

The six-sentence abstract

A reviewer with four other papers due the same night decides from your abstract how carefully to read the rest. Give them six sentences, each doing exactly one job.

1 Context	2 Gap	3 Contribution	4 Approach	5 Evidence	6 So what
-----------	-------	----------------	------------	------------	-----------

About 170 words in all. Each block is that sentence's share.

THE FORMULA

One job per sentence. Fill the brackets, dodge the trap.

1 Context ~20 words Where are we, and why does it matter now?	[<i>Task or system</i>] is now the standard way to [<i>do what</i>], and [<i>property</i>] decides whether it works.	Trap Opening with the history of your field. “Deep learning has revolutionised...” tells a reviewer nothing.
2 Gap ~30 words What is missing, broken or never measured?	Yet existing [<i>methods or benchmarks</i>] assume [<i>assumption</i>], so [<i>failure or question</i>] is never [<i>tested or solved</i>].	Trap “Little work has been done on X.” Name the assumption that is wrong, not the amount of work.
3 Contribution ~30 words What exactly did you find or build?	We show that [<i>finding</i>]. Or: We introduce [<i>artifact</i>] that [<i>does what, for whom</i>].	Trap “We propose a novel framework.” Novel is the reviewer's verdict to give, not yours.
4 Approach ~35 words How, in one sentence that proves it is real?	We [<i>key idea</i>] by [<i>mechanism</i>], and evaluate it on [<i>datasets or setting</i>] against [<i>strongest baseline</i>].	Trap Listing every component. Give the one idea that makes it work; the rest belongs in Section 3.
5 Evidence ~30 words What number would convince a sceptic?	On [<i>benchmark</i>], [<i>method</i>] improves [<i>metric</i>] from [<i>A</i>] to [<i>B</i>], [<i>N</i>] points over [<i>named baseline</i>].	Trap “Outperforms many existing methods.” No number and no named baseline means no credibility.
6 So what ~25 words What should the reader now believe or do?	This suggests [<i>claim no bigger than your evidence</i>]. We release [<i>code or data</i>] at [<i>link</i>].	Trap Claiming more than sentence 5 shows. Overclaim here and Reviewer 2 will quote it back to you.

WORKED EXAMPLE

The same paper, before and after.

BEFORE · LOW SCORES

~~In this paper we propose a novel framework which achieves good results on several benchmarks and outperforms many existing methods.~~

Why it fails

- No field: which models, which task?
- No gap: what was wrong before?
- No number: how good is “good”?
- No baseline: better than whom?
- It could describe almost any paper.

AFTER · ACCEPTED

① Retrieval-augmented models are now the standard way to ground a language model's answers in external documents. ② Yet existing benchmarks pair every question with a passage that supports the answer, so no one measures what happens when the retrieved passage contradicts it. ③ We show that models degrade by 14 points in exactly this case—a failure mode three prior benchmarks miss because they never test it. ④ We build a 2,400-item diagnostic set of contradicting passages and a fine-tuning recipe that teaches the model to flag the conflict before it answers. ⑤ The recipe recovers 11 of those 14 points, with no loss when the passage agrees. ⑥ Contradiction should be a standard check for any retrieval-augmented system; we release the set and the recipe.

Read only sentences 3 and 5 and you still know what the paper found. That is the test.

WORD SWAPS

Replace the phrase, win back the point.

INSTEAD OF

In this paper we propose

a novel framework

achieves good results

on several benchmarks

outperforms existing methods

extensive experiments

robust

to the best of our knowledge, the first

has attracted significant attention

WRITE

We show that ...

the one idea that is new, by name

improves F1 from 61.2 to 68.9

on three benchmarks: A, B and C

beats the strongest baseline, X, by 4 points

4 datasets, 3 model sizes, 5 seeds

holds when Z varies by up to W

a gap so precise it is obviously first

why it matters now, in one clause

ORDER OF WORK

Write the abstract last.

- 1 Figures and tables — they are your evidence
- 2 Method
- 3 Results
- 4 Introduction, ending in your contribution list
- 5 Related work and conclusion
- 6 Abstract, then the title

An abstract written first describes the paper you hoped to write. Written last, it describes the one you did.

SELF-CHECK

Read it once, cold. Tick every box.

- ☐ A reader from the next field over could say what you did.
- ☐ Sentence 2 names an assumption someone could disagree with.
- ☐ There is at least one number, and the baseline it beats is named.
- ☐ Sentence 6 claims nothing that sentence 5 does not show.
- ☐ No “novel”, “robust” or “state of the art” without a number beside it.
- ☐ No citations, equations or undefined abbreviations.
- ☐ Inside the venue’s word limit. Check this year’s call for papers.
- ☐ It promises exactly what your introduction’s contribution list delivers.
- ☐ The abstract on the submission form matches the one in the PDF.

DESK-REJECT TRIPWIRES

Rejected before a reviewer reads a word.

- ✗ Over the page limit, or the template’s margins or font changed.
- ✗ Anonymity broken: author names in the PDF metadata, “our previous work [12]”, a code link with your username.
- ✗ A required checklist or AI-use disclosure missing.
- ✗ Last year’s template, or another venue’s.
- ✗ The same paper under review somewhere else.

PREDATORY VENUES

Five signs a journal wants your fee, not your paper.

- ✗ Acceptance promised, or a review in days rather than weeks.
- ✗ Payment asked for before any peer review has happened.
- ✗ An unsolicited email praising work you have not published.
- ✗ A scope covering engineering, medicine and management at once.
- ✗ Indexing claims you cannot confirm on the indexer’s own site, only on the journal’s.

THROUGH THE REVIEWER’S EYES

Score names follow the NeurIPS review form; other venues vary.

Every score on the form is decided, a little, by the abstract.

SOUNDNESS

Do the experiments support the claim?

Answered by sentences 4 and 5

A number and a named baseline head off “the evaluation is weak” before the reviewer thinks it.

CONTRIBUTION

Does this matter, and is it new?

Answered by sentences 2 and 3

A precise gap makes the novelty obvious without you ever writing “novel”.

PRESENTATION

Could I follow it in one read?

Answered by all six

One job per sentence is what “clearly written” means to a reviewer on their fortieth paper.

ANOTHER FIELD ENTIRELY

The six jobs do not care what you study.

BEFORE · LOW SCORES

~~This paper proposes a novel technique for improving the performance of rooftop solar systems, which gives good results and outperforms conventional methods.~~

Why it fails

- “Performance” of what — efficiency, stability, cost?
- No gap: what do current methods get wrong?
- No number, no standard, no named comparison.
- Swap three words and it is a different paper.

AFTER · ACCEPTED

1 Rooftop solar inverters must disconnect within two seconds of a grid fault, and that margin is what keeps a distribution feeder stable. 2 Yet anti-islanding conformance is tested one inverter at a time, so nobody measures what happens when thirty of them on the same feeder trip together. 3 We show that simultaneous tripping produces a voltage overshoot on reconnection that single-unit testing cannot predict. 4 We model a 415 V feeder carrying 30 inverters in real-time hardware-in-the-loop, and compare it against the single-unit IEEE 1547 procedure. 5 Overshoot reaches 12.4% against a 5% limit, and a staggered reconnection schedule brings it back to 4.1%. 6 Conformance testing should specify aggregate behaviour, not only single-unit behaviour; we release the feeder model and the schedule.

Different field, different evidence, identical skeleton. Sentence 2 still names an assumption someone could argue with.

IF YOUR JOURNAL WANTS HEADINGS

Structured abstracts: the same six jobs, relabelled.

HEADING	SENTENCES	WHAT GOES UNDER IT
Background	1 and 2	The setting, then the assumption nobody has tested.
Methods	4	Design, data or participants, and what you compared against.
Results	5	The numbers, with effect size and precision. No interpretation.
Conclusions	3 and 6	What you now claim, and it may not exceed Results.

READ THE GUIDELINES FIRST

Medicine, life sciences and many engineering journals accept *only* this form, and the headings are fixed — often 250–300 words against an unstructured abstract’s 150–200.

Some add Objective, Design, Setting or Limitations. Use exactly the headings the call for papers lists, in exactly that order. A structured abstract with invented headings is a desk reject.

KEYWORDS

Three to six words that decide who reviews you.

- Most submission forms make you pick from the venue’s own subject list before you type anything of your own. Open that list first — it is how papers are routed to an area chair, and a wrong pick sends yours to someone who does not work on it.
- Take the rest from sentences 2, 3 and 4:
- ✓ One that names the **field**, so the right chair claims it.
 - ✓ One that names the **task or application**.
 - ✓ One that names your **method or artifact**.

- ✗ Words already in your title. You have that ground; take more.
- ✗ An acronym only your lab uses.
- ✗ Anything so broad it describes a third of the conference.

From the solar abstract above:

distributed generation · anti-islanding · inverter control · feeder stability · IEEE 1547

YOUR TURN

Write the six now. Read them cold tomorrow, then tick page 2.

1 Context ~20 words

Where are we, and why does it matter now?

2 Gap ~30 words

What assumption is wrong, untested or unmeasured?

3 Contribution ~30 words

What exactly did you find or build?

4 Approach ~35 words

The one idea that makes it work, and what you compared against.

5 Evidence ~30 words

The number that would convince a sceptic, and the baseline it beats.

6 So what ~25 words

What should the reader now believe — no more than sentence 5 shows?

BEFORE YOU CALL IT DONE

- ☐ Read it aloud once. Any sentence you stumble on is doing two jobs.
- ☐ Count the words. Over your venue's limit, cut sentence 1 first.
- ☐ Give it to someone outside your field and ask what you did.

BRING THIS PAGE TO DAY 2

Session 2.2 rewrites these six sentences live — yours, not ours — against the formula and against each other's. Turning up with the page already filled in is the difference between watching and working.

TITLES

Your title is sentence 3, shortened.

THE FINDING

Retrieval-Augmented Models Break When the Passage Disagrees

Use it when the result itself is the surprise.

THE ARTIFACT AND ITS JOB

A Diagnostic Set for Conflicting Evidence in Retrieval-Augmented QA

Use it when you built something other people will reuse.

THE QUESTION

Do Retrieval-Augmented Models Notice When the Evidence Disagrees?

Use it only when the paper answers it cleanly.

✗ “Towards...” promises less than you did. ✗ “A novel...” is the reviewer’s word. ✗ **Acronym first:** “CARE:” means nothing yet. ✗ **Over 15 words.**

BY PAPER TYPE

The six jobs stay the same. Sentences 3 and 5 change shape.

PAPER TYPE	SENTENCE 3 CLAIMS	SENTENCE 5 PROVES
New method	We introduce [X], which [does Y].	The gain on named benchmarks over the strongest baseline.
Dataset or benchmark	We release [X], which tests [Y] that existing sets miss.	Its size and coverage, and how far current models fall short on it.
Empirical study	We show that [X] happens when [Y].	The effect size, across how many models, datasets or settings.
Negative result	We find that [X], widely assumed, does not hold when [Y].	The comparison, and why it is a fair one.
Theory	We prove [X] under [assumptions].	The bound or guarantee, and when it is tight.
Survey	We organise [N] papers on [X] by [Y].	The taxonomy, and the open problems it exposes.

AI AS A CRITIC, NOT A GHOSTWRITER

Two prompts to run before anyone else reads it.

1 Label the jobs

Here is my abstract. For each sentence, say which job it does: context, gap, contribution, approach, evidence or implication. List any job that is missing or done twice. Do not rewrite anything.

[paste abstract]

2 Be Reviewer 2

You are a sceptical reviewer at [venue]. Using only this abstract, write the three strongest objections you would raise and quote the words that triggered each. Do not suggest fixes.

[paste abstract]

- ✗ Never keep a number the model wrote. Every figure comes from your results.
- ✗ Never keep a citation you have not opened and read.
- ✓ Check the venue’s AI-use policy, and disclose the way it asks.
- ✓ Check with your supervisor before pasting unpublished work into a public tool.

THE FULL MASTERCLASS

This sheet is one idea from twelve sessions.

Two live evenings on Zoom: finding a gap that is actually open, designing the experiment that answers it, writing the paper and answering Reviewer 2 — taught on one real paper’s draft, reviews, rebuttal and accepted version.

SAT 24 & SUN 25 OCT 2026

7–9 PM IST · Live on Zoom

₹ 499 early bird to 11 October

Then ₹999 · 60 seats · refund before 21 October

Reserve at firstauthor.in