Your First Research Paper

Your First Research Paper

From Project to Submission,
Peer Review, and Publication

Gaurav Tiwari

gauravtiwari.org

Copyright © 2026 Gaurav Tiwari.

All rights reserved. No part of this book may be reproduced in any form without written permission from the author, except for brief quotations in reviews.

First edition, 2026.

gauravtiwari.org

To everyone sitting on a finished project
and a blank document.

Contents

Preface

Somebody has told you your project might be publishable.

It's a nice thing to hear and a strange place to be left. You have a result, a supervisor who is busy, a set of scripts, three hundred figures, and no idea what happens next. Nobody has ever shown you a paper being made. You've read hundreds of them and never seen one written.

This book shows you that.

It assumes you have work that might be worth publishing, and it takes you from there to a submitted manuscript, through peer review, and out the other side. Choosing a venue. What a contribution actually is. How to write an introduction that a referee will believe. Who counts as an author. What to do when three reviewers disagree with each other and one of them is annoyed.

A word about who I am to write it. I'm not a full-time academic, and this isn't a book about building a research career. It's a book about the mechanics of a first paper, written by somebody who has spent eighteen years explaining technical work to people who need to act on it, and who has watched a lot of good undergraduate and master's projects die in a drawer because nobody told the student what to do next.

Where the process varies by field, and it varies enormously, I say so rather than pretending my corner of it is universal. Mathematics and physics preprint everything; parts of medicine and psychology have registration requirements this book only gestures at; computer science runs on conferences with hard deadlines and rebuttal periods. Where a practice is field-specific, it's labelled.

Two things this book will keep saying.

The paper is not the project. A project is everything you did. A paper is one claim, supported. Most first drafts fail because the writer tried to report the project rather than argue the claim, and cutting is the single most useful revision skill there is.

Rejection is the normal case. Good papers get rejected. First papers get rejected more. A rejection is information about fit and about what needs strengthening, and the students who publish are not the ones who never get rejected, they're the ones who resubmit within a month.

If you're at the very start, sitting with a project and a supervisor's vague encouragement, begin with Chapter 3, which asks the honest question about whether this is a paper yet, and tells you what to do if the answer is no.

Gaurav Tiwari
gauravtiwari.org

Part I

What a Paper Is

1

Who Reads a Paper, and Why

Before you write a word, it helps to know who this document is for, because it isn't who students assume.

Almost nobody will read your paper the way you're about to write it: from the beginning, carefully, to the end.

1.1 The four readers

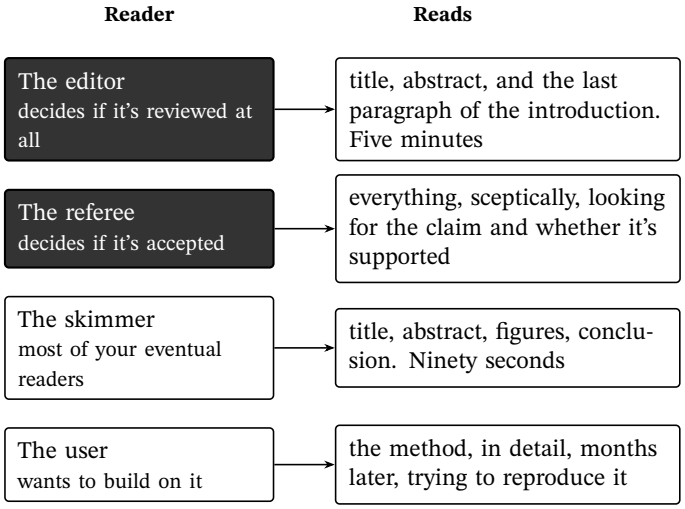

Figure 1.1: The four people who will encounter your paper, and how much of it each one reads. Write for the top two and the fourth is satisfied.

Two consequences fall straight out of that table, and they shape the whole book.

The title, abstract and introduction do most of the work. They're all the editor sees before deciding whether to send it out, and all most readers ever see. Students spend their time on the results section and write the abstract in ten minutes at the end.

The referee is the reader you write for. They're the only person who reads it all, they're professionally sceptical, and they're doing it unpaid on a Sunday. Everything in Part III is about making their job possible.

> *Write for a tired, competent, sceptical reader who is looking for a reason to stop. Every ambiguity you leave is a place they can stop.*

1.2 What the referee is actually asking

Not "is this good?" Four specific questions, and a good paper answers all four explicitly.

1. **What is the claim?** One sentence. If they can't find it, that's usually the review right there.
2. **Is it new?** Compared with what, specifically, and cited.
3. **Is it supported?** Do the results actually establish the claim, or something weaker?
4. **Is it correct?** The details: the derivation, the statistics, the controls, the assumptions.

Notice the order. Correctness comes last, and papers are rejected on questions 1 to 3 far more often than on 4. A technically flawless paper whose contribution is unclear gets rejected; a paper with a fixable error and a clear, novel, supported claim gets a revision.

FROM THE OTHER SIDE OF THE DESK

The most common thing I've seen sink a first submission isn't a mistake. It's a referee writing "I could not determine what is new here". That's question 2, and it's answerable in the introduction with two sentences the author knew perfectly well and never wrote down.

1.3 The paper is not the project

The single most important idea in this book, so it goes early.

Your project is everything you did over six or twelve months: the false starts, the pipeline you built, the parameter sweep, the three approaches that failed, the literature you read.

Your paper is *one claim, supported*. Most of the project doesn't appear.

In the project	**In the paper**
Three approaches, two abandoned	The one that worked, plus one sentence on why the others were rejected if it's informative
Six months of pipeline building	A method section, and a link to the code
Everything you read	The dozen references that position the claim
Every figure you generated	Four
Your learning curve	Nothing

Students find this brutal, and it's the difference between a manuscript and a report. A supervisor saying "it's too long" almost always means "it's still the project".

1.4 What this book covers, and what it doesn't

Covers: deciding whether you have a paper, choosing a venue, writing each section, authorship, data and code, submission, peer review, responding to referees, and what to do when it's rejected.

Doesn't cover: how to do research, how to build an academic career, grant writing, or the specifics of any one field's conventions beyond flagging where they differ.

Assumes: you have results, and a supervisor or collaborator who is nominally involved. If you don't have a supervisor, the book still works, and Chapter 13 has a note about the extra care needed.

1.5 Field differences, stated once

Mathematics	Journals; preprints on arXiv are universal; review is slow, often a year; single-author papers are normal
Physics	Journals plus arXiv; review months; large author lists in some subfields
Computer science	Conferences are primary, with hard deadlines and rebuttal periods; journals secondary
Engineering	Journals and conferences both count; industry co-authors common
Life sciences	Journals; preprints growing; strict authorship and ethics conventions

Ask your supervisor which model your subfield follows before you plan anything, because the conference model in particular changes the whole calendar.

DO THIS NEXT

Take a paper you admire in your area. Read only the title, the abstract, and the last paragraph of the introduction. Can you state its claim in one sentence? That's what an editor gets, and it's the standard your own paper will be held to.

2

The Anatomy of a Paper

Every paper has the same skeleton, and each section answers exactly one question.

Knowing which question belongs where prevents most structural problems before they happen.

2.1 The sections, and their one question each

Section	The question it answers
Title	What is this about, and is it for me?
Abstract	What did you do and what did you find?
Introduction	Why does this matter, what's missing, and what do you claim?
Related work	Who else has done what, and how is yours different?
Method	What exactly did you do, in enough detail to repeat?
Results	What happened?
Discussion	What does it mean, and what are the limits?
Conclusion	What's the takeaway, and what's next?
References	Where did the prior work come from?

The commonest structural error is answering a question in the wrong section: interpretation in results, method detail in the introduction, new material in the conclusion. Each one costs you, because a referee looking for something in its usual place and not finding it concludes it isn't there.

> *Before writing any section, write its one question at the top of the page in pencil. Then check every paragraph answers that question and no other.*

2.2 The shape of the argument

Underneath the sections, a paper is a single argument with four moves.

Move 3 is the one first papers most often omit. The authors know their claim so well that they never write it down as a sentence, and the referee has to infer it. Referees asked to infer a claim usually infer a weaker one than you meant.

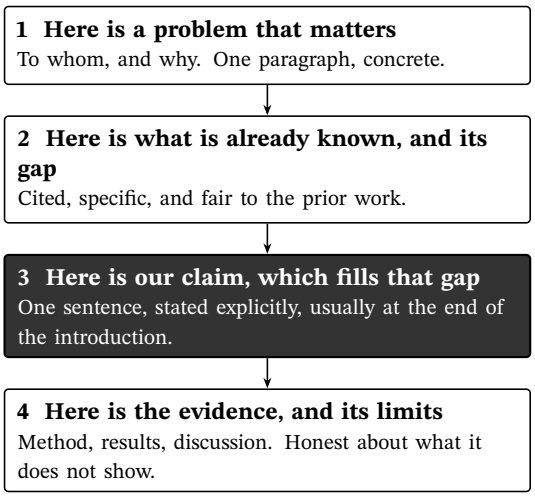

Figure 2.1: The four moves. Everything in a paper serves one of them, and a paragraph that serves none should be cut.

2.3 The IMRaD variant, and where it doesn't apply

Introduction, Methods, Results, and Discussion is the default in most experimental sciences and much of engineering.

It doesn't fit everywhere:

Mathematics	Introduction, preliminaries, results with proofs, then often a section of examples or applications. No "methods" section
Theoretical CS	Similar: definitions, theorem statements, proofs
Systems / ML	Introduction, related work, method, experiments, ablations, discussion
Case-based work	Introduction, background, the case, analysis, implications
Review papers	Entirely different: the structure is the taxonomy you impose on the field

Whatever your field does, the four moves are still there. They're just distributed differently.

2.4 Length, and where it goes

For a first paper, in most journals, six to ten pages in the venue's format. Here's roughly where it goes.

Section	Share	Note
Abstract	3%	Written last
Introduction	15%	The most-read part after the abstract
Related work	10%	Positioning, not a literature survey
Method	20%	Enough to repeat
Results	25%	Including figures
Discussion	20%	Where the thinking shows
Conclusion	5%	Short
References	2%	

If your related work is 30% of the paper, you've written a literature review with a result attached. If your method is 40%, you've written a technical report. Both are common first-draft shapes and both are fixed by cutting rather than by adding.

FROM THE OTHER SIDE OF THE DESK

An experienced referee can tell within a page whether a manuscript is a paper or a converted thesis chapter. The tell is proportion: a thesis chapter has a long background, an exhaustive method, and a thin discussion, because those are the proportions a thesis wants.

2.5 Converting a thesis chapter

Which many first papers are, so it's worth naming the surgery.

Thesis chapter	Paper
Background chapter, 20 pages	Related work, 1.5 pages, positioning only
Every method detail	Enough to repeat; the rest to a supplement
All results	The ones that support the claim
Discussion of everything	Discussion of this claim
Written for an examiner who must see your competence	Written for a referee who must see your contribution

That last row is the whole difference. A thesis proves you can do research. A paper contributes a result. The audiences want opposite things, and text that persuades one bores the other.

DO THIS NEXT

Take the four moves and write one sentence for each about your own work, on a single sheet. If move 3, the claim, takes you more than a sentence or needs an "and", you may have two papers, or you may not yet have one. Chapter 3 is next for a reason.

3

Is This Publishable?

The honest question, asked early, before you spend three months on a manuscript.

3.1 The four tests

A piece of work is publishable if it passes all four. Answer them in writing, in one sentence each.

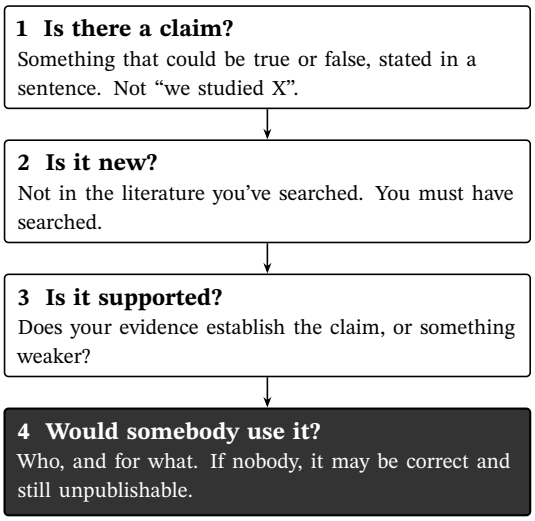

Figure 3.1: The four tests. Failing any one of them is fixable; the fix is different in each case.

3.2 What each failure looks like, and what to do

No claim. You have results and no assertion. Very common with project work: you built something, or characterised something, and never asked what it shows.

Fix: write down the three most interesting things in your data. For each, write the sentence a reader would take away. One of them is usually a claim, and the paper is then about that one thing.

Not new. Somebody published it, often five years ago in a journal you haven't read.

Fix: first, be certain. Search properly, as in Chapter 5. If it's genuinely done, ask whether your version differs in a way that matters: a different regime, a stronger result, a simpler proof, a real-world dataset where theirs was simulated. Those are legitimate contributions if you claim them accurately.

Not supported. You have a claim and the evidence establishes something narrower. This is the most common referee objection and the most fixable.

Fix: weaken the claim to what you can support. A narrow claim, solidly established, is a paper. A broad claim, thinly supported, is a rejection.

Nobody would use it. Correct, new, supported, and of no interest.

Fix: the hardest one, and sometimes the answer is that this is a good project and not a paper. Sometimes the fix is framing: the same result aimed at a different community is interesting to them.

> *The most common repair to a first paper is to make the claim smaller. Not weaker writing, a narrower assertion, fully supported. Referees reward that and punish the opposite.*

3.3 Types of contribution that count

Students think "new result" means a discovery. The category is wider.

New result	The obvious one
New method	A better way to do something known
New application	A known method applied where nobody had, with the adaptation that made it work
Replication or refutation	Especially of something widely believed. Under-valued and genuinely useful
Negative result	"This approach does not work, and here is why" saves other people months. Harder to place, and some venues want them
New dataset or benchmark	With the documentation that makes it usable
Simplification	A shorter proof, a cheaper algorithm, a smaller apparatus
Systematisation	A review or taxonomy that organises a confused literature

Undergraduate and master's projects most often produce rows 3, 4, 6, and 7. Those are publishable, in the right venue, if claimed honestly.

FROM THE OTHER SIDE OF THE DESK

Framing matters more than students expect. The same work described as "we applied a standard method to a new dataset" and as "we show the standard method fails on realistic data, and identify why" can be a desk reject and an accept. The second isn't spin: it's identifying which part is the contribution.

3.4 Ask your supervisor these four questions

Specifically, rather than "is this publishable?", which gets you a shrug.

1. "What would you say the claim is, in one sentence?" If their sentence differs from yours, that's the conversation.
2. "What's the closest published work?" They usually know, and it saves you a week.
3. "What would a referee attack first?" This is the most useful question in the book and almost nobody asks it.
4. "Which venue, and why that one?" Chapter 4.

3.5 When the answer is not yet

Three common situations, with what to do.

One more experiment would do it. Then do it, if it's weeks rather than months. A paper that waits eight weeks for a control is much better than one that goes out without it and gets rejected for the same reason.

It needs a comparison with an existing method. Very often required, and often the missing piece. Budget for it.

It's a good project and not a paper. This happens, and it's not a failure. A well-written project report, a public repository with a good README, or a blog post that explains the negative result all have value. Not everything is a paper, and pretending otherwise wastes a term.

DO THIS NEXT

Write the four tests on one page, with a one-sentence answer to each about your own work. Take it to your supervisor and ask their version of test 1 and question 3 above. Thirty minutes, and it decides whether the next three months are a manuscript or something else.

Part II

Before You Write

4

Choosing the Venue

Choose the venue before you write, not after.

The venue determines the length, the format, the expected level of novelty, the audience's background, and how much you have to explain. A manuscript written for nobody in particular has to be rewritten for whoever eventually gets it.

4.1 How to build a shortlist

Not by browsing rankings. By looking at your own reference list.

1. Take the ten most relevant papers you cite.
2. Write down where each was published.
3. The two or three venues that appear more than once are your shortlist.

That's it, and it works because the venues publishing the conversation you're joining are the venues whose readers care about your result.

> *Your reference list is your venue shortlist. If none of the work you build on was published there, the editor will wonder why you sent it, and they will be right to.*

4.2 The three-tier approach

Pick three: a stretch, a target, and a safety. Decide all three now, so that a rejection is a redirection rather than a crisis.

Stretch	The best venue where the work is plausibly in scope. Submit here first if the work is genuinely strong and you can afford the time
Target	Where your references cluster. The realistic home
Safety	A solid, indexed venue that publishes work like yours. Not a predatory one

For a first paper from an undergraduate or master's project, starting at the target is usually right. Starting at the stretch costs three to six months for a low probability, and time matters more than prestige when you're about to leave for a job or a PhD.

4.3 What to check before committing

Scope statement	Read the actual “aims and scope”. Half of desk rejections are scope mismatches
Recent issues	Does anything look like yours? If nothing has for two years, it’s the wrong venue
Length limits	Some are strict, and they shape the paper
Format	Two-column, template, reference style
Review model	Single blind, double blind, open. Double blind changes how you write about your own prior work
Time to first decision	Often published. Ranges from six weeks to a year
Open access and fees	Some charge substantial article processing charges. Check before submitting, not after acceptance
Indexing	Is it in the databases your field uses

That fees row catches students. An acceptance letter followed by a £2,000 invoice is a genuinely awful surprise, and many institutions have funds for it but need notice.

4.4 Conferences, if that’s your field

In computer science and much of engineering, conferences are primary. Different rules:

- **Hard deadlines.** The calendar is fixed a year out and there is no extension. Work backwards from the date.
- **Page limits are strict**, usually with a separate limit for references and appendices.
- **Rebuttal periods**, typically a week, where you answer reviewers in a few hundred words. Chapter 20 applies in compressed form.
- **Acceptance rates are published**, often 15–25% at the top venues. Plan accordingly.
- **Workshops** are an excellent first venue: lower stakes, real feedback, and often a route to a fuller paper later.

For a first paper, a workshop at a good conference is frequently a better choice than the main track, and supervisors often forget to suggest it.

4.5 Predatory venues, and how to spot one

Real, common, and they target exactly your situation: an inexperienced author with a finished manuscript.

Warning sign	What it means
Unsolicited email praising your "esteemed research"	They emailed thousands of people
Promised review in 72 hours	There is no review
Scope covering everything from nanotech to pedagogy	It's a fee collector
Fees revealed only after acceptance	
Editorial board members who don't know they're on it	Check two, by email
Not indexed anywhere your field uses	

The test that settles it: ask your supervisor or a librarian, "have you heard of this?" Publishing in a predatory venue is worse than not publishing, because the work is then effectively burned for a real venue.

FROM THE OTHER SIDE OF THE DESK

Editors can tell within a paragraph whether an author read the scope statement. A submission that opens by explaining why it fits this particular journal, in one sentence, starts from a better place than one that reads as though it was sent to five venues simultaneously. Which, incidentally, you must not do.

4.6 One at a time

Simultaneous submission to more than one venue is prohibited essentially everywhere and is treated as serious misconduct.

Preprints are a different matter and usually allowed; Chapter 16 covers that.

DO THIS NEXT

List the venues of your ten key references, and pick your three tiers today. Then read the target venue's scope statement and one recent issue, and write one sentence saying why your paper belongs there. That sentence goes in your cover letter (Appendix A).

5

Reading the Conversation You're Joining

A paper is a turn in a conversation that started before you arrived.

You can't take a good turn without knowing what's already been said, and the referee who rejects your paper for missing an obvious reference is telling you that you didn't listen.

5.1 What a literature search is for

Not to prove you read a lot. Three specific jobs:

1. **Find out whether your result already exists.** If it does, better to know now.
2. **Find the dozen papers that position your claim.** Those are your references.
3. **Learn the vocabulary.** Every subfield names things its own way, and using the wrong word marks you as an outsider.

Job 3 is the one students underestimate. If the field calls it "domain adaptation" and you call it "retraining on new data", your paper won't be found by search, and it will read as though you don't know the area.

5.2 The three-pass search

Searching once and stopping is the usual mistake. Do it in three passes, each with a different move.

Pass 3 is the one that saves you. It's where you find the 2024 paper that did your experiment, and it's much better to find it in week one than in a referee report in month eight.

> *Search until saturation, not until you have enough references. Saturation is when new searches return papers you've already seen. Anything less and you're guessing about novelty.*

5.3 Getting the search terms right

Your first search terms are your own words for the problem, and they're usually wrong.

The repair is quick. Find one paper that's definitely relevant. Look at its keywords, its title, and the terms it uses in the abstract. Search again with those. Repeat once. Two iterations usually gets you to the field's actual vocabulary.

Also worth trying:

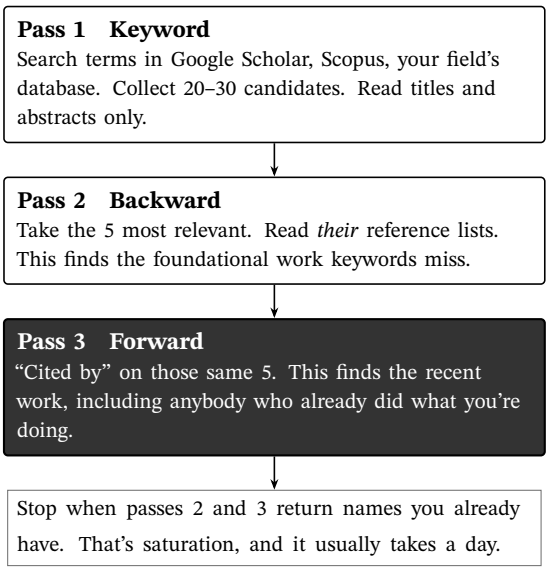

Figure 5.1: The three-pass search. Each pass uses a different mechanism, so each finds papers the others miss.

- The same idea named in a neighbouring field. Statistics, physics and computer science have three names for a surprising number of the same objects.
- The method name plus your application domain.
- The name of the person whose work yours is closest to, on its own. Their recent papers are the front line.

5.4 What to record, and how

Not PDFs in a downloads folder. For each paper you'll actually use, one short note:

Citation	Full, in your reference manager, right away
Claim	Their claim in one sentence, in your words
Relation	How yours differs: newer data, different regime, stronger result, contradicts it
Quote	One quotable line with a page number, if you'll need it

That "relation" line is the whole related-work section, written in advance, one paper at a time, while you still remember. Writing related work from a pile of unread PDFs six months later is the reason that section is so often bad.

Use a reference manager. Zotero, JabRef, Mendeley, whichever. Not a Word document with pasted citations, which will cost you an entire day at submission when the venue wants a different style.

FROM THE OTHER SIDE OF THE DESK

The reference lists I see from first-time authors have a characteristic shape: everything is either from the last two years or from a famous 1970s paper, with nothing in between. It reads as though the author searched once, sorted by date, and stopped. The middle

decade is where most of the directly relevant work usually is.

5.5 Reading efficiently

You cannot read 30 papers properly. Read them in layers.

Layer	Time	What you get
Title, abstract, figures	3 min	Whether it's relevant at all
Introduction, conclusion	10 min	The claim and how it's positioned
Full read	1–2 h	Method detail, enough to compare with yours

Of 30 candidates, roughly 20 stop at layer 1, 8 at layer 2, and 2 to 5 get a full read. Those last few are the ones your work is really in dialogue with.

5.6 When you find your result already published

It happens, and it feels like the end of the world for about a day.

Then check carefully, because "the same" is rarely the same:

- Different assumptions, or a narrower setting than yours?
- Simulated where you have real data, or one dataset where you have five?
- An existence result where you have a construction?
- Published without code, so nobody has reproduced it?

Any of those can still be a paper, provided you cite them prominently and state exactly what you add. The one thing you must not do is fail to cite them, which turns an honest overlap into something much worse.

DO THIS NEXT

Do all three passes this week and stop at saturation. Then write the one-line "relation" note for your ten closest papers. When you get to Chapter 9, the related-work section will already be drafted.

6

Planning the Paper Before You Write

Half a day of planning saves three weeks of rewriting.

The alternative, which almost every first-time author tries, is to start typing at the introduction and see where it goes. That produces a draft whose shape is the order you did the work in, and the order you did the work in is never the order a reader needs.

6.1 The one-page plan

Before any prose, fill in one page. Handwritten is fine.

Line	What goes there
Claim	One sentence. The thing you assert
Audience	Who specifically. Name the venue
Evidence	The 3–5 results that establish the claim
Figures	The 3–5 figures, each with the sentence it proves
Limits	What it does not show. Two lines
Prior work	The 5 papers you must position against
Takeaway	What a reader should do differently

If you can't fill the claim line in one sentence, you're not ready to write. Go back to Chapter 3.

If the evidence line has eight items, you have two papers, or you have one paper plus a supplement.

> *The one-page plan is the paper. Everything after it is expansion. If the plan is confused, no amount of good sentence-writing will fix the manuscript, and you'll spend the fix trying.*

6.2 Plan the figures first

This is the highest-leverage habit in the book, so it gets its own section.

Most readers see your figures and nothing else. So design them before you write, and design them as a sequence that tells the story on its own.

Sketch all four on paper, with axes labelled and the caption written, before you generate a single plot. Two things happen. You discover which experiment is missing, while there's still time to run it. And you find out that two of your six planned figures say the same thing.

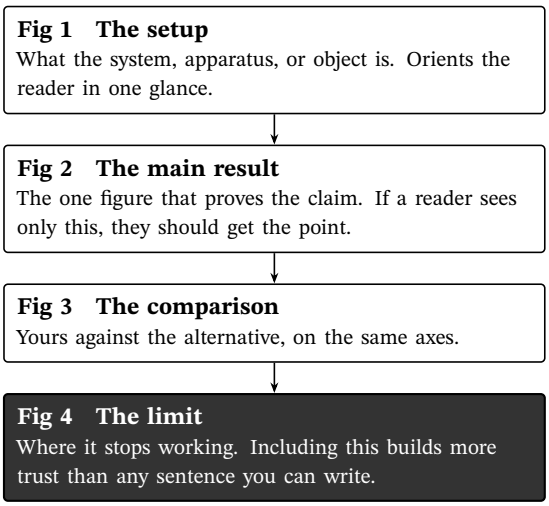

Figure 6.1: Figure-first planning. The figures carry the argument; the text explains the figures. Doing it the other way round produces figures that decorate rather than prove.

6.3 The reverse outline, for text

Once the figures exist, outline in sentences rather than headings.

Write one sentence per intended paragraph, in order, on a single page. Nothing else. Then read the page top to bottom and ask:

- Does each sentence follow from the one before?
- Is anything asserted before the reader has what they need to believe it?
- Is anything there twice?
- Can you cut three sentences and lose nothing?

Reordering a page of sentences takes ten minutes. Reordering eight pages of finished prose takes two days, and you'll resist doing it because the prose feels expensive. That's exactly why you outline first: sentences you haven't polished are cheap to throw away.

6.4 The order to write in

Not the order the reader reads.

1	Figures and tables	The evidence, finished
2	Method	Easiest. You just describe what you did
3	Results	Written against the figures
4	Discussion	Where the thinking happens
5	Introduction	Now you know what you're introducing
6	Related work	Now you know what to position against
7	Conclusion	Short
8	Abstract	Last, always
9	Title	Last, or nearly

Writing the introduction first is the classic trap. You spend two days on a framing that no longer fits by the time the discussion is done, and then you're reluctant to change it because it cost two days.

6.5 A realistic schedule

For a first paper alongside other work, from finished results to submission:

Stage	Time	Note
Plan and figures	1 week	Including the extra run you'll discover
First full draft	3 weeks	Method, results, discussion first
Supervisor read	1–2 weeks	Their calendar, not yours. Ask early
Revision	1 week	
Second read	1 week	Someone outside the project
Polish and format	1 week	Longer than you think

That's about eight weeks of elapsed time, and it assumes the results are done. Students routinely plan for three. The two supervisor reads are the part that can't be compressed, and the way to protect them is to give notice: tell them in week one when the draft will arrive.

FROM THE OTHER SIDE OF THE DESK

The single best predictor of whether a student's paper gets submitted isn't the quality of the work. It's whether they booked the supervisor's reading time in advance. Drafts that arrive unannounced in a busy month sit for six weeks, and the student concludes their supervisor isn't interested.

DO THIS NEXT

Fill in the one-page plan today, then sketch the four figures on paper with captions. Take both to your supervisor in the same meeting and ask which figure is missing. That meeting is worth more than the next two weeks of writing.

Part III

Writing It

7

The Title and the Abstract

These two are read a hundred times more often than anything else you write, and they're usually written last, in an hour, when you're tired.

Give them a day.

7.1 What a title has to do

Two jobs. Tell a scanning reader what the paper is about, and be findable by someone searching for it.

Do	Don't
Name the object and the finding	Ask a question the paper answers
Use the field's search terms	Use a pun or a colon-joke
8–15 words	25 words with four sub-clauses
Be specific enough to distinguish it	"A Study of…"

BEFORE

A Study of Machine Learning Approaches for Data Classification

AFTER

Gradient Boosting Outperforms Deep Networks on Small Tabular Clinical Datasets

The second names the method, the comparison, the domain and the finding. Somebody searching "tabular deep learning clinical" finds it. Somebody scanning a table of contents knows instantly whether to stop.

Two title styles are both acceptable. A *descriptive* title names the topic ("Thermal Stability of Perovskite Films Under Humid Cycling"). A *declarative* title states the finding ("Humid Cycling Degrades Perovskite Films Through Grain-Boundary Ingress"). Declarative titles are more useful to readers and slightly more exposed if the finding is later qualified. Mathematics and theoretical work almost always use descriptive.

> *Write the title after the abstract, and the abstract after the paper. A title written first describes the paper you intended to write.*

7.2 The abstract: five sentences

An abstract is not a summary of the paper. It's a compressed version of the argument, and it has a fixed shape.

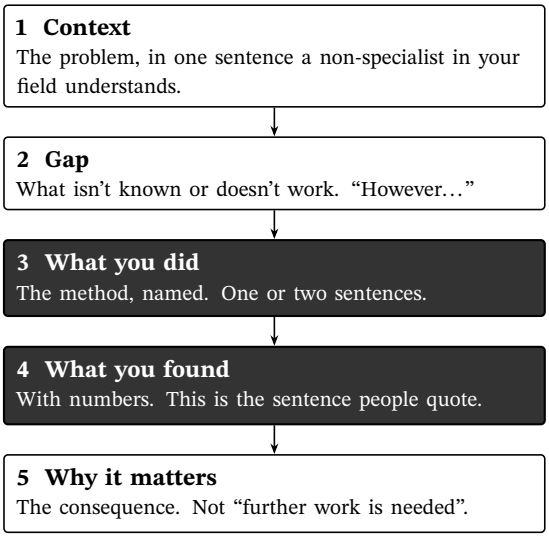

Figure 7.1: The five-sentence abstract. Most 200-word abstracts are exactly these five sentences with a little elaboration on sentences 3 and 4.

Sentence 4 is the one that must contain numbers. "Accuracy improved substantially" is worth nothing; "accuracy improved from 71% to 88% on the held-out cohort" is the reason someone reads on.

7.3 An abstract, before and after

BEFORE

> Machine learning has become increasingly important in many fields. In this paper, we investigate the application of various machine learning techniques to the problem of classification. We conducted extensive experiments on several datasets. The results show that our approach performs well. Future work will explore additional directions.

Five sentences and no information. There's no problem, no gap, no named method, no number, and no consequence. An editor reads this and desk rejects, not from malice but because nothing in it says the paper has a result.

AFTER

> Clinical prediction models are often trained on datasets of a few thousand rows, where deep networks are widely assumed to underperform. However, the comparison has rarely been made on real clinical data with proper hyperparameter budgets for both families. We benchmarked gradient boosting against three deep architectures on 14 clinical tabular datasets (n = 800–12,000), with equal tuning budgets. Gradient boosting won on 12 of 14 datasets, with a mean AUROC advantage of 0.041, and trained 22 times faster. For datasets of this size, the tuning effort spent on deep architectures is better spent on feature construction.

Same five moves. Named datasets, sample sizes, a controlled comparison, a number, and a recommendation someone can act on.

7.4 Rules that hold nearly everywhere

- **No citations** in the abstract, unless the venue's style demands it.
- **No abbreviations** you haven't defined, and prefer to define none.
- **No "we will show"** or "this paper is organised as follows". Say the thing.
- **Self-contained.** It's read separately from the paper, in search results and databases.
- **Past tense for what you did**, present for what's true.
- **Respect the word limit exactly.** Some submission systems truncate.

7.5 Structured abstracts

Some venues, particularly in medicine and parts of engineering, require labelled sections: Background, Methods, Results, Conclusions. The five moves map onto those directly. Sentences 1 and 2 become Background, 3 becomes Methods, 4 becomes Results, 5 becomes Conclusions.

7.6 Keywords

Five or six, and they matter more than the two minutes they usually get. Pick them so that:

- at least three are standard terms in your field, taken from other papers rather than invented;
- at least one names the application domain;
- none simply repeat a word already in the title, since that wastes a slot in the index.

FROM THE OTHER SIDE OF THE DESK

An editor's decision to send a paper out for review takes about four minutes. They read the title, the abstract, and the last paragraph of the introduction. If those three don't establish that there's a claim, that it's new, and that it fits the journal, the paper comes back without a review. That's not a judgement on the work. It's a judgement on three paragraphs.

DO THIS NEXT

Write your abstract as exactly five sentences, one per move, with a number in sentence 4. Then show it to someone in an adjacent field and ask them what you found. If they can't tell you, sentence 4 isn't doing its job.

8

The Introduction

The introduction is where papers are won and lost.

It's the last thing the editor reads before deciding, the first thing the referee reads properly, and the place where the claim either appears or doesn't.

8.1 The five-paragraph introduction

There's a standard shape, and it works. Use it until you have a reason not to.

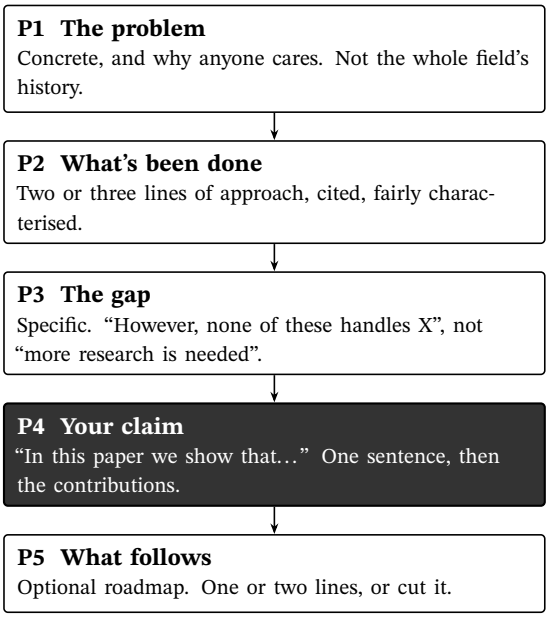

Figure 8.1: The five paragraphs. Paragraph 4 is the one first papers most often leave out, and it's the one the whole review turns on.

The whole thing is one page in most formats. If your introduction is three pages, the extra two are related work in the wrong place.

8.2 Paragraph 1: start concrete

The commonest opening in student papers is the widening-gyre sentence: "In recent years, machine learning has attracted increasing attention in many fields."

Nobody has ever read that sentence and learned anything. Start one level narrower and one level more specific.

BEFORE

In recent years, energy storage has become increasingly important due to growing global demand for renewable energy sources.

AFTER

Grid-scale lithium-ion installations lose 8–12% of capacity in their first two years, and the loss is dominated by a mechanism that only appears above 35°C. Most cycling protocols test at 25.

Numbers, a mechanism, a specific mismatch. Two sentences in, the reader knows what the paper is about.

> *Delete your first sentence and read what's left. If the paragraph is better, the sentence was throat-clearing. This is true about 70% of the time.*

8.3 Paragraph 3: the gap has to be real

A gap is not "nobody has done exactly my thing". That's true of everything and impresses no one. A real gap is one of:

Untested assumption	Everyone assumes X; nobody has checked
Unexplored regime	It works at low temperature; nobody knows about high
Conflicting results	Two groups disagree and nobody resolved it
Missing comparison	Three methods exist and have never been compared on the same data
Practical barrier	The method works but needs data nobody has

Whichever it is, name it in a sentence beginning "However" or "What remains unclear is". Referees look for that sentence.

And be fair to the prior work while you do it. "Smith et al. ignored temperature effects" is a bad sentence for two reasons: it's usually untrue (they scoped their study), and Smith is quite likely your referee. "Smith et al. characterised the mechanism at 25°C; the behaviour above 35°C has not been reported" says the same thing and is both accurate and safe.

8.4 Paragraph 4: state the claim, then list the contributions

Write the claim as a sentence starting "We show that" or "We demonstrate" or "This paper establishes". Explicitly. Then a short list.

> We show that grain-boundary ingress, not bulk decomposition, dominates degradation above 35°C. Specifically:
>
> - We report cycling data for 24 cells across four temperatures, with per-cell impedance spectra (Section 3).
> - We show that the Arrhenius fit used in current lifetime models breaks down above 35°C, over-predicting life by 40% at 45 (Section 4).

- We propose a two-mechanism model that fits all four temperatures with two parameters, and release the fitting code (Section 5).

Three things that list does. It gives the referee a checklist they can tick, it tells them where to look, and it forces you to be honest, because a contribution you can't point at a section for isn't one.

Keep it to three or four items. Seven bullets means each is small.

FROM THE OTHER SIDE OF THE DESK

When I read a manuscript, I look for one sentence: the one that starts "we show that". If I can't find it in the introduction, I write the review around that absence, because every subsequent judgement depends on knowing what's being claimed. Authors are always surprised. They knew what they were claiming. They just never wrote it in a sentence.

8.5 Things to keep out of the introduction

- **Method detail.** It has a section.
- **Results in numbers**, beyond a single headline figure.
- **A history of the field** from 1950.
- **Definitions** of standard terms your readers know.
- **Apology.** "Due to time constraints, we were only able to…" belongs in the limitations, phrased differently.

8.6 On the roadmap paragraph

"Section 2 reviews related work. Section 3 describes our method…"

Some venues expect it, most readers skip it, and it costs you a quarter of a page. Include it if the paper's structure is unusual, cut it if the structure is standard. In a conference paper with a hard page limit, cut it first.

DO THIS NEXT

Take your draft introduction and mark each paragraph P1 to P5. If two paragraphs have the same label, merge them. If P4 is missing, write it now, in one sentence beginning "we show that". That sentence is your paper.

9

Related Work Without the Laundry List

The most commonly bad section in a first paper, and the easiest to fix.

Its job is not to prove you read the literature. Its job is to place your claim in the field so a reader can see exactly what's yours.

9.1 The list versus the map

Here's the shape almost every first draft has:

BEFORE

> Smith et al. [1] proposed a method for X. Jones et al. [2] developed an approach using Y. Kumar and Lee [3] applied Z to the same problem. Wang et al. [4] extended [2] with additional features. Chen [5] studied the effect of temperature.

Five sentences, five papers, no organisation, and crucially, no relationship to your work. A referee reads that and learns nothing except that you have a reference manager.

AFTER

> Existing lifetime models fall into two families. **Empirical Arrhenius fits** [1, 3, 5] treat degradation as a single thermally activated process, which fits well below 30°C but has only been validated in that range. **Physics-based multi-mechanism models** [2, 4] resolve separate pathways but require impedance data at each temperature, which routine testing does not collect. Our approach sits between them: two mechanisms, but fitted from capacity data alone.

Same five references. Now they're organised into families, each family has a stated limitation, and the last sentence positions your work. That's a map.

> *Group the literature into two to four families, characterise each in a sentence, name what each family can't do, then say where you sit. Never a paper-per-paragraph walk.*

9.2 The three-column table you write first

Before drafting the section, make this table. It never goes in the paper; it makes the section write itself.

Paper	What they did	Why yours differs
Smith 2021	Arrhenius fit, 15–30°C	Doesn't cover our regime
Jones 2022	Two-mechanism, needs EIS	We fit from capacity alone
Kumar 2023	Same regime as ours	Simulated cells; ours are physical

Column 3 is the section. If you can't fill column 3 for a paper, ask why it's cited. Sometimes the answer is that it shouldn't be, and cutting it strengthens the section.

9.3 Where the section goes

Two conventions, and it depends on your field.

After the introduction	Standard in CS, engineering. A separate Section 2. Good when the field needs mapping before your method makes sense
Woven into the introduction	Standard in mathematics, physics, much of the life sciences. No separate section; positioning happens in P2 and P3
Near the end	Occasionally, when the comparison only makes sense after readers know your method. Uncommon; check the venue does it

Look at three recent papers in your target venue and do what they do.

9.4 Length and proportion

Ten percent of the paper. In a ten-page paper, one page.

If yours is three pages, you've written a literature review. That's a different genre with its own venues, and a referee will say so.

The reason it bloats is thesis inheritance: the background chapter was 20 pages because the examiner wanted evidence you'd read the field. The referee wants the opposite, which is evidence you can identify the six papers that matter.

9.5 How to cite

- **Cite what you read.** Not what you found in someone else's reference list. Citing a paper you haven't opened is how errors propagate for decades.
- **Cite the original**, not the review that mentions it, when you're attributing an idea.
- **Cite recent work**, especially anything from the last eighteen months. A reference list ending two years ago says the search was stale.
- **Cite your referee's paper** when it's genuinely relevant. Not flattery, just accuracy: the editor picks referees who work in the area, so the relevant people probably are the reviewers.
- **Don't over-cite yourself.** Two or three self-citations in a first paper is fine; eight looks like padding, and under double-blind review it also breaks anonymity.

FROM THE OTHER SIDE OF THE DESK

Referees notice one thing above all in related work: whether their own closest paper is cited, and whether it's characterised correctly. Getting somebody's contribution wrong in print annoys them more than not citing them at all. If you're describing a paper in a sentence, make sure the sentence is one the author would agree with.

9.6 Double-blind review changes the wording

Under double blind, you can't write "in our previous work [7] we showed". That identifies you.

The standard fix is third person: "Prior work [7] showed". Cite normally, just don't claim it. And check whether the venue wants prior work anonymised in the submission version; some do, most don't.

Similarly, avoid "our group's dataset" and named internal tools that identify the lab. This trips up nearly everyone the first time.

DO THIS NEXT

Build the three-column table for your ten key references. Then write the section as: one paragraph per family, ending with one paragraph that says where you sit. If it runs past a page, you've got a family too many.

10

The Method Section

The easiest section to write and the one most likely to sink a resubmission.

Easy, because you're describing what you did and you were there. Dangerous, because the referee's job is to decide whether they believe it, and every detail you omit is a reason to doubt.

10.1 The standard: could a competent stranger repeat it?

Not you in six months. A competent person in your field, with your paper and nothing else.

That standard resolves nearly every "should I include this?" question.

Include	Leave out
Anything that changes the result if changed	Anything that wouldn't
Versions, settings, quantities, durations	The brand of your laptop
Why you chose this design	Your search for the design
What you excluded and why	The failed pilot, unless it's informative

The second column has one exception worth naming: *exclusions*. Every excluded sample, dropped run, or filtered data point must be stated with its rule and its count. This is where honest work most often looks dishonest, purely because the numbers don't add up and nobody said why.

> *Every number in your results must be traceable to a stated procedure. If a referee can't get from your method to your results section without guessing, you'll get a major revision, and you'll deserve it.*

10.2 What a method section contains

Order varies by field, but the content is stable.

Block 4 deserves the emphasis. "We used a learning rate of 0.001" is half an answer. "We used a learning rate of 0.001, selected by grid search over $\{10^{-4}, 10^{-3}, 10^{-2}\}$ on the validation split" is the whole one, and it's the difference between a tuned result and an arbitrary one.

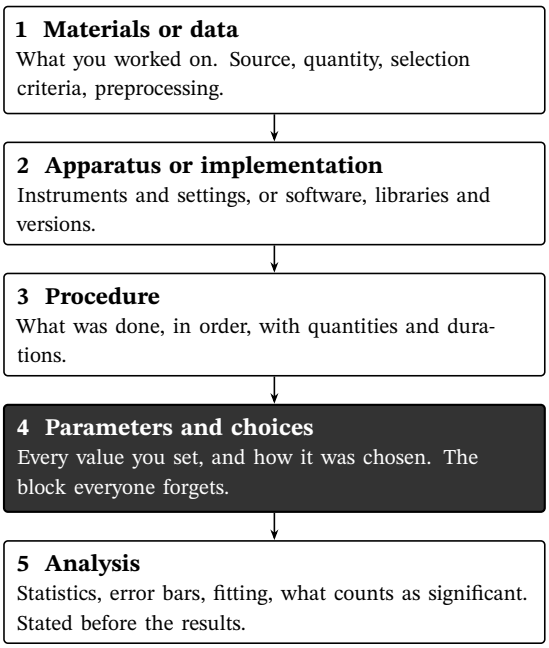

Figure 10.1: The five blocks of a method section. Missing block 4 is the single most common cause of "insufficient detail to reproduce".

10.3 Tense, voice, and the passive

Past tense for what you did. Present tense for what is true.

> "Samples *were annealed* at 400°C for two hours."
> "The model *assumes* isotropic diffusion."

On the passive voice: it's fine here, and here only. "We annealed" and "samples were annealed" are both acceptable in a method section, and fields differ. Elsewhere in the paper, prefer the active voice, because "it was observed that" hides who observed it.

Check the venue. Some journals now explicitly ask for the active voice throughout.

10.4 Being honest about what you don't know

Sometimes a parameter was inherited from a previous student, or a setting came from a lab default nobody has questioned in a decade.

Say so. "The annealing schedule follows the laboratory's standard protocol [ref], and was not varied in this study" is a legitimate sentence. It's honest, it's reproducible, and it pre-empts the obvious question.

What you must not do is invent a rationale after the fact. If you tuned until it worked, describe the tuning. Referees have done the same thing and they recognise the smell of a story constructed backwards.

10.5 Supplements and where to put detail

Most venues allow supplementary material, and it exists for exactly this problem: detail that's necessary for reproduction and destructive to readability.

Main text	Supplement
Everything needed to understand the result	Everything needed to repeat it exactly
The design, and why	Full parameter tables
Sample size and selection rule	Per-sample listings
The statistical approach	Full model output, diagnostic plots
	Instrument calibration records

The rule for the split: a reader should be able to *follow* the paper without the supplement, and *repeat* the work with it. And put the code and data links here too, with Chapter 14 covering what that involves.

FROM THE OTHER SIDE OF THE DESK

The referee question I've seen sink the most resubmissions is some version of "how many samples were excluded, and by what rule?" It's asked because the numbers in Table 2 don't reconcile with the sample count in the method. The work was almost always fine. The reporting wasn't, and by the time it's asked, the referee has started to wonder what else is missing.

10.6 The reconciliation check

Do this before submitting, and it takes twenty minutes.

1. Write down every sample size, count, and n that appears anywhere in the paper, including in figure captions.
2. Check they add up: collected, minus excluded, equals analysed.
3. Check the excluded count has a stated rule.
4. Check every parameter in the results has a value in the method.

Most first submissions fail at least one of those four, and every failure is free evidence for a sceptical referee.

DO THIS NEXT

Give your method section to someone in your lab who didn't do the work, and ask them to list what they'd need to repeat it. Everything on their list that isn't in the section goes in, or goes in the supplement. Twenty minutes of their time, and it's the single highest-yield check in this book.

11

Results, Figures, and Tables

The results section reports what happened. It does not explain what it means.

Keeping those apart is the discipline of the section, and mixing them is the most common structural fault in a first paper.

11.1 Report, don't interpret

BEFORE

Capacity fade was much worse at 45°C, which clearly shows that the Arrhenius model is inadequate and suggests that a second degradation mechanism becomes active at elevated temperature.

AFTER

Capacity retention after 500 cycles was 91.2% at 25°C and 68.4% at 45°C (Fig. 3). The Arrhenius fit calibrated on the 25–35°C data predicted 81% at 45°C, over-predicting retention by 12.6 percentage points.

The second states numbers and a discrepancy. The interpretation ("a second mechanism") is a claim about causes, and it belongs in the discussion, where the reader expects argument rather than observation.

Why it matters: a referee who agrees with your numbers but not your interpretation can accept a paper that separates them. If they're entangled, doubting the interpretation means doubting the results.

> *Results answer "what happened". Discussion answers "what does it mean". A results paragraph containing "because", "suggests", or "demonstrates that" has usually crossed the line.*

11.2 The order of results

By argument, not by chronology. Nobody needs to know which experiment you ran first.

A workable default:

1. The headline result that establishes the claim.
2. Supporting results that strengthen it.
3. Controls, and results that rule out alternative explanations.
4. Where it breaks down.

Point 4 belongs in the results, stated flatly. Burying a failure in the discussion looks like concealment; reporting it in its place looks like competence.

11.3 Figures: the rules that matter

Your figures will be looked at by more people than will read your sentences. Give them the time that implies.

One point each	A figure making three points makes none
Axes labelled with units	Every time. No exceptions
Readable at print size	Reduce to column width and read it. Most journals print at 8.5 cm
Legible in greyscale	Many readers print in black and white. Use markers and line styles, not colour alone
No chartjunk	3D bars, gradients, shadows, and gridlines that don't help all cost credibility
Error bars defined	In the caption: standard deviation, standard error, or confidence interval, and n
Consistent across figures	Same colour for the same condition in every plot

That greyscale point catches people. Roughly one in twelve men has some degree of colour vision deficiency, and red-green is the common one, so a red-versus-green comparison is unreadable for a meaningful share of your audience even in colour.

11.4 Captions do real work

A caption should let a figure stand alone. Three parts:

> **Figure 3.** *[What it shows]* Capacity retention against cycle number at four temperatures. *[How to read it]* Solid lines are measured means over six cells; shaded bands are ± 1 standard deviation. The dashed line is the Arrhenius prediction calibrated on the two lowest temperatures. *[What to conclude]* The prediction tracks the data to 35°C and diverges above it.

That third part is the one students omit, and it's the one that makes the figure useful to a skimmer. A caption that ends without telling the reader what to conclude has left the most important job undone.

11.5 Tables

Use a table when the reader needs exact values, and a figure when they need a pattern. Never both for the same data.

- **No vertical rules.** Horizontal rules at the top, under the header, and at the bottom. That's the convention in nearly every venue and it's genuinely more readable.
- **Align numbers on the decimal point.**
- **Consistent significant figures** down a column, and no more than your measurement justifies. Reporting 87.4213% from six samples claims a precision you don't have.
- **Units in the header**, not repeated in every cell.
- **Bold the best value** in a comparison, and say in the caption that bold means best.

11.6 Statistics, stated plainly

Whatever your field's conventions, three things must be unambiguous:

What test	Named, with the assumption it requires and whether you checked it
Sample size	Per group, in the text or the caption. Not just the total
Effect size	Not only the p-value. A significant difference of no practical consequence is worth saying out loud

And a caution: if you ran twenty comparisons and reported the one that reached significance, say that you ran twenty. Referees increasingly ask, and correcting for multiple comparisons is expected in most fields now.

FROM THE OTHER SIDE OF THE DESK

The fastest way to lose a referee's confidence is a figure whose axis is unlabelled or whose n is unstated. It's a small thing, and that's precisely why it reads as carelessness. If the obvious things aren't checked, the reviewer starts checking the non-obvious ones much harder.

11.7 The five-minute figure audit

Before submission, for every figure:

1. Print it at final size. Can you read the smallest text?
2. Print it in greyscale. Can you still tell the conditions apart?
3. Cover the text and read only the caption. Do you know what to conclude?
4. Is every axis labelled with units?
5. Is n stated, and is it the same n as in the method?

DO THIS NEXT

Run the five-minute audit on your main figure today. Then hand it, with only its caption, to someone who hasn't seen the work and ask what it shows. Whatever they say is what the figure actually communicates, regardless of what you meant.

12

Discussion, Limitations, and Conclusion

The discussion is where a paper stops reporting and starts thinking.

It's also the section where experienced readers decide whether the author understands their own work, which is why a thin discussion attached to good results is such a common and such a fixable weakness.

12.1 What the discussion does

Four jobs, in roughly this order.

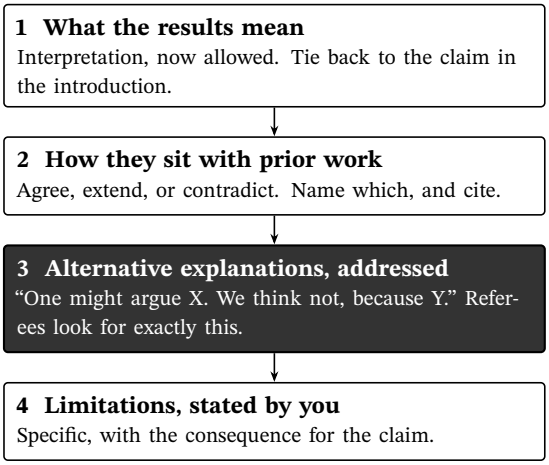

Figure 12.1: The four jobs of a discussion. Job 3 is the one that converts a sceptical referee, and the one first papers most often skip.

Job 3 is worth dwelling on. If you can anticipate the referee's objection and answer it in the paper, you have converted a rejection into a sentence. If you can anticipate it and *can't* answer it, you've at least shown you know it's there, which is a much better position than appearing not to have thought of it.

> *Write down the three objections a hostile expert would raise. Address all three in the discussion. This is the highest-return hour you will spend on the manuscript.*

12.2 Contradicting prior work, carefully

If your result disagrees with a published one, that's interesting, and you have to handle it properly.

BEFORE

Our results contradict those of Kumar et al. [3], whose experimental setup was likely flawed.

AFTER

Our retention figures at 45°C are 9 points below those reported by Kumar et al. [3]. The two studies differ in cell format (pouch versus cylindrical) and in charge rate (1C versus 0.5C). Either could account for the difference; distinguishing them would require cycling both formats at matched rates, which we did not do.

The second names the discrepancy, offers candidate explanations, and says what would settle it. It's more useful, it's more honest, and it doesn't make an enemy of a likely referee.

12.3 Limitations: how to write them

Every paper has limitations. Stating them well makes you more credible, not less, and the way to do it badly is either extreme.

Too vague	"Our study has some limitations. Future work should address them." Says nothing
Too abject	"Due to severe time constraints we were unable to do anything properly." Invites rejection
Right	Specific limitation, its consequence for the claim, and what would resolve it

> All cells came from a single manufacturing lot, so lot-to-lot variation is not captured. Our absolute retention figures may therefore not transfer to other lots, though the temperature dependence, which is our claim, is unlikely to be lot-specific. Replication across three lots would settle this.

Notice the middle sentence. It scopes the limitation to the part of the paper it actually threatens, and it protects the claim. A limitation without that sentence lets the referee decide how much damage it does, and referees are not generous with that decision.

12.4 Future work, without the padding

"Future work will explore other datasets" is filler. Everyone knows future work exists.

Write two or three concrete next steps that follow from what you found, and only those. The test is whether somebody could actually start it on Monday.

> Two extensions follow directly. First, the two-mechanism fit should be tested on cells with different cathode chemistries, where the grain-boundary pathway may not dominate. Second, the crossover temperature we observe near 35°C is a fitted parameter; measuring it directly by impedance would test whether it corresponds to a physical transition.

12.5 The conclusion

Short. In a ten-page paper, one paragraph, occasionally two.

Three things, and nothing else:

1. Restate the claim, now with the evidence behind it.
2. State what it changes for someone working in the area.
3. One line on what's next, if it hasn't been said.

Two prohibitions. **No new results in the conclusion**, ever. And **no new claims** that the paper hasn't supported. Both make a referee go back and re-read suspiciously, which is the last thing you want on the final page.

FROM THE OTHER SIDE OF THE DESK

There's a tell that separates a first paper from a fifth. In the first, the discussion restates the results in different words. In the fifth, the discussion argues with itself: it raises the objection, weighs it, and concludes. Referees read that argument as evidence the author knows what their result does and doesn't establish, and they trust the rest of the paper more because of it.

12.6 Checking discussion against introduction

A quick structural check, and it catches a surprising number of problems.

Read the last paragraph of your introduction and the first paragraph of your discussion, one after the other, ignoring everything between.

- Do they refer to the same claim?
- Does the discussion answer the question the introduction asked?
- Has the claim quietly grown between the two?

That third one is common. The introduction claims something modest, the discussion claims something larger, and the referee notices even when the author doesn't. If they've drifted, the introduction is almost always the one to change, because the discussion is closer to what the evidence supports.

DO THIS NEXT

Write your three hostile objections on a card before drafting the discussion. Address each in a named paragraph. Then run the introduction-to-discussion check above. Half a day, and it's the difference between "accept with minor revisions" and "major revisions".

Part IV

Authorship, Data, and Honesty

13

Authorship, and the Conversation Nobody Has

Authorship disputes are the most common serious conflict in research, and almost all of them come from the same cause: nobody discussed it until the paper was nearly finished.

Have the conversation in week one, when it's easy.

13.1 What authorship means

Authorship means you contributed substantially to the work *and* you take responsibility for it. Both halves.

Most fields use some version of four criteria. You should be an author if you:

1. contributed substantially to the conception, design, data collection, or analysis;
2. drafted the paper or revised it critically for intellectual content;
3. approved the final version;
4. agree to be accountable for the work.

The strict reading requires all four. In practice fields vary, and the common failure isn't the criteria, it's that nobody applied them.

13.2 Who is not an author

Contribution	Proper credit
Funded the lab	Acknowledgement, funding statement
Provided a sample or dataset	Acknowledgement, unless they shaped the study
Ran a routine measurement	Acknowledgement
Head of department	Nothing, unless they did the work
Proofread it	Acknowledgement
Gave useful advice over coffee	Acknowledgement

Adding a senior name that didn't contribute is *gift authorship*. Leaving off someone who did is *ghost authorship*. Both are misconduct, both are common, and as a student you may be pressured into the first. Chapter 15 says more about what to do if you are.

13.3 Author order, which varies enormously

There is no universal convention, and assuming your field's is universal causes trouble in collaborations.

Mathematics	Alphabetical, almost always. Order carries no information
Physics	First author did the work; last is often the group leader; huge collaborations use alphabetical
Life sciences	First is the main contributor, last is the senior author, middle roughly by contribution
Computer science	First did the most; order roughly by contribution; last often the supervisor
Engineering	Mixed. Ask

For a first paper from your project, you are usually first author and your supervisor is usually last. If that's not what's being proposed, ask why, politely and early.

> *Agree the author list and the order in writing, at the start, and revisit it if the contributions change. An email saying "just to confirm what we discussed" is enough, and it prevents almost every dispute.*

13.4 How to have the conversation

Students dread this. It's easier than you think, because you're not negotiating, you're asking a procedural question.

Something like:

> "Before I start drafting, can we agree the author list? My understanding is that it'd be me, then Priya for the imaging, then you. And I'd put Dr. Rao in the acknowledgements for the sample prep. Does that match how you see it?"

Concrete, specific, and it invites correction rather than confrontation. Send a one-line email afterwards recording what was agreed.

If it changes later, because someone runs a new experiment or drops out, raise it then rather than at submission.

13.5 Contribution statements

Many venues now require a CRediT-style statement naming who did what: conceptualisation, methodology, software, validation, formal analysis, investigation, data curation, writing (original draft), writing (review and editing), visualisation, supervision, funding acquisition.

This is good for you. A first author who wrote the draft, built the software, and ran the analysis has that on record, which matters for PhD applications and fellowships in a way that position in a list doesn't.

Draft it yourself and circulate it. It also surfaces disagreements about contribution while they're still fixable.

13.6 The corresponding author

Handles submission, correspondence with the editor, and post-publication queries. Often the supervisor, and there's an argument for it being you.

If you're about to leave the institution, make sure the listed email is one you'll still read, or that a co-author's is. Papers get queries years later, and an unreachable corresponding author is a real problem.

FROM THE OTHER SIDE OF THE DESK

The disputes that get ugly are almost never about who did the work. They're about a promise that was implied and never stated: the student who assumed they'd be first author because it was their project, the postdoc who assumed they'd be included because they set up the rig. Five minutes of explicit conversation in month one would have prevented every one I've seen.

13.7 If you're not on a paper you should be

It happens, more often to students than to anyone else.

1. **Assume a mistake first.** Ask directly and without accusation: "I wanted to check about the authorship on the X paper. I'd understood I'd be included for the analysis. Has something changed?"
2. **Put it in writing** if the first conversation doesn't resolve it. Describe your contribution factually against the criteria.
3. **Escalate to a neutral party** if it still doesn't. Every institution has a research integrity officer, and a graduate school will usually have someone whose job includes exactly this.
4. **Keep records throughout.** Lab notebooks, commit histories, dated drafts, and emails. Not because you expect a fight, but because they make the factual question unarguable if there is one.

13.8 If you have no supervisor

Independent work happens: after graduation, outside an institution, or alongside a job.

It's harder but not impossible. Two cautions. Some venues require an institutional affiliation; list yourself as independent, which is accepted. And you lose the informal read that catches obvious problems, so buy that back deliberately by asking two people in the field to read a draft. Most researchers will read ten pages for a polite stranger who asks well.

DO THIS NEXT

Send the authorship email this week, before you start drafting. List the proposed authors, the order, and who goes in the acknowledgements, and ask whether it matches everyone's understanding. It's the cheapest insurance in research.

14

Data, Code, and Reproducibility

More and more venues require you to share data and code, and more and more referees check.

The good news is that the work involved is small if you do it as you go and enormous if you leave it to the week of submission.

14.1 Why it's worth doing properly

Not only compliance. Three self-interested reasons.

- **Papers with available code get cited more.** Consistently, across fields.
- **Referees trust it more.** A referee who can see the analysis script asks fewer sceptical questions about the analysis.
- **Your future self needs it.** The revision arrives eight months later, the referee asks for one more plot, and you cannot remember which script made Figure 3.

That last one is the argument that convinces people. Every researcher has lost a week to it once.

14.2 What a shareable repository contains

Not your working directory. A cleaned version, and it takes an afternoon.

The environment file is the difference between code that runs and code that produces a stack trace on somebody else's machine in 2029. It takes one command to generate.

> *The test for a repository: clone it onto a machine that has never seen your project, follow the README, and reproduce one figure from the paper. If you can't, nobody can.*

14.3 Where to put it

GitHub is fine for working code, and it isn't an archive, because repositories can be deleted or renamed and the URL then dies.

For the version cited in the paper, use an archive that issues a DOI:

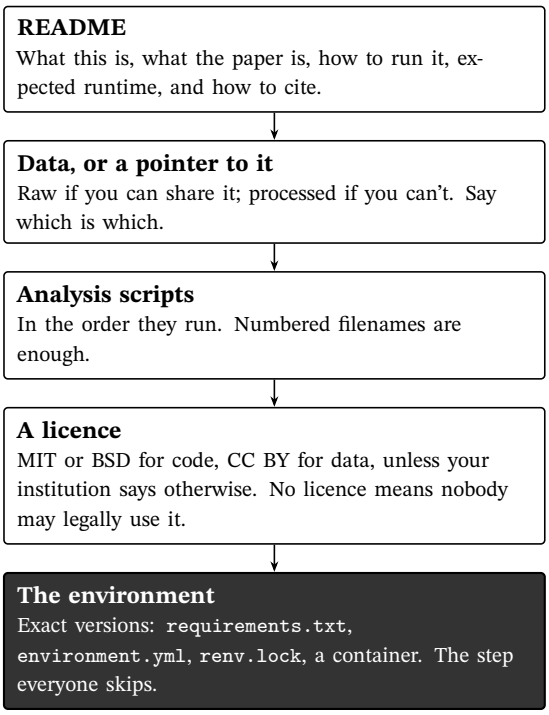

Figure 14.1: The five things a usable repository has. Missing the last one is why most shared code doesn't run for anybody else.

Zenodo	General purpose, free, integrates with GitHub to archive a release and mint a DOI
Figshare	General purpose, common in the life sciences
Dryad	Data, particularly biology
OSF	Whole projects, preregistration, common in psychology
Field archives	Many fields have one. Ask, and prefer it if so

Workflow that takes ten minutes: tag a release on GitHub, connect the repository to Zenodo, let it archive the release, and cite the DOI in the paper. Development continues on GitHub; the citation points at a frozen copy.

14.4 When you can't share the data

Common and legitimate: human subjects, patient records, commercial partners, export-controlled material, third-party data you licensed.

Say so explicitly, in a data availability statement, and share what you can.

> The clinical records analysed here cannot be shared publicly under the terms of the ethics approval (ref. 2025/117). De-identified summary statistics, the full analysis code, and a synthetic dataset with matched structure are available at [DOI]. Requests for access to the underlying data may be made to the data controller at [address].

That's a good statement. It names the restriction, gives the reference, shares the code regardless, provides a synthetic substitute so the code is runnable, and tells people how

to ask. “Data available on reasonable request” on its own is increasingly treated as a non-answer, and studies that have tested it find most such requests go unanswered.

14.5 Reproducibility for non-computational work

The principle transfers even without code.

- **Instrument settings and calibration records** in the supplement.
- **A step-by-step protocol**, more detailed than the method section, deposited on a protocols repository.
- **Raw traces or images**, not only the processed ones.
- **Material sources**, with supplier and catalogue number.
- **The spreadsheet**, if analysis was done in one, with formulas intact.

That last one causes discomfort, and it shouldn’t. A spreadsheet with visible formulas is more reproducible than one where only the numbers were copied into the paper.

FROM THE OTHER SIDE OF THE DESK

The reproducibility failure I see most often isn’t fraud or even carelessness. It’s a figure produced by a script that was edited after it ran. The numbers in the paper are real; the script in the repository produces different ones; nobody can tell why. Tag a commit when you generate the final figures, and record which commit made which figure.

14.6 Habits that make this cheap

Adopted during the project, these cost nothing. Adopted at submission, they cost a week.

Version control from day one	Even alone. Even for a thesis
One script per figure	Named for the figure it makes
Never edit raw data	Read it, write derived files elsewhere
Seed your random numbers	And record the seed
Write the README first	Then keep it current
A `make figures` command	One command, all figures, from raw data

The last row is the whole thing in one line. If a single command regenerates every figure from raw data, the paper is reproducible almost by definition.

DO THIS NEXT

This week, make a clean copy of your project in a fresh folder, write the README, freeze the environment, and regenerate one figure from raw data. Whatever breaks is what would have broken for a referee, and it’s much cheaper to find now.

15

Honesty: Plagiarism, Ethics, and AI Tools

Most research misconduct isn't committed by villains. It's committed by people under pressure who took a shortcut they told themselves was small.

Knowing where the lines are, before you're tired and behind schedule, is most of the protection you need.

15.1 Plagiarism, including the kinds you didn't know about

Type	What it is
Direct	Copying text without quotation and citation
Mosaic	Rewording sentences from a source, keeping the structure, citing the source. Still plagiarism
Self	Reusing your own published text without saying so. Yes, this counts
Idea	Using someone's idea or design without credit, even in your own words
Figure	Redrawing a figure without attribution

Mosaic is the one that catches honest students. You read a paragraph, put it in your own words, cite the paper, and believe you've done it properly. If the sentence order and argument structure are theirs, you haven't.

The fix is mechanical: read the source, close it, wait, and then write from your understanding. If you can't write it without the source open, you don't understand it well enough to summarise it yet.

> *If you can't state something in your own words with the source closed, either quote it properly or go and understand it. Those are the only two options.*

15.2 Self-plagiarism and thesis reuse

Reusing your own conference paper in a journal version, or your thesis chapter in a paper, is normal and expected. It's also something you must declare.

- **Conference to journal:** most venues require a stated percentage of new material, often 30%. Cite the conference paper and say in the cover letter what's new.

- **Thesis to paper:** generally fine, since theses aren't usually considered prior publication, but check the venue's policy and mention it in the cover letter.
- **Method sections:** reusing your own method description across your own papers is the grey area most people live in. Cite the earlier paper and rewrite rather than paste.

15.3 Data honesty, and the pressures on it

The bright lines, stated plainly:

- Never invent data. Obviously.
- Never alter an image beyond uniform adjustments applied to the whole image, and say what you did.
- Never drop an inconvenient data point without a rule that was decided in advance and is stated in the paper.
- Never run analyses until one is significant and report only that.
- Never present an exploratory finding as though it were the hypothesis you started with.

Those last three have names now: *p*-hacking, selective reporting, and HARKing (hypothesising after results are known). They're widespread, they often feel like tidying rather than cheating, and referees are increasingly alert to all three.

The honest version of an exploratory finding is easy to write: "This comparison was not planned; it emerged during analysis and should be treated as hypothesis-generating." That sentence costs you nothing and protects everything.

15.4 Ethics approvals

If your work involves people, animals, human tissue, or personal data, you need approval, and you need it *before* collecting, not before publishing.

State it in the method: the committee, the reference number, and the consent procedure. Most venues will not review a paper without it, and no committee can grant retroactive approval, so data collected without it usually cannot be published at all.

If you're a student who collected data without realising approval was needed, tell your supervisor immediately. It's occasionally recoverable and never recoverable by waiting.

15.5 Conflicts of interest

Declare funding sources, commercial relationships, patents, and anything that could reasonably look like an interest.

Declaring one is not an admission of bias. Failing to declare one, and having it surface later, is what damages people.

15.6 AI writing tools, and where the line is

This one is new enough that policies vary, so check your venue's. The current consensus across major publishers is reasonably consistent, and it draws the line in a sensible place.

Generally accepted	Generally not
Grammar and language editing	Generating results or data
Improving the clarity of text you wrote	Generating text you present as your own analysis
Translation	Writing sections you haven't verified
Code assistance you then verify	Producing references without checking they exist
Summarising papers you also read	Listing an AI system as an author

Three practical rules that keep you safe regardless of the policy details:

1. **You are accountable for every sentence.** If it's wrong, it's your error, and "the tool wrote it" is not a defence anywhere.
2. **Check every reference exists and says what you claim.** Fabricated citations are the characteristic failure mode, and they're trivially detected.
3. **Disclose use where the venue asks.** Many now have a statement field. Fill it in honestly.

No major publisher permits an AI system as an author, on the reasoning that authorship requires accountability and a tool cannot be accountable.

FROM THE OTHER SIDE OF THE DESK

The retractions I've watched happen to early-career researchers almost never began with a decision to deceive. They began with a deadline, a dataset that nearly worked, and a small adjustment that seemed defensible at the time. The protection isn't virtue. It's writing down the analysis plan before you look at the data, so the small adjustment becomes visible as a change rather than feeling like a judgement call.

15.7 If you're pressured

It happens: a supervisor who wants a name added, a co-author who wants an outlier removed, a lab that wants a result framed more strongly than the data supports.

1. **Ask for the reasoning in writing.** Often that alone ends it, and where it doesn't, you have a record.
2. **Propose the honest version.** "Could we report both, with and without the outlier?" resolves a surprising number of these, because the honest version is frequently fine.
3. **Talk to someone outside the group.** Every institution has a research integrity officer, and the conversation can be informal and confidential.
4. **Don't sign off on something you believe is wrong.** You're accountable as an author, and that accountability outlives the project by decades.

This is genuinely hard when the person has power over your future. It's still the right answer, and asking for advice early, before anything is submitted, keeps the options open.

DO THIS NEXT

Before you analyse anything, write down your analysis plan: what you'll test, what your exclusion rules are, and what counts as support for your claim. One page, dated. It costs an hour and it makes every later decision a documented one.

Part V

Submission, Review, and After

16

Preprints

A preprint is your manuscript, posted publicly, before or during peer review.

In some fields it's universal and in others it's still unusual, so the first question is what your field does.

16.1 What you get

- **A timestamp.** Public evidence of what you had and when, which matters if someone else is close to the same result.
- **A citable link now**, rather than in fourteen months. Useful for job and PhD applications while the paper is under review.
- **Readers**, and occasionally useful comments from people who would never have refereed it.
- **A DOI**, so the work is findable and countable.

16.2 What you risk

- **It's public with its errors in.** You can post a new version, and the old one stays visible.
- **A few journals still refuse** papers with preprints. Rare now, and worth checking rather than assuming.
- **It weakens double-blind review** in practice, since a referee can search a distinctive title. Most venues accept this; some ask you not to post during review.
- **Being scooped**, which is the fear people raise. In practice the timestamp protects you more than the exposure endangers you.

16.3 Field norms, which vary a lot

Mathematics	arXiv, essentially universal. Not posting is the unusual choice
Physics	arXiv, universal in most subfields
Computer science	arXiv, very common, especially in machine learning
Engineering	arXiv or engrXiv; growing, not yet default
Life sciences	bioRxiv, medRxiv; now common and increasingly expected
Chemistry	ChemRxiv; slower uptake

Ask your supervisor before posting, both because norms differ and because they may have a reason to wait, such as a patent filing.

> *Post the preprint at the moment you submit to the journal, not before you have a submittable manuscript. Preprints are public; a rough draft posted publicly is a rough draft everybody can find.*

16.4 Practicalities

- **Check the journal's policy first.** Nearly all allow it; some restrict which version you may post after acceptance.
- **Post the submitted version**, not a half-finished one.
- **Choose a licence.** CC BY is the usual default. Check whether your funder mandates one.
- **Update it** when the paper is revised, and again on acceptance if the publisher allows.
- **Link the published version** once it exists. Preprint servers do this automatically if you give them the DOI.

16.5 Versions, and their names

Publishing has three, and knowing the words saves confusion later.

Preprint	Before peer review. Yours to post
Accepted manuscript	Peer reviewed, before the publisher's typesetting. Usually postable after an embargo
Version of record	The publisher's final typeset version. Usually theirs, unless the paper is open access

Most funder open-access mandates are satisfied by depositing the accepted manuscript in a repository, which costs nothing. Article processing charges buy you the version of record being open, which is a different thing.

FROM THE OTHER SIDE OF THE DESK

The scooping story people tell each other usually runs: "I posted a preprint and a bigger group copied it." The story I've actually seen play out is the reverse: two groups discovered they were working on the same problem, one had a public timestamp, and it was the preprint that made the priority obvious and the eventual co-citation natural.

DO THIS NEXT

Ask your supervisor whether your subfield posts preprints, and if it does, plan to post on the day you submit. Check the target venue's policy while you're checking its scope statement, so both questions are answered at once.

17

The Submission Package

Submission day takes about four hours if you're prepared and two days if you're not.

Most of the difference is knowing what you'll be asked for before the system asks.

17.1 What you'll need

Manuscript	In the venue's format, usually PDF, sometimes source
Figures	Often as separate files, at specified resolution
Cover letter	Short. Appendix A
Abstract	Retyped into a form field, which will strip your formatting
Keywords	
Author details	Full names, affiliations, emails, ORCIDs
Suggested referees	Three to five, often required
Statements	Data availability, funding, conflicts, ethics, AI use
Supplement	Separate files

The ORCID line is worth flagging: register for one now if you haven't. It takes two minutes, many venues require it, and it permanently distinguishes you from every other researcher with your name.

17.2 Formatting to the venue

Download the template and use it. Not something close to it.

- **Reference style** exactly as specified. Your reference manager does this in one click if you've been using one.
- **Line numbers and double spacing** if asked for. Referees like line numbers, and some journals reject submissions without them.
- **Figures at the required resolution**, usually 300 dpi for photographs and vector or 600 dpi for line art.
- **Length within the limit**, including or excluding references as the venue specifies.
- **Anonymised** if double blind: no names, no affiliations, no acknowledgements, no "our previous work", and check the PDF metadata, which is where people get caught.

> *Formatting non-compliance gets a manuscript returned before an editor ever reads it. It's the cheapest possible rejection to avoid and one of the most common.*

17.3 Suggesting referees

Most journals ask for three to five, and students often treat it as a formality. It isn't; editors frequently use one or two of your suggestions.

Good suggestions	**Not eligible**
Authors of papers you cite	Your supervisor or co-authors
People working on the problem now	Anyone at your institution
Spread across countries and groups	Recent collaborators
Senior enough to be credible	Your friends

Give full names, institutional emails, and one line on why each is appropriate. Institutional emails, not Gmail, because editors check.

You may also usually request that specific people are *not* used, and you need a reason: a live dispute, a competing group working on the same result, a personal conflict. Use it sparingly, and one or two names at most, since a long exclusion list reads badly.

17.4 The cover letter

Half a page. Its job is to help the editor decide to send it out.

Four things: what the paper claims, why it fits this venue, that it's original and not under review elsewhere, and any declarations (a related conference paper, a preprint, a conflict).

No summary of the whole paper. The editor has the abstract. Appendix A has templates you can adapt.

17.5 The pre-submission checklist

Do this the day before, not the morning of.

Block 6 is not a formality. Submitting without a co-author's approval is a genuine problem, and an email saying "submitting this version tomorrow unless I hear otherwise" takes thirty seconds.

17.6 The submission system itself

They are, without exception, unpleasant. Budget three hours.

- Create the account and start the submission a day early, so you discover what it wants while there's still time.
- It will make you retype the abstract, the title, and every author's details.
- It will strip your formatting from the abstract field. Check the proof it generates.
- It will generate a merged PDF. **Read the whole thing** before approving. Figures get scrambled, fonts get substituted, equations break.
- Save the manuscript number it gives you.

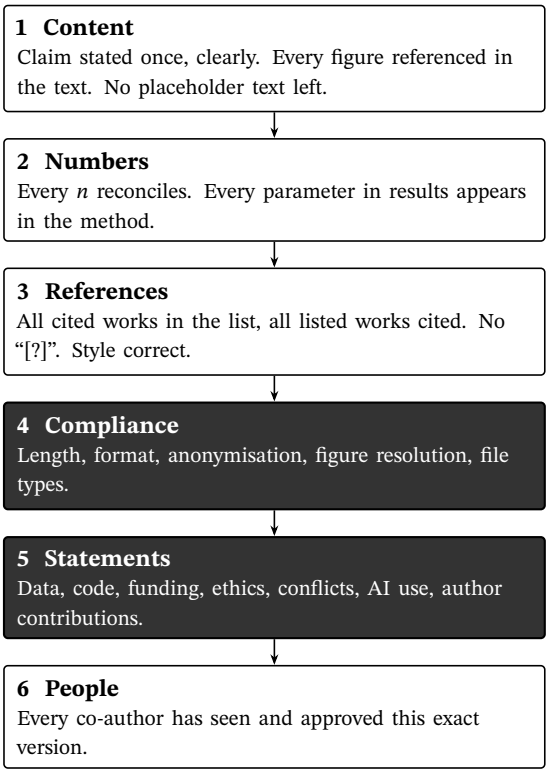

Figure 17.1: The six-block pre-submission check. Blocks 4 and 5 are the ones that get skipped and the ones that cost most when they're wrong.

FROM THE OTHER SIDE OF THE DESK

Editors see the same avoidable failures every week: the reference style ignored, the anonymised version with the author's name in the PDF properties, the figure that's illegible at print size, the abstract pasted with markdown asterisks still in it. None of them says anything about the research. All of them shape the first impression before a referee is even chosen.

17.7 After you hit submit

Two things, then stop.

Post the preprint if you're posting one. Write the manuscript number, the submission date, and the expected decision time in a file where you'll find it.

Then leave it alone. The status will say "with editor" for weeks and refreshing it changes nothing. Chapter 18 covers what's happening behind that status, which helps with the waiting.

DO THIS NEXT

Work through the six-block check the day before you submit, and send the "submitting tomorrow" email to your co-authors at the same time. Then do the submission itself in the morning, unhurried, with the venue's author guidelines open in another window.

18

What Actually Happens in Peer Review

The status bar says “under review” for four months and tells you nothing.

Here’s what’s happening behind it, because knowing makes the wait considerably less unpleasant.

18.1 The pipeline

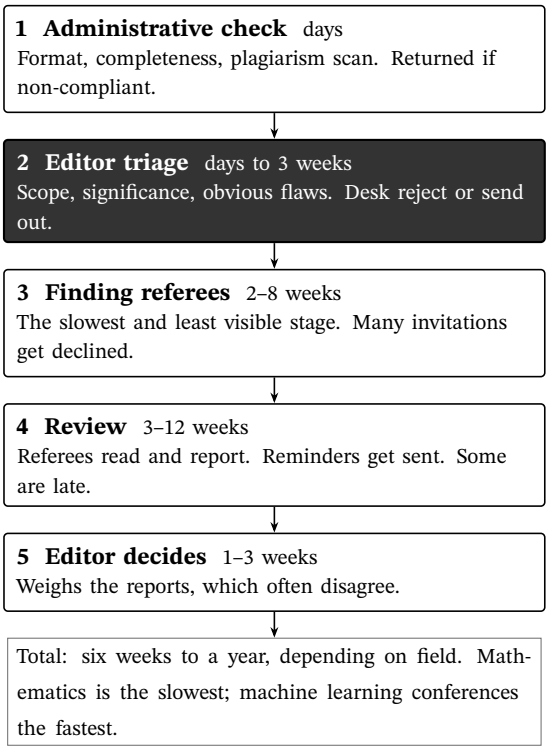

Figure 18.1: The path a manuscript takes. Most rejections happen at stage 2, before any referee sees the paper.

Stage 3 is why the wait is so long and so silent. An editor may invite eight people to get two acceptances, each invitation waits a week for a reply, and there is nothing to report while it happens.

18.2 Who your referees are

Working researchers in your area, usually two or three, occasionally one.

Unpaid, doing it in evenings and weekends on top of their own work, and typically reviewing somewhere between five and twenty papers a year. A careful review takes three to six hours.

That last fact explains a great deal. A referee who can't find your claim in ten minutes will not spend two hours hunting for it. A referee whose figure is illegible will not email to ask for a better one. Every piece of friction in your manuscript is friction for somebody working for free at eleven at night.

> *Write for a tired volunteer expert. That single image explains nearly every piece of advice in this book.*

18.3 The decisions, and what they mean

Decision	What it means
Accept	Essentially never on first submission. Don't expect it
Minor revision	It's in, subject to small fixes. Excellent outcome
Major revision	Real work needed, often new experiments, but they want it. A good outcome
Reject and resubmit	Sometimes a disguised major revision, sometimes a polite no. Read the editor's letter carefully
Reject	Not for this venue. Chapter 21
Desk reject	Rejected without review. Usually scope, sometimes obvious flaws

Major revision is not a failure. For a first paper it's a good result, and most published papers went through at least one round of it.

18.4 Rejection rates, for calibration

Selective journals reject 70–90%. Top conferences accept 15–25%. Solid, respectable journals reject perhaps 40–60%.

So a first-submission rejection is the single most likely outcome everywhere, and it says almost nothing about the work. Every researcher you admire has a folder of them.

18.5 Why reports disagree

They often will, and it's disconcerting the first time. Referee 1 finds the method sound and the framing weak; referee 2 finds the framing fine and wants three more controls.

This isn't a system failure. They're different people with different expertise reading for different things, and the editor's job is precisely to weigh that. What matters for you is the editor's letter, which tells you which points they consider binding.

18.6 How long to wait before asking

Under 6 weeks	Wait. Nothing has gone wrong
6–12 weeks	Normal for most journals. Still wait
3 months	A polite enquiry is reasonable
6 months	Enquire, and ask for an expected date
9 months, no reply	Escalate to the editor-in-chief; consider withdrawal

The polite enquiry is short:

> Dear Dr. [Editor],
>
> I'm writing to ask about the status of manuscript [number], "[Title]", submitted on [date]. I'd be grateful for any indication of the expected timeline.
>
> Thank you for your time.

That's the whole email. Not a complaint, not a list of how long it's been, and it usually gets a reply within a week.

Withdrawing to submit elsewhere is your right, and it resets the clock to zero, so it's rarely worth it before nine months unless something else has changed.

FROM THE OTHER SIDE OF THE DESK

The thing that surprises first-time authors most is how much of the delay has nothing to do with them. A manuscript can sit for two months because one referee agreed and then had a difficult term, and the editor is reluctant to start again with somebody new. Nobody is evaluating you during that time. Nobody is thinking about your paper at all, which is oddly reassuring once you accept it.

18.7 What to do while waiting

Not refreshing the portal.

Start the next thing. Write the talk. Post the preprint if you haven't. Put the manuscript out of your mind entirely, because you'll need fresh eyes when the reports come, and rereading your own paper anxiously for four months guarantees you won't have them.

DO THIS NEXT

Record the submission date and add a calendar reminder for twelve weeks out. If nothing has happened by then, send the three-line enquiry above. Between now and then, don't check the portal more than once a month.

19

Reading the Decision Letter

The email arrives. Your chest tightens. You read it badly.

Everyone does this. Here's how to read it properly, which is mostly a matter of when.

19.1 The three-day rule

Read the decision line. Note whether it's a revision or a rejection.

Then close it for three days.

Not because the reports will change, but because you will. On day one you read "the authors appear not to understand the underlying mechanism" as an attack on your competence. On day four you read it as a sentence that tells you exactly which paragraph to rewrite.

> *Never reply to referees within 48 hours. Nothing good has ever come from a same-day response to a review, and the delay costs you nothing in a process measured in months.*

19.2 Read the editor's letter first, and twice

Students skip straight to the reports. The editor's letter is more important, because the editor decides.

Look for:

The decision word	Minor, major, reject and resubmit, reject
Which points are binding	"Please pay particular attention to Reviewer 2's concerns about…"
What they've disregarded	Sometimes explicit: "you need not address point 4"
A deadline	Often 60 or 90 days. Extensions are usually granted if you ask before it expires
Their own view	Some editors add substantive comments. Weight these highest

If the editor names two of six concerns as essential, those two are the paper's fate. The other four still need answering in your response, and they need less rebuilding.

19.3 Sorting the comments

Print the reports. Go through with three colours.

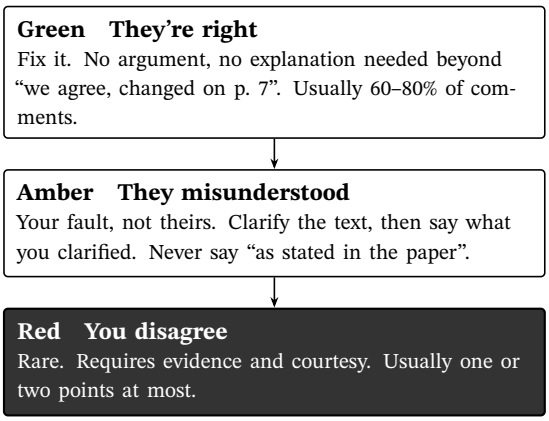

Figure 19.1: Sorting every comment into one of three bins before you plan any work. Most comments are green, and most of the anxiety comes from the few that aren't.

The sorting matters because it converts an overwhelming document into a task list. Twenty-two comments feels impossible. Sixteen green, four amber and two red is a fortnight of work.

And notice what amber means. If a competent expert misread your paper, the paper is unclear. That's true even when their reading is one nobody else would have made, because your next reader might be them.

19.4 Decoding referee language

Reviews are written in a register. Some translations:

They write	They mean
"The novelty is unclear"	I couldn't find your claim
"The authors may wish to consider"	Do this
"It would strengthen the paper if"	Do this
"I am not convinced that"	Your evidence doesn't support this claim
"Minor point"	Genuinely minor, but fix it
"The English requires attention"	Get it edited, and don't take it personally
"This is an interesting study, however"	The however is the review
"The authors should compare with [X]"	Usually X is theirs, and it's usually a fair request
"Insufficient detail to reproduce"	Chapter 10

Appendix D has a longer list.

19.5 Hostile and unhelpful reviews

They exist. A review that's dismissive, sarcastic, or obviously written in fifteen minutes is genuinely unpleasant to receive, particularly the first time.

Three things to know.

It's not about you. Referees are anonymous, tired, and occasionally graceless, and the tone rarely correlates with the substance.

Extract the substance anyway. Even a bad review usually contains one real objection, and finding it is the professional move.

Tell the editor if it crosses a line. A review containing personal remarks, or one that's clearly uninformed about the area, can be raised with the editor factually: "Reviewer 2 appears to have understood the paper to be about X; we wondered whether the report could be weighted accordingly." Editors know which of their referees are difficult.

FROM THE OTHER SIDE OF THE DESK

The most useful reframe I know: the referee spent hours on your paper for nothing. Even the harsh ones. They are the only person besides your supervisor who will ever read it that closely, and their objection is a free preview of what every future reader will think. You can have it now, privately, or later, in public.

19.6 Deciding whether to revise

For a major revision, sit down with the sorted list and ask three things.

Is the work feasible in the time given? Six extra weeks of experiments inside a 90-day deadline is fine; six months isn't, and you should ask for an extension rather than rush it.

Does the revision improve the paper? Usually yes. Occasionally a referee wants a different paper, and that's worth naming.

Is the venue still right? If the reports say the contribution is too small for this journal, and the editor agrees, a well-argued revision may still lose. Chapter 21 covers moving on.

For most first papers with a major revision, the answer is: revise, take the time, do it thoroughly. The paper that comes back is usually substantially better than the one you sent.

DO THIS NEXT

When the decision arrives, read only the decision line, then close it for three days. On day four, print everything, sort every comment green, amber or red, and count them. Then plan the work from the counts, not from the feeling.

20

Responding to Referees

The response letter is a document most students have never seen, and it carries as much weight in the second decision as the manuscript does.

Write it well and a borderline paper gets accepted. Write it badly and a good revision gets another round.

20.1 What you submit

Three things:

1. The revised manuscript.
2. A marked-up version showing changes, if the venue asks.
3. The response letter, answering every comment.

"Every comment" is literal. Skipping one, even a trivial one, reads as carelessness and invites the referee to check whether you skipped others.

20.2 The structure that works

Point by point, in the referee's order, with their text quoted.

> **Reviewer 2, Comment 3.** *"The authors do not report how many cells were excluded from the analysis, or on what basis."*
>
> **Response.** We agree, and this was an omission. Four cells of 28 were excluded: three for a failed initial capacity check (below 95% of nameplate) and one for a thermocouple failure at cycle 140. The rule was set before cycling began. We have added this to Section 3.2 (p. 6, lines 142–149) and updated the sample counts in Table 2 accordingly.

Four things in that response. Agreement, the specific fact, the assurance that the rule was pre-specified, and where to find the change. A referee can verify it in thirty seconds.

> *Every response ends with a page and line number. A referee who has to hunt for your change will assume you didn't make it.*

20.3 Formatting the letter

- **Quote each comment** in full, in italics or a distinct typeface, then respond below it.
- **Number everything** to match the reports.
- **Open with a short paragraph** thanking the referees and summarising the main changes in three or four lines.
- **Use the referees' own numbering**, even when it's chaotic.
- **Quote the new text** where a change is short. It saves the referee a lookup.

Length: it's normal for the response letter to be as long as the paper. That's fine, and it's expected.

20.4 Language for the three bins

Different moves for green, amber and red.

Green (they're right). Short and clean.

> We agree. The sentence has been rewritten as follows: "[new text]" (p. 4, l. 87).

Don't over-explain a fix. Don't apologise repeatedly. Don't argue that it was nearly right.

Amber (they misunderstood). The trap here is the phrase "as stated in the paper". It's technically true and it reads as an accusation that the referee didn't read carefully, which is a poor idea when they're about to read it again.

BEFORE

As stated in Section 3, the samples were annealed before measurement. The reviewer appears to have missed this.

AFTER

We take the point that this was not clear. The samples were annealed before measurement, but the ordering of the paragraphs made this easy to miss. We have moved the annealing step to the start of Section 3.1 and added a summary of the full sequence as Table 1 (p. 5).

The second treats a misreading as a writing problem, which it usually is.

Red (you disagree). Rare and it needs care. Four moves, in order: acknowledge the concern, give your reasoning with evidence, offer something, and leave the decision open.

> We understand the concern that a single manufacturing lot limits generality. We looked at this: the temperature dependence we report is consistent with the three-lot data of Nakamura et al. [8, Fig. 4], where absolute retention varied by 6 points across lots while the activation behaviour did not. Obtaining three lots ourselves would take approximately eight months. We have added a paragraph to the limitations stating the constraint explicitly (p. 11, l. 304) and softened the generality claim in the abstract. We would of course be glad to reconsider if the editor regards multi-lot data as essential.

No defensiveness, evidence for the position, a concrete concession, and the editor left holding the decision. That's the register.

20.5 What not to do

Ignore a comment	Reads as evasion. Answer all of them
Argue at length	Two paragraphs maximum, even on a red
Say “the reviewer is wrong”	Say “we respectfully disagree, for the following reason”
Get sarcastic	They will read it, and they may referee your next one
Claim a change you didn’t make	Catastrophic if noticed, and it will be
Rewrite the whole paper	Unless asked. Large unrequested changes force a full re-review

20.6 Handling contradictory referees

Referee 1 wants the theory section expanded; referee 2 wants it cut.

Say so, plainly, pick one, and explain the choice.

> Reviewer 1 asked for an expanded derivation, while Reviewer 2 suggested the theoretical material be shortened in favour of the experimental results. We have taken a middle course: the main steps of the derivation now appear in Section 2 (p. 3), with the full working moved to Appendix B. We are happy to adjust further at the editor’s direction.

Editors expect this and are usually happy to arbitrate when asked directly.

FROM THE OTHER SIDE OF THE DESK

A good response letter does something that isn’t obvious: it makes the referee’s second read fast. They have your letter, each of their points quoted, and a line number for each change. That’s a twenty-minute job instead of a two-hour one, and a referee who can verify your revision quickly is far more likely to recommend acceptance.

20.7 The resubmission itself

- **Meet the deadline**, or ask for an extension before it expires. They’re nearly always granted and rarely offered after the fact.
- **Check the changes are actually in** the file you’re uploading. Version confusion at this stage is common and embarrassing.
- **Re-run the reconciliation check** from Chapter 10; revised numbers often break it.
- **Have a co-author read the response letter** before it goes. Tone is hard to judge on your own work.

Appendix C has a full response-letter template you can fill in.

DO THIS NEXT

Draft your response letter before you make the changes. Writing the response first tells you exactly what work the revision requires, and it stops you making changes nobody asked for.

21

Rejection, Resubmission, and After Publication

Rejection is the most likely outcome of any single submission, so it's worth having a procedure rather than a feeling.

21.1 The three-day rule again

Read the decision. Close it. Come back on day four.

Then read the reports properly, because a rejection at one venue is usually a free review for the next, and it's the only free review you'll get.

21.2 Work out which rejection it was

They're not the same thing, and the right next move differs.

Kind	What to do
Scope	Wrong venue, no criticism of the work. Send it to a right one, almost unchanged, this week
Significance	"Sound but not important enough for this journal". Go down a tier; the work is fine
Fixable flaw	A missing control, an unaddressed alternative. Fix it, then resubmit elsewhere
Fundamental flaw	The claim isn't supported and can't be with this data. The hard one. Rescope to what you can support
Taste	Two referees disagreed and the editor sided against you. Try elsewhere

Most rejections of competent student work are the first two, and both mean send it somewhere else quickly.

> *A paper rejected on scope or significance should be resubmitted within two weeks. Every week it sits in a folder is a week of momentum lost, and the usual reason it sits is that nobody wanted to reopen the file.*

21.3 Before you resubmit elsewhere

Don't just change the template and press send. Two hours of work substantially improves the odds.

1. **Fix everything the referees were right about.** Free improvements, and the same referee may see it again at the next venue.
2. **Reformat to the new venue**, properly.
3. **Rewrite the introduction** for the new audience. A different readership needs different framing.
4. **Update the cover letter**, including, if you choose, a brief note that the paper was previously reviewed elsewhere and improved accordingly. Optional, honest, and some editors appreciate it.

Referee overlap is real and is a reason to make the fixes. Fields are small, and sending an unchanged manuscript that a referee already criticised is a bad look for you and irritating for them.

21.4 When it's the fundamental kind

Occasionally the reports agree that the claim isn't supported and no revision will fix it with the data you have.

Three options, in order of how often they're right.

Narrow the claim. What *does* the data support? Often a smaller, sound paper is sitting inside the failed larger one, and it will be published.

Get the missing piece. If it's one experiment and you still have access to the equipment, do it.

Put it out another way. A preprint, a technical report, a well-documented repository, a blog post. A negative or inconclusive result that's publicly available still has value, and it's still citable.

What you should not do is put it in a drawer and treat the whole thing as wasted. The single most common outcome for a student paper is that it never gets resubmitted after the first rejection, and that's a shame every time.

FROM THE OTHER SIDE OF THE DESK

Every researcher whose work you admire has a rejection folder. Papers that went on to be highly cited were rejected first, sometimes several times. The difference between a published paper and an unpublished one is very often nothing but persistence: somebody kept sending it out, and somebody else didn't.

21.5 Acceptance, and what happens next

The email arrives and it's genuinely a good day. Then there's admin.

Copyright form	Read it. Know what rights you keep, particularly on posting the accepted manuscript
Open access choice	Decide whether to pay, and check funder mandates and institutional agreements first
Proofs	Arrive in days to weeks, usually with a 48-hour deadline
Publication	Online first, then assigned to an issue

21.6 Reading proofs

The last chance to fix anything, and people rush it.

- **Check every equation**, character by character. Typesetting breaks maths more than anything else.
- **Check every figure** rendered correctly and is in the right place.
- **Check the author names, affiliations, and funding statement.**
- **Check the references** against your originals. Typesetters introduce errors.
- **Answer every query** the typesetter has flagged.
- **Only correct errors.** Rewriting at proof stage may incur a charge and will annoy everyone.

Two people should read the proofs, and one should be a co-author who hasn't been staring at it.

21.7 After it's published

The paper exists. A few things worth doing in the first week, and they take an hour in total.

- **Update the preprint** with the published DOI.
- **Deposit the accepted manuscript** in your institutional repository, which is usually required by your funder.
- **Update ORCID, Google Scholar**, and your CV.
- **Tell people.** A short post, an email to the people who helped, a mention at your next group meeting. Not self-promotion; papers nobody hears about get read by nobody.
- **Answer emails** about it. Somebody asking for your data or your code is the best possible outcome, and they'll cite you.

21.8 And the next one

The second paper is easier than the first by a large margin, because most of what made the first hard was not knowing the process.

You know it now. You know what a claim sentence looks like, what a method section owes a stranger, what a referee is asking, and what a response letter does. None of that has to be learned again.

Write the next one sooner than feels comfortable. The gap between a first and second paper is where a lot of people stop, usually for no better reason than that the first one took so long that it felt like an achievement to be finished with rather than a thing to repeat.

DO THIS NEXT

If a paper of yours is sitting rejected in a folder right now, take an hour this week: sort the referee comments, fix the green ones, pick the next venue, and put it in the diary for a fortnight's time. That hour is worth more than the next month of new work.

Part VI

Toolkit

A

Cover Letter Templates

Half a page. Adapt, don't paste.

A.1 Standard journal submission

> Dear Professor [Editor's name],
>
> We are submitting "[Title]" for consideration as a [research article / letter / short communication] in [Journal].
>
> [One or two sentences: the claim.] Grid-scale lithium-ion cells are routinely qualified using Arrhenius lifetime models calibrated below 30°C. We show that this calibration over-predicts retention by 12.6 percentage points at 45°C, and we identify the mechanism responsible.
>
> [One sentence: why this journal.] Because [Journal] has published the principal work on cell-level degradation modelling ([ref], [ref]), we believe the result will be of direct interest to its readers.
>
> The manuscript is original, has not been published elsewhere, and is not under consideration by another journal. All authors have approved the submission. [A preprint is posted at [DOI]. / Analysis code and data are available at [DOI].]
>
> We have suggested [n] potential reviewers in the submission system.
>
> Thank you for your consideration.
>
> Yours sincerely,
> [Name], on behalf of the authors
> [Affiliation, email]

A.2 Conference submission

Most conferences don't want a cover letter. Where a field is provided, keep it to three lines:

> This paper shows that [claim]. It is original work, not under review elsewhere, and not an extension of a previously published paper by the same authors. [A preprint is posted at [DOI].]

A.3 Extended version of a conference paper

The declaration matters here, and vagueness about it causes real trouble.

> This manuscript extends our conference paper [full citation], presented at [venue] in [month year]. The extension comprises: the two-mechanism model of Section 5, which is entirely new; cycling data for two additional temperatures (Section 3.3); and a full comparison against three alternative lifetime models (Section 6). Approximately 45% of the material is new, and the conference paper is cited as [ref].

A.4 Resubmission after rejection elsewhere

Optional, and some editors appreciate the candour.

> This manuscript was previously reviewed at [Journal], where it was declined on grounds of scope. We have taken the referees' technical comments into account: [one line on the main improvement]. No part of the work has been published.

A.5 Requesting an extension on a revision

Send it *before* the deadline, not after.

> Dear Dr. [Editor],
>
> Thank you for the opportunity to revise manuscript [number]. Addressing Reviewer 2's request for [the additional measurement] requires access to [instrument], which is available to us from [date].
>
> Would it be possible to extend the revision deadline from [date] to [date]? We expect to submit the revised manuscript by then.
>
> Thank you for your understanding.

A.6 Status enquiry

Three lines, no complaint. After roughly three months.

> Dear Dr. [Editor],
>
> I am writing to ask about the status of manuscript [number], "[Title]", submitted on [date]. I would be grateful for any indication of the expected timeline.
>
> Thank you for your time.

A.7 Six rules for all of them

1. Address the editor by name. It's on the journal's website.
2. Half a page, and never more than one.
3. Don't summarise the paper; the editor has the abstract.
4. State the claim in one sentence.
5. Say why this venue, specifically.
6. Declare everything: preprints, related conference papers, conflicts.

B

The Submission Checklist

Print this. Work through it the day before you submit.

B.1 The claim

- ☐ The claim appears as one explicit sentence in the introduction, beginning "we show that" or equivalent.
- ☐ The same claim appears in the abstract.
- ☐ The same claim appears in the conclusion, and hasn't grown.
- ☐ The contributions are listed, with section pointers.
- ☐ Every claim in the paper is supported by something in the paper.

B.2 Structure

- ☐ Each section answers only its own question (Chapter 2).
- ☐ Related work is grouped into families, not listed paper by paper.
- ☐ Related work is roughly 10% of the paper, not 30%.
- ☐ Results report; discussion interprets.
- ☐ Limitations are stated, specifically, with their consequence for the claim.
- ☐ The three hostile objections are addressed.
- ☐ No new material in the conclusion.

B.3 Numbers and reconciliation

- ☐ Every sample size in the paper reconciles: collected minus excluded equals analysed.
- ☐ Every exclusion has a stated rule and a count.
- ☐ Every parameter mentioned in the results has a value in the method.
- ☐ Statistical tests are named, with their assumptions.
- ☐ Effect sizes are given, not only p-values.
- ☐ Significant figures are consistent and justified by the measurement.

B.4 Figures and tables

- ☐ Every figure and table is referenced in the text.
- ☐ Every axis is labelled, with units.
- ☐ Every caption says what to conclude, not only what is plotted.
- ☐ Error bars are defined in the caption, with n.
- ☐ Each figure is legible printed at final size.
- ☐ Each figure is legible in greyscale.
- ☐ Colours mean the same thing across every figure.

- ☐ Tables have no vertical rules; numbers align on the decimal point.
- ☐ Figure files are at the venue's required resolution.

B.5 References

- ☐ Every citation in the text is in the list.
- ☐ Every entry in the list is cited in the text.
- ☐ No "[?]", no "TODO", no broken links.
- ☐ Reference style matches the venue exactly.
- ☐ Every reference has been read, at least at layer 2.
- ☐ Every paper described in a sentence is described in a sentence its author would accept.
- ☐ The list includes work from the last eighteen months.

B.6 Compliance

- ☐ The venue's template is used, not something similar.
- ☐ Length is within the limit, counted the way the venue counts.
- ☐ Line numbers and spacing as required.
- ☐ If double blind: no names, no affiliations, no acknowledgements, no "our previous work", and the PDF metadata is clean.
- ☐ File types and naming as specified.

B.7 Statements

- ☐ Data availability, saying what is available and where, or why it isn't.
- ☐ Code availability, with an archived DOI rather than only a repository URL.
- ☐ Funding, with grant numbers.
- ☐ Ethics approval, with committee and reference number, if applicable.
- ☐ Conflicts of interest, or an explicit statement of none.
- ☐ AI tool use, if the venue asks.
- ☐ Author contributions, in the venue's format.

B.8 People

- ☐ The author list and order are agreed, in writing.
- ☐ Every co-author has seen and approved *this exact version*.
- ☐ Everyone in the acknowledgements has been asked.
- ☐ Suggested referees are named, with institutional emails, and none is a collaborator or colleague.
- ☐ The corresponding author's email will still work in two years.

B.9 The final read

- ☐ Read the merged PDF the system generates, end to end, before approving it.
- ☐ Check equations survived the conversion.
- ☐ Check figures are in the right order and the right way up.
- ☐ Record the manuscript number and submission date somewhere you'll find them.

C

Response-to-Referees Template

Fill in the brackets. Keep the structure.

C.1 The opening

> Dear Professor [Editor],
>
> Thank you for the opportunity to revise manuscript [number], "[Title]". We are grateful to both reviewers for careful reports that have improved the paper substantially.
>
> The principal changes are: (i) we now report the exclusion criteria and counts in full (Section 3.2); (ii) we have added the comparison against [method] requested by Reviewer 1 (Section 5.2 and new Figure 6); and (iii) the generality of the claim has been narrowed in the abstract and conclusion in line with Reviewer 2's concern about single-lot data.
>
> Our point-by-point responses follow. Reviewer comments are in italics. Page and line numbers refer to the revised manuscript with changes marked.
>
> Yours sincerely,
> [Name], on behalf of the authors

Three or four lines of principal changes. That paragraph is what the editor reads.

C.2 The point-by-point body

> **Reviewer 1**
>
> *General comment: "This is a careful study, but the novelty relative to [ref] is not established."*
>
> **Response.** We agree that this was not clear, and we have addressed it directly. [Ref] measured the same quantity below 30°C using a single mechanism model; our contribution is the behaviour above 35°C, where that model fails, and the two-mechanism fit that covers both regimes. We have added a paragraph to the introduction making this explicit (p. 2, l. 48–56) and a row to Table 1 comparing the two studies directly.
>
> *Comment 1.1: "The exclusion criteria are not stated."*
>
> **Response.** We agree; this was an omission. Four cells of 28 were excluded, three for a failed initial capacity check and one for a thermocouple failure. The rule was fixed before cycling began. Added at p. 6, l. 142–149; sample counts in Table 2 updated accordingly.
>
> *Comment 1.2: "Figure 3 is difficult to read in print."*
>
> **Response.** Corrected. Figure 3 has been redrawn at column width with larger type, distinct line styles as well as colours, and the legend moved below the axes (p. 8).

C.3 Language for a disagreement

Comment 2.3: “A single manufacturing lot cannot support a general claim.”

Response. We understand the concern and think it is well placed with respect to absolute retention. Our claim, however, concerns the temperature dependence rather than the absolute values, and the three-lot data of [ref, Fig. 4] show absolute retention varying by 6 points across lots while the activation behaviour does not. Obtaining three lots ourselves would require approximately eight months. We have (i) added an explicit statement of the constraint to the limitations (p. 11, l. 304–311), and (ii) narrowed the abstract to claim the temperature dependence rather than the absolute figures. We would of course reconsider if the editor regards multi-lot data as essential for publication.

Four moves: acknowledge, reason with evidence, concede something concrete, leave the decision with the editor.

C.4 Language for a misunderstanding

Comment 2.5: “The authors do not appear to have annealed the samples.”

Response. We take the point that this was not clear. Samples were annealed at 400°C for two hours before measurement, but this was described mid-paragraph in Section 3.1 where it was easy to miss. The annealing step now opens Section 3.1 (p. 5, l. 104), and the full preparation sequence is summarised in new Table 1.

Never “as stated in the paper”.

C.5 Language for contradictory reviewers

Note on Reviewers 1 and 2. Reviewer 1 asked for an expanded derivation in Section 2, while Reviewer 2 suggested shortening the theoretical material in favour of the experiments. We have taken a middle course: the main steps now appear in Section 2 (p. 3) with the full working in Appendix B. We are happy to adjust further at the editor’s direction.

C.6 The closing

We hope the revised manuscript addresses the reviewers’ concerns. We are grateful for the time both reviewers gave to it, and we are glad to make further changes if required.

C.7 Checks before it goes

- ☐ Every comment answered, including the trivial ones.
- ☐ Every response ends with a page and line number.
- ☐ Every claimed change is genuinely in the uploaded file.
- ☐ No sarcasm anywhere, and no “as stated in the paper”.
- ☐ Disagreements run to two paragraphs at most.
- ☐ A co-author has read the letter for tone.
- ☐ Sample-size reconciliation re-run after the revisions.

D

Referee Phrases, Translated

Peer review has a register. Here's what the standard phrases usually mean, and what to do about each.

D.1 About the contribution

They write	They mean, and what to do
"The novelty is unclear"	I couldn't find your claim. Add the "we show that" sentence and a contributions list
"Incremental"	Real but small for this venue. Either argue the size, or go down a tier
"The contribution over [ref] is not established"	Add a direct comparison, ideally a table row
"Of limited interest to this journal's readership"	Scope. Change venue
"The paper reads as a technical report"	Method-heavy, claim-light. Rebalance

D.2 About the evidence

"I am not convinced that"	Your evidence supports something weaker. Narrow the claim or add evidence
"The sample size is small"	Report a power analysis, or state the limitation and soften the claim
"No comparison with existing methods"	Usually fair, usually required. Do it
"The control is missing"	Add it, or explain precisely why the design doesn't need one
"The statistical treatment is inadequate"	Name your test, state its assumptions, report effect sizes
"Insufficient detail to reproduce"	Chapter 10. Add parameters, versions, exclusions
"How were these excluded?"	Your numbers don't reconcile. Fix and state the rule

D.3 About the writing

“The English requires attention”	Get it edited. Not a comment on you
“The paper is difficult to follow”	Structural. Reverse-outline it
“The introduction is too long”	Related work has migrated into it
“Figure X is illegible”	Redraw at column width
“The abstract does not reflect the paper”	Usually the claim grew somewhere. Reconcile them
“Terminology is inconsistent”	You used two words for one thing. Search and standardise

D.4 Politeness formulas

“The authors may wish to consider”	Do it
“It would strengthen the paper if”	Do it
“I would encourage the authors to”	Do it
“A minor point”	Genuinely minor. Still do it
“Perhaps beyond the scope of this work”	Optional. Say why it’s out of scope
“This is an interesting study, however”	Everything after “however” is the review
“I have only one substantive concern”	That concern is the decision

D.5 Signals about the outcome

“I would be happy to review a revised version”	They expect it to be accepted eventually. Good news
“The following would need to be addressed before publication”	A checklist. Work through it
“I do not believe this can be addressed by revision”	They’re recommending rejection
“I defer to the editor on suitability”	They think it’s sound but possibly wrong-venue
“I am not an expert in [X]”	That part will be weighted less. Don’t assume it’s unchecked

D.6 And in the editor's letter

"Please pay particular attention to"	These points decide it
"You need not address point N"	Genuinely. Say thank you and skip it
"A revised version is not guaranteed acceptance"	Standard wording. Not a signal
"We can consider a substantially revised manuscript as a new submission"	Reject and resubmit. The clock restarts
"We receive many more submissions than we can publish"	Desk reject on significance, not on flaws

D.7 Two rules for reading any of it

Treat every criticism as a writing problem first. If an expert misread your paper, the paper permitted the misreading. That's true even when you're sure the text is clear.

Separate tone from substance. A rude review and a kind review containing the same objection require exactly the same fix. Extract the objection, discard the register, and answer in the register you'd like to have received.

E

An Eight-Week Plan

From finished results to submitted manuscript, assuming the paper is not your only responsibility.

Adjust the calendar, keep the order.

E.1 Week 0: before anything

- ☐ Agree the author list and order, in writing (Chapter 13).
- ☐ Apply the four tests. Confirm you have a claim (Chapter 3).
- ☐ Pick three venues: stretch, target, safety (Chapter 4).
- ☐ Tell your supervisor when the draft will arrive, and book their reading time.

E.2 Week 1: plan and figures

- ☐ Fill in the one-page plan (Chapter 6).
- ☐ Sketch the four figures on paper, with captions.
- ☐ Take both to your supervisor. Ask which figure is missing.
- ☐ Run the missing analysis or experiment if it's small.
- ☐ Finish the three-pass literature search (Chapter 5).
- ☐ Write the one-line "relation" note for your ten key papers.

E.3 Weeks 2–4: first draft

Write in this order, not the reader's order.

- ☐ Figures and tables, finished properly.
- ☐ Method. Give it to a lab-mate and ask what's missing.
- ☐ Results. Report only; no interpretation.
- ☐ Discussion. Write the three hostile objections first, then address them.
- ☐ Introduction, in five paragraphs.
- ☐ Related work, from the three-column table.
- ☐ Conclusion. One paragraph.
- ☐ Abstract, in five sentences, with a number in sentence 4.
- ☐ Title.

E.4 Week 5: supervisor read

Their calendar, not yours, which is why you booked it in week 0.

While you wait:

- ☐ Build the shareable repository (Chapter 14): README, environment file, licence, one figure regenerated from raw data.

- ☐ Run the reconciliation check on every n in the paper.
- ☐ Draft the cover letter.
- ☐ Draft the data, funding, ethics, conflicts and contribution statements.

E.5 Week 6: revise

- ☐ Work through your supervisor's comments.
- ☐ Run the introduction-to-discussion consistency check.
- ☐ Run the five-minute figure audit on every figure.
- ☐ Reverse-outline the paper: one sentence per paragraph, one page. Fix anything out of order.

E.6 Week 7: second read

Someone outside the project. A friend in an adjacent field is ideal, because they will find the places where you assumed knowledge.

Ask them three questions rather than "what do you think?":

1. What's the claim?
2. Where did you get lost?
3. What would you attack?

E.7 Week 8: format and submit

- ☐ Reformat to the venue's template exactly.
- ☐ Reference style correct; every "[?]" resolved.
- ☐ Anonymise if double blind, including the PDF metadata.
- ☐ Archive the repository and mint a DOI.
- ☐ Work through Appendix B in full.
- ☐ Email co-authors: "submitting this version tomorrow unless I hear otherwise."
- ☐ Submit in the morning, unhurried. Read the merged PDF before approving.
- ☐ Post the preprint, if your field posts them.
- ☐ Record the manuscript number and date. Set a reminder for twelve weeks.

E.8 What happens after week 8

Weeks 9–20	Silence. Start the next thing. Don't check the portal more than monthly
Week 12	If nothing, send the three-line enquiry
Decision	Read the decision line, close it for three days
Day 4	Sort every comment green, amber, red. Count them
Revision	Draft the response letter *before* making the changes
Resubmit	Within the deadline, or ask for an extension before it expires

E.9 Where the eight weeks actually go wrong

Three failure modes, all avoidable.

The supervisor read wasn't booked. The draft lands unannounced in a busy month and sits for six weeks. Fix it in week 0.

The figures were left until the end. Then a missing experiment surfaces in week 6, when there's no time to run it. Fix it in week 1.

Formatting was assumed to take an afternoon. It takes two days the first time, and the submission system takes three hours on top. Budget the whole of week 8.

F

Glossary

The vocabulary of publishing, which nobody teaches and everybody assumes you know.

F.1 The process

Desk reject	Rejected by the editor without going to referees. Usually scope or significance
Triage	The editor's initial decision on whether to review at all
Single blind	Referees know who you are; you don't know them
Double blind	Neither side knows. Requires anonymising your manuscript
Open review	Reports are published, sometimes signed
Rebuttal	A short author response during review, standard at CS conferences
Major revision	Substantial work required, but they want it
Minor revision	Small fixes. Essentially an acceptance
Reject and resubmit	A new submission with a new clock. Sometimes a disguised major revision
Camera ready	The final version you supply for typesetting
Proofs	The typeset version, for you to check for errors only
Online first	Published online before assignment to an issue

F.2 Versions and access

Preprint	Your manuscript, posted publicly before peer review
Accepted manuscript	Peer reviewed, before publisher typesetting. Usually yours to deposit
Version of record	The publisher's final typeset version
Gold open access	Free to read on the publisher's site, usually paid for by an APC
Green open access	Free via a repository copy, usually the accepted manuscript, sometimes after an embargo
APC	Article processing charge. Can be substantial. Check before submitting
Embargo	The period after publication before you may post the accepted manuscript
CC BY	A licence permitting reuse with attribution. Required by many funders

F.3 Identifiers and metrics

DOI	A permanent identifier for a document. Cite these, not URLs
ORCID	A permanent identifier for you. Free, two minutes, increasingly required
Impact factor	Mean citations per paper for a journal over two years. A journal-level number, widely misused to judge individual papers
h-index	You have h papers with at least h citations each
Altmetrics	Attention measures: news, social media, policy documents

F.4 Integrity terms

Gift authorship	Adding someone who didn't contribute. Misconduct
Ghost authorship	Omitting someone who did. Misconduct
Salami slicing	Splitting one study into several thin papers
Self-plagiarism	Reusing your own text without disclosure
p-hacking	Analysing repeatedly until something reaches significance
HARKing	Presenting a post-hoc finding as an a-priori hypothesis
Preregistration	Publicly recording your plan before collecting data
Corrigendum	A correction by the authors
Erratum	A correction by the publisher
Retraction	Withdrawal of a published paper
Expression of concern	A published flag that questions have been raised
CRediT	The standard taxonomy for contribution statements

F.5 Places to put things

arXiv	Preprints: maths, physics, computer science, quantitative biology
bioRxiv, medRxiv	Preprints: life sciences and clinical research
ChemRxiv	Preprints: chemistry
Zenodo	Archives code and data, mints a DOI, integrates with GitHub
Figshare, Dryad	Data archives
OSF	Projects, preregistration, materials
Institutional repository	Where your accepted manuscript usually must be deposited

F.6 Ten things worth remembering

If you retain nothing else from this book.

1. The paper is not the project. One claim, supported.
2. Write the claim as a sentence beginning "we show that".
3. Choose the venue before you write.
4. Plan the figures before you write the text.
5. Write the abstract last and the introduction second-to-last.
6. Every *n* in the paper must reconcile.
7. State your limitations before a referee does.
8. Agree authorship in writing, in week one.
9. Never respond to a review within 48 hours.
10. A rejection is a free review. Fix what's fair, and send it out again within a fortnight.

A Note from the Author

Thank you for reading. Books like this one travel by word of mouth among students, so if it saved you a term of guessing, two minutes of your help matter more than you'd think: leave a short, honest review wherever you bought it, or pass your copy to whoever in your group is sitting on finished results and hasn't started writing.

The appendices are working tools rather than reading. The submission checklist, the response-letter template and the eight-week plan are meant to be photocopied and used with a pen. You'll find printable companions and more of my writing for students at **gauravtiwari.org**, where corrections and questions are welcome.

And if your first paper does get accepted, I'd genuinely like to hear about it. Tell me which part of the process turned out to be harder than this book made it sound, and the next edition will be better for it.

Also by Gaurav Tiwari

How to Read Mathematics

Definitions, Theorems, Proofs, and Papers for Students

The reading half of the skill this book assumes. A six-step drill for opening any definition, three ways to read a proof, a stuck protocol that always ends in an action, and honest reading speeds for research papers.

How to Write Mathematics

Proofs, Papers, and Problem Sets for Students

Writing solutions that collect the marks the work deserved: proof makeovers, full-credit problem sets, and the grader's-eye view of partial credit.

How to Write Physics and Engineering

Lab Reports, Derivations, Design Documents, and Your Final-Year Project

The report side of a STEM degree, from the four-sentence abstract to the final-year thesis, including uncertainty written as prose rather than arithmetic.

How to Study STEM

Problem Sets, Practicals, and Exams Without the All-Nighter

The week-by-week system underneath everything else: how to work a problem set, what to do when you're stuck, and how to keep a term from collapsing into the fortnight before finals.

The STEM Exam Handbook

Timing, Triage, and Partial Credit in Maths, Physics, and Engineering Exams

The four-minute triage, how method marks really work, salvage writing, and ten annotated scripts where every lost mark is traced to its cause.

Other titles

- *LaTeX for Students*
- *LaTeX for Theses and Dissertations*
- *The Student Research Workflow*
- *Indian Polity: Complete Notes for UPSC Prelims & Mains*

All titles are available on Amazon. Find the full, current catalog at **gauravtiwari.org**.