

Writing Advice

Gerard J. O'Reilly

Last updated: October 2025

1 Tenses

When writing a scientific article, the tense in which the text is presented is important. It is easy to get mixed up between past and present tense, especially if you write different parts of the article at different times. This can influence your train of thought, and you may find yourself describing methods in the present tense (e.g., A test specimen is created and subjected to several cycles of a loading protocol). However, after taking a break and returning to your writing, you might start thinking of the experiment differently like something you completed in the past and begin using the past tense instead (e.g., The specimen behaved very well and showed good ductile behaviour). The result is a piece of text that may describe the activities clearly but contains inconsistent tenses, which can be confusing and is grammatically incorrect.

To avoid this, there is a simple set of rules that can be followed:

1. Actions or observations that were carried out in the past should be written in past tense. For example:
 - The test specimen *was* constructed and placed in the loading rig.
 - A standard loading protocol *was* applied and the specimen *exhibited* ductile behaviour.
 - A numerical model *was* constructed using OpenSees.
 - A survey *was* carried out.
2. Anything that is presented in the article itself should be written in the present tense. For example:
 - A numerical model was constructed using OpenSees and *is* shown in Figure 1.
 - A survey was carried out and the data *is* listed in Table 2.
3. Anything that is generally true in the broader context of the paper should be written in the present tense. For example:
 - Steel moment frames *are* a common lateral force resisting system.
 - Incremental dynamic analysis *is* a popular analysis method for quantifying building response.

2 Active versus passive voice

Another common theme in technical papers today is the choice between using the active voice (e.g., We built a model) or the passive voice (e.g., A model was built). While both are valid ways of conveying the same information, the passive voice is traditionally considered more appropriate for technical writing. However, in recent years, the active voice has gained popularity in some research communities, particularly in North America, due to its simplicity and less formal tone. I admit that the active voice is generally easier to read, but it's important to note that the passive voice remains the more conventional choice in technical contexts. Ultimately, it comes down to personal preference. Many grammar-checking tools will encourage the use of the active voice, but whichever style you choose, it's important to remain consistent. Personally, I prefer to use the passive voice, noting the irony that I wrote this text in the active voice.

3 Acronyms, abbreviations and symbols

In general, acronyms and abbreviations can be very useful for helping readers follow the text and avoiding overly repetitive phrasing. For example, the term reinforced concrete is almost synonymous with the acronym RC. That said, defining too many acronyms can become overwhelming and may create confusion due to the number of terms the reader must keep track of.

As for when and where to define an acronym, a simple rule is: define it at its first mention. So, the first time you write reinforced concrete, you should introduce it as *reinforced concrete* (RC). From that point onward, you can use RC consistently. Needless to say, there's no point in defining an acronym if you don't plan to use it regularly. If the term appears only once or twice more in the document, it may not warrant an acronym, unless it refers to a well-known method, organisation, or concept (e.g., American Society of Civil Engineers (ASCE)).

In some cases, especially in theses or dissertations, authors include a complete list of symbols and acronyms at the beginning of the document. While this is certainly useful, it's not always practical or required. Therefore, the advice here applies to documents where such a list has not been implemented.

On the issue of symbols, these should be introduced and explained fully in the text by saying something like: Newton's second law states that the force, F , is given by:

$$F = ma \tag{1}$$

where m is the mass of the object and a is the acceleration. Defining the symbol F above should come directly after mentioning it (i.e., the force, F , or the mass, m), where I tend to separate these with commas. Other use parentheses (i.e., the force (F) or the mass (m)), which can also be fine but again, a personal preference. Symbols should also always be written in math font and not normal font (i.e., F and not F).

Another common misconception is capitalising the words when defining an acronym, such as writing *Reinforced Concrete* (RC) instead of *reinforced concrete* (RC). This is unnecessary and should generally be avoided unless a proper noun is involved. Otherwise, it is an indicator that your text is AI-generated.

A further point concerns plural forms. If you define an acronym in the singular form (e.g., moment-resisting frame (MRF)) and later refer to it in the plural (e.g., moment-resisting frames (MRFs)) there is no need to redefine it. Simply using MRFs is sufficient and understood.

An exception to these conventions applies to figure captions and axis labels. Excessive use of acronyms in figures can reduce clarity. Since the goal of a figure is to be as self-explanatory as possible, it is often better to write out the term fully (e.g., engineering demand parameter (EDP)) on axis labels or in figure captions. Of course, this must be balanced with avoiding clutter or overly long labels, so some ad hoc judgment is sometimes necessary.

Finally, a key consideration is where acronyms should be defined. A well-structured technical article typically includes three main parts: the abstract, the main body, and the summary and conclusions. Each of these sections should be as self-contained as possible, allowing readers to understand the content without needing to refer elsewhere. Therefore, if an acronym is used in any of these sections, it should be defined there—even if it has already been defined elsewhere in the document. For example, if an acronym appears in the abstract, it should be redefined the first time it appears in the main body, and again in the conclusions if used there. While this may seem redundant, it improves accessibility for readers who may only skim the abstract or conclusions without reading the full article.

4 Summary and conclusions

Another important aspect of academic writing is how to effectively summarise and conclude an article. In modern academic publishing, this section plays a crucial role, as many readers will not take the time to read the entire paper and thoroughly digest the findings and arguments presented in the main body. More often, researchers will simply read the abstract, possibly the introduction, and the summary and conclusions. Therefore, this final section should do precisely that: summarise the work carried out and clearly present the main findings.

A strategy I’ve used for many years, which is common among many researchers and not something unique to me, is to divide this section into two distinct parts. The first part consists of a single paragraph, no longer than half a page, which briefly summarises the work presented in the paper. This is then typically followed by the sentence: “*Based on this, the following conclusions can be drawn:*” and a concise set of bullet points. These bullet points should clearly state the key takeaways, which are what the reader should remember after going through 20 or so pages of detailed work and discussion.

In fact, when considering whether a piece of research is worth developing into a journal paper, it can be helpful to try writing this section first. If you find that the bullet points are few, vague, or simply restate the obvious, it’s usually a sign that the paper may not add much value or generate significant interest. This doesn’t necessarily mean it won’t be accepted for publication, but it may indicate that it’s not the kind of paper others will find especially useful.

Here is an example of this approach from O’Reilly and Shahnazaryan [2024] with the summary, followed by the bullet point conclusions:

This paper presented a comparative study of two methods to develop vulnerability functions for regional seismic risk modelling. These methods were based on analytical methods and involved several key assumptions to make them widely applicable and computationally highly efficient. The first of these methods was based on the fragility analysis of displacement-based demands on an equivalent single-degree-of-freedom (SDOF) model, which, in tandem with loss ratios, returns vulnerability functions. A second and alternative approach described here utilises an equivalent SDOF model to estimate both displacement- and acceleration-based demands at each building level. Combining these demands with storey loss functions (SLF)s leads to the estimation of

vulnerability at each storey and, ultimately, the vulnerability function of the building. The two methods were described in detail, and a comparative case study example was presented to scrutinise their similarities and differences. Some final remarks on the wider implications and possible future directions were also presented.

Based on the work presented here, the following conclusions can be drawn:

- The SLF-based approach is a more detailed and robust methodology compared to the widely-used fragility and damage-to-loss ratio-based approach. The SLF-based approach is more detailed but does not carry the burden of excessive extra computation or effort. In fact, when simplified, the SLF-based approach can be reduced to give essentially the same result as the fragility-based approach.
- The SLF-based approach directly considers both peak storey drift (PSD)- and peak floor acceleration (PFA)-based demands at all storeys of the building. This breakdown allows for a more detailed consideration of non-structural elements and identifies which performance groups at which building location contribute most to the expected loss.
- The issue of collapse and demolition are handled more directly, which can be useful for decision-making in issues beyond economic losses, such as estimating casualties or identifying when retrofitting interventions may be futile.
- While SLFs are needed for different building taxonomies, the generalised format may be a prudent strategy given that a physics-based approach to modelling the building component fragility can be adopted, followed by a rational approach to estimating the expected repair costs.
- The uses of this approach extend beyond vulnerability modelling and can also aid insurance companies in deciphering and establishing a relative scale of expected damage and loss distribution within a building to determine more reasonable tariffs.

Overall, this extension of the existing to the proposed approach is without much additional computational cost and can represent an enhanced level of quality in the output decision variables used in regional seismic risk assessment.

5 Figures and graphics

When presenting results or data, plotting figures is almost always a very effective way to convey information. While there are many guides on how to create good versus poor plots (see Jack Baker's YouTube video with some excellent pointers on making effective figures), one aspect that is sometimes overlooked is the relative dimensions and sizing of figures.

This may seem trivial, as once you define and create a figure, you can simply scale and crop it in your word processor to make it appear correctly in the final document. While this approach works in many cases, it can lead to problems when publishing through journals or dealing with typesetting staff. Many of us have experienced the situation where a manuscript is accepted, and we eagerly await the proofs, only to discover that the figures are awkwardly scaled, either too large or too small, and not at all as originally submitted. What has likely happened is that the publisher's typesetters imported your original figure files without adjusting their size or formatting. As a result, the figures may end up barely legible and visually unappealing.




Figure 1: Example plot with no scaling

To avoid this, it is helpful to be a bit more prepared and provide figures in PDF format that are sized appropriately for 100% scaling—rather than relying on arbitrary scaling (e.g., “eyeballed” adjustments in Microsoft Word). This ensures that line weights and font sizes appear clean and professional in the final version.

A good figure dimension for a single plot is around 6 cm tall by 8 to 10 cm wide, with a font size of approximately 8 pt to 10 pt maximum. Slightly smaller font sizes like 6 pt for axis ticks and legend text can help the axis labels stand out and reduce visual clutter. While there’s no fixed rule for these numbers, I’ve found them to work well in practice, but the dimensions do approximate the Golden Ratio, which is somewhat reassuring.

Figure 1 is an example of such an approach from O’Reilly et al. [2022], where the plot was prepared in Python and the source code as given as shown below. Here it can be seen how the dimensions are initially set to about 9 cm by 6 cm (Python works in imperial units, hence 2.36 inches is around 6 cm, and 1.5 times 6 cm is 9 cm). The default font size is set to 8 pt, where the legend font is set to 6 pt. \LaTeX syntax is used to display symbols in the axis labels and the figure is saved as a vector-based PDF.

```
plt.figure(figsize=(1.5 * 2.36, 2.36))
plt.rcParams.update({'font.size': 8})

# Plotting functions to display data

plt.xlabel(r'Displacement,  $\Delta$  [m]')
plt.ylabel(r'MAFE,  $\lambda_{\Delta}$ ')
plt.yscale('log')
plt.ylim([1e-5, 1e-1])
plt.xlim([0, 1])
plt.legend(ncol=2, fontsize=6, frameon=False)
plt.savefig('plots/Demand-Hazard-Curve-Case-Study.pdf', bbox_inches='tight')
plt.close('all')
```

Although I no longer use Matlab and have since passed to Python, I was an avid user during my PhD days, and had several tools and shortcuts that helped make nice plots and data presentation. A GitHub repository is available online with much of what was described above presented there also.

In addition, when creating graphics, I often use Inkscape, which is a free online software.

One of the first steps I take is to enable a grid and place an A4 sheet behind the drawing area for reference. This helps guide the image creation relative to the actual size of an A4 page, so that the final output is neither too small nor too large and does not require any of the ad hoc adjustments previously mentioned.

References

Gerard J O'Reilly and Davit Shahnazaryan. On the utility of story loss functions for regional seismic vulnerability modeling and risk assessment. *Earthquake Spectra*, 40(3):1933–1955, aug 2024. ISSN 8755-2930. doi: 10.1177/87552930241245940. URL <https://journals.sagepub.com/doi/10.1177/87552930241245940>.

Gerard J. O'Reilly, Hiroshi Yasumoto, Yoshitaka Suzuki, Gian Michele Calvi, and Masayoshi Nakashima. Risk-based seismic design of base-isolated structures with single surface friction sliders. *Earthquake Engineering & Structural Dynamics*, 51(10):2378–2398, aug 2022. ISSN 0098-8847. doi: 10.1002/eqe.3668. URL <https://onlinelibrary.wiley.com/doi/10.1002/eqe.3668>.