quote marks, such as some kinds of elision (e.g., ‘du’ for ‘de le’ in French), the use of ‘an’ before vowels, etc. Thus one says:
 - a. An “Iphigenia in Brooklyn” production will be held tomorrow night.
 - b. On parle du “Sacre de Printemps.”
3. Don’t think that there are perfectly precise rules for quotation (cf. B.1). People vary, for example, in which of the following they find acceptable:
 - a. ‘Boston’ has six letters.
 - b. ‘Boston’ is capitalized.
 - c. The six-letter word ‘Bos-ton’ is hyphenated
 - d. ‘Boston’ is smudged.

My own view is that all of these four are either OK as they stand, or could be OK in appropriate contexts. But don’t imagine that there is an absolute fact of the matter; the legitimacy of these and a myriad other constructions is a matter of judgment and of taste.

IV • Vocabulary

- A. **On ‘comprise’ vs. ‘constitute’:** The traditional rule is that parts *constitute* a whole; a whole *comprises* its parts. According to this, it would be wrong to say that “these three movements comprise the symphony.” However some people feel that in the passive voice, ‘comprised of’ can be used in the sense of ‘constituted of.’ Thus “poems are often comprised of several verses” would be acceptable to some people.
- B. **On “its” vs. “it’s”:** See C.2, above.
- C. **On ‘effect’ vs. ‘affect’:** This pair is especially difficult, since both words occur both as nouns and as verbs.
 1. *Noun:* An ‘effect’ is something brought about by a cause or an agent—i.e., a result. (“The effects of the storm were felt for years thereafter.”) An ‘affect’ is a feeling or emotion—as distinguished

from a thought or action. (“Lesley had a habit of entering a room projecting a strong affect, which tended to put people off.”)

2. *Verb*: To ‘effect’ is to bring into existence. (“She said that she would effect the changes by reorganizing people’s offices.”) To ‘affect’ is to have an influence on, or effect a change in. (His joining the committee affected everyone negatively.)

What is confusing is that the a-verb (‘affect’) is most closely related to the e-noun (‘effect’). If you *affect* something, you will have had an *effect* on it.

- D. **On ‘ensure’ vs. ‘insure’**: It is best to use ‘ensure’ to mean “to make sure or certain”; and ‘insure’ to mean “to cover with insurance” (i.e., to protect against loss). So: “They insured their house and car, in order to ensure that they would never again be driven back into poverty.”
- E. **On ‘accept’ vs. ‘except’**: To ‘accept’ is to receive or admit something offered, or to regard as correct. To ‘except’ is to leave out or exclude. (“Although he typically accepted papers up until a week after they were due, he excepted the last week of classes, when they had to be submitted on time.”)
- F. **On ‘distinguish’ and ‘discriminate’**: It is common to write “distinguish *between* α and β ” but the simpler “distinguish α and β ” suffices. Similarly for “discriminate *between* α and β ”; it is enough to say “discriminate α and β .” It is likely that the “between” construction developed in the nominative case, where it is perfectly acceptable (virtually necessary) to say “a distinction *between* α and β .” In the verbal form, however, avoid ‘between.’

V • The Writing Process

A huge amount can (and has) been written on how to write. Some excellent resources are available on the University’s writing web site.⁶ Here I will make just a few comments—again dealing with issues I encounter particularly frequently.

- A. **Give it time**: This one is a no-brainer, but worth repeating. Write what you consider to be a fully completed, “let go” of it for a few days, and then re-read it—and edit as appropriate. It is very hard for anyone (including me) to tell, when up close and involved with a text, how it will come across to someone on the outside—someone who is not embroiled in its every detail, assumption, and nuance.⁷

An obvious moral: The best writing is rarely writing completed just before a deadline.

- B. **Favourite passages**: Be especially suspicious of words, phrases, sentences, passages, or even whole sections that you particularly like. Some of the best—but toughest—writing advice I ever received was to identify those passages in a draft I was particularly fond of, and strike them all out.

The problem is that emotions can so easily outrun reason, especially as regards one’s judgment about the products of one’s own labours. Revising a draft is little like serving as a judge at an athletic competition where you yourself are the performer. Not only would that situation be banned; it is common wisdom that, if there is an athlete of whom you are particularly fond, your judgment is likely to be skewed. Why should the situation be any different, when what is being judged is a string of your own hard-won words?

⁶<http://www.writing.utoronto.ca/>

⁷I am perfectly capable of striving mightily to write a passage or section, complete something I am very happy with, and then—when looking at it several days later—realizing that it is just terrible.

C. **Detachment:** Many people say they “think by writing”—and so write a draft as a way of formulating their views and arguments. That may be fine (though see E.4, below, on skeletons), but it presents a special hazard. It is much more difficult to think than to edit—and so there is a huge temptation, as soon as one has a draft (of any length: a paragraph or two, a section, a paper), to revise it rhetorically and stylistically, rather than wrestling with whether it actually presents the proper (or best) conceptual argument.

It is tremendously important, therefore, to be *detached* from your own writing—not to cling to it, or to feel that anything you have written is sacred or precious or too valuable to discard. Words are easy, if thoughts are clear. Remember, too, that passion and dispassion are not opposites. Strive for both!⁸

D. **Skeletons:** In part to deal with the above issues, I recommend (especially for my graduate students) the development of what I call *skeletons*—short, telegraphic summaries of the stuff and substance of a paper, stripped as much as is humanly possible of all issues of presentation, rhetoric, style, etc. Skeletons look rather like outlines, but they are different in a crucial way: instead of having *noun phrases* as entries on their various levels of hierarchical structure, they have *full sentences*—sentences that say, in the most distilled, abbreviate way possible, what point or claim is being made at that point.

Skeletons are non-standard; both the word and the practice are my own. They are much more different from outlines than a cursory examination of one might suggest. It is a form that I have developed over many years of wrestling with my own writing—and one I especially recommend to students that I supervise.

A sample skeleton, with some explanatory description, is included, below, as an Appendix.

VI • Further Information

Inevitably, these notes reflect one individual’s idiosyncratic set of judgments. Whether or not you agree with all of them, one moral to take away is that they illustrate the kind of care with which you should think about your own writing.

- A. For a wide range of more standard resources, see the University of Toronto’s writing web site:
 1. <http://www.writing.utoronto.ca>
- B. For some specific authoritative—and comprehensive—sources, I especially recommend:
 1. The usage notes in the *American Heritage Dictionary*, 3rd edition, compiled and prepared by Geoff Nunberg, mentioned in notes B.2 and B.5, above;
 2. (Henry Watson) Fowler’s *Dictionary of Modern English Usage*; and
 3. The University of Chicago Press’s *The Chicago Manual of Style*.
- C. For issues of vocabulary, I recommend not only the American Heritage (above) as a general dictionary, but also the following, which instead of saying what words mean, spells out the distinctions between words that are close in meaning:
 1. The *Merriam-Webster Dictionary of Synonyms: A Dictionary of Discriminated Synonyms with Antonyms and Analogous and Contrasted Word*, 1968; aka *Webster’s New Dictionary of Synonyms*, 1978.

This book is unfortunately no longer in print. A similarly-named volume, the *Merriam-Webster Dictionary of Synonyms & Antonyms*, is available, as is a *Merriam-Webster’s Dictionary and Thesaurus*. The latter two, however, very different books—merely abbreviated dictionaries and word lists amplified

⁸I enjoin my doctoral students to write a dissertation exhibiting passion, dispassion, and compassion.

with lists of other words with similar and opposing meanings. The older *Dictionary of Synonyms* was radically different, much longer, and much more interesting. It has extensive discussions of the subtleties, nuances, and particular connotations of each members in a group of similar words. Used copies are available on Amazon for just \$10 or so (though unused ones sell for hundreds of dollars).



Appendix – Skeletons

By a **skeleton** I mean a short distilled summary (one to just a very few pages) of the claims and arguments to be made in a paper. The aim is to present the argument in *logical* rather than *rhetorical* order, to avoid introductions, and just to focus on the “guts” of what is to be claimed. Grammatically, a skeleton need not (probably should not) be well-formed prose. Substantively—this is the critical part—each leaf node should be a sentence or a claim, not simply a noun phrase of the sort that might be found in an outline (such as “background,” “history of the topic,” “Derrida’s notion of ‘différance,’” “Llewellyn’s critique,” “epigenetics in the loon,” etc.).

Developing skeletons is *diabolically difficult*. Some years ago, Joel Walmsley, a doctoral student in philosophy at U.of.T, wrote a great PhD on the notion of emergence, for which I served as supervisor.⁹ We met regularly—every week or so. He brought to each meeting a new version of a skeleton of his argument. We would wrestle with it, question its assumptions, uncover inconsistencies, recognize holes, tear it apart, write all over it, and in general plot ways to improve it. After *two full years*, he finally developed a version (still just 2 pages!) that we both felt was sound. He then sat down and wrote his dissertation straight, from start to finish, in about 3 months. The result was unanimously accepted by his committee *without modifications*—an extremely rare accomplishment. Moral: all of the work was in developing the argument, reflected in the skeleton.

Some of you will argue that you can only “think” by writing—i.e., that it is only in the process of writing that your ideas become clear. Fair enough. I know many people work that way (sometimes I even do myself). But do not confuse those drafts with a skeleton. (I often write a skeleton *after* completing a first draft, which draft I then throw away). The aim of developing the skeleton is to distill the substance of one’s thoughts, however they have been developed, in order to focus on and subject to critical scrutiny what is core and essential to the argument. Laying out the skeleton has psychological advantages, too: it helps one shed the emotional attachment that one almost invariably attaches to one’s prose, freeing you up to concentrate on issues of intellectual substance. Prose can be elegant and compelling; but it can also be unbelievably distracting. Only once one has an adequate grip on what is to be said is it time for exegetical prowess and rhetorical flourish.

Remember: a skeleton is not an *outline*!

What follows is sample skeleton, to illustrate the form. The (lettered) endnotes are for purposes of this presentation, rather than being part of the skeleton *per se*.

Example Skeleton – Formal Symbol Manipulation^a

I • Project

A. Background

1. “Computational theory of mind” (CTOM): mind \approx computer^b
2. CTOM variously interpretable (different people: different allegiances—some an admixture)
 - a. *Empirical*: mind \approx computer, *whatever computers are like* (“computation in the wild” \equiv CITW)
 - b. *Theoretical*: mind explicable by a particular θ_i (computing), for various θ_i .^c Label: CTOM $_i$.
3. Therefore multiple ways to challenge CTOM, falling into two broad classes:
 - a. *Human side*: (i) mind \neq computer as in CITW (emp.); or (ii) $\neg\theta_i(\text{mind})$,^d for a particular θ_i (theory)
 - b. *Computer side*: $\neg\theta_i(\text{CITW})$

B. Candidate θ_i s (call them *construals* of computation)

1. θ_1 : *Formal symbol manipulation* (FSM)—primarily: philosophy of mind

⁹Joel obtained his PhD in 2005, and now teaches in the philosophy department at the University College, Cork.

2. θ_2 : *Information processing*—primarily: cognitive science, biology, “public”
3. θ_3 : *Digital state machine*—assumed by everyone?
4. θ_4 : (Realisation of an) *effectively computable function*—primarily: computer science

C. Project

1. Choose θ_1 (FSM) as candidate θ_i , and adopt strategy I.A.3.b: i.e., ask whether FSM(computers)
2. If \neg FSM(computers), then hypothesis that FSM(mind) (i.e., CTOM_{FSM}) is substantially weakened.

II • Formal symbol manipulation—prefatory analysis

A. FSM · ambiguity #1

1. Two readings: “((formal symbol) manipulation)” vs. “(formal (symbol manipulation))”
2. Choose latter: genuine symbols, with (at least potentially) real semantics, *formally manipulated*
3. This is what underwrites logic [[does this need defense?]]

B. FSM · ambiguity #2: positive and negative characterisations

1. Positive (FSM+): symbols manipulated in virtue of *syntax or grammatical shape*
2. Negative (FSM−): manipulated *independent of semantics*

C. Equivalence hypothesis (EH)

1. Assumption: FSM+ and FSM− are extensionally equivalent
2. Implies: *semantical* properties and *syntactic/grammatical* properties do not overlap
3. NB: is EH true? will come back to this [[do I ever do this?]]

III • Positive construal: FSM+

A. Intro

1. Need to know what syntactic and grammatical properties are.
2. Exemplars (logic, Turing machines, formal languages, computers): always demonstrated *ostensively*
3. General problem: what are they *in general*? Rarely theorized (ever?). Worrying.

B. Intuitions

1. Theorists— Δ^c articulations, perhaps different underlying ideas?
 - a. Cf. Turing’s original (1936–37) article: *what can be done*
 - b. Cf. Haugeland’s “Semantic Engines”: *read/write*—anything something can react to
 - c. Cf. Gandy’s “Principles of Mechanism”: *mechanical principles*
2. Where do these ideas come from? How do they relate?
 - a. Historically: canonical “formal language” definitions motivated by (refer to?) *mechanical realizability*
 - b. Conceptually: any (recursively) specifiable way of demonstrating that something “works,” *mechanically*, can count as a “syntactic” property
 - c. Argument
 - i. Hard to argue: mechanically implementable machine \neg computer because works via \neg *syntactic* physically effective property. So ‘mechanizable’ seems important, even definitive?
 - ii. Conversely, hard to argue: machine w/ well-defined but non-implementable “syntactic” property (e.g., “true”) still computationally OK. ‘Syntactic’ seems *subservient* to mechanizable.
3. Conclusion
 - a. ‘Syntactic’: no more general nor restrictive than “mechanically effective”
 - b. Doesn’t imply that ‘syntactic’ *means* “effective”—‘syntactic’ might require being *both*:
 - i. Property in virtue of which the system *works*, mechanically; and
 - ii. Property in virtue of which it *means something*, compositionally.

C. Conclusion (surprising!)

1. FSM+: “works in virtue of syntactic properties” means (no more than) “work in virtue of properties it works in virtue of”—i.e., vacuous! So yes, FSM+ true of computation—but only because content-free!

2. FSM+ true of mind (i.e., $CTOM_{FSM+}$)? Yes, unless dualist!
3. But boring; no intellectual interest. Imposes no restriction; provides no insight into how mind works.

IV • Negative construal: FSM–

- A. FSM– (“Independent of semantics”): two readings:
 1. FSM–1: *Defined* independent of semantics
 2. FSM–2: *Works* independent of semantics
- B. FSM–1
 1. Analysis
 - a. If defined independent of semantics, then (cf. §III, above) no positive substance to add to it.
 - b. So “FSM–1” means (combine w/ III.C.1) “works how it works” (evacuating ‘symbol’ of all meaning)
 - c. Hence: this construal should be recast “**stuff manipulation**”. Empty!
 2. Truth
 - a. FSM–1 true of CITW? Yes—but again, vacuous.
 - b. FSM–1 true of (non-dualist) mind, too—again, vacuously. I.e., $CTOM_{FSM-1}(\text{people})$: true but boring!
- C. FSM–2
 1. *Works* independent of semantics
 - a. Requires semantics, this time—for manipulation to be independent of
 - b. Cf. temperance union: live *independent of alcohol* (\neg like luminiferous ether)
 2. Analysis
 - a. True of all “exemplar” cases—logic, arithmetic, etc.
 - b. Critical insight: true because *semantic domain* (numbers, math, etc.) *causally inaccessible*
 - c. Critical test: when semantic domain causally accessible, can it be used (still be computational?)
 - d. Yes! \Leftarrow this will be FSM–2’s Achilles’ heel
 3. Counterexample (to FSM–2)
 - a. “Length(‘ABCDEFG’) \Rightarrow 7”
 - b. Result: *numeral* (\neg number!), *denoting* a number. Call numeral “hepto”; so *hepto* denotes 7.
 - c. What played a causal role in engendering that’s being the answer?
 - d. The *length*: the *number* (seven) of letters; \neg the numeral (hepto) of numbers.
 - e. So “Length(‘ABCDEFG’) \Rightarrow 7” fails FSM–2
 4. Conclusion
 - a. Is “Length(‘ABCDEFG’) \Rightarrow 7” a computation? Yes!
 - b. Hence: FSM–2(CITW) is false
 - c. $CTOM_{FSM-2}$ probably false, too (since we, qua people, can also count up inner states)

V • Conclusion

- A. As a θ_c (computers), FSM is either
 1. Vacuous—if taken to mean FSM+
 2. Vacuous—if taken to mean FSM–1
 3. False—if taken to mean FSM–2
- B. Hence
 1. FSM fails as a theory of computing
 2. Likely fails as a theory of mind, too! QED (◆)

Notes

^aA paper that presented the ideas summarized in this skeleton would likely be 30–40 pages (c. 6,000–10,000 words).

^bA cryptic sentence. If written out in English, this line and the three immediately below it would read roughly as follows:

“What is known as the ‘computational theory of mind’ (which I will refer to as ‘CTOM’) claims that the mind is, or anyway is like, a computer. This claim can be seen to consist of two at least partially separable theses. Empirically, the CTOM can be interpreted as a thesis that minds are like computers, in whatever way *computers in the world* are computers. I.e., on an empirical reading the CTOM is essentially an ostensive thesis, which: (i) points to two classes of extant entities in the world—*minds*, on the one hand, and *computers*, on the other; and then (ii) claims, of the former, that they are (substantively or constitutively) the same as the latter. Theoretically, the CTOM can be interpreted as making a theoretical claim—i.e., as being framed with implicit reference to a theory about what computers are like, a theory I will call θ_c . On this interpretation, the CTOM would imply (might be entertained or defended because of this implication) that that theory was also substantively or constitutively true of minds—i.e., that $\theta_c(\text{minds})$, in the same way as $\theta_c(\text{computers})$. Given that there are different theories θ_c of what computers are like (indexed by different subscripts ‘c’), there are correspondingly many different theoretical interpretations of the CTOM, which I will refer to as CTOM_c .

There are details in this formulation that would need to be worked out (I am not claiming that the foregoing paragraph is particularly good). The advantage of the skeleton is exactly that it abstracts away from some of these details, in order to foreground the issues most centrally at stake.

^cI use ‘ θ ’ for *theory*; ‘ θ ical’ for *theoretical*, etc.

^d \neg is widely used for ‘not’ in logic.

^eI use ‘ Δ ’ for *different* and its variants (*differently*, *difference*, etc.)

