

Submission Criteria and Review Rubric

Student AI Research Collective (SAIRC)

Version 1.2 · September 2026

This is the document our reviewers work from. We publish it so that nobody has to guess what we are looking for, and so that a rejection can be checked against something written down. If you think a review of your work departed from what is below, say so and we will look again.

1 What we accept

Any original work in artificial intelligence or machine learning, or a literature review of existing work in those areas. Subfields we regularly publish in include deep learning, natural language processing, computer vision, reinforcement learning, interpretability, AI safety, robotics, multi-modal systems, and theory.

There is no affiliation requirement. We do not ask where you go to school, and the review is carried out without knowledge of the author's institution where that is practical.

1.1 Formats

- **Poster.** One page, mostly visual. Suited to early or exploratory work.
- **Short paper.** Up to five pages excluding references and appendices. Suited to a single focused experiment or analysis.
- **Long paper.** Up to fifty pages excluding references and appendices. A complete project with methodology and results written out.

1.2 Project types

- **Review.** A critical reading of one or more existing papers. Place the work in context, state what it does well and where it is weak, and give your own assessment.
- **Research.** Your own question or experiment. It does not have to beat the state of the art, and most of what we publish does not.

2 The rubric

A reviewer scores each criterion as *meets*, *needs work*, or *does not meet*. Any criterion marked *does not meet* sends the submission back with notes rather than rejecting it outright, except where noted in Section 3.

Criterion	Meets	Needs work
Question	The paper states what it is trying to find out, and why that is worth knowing.	The topic is clear but the specific question is not, or the question is too broad to answer with the work presented.
Method	Described in enough detail that a reader with the same tools could follow it. Hyperparameters, data sources, and evaluation are all stated.	Key steps are missing, or the description would not let anyone reproduce what was done.
Evidence	The results shown actually bear on the question asked. Baselines and controls are present where the claim needs them.	Results are reported but do not test the stated claim, or there is no baseline to compare against.
Claims	Conclusions stay within what the results support. Limitations are stated.	Conclusions overreach, or the limitations section is absent when the work plainly has them.
Attribution	Prior work is cited. Contributions of collaborators and of any AI tools are disclosed.	Relevant prior work is missing, or the paper does not say who did what.
Presentation	Readable. Figures are labelled, tables have units, and the structure is navigable.	Unlabelled figures, missing axes, or a structure that makes the argument hard to follow.

Note what is not on this list: novelty, significance, and whether the result is positive. A careful negative result is publishable here. A replication is publishable here. Work that reaches a modest conclusion honestly is worth more to us than work that reaches a large one loosely.

3 Integrity requirements

These are not scored. Failing any of them stops the submission.

1. The work is yours, and everyone who contributed is listed as an author.
2. Anything taken from elsewhere is cited, including text, figures, code, and data.
3. The data is real. Fabricated or selectively edited results end the review and we will not accept further submissions from the same author.
4. If a language model wrote or substantially edited part of the paper, say which part. Using AI tools is permitted and common. Presenting their output as your own reasoning is not.
5. You have whatever permission you need to publish, from a school, an employer, or a mentor.
6. Work involving human subjects describes how consent was obtained.

4 Process

1. You submit a PDF through the form at sairc.vercel.app/submit.
2. A student reviewer is assigned. We aim to avoid assigning anyone who knows the author.

3. The reviewer works through the rubric above and writes notes.
4. You receive one of three outcomes: publish, revise and resubmit, or decline. All three come with the reviewer's notes.
5. Median time from submission to reply is nine days. Exam periods make it longer, and we will tell you if it will.

5 What our review is not

We want to be exact about this, because it affects how you should describe a SAIRC publication.

SAIRC is run by high school students. Our reviewers are students. What we provide is a careful read by a person who knows the area, checked against the rubric above. It is not formal academic peer review, it is not anonymous in both directions, and it does not involve subject-matter experts with doctorates in the field.

You may accurately say that your work was published in a student-run open-access journal and that it was reviewed by a student editor. You should not describe it as peer reviewed in the sense that an academic journal means that phrase. Anyone reading your CV who knows the difference will spot the overstatement, and it costs you more than the phrase is worth.

6 Common reasons work comes back

In order of how often we write them:

1. No baseline. A number on its own does not show that a method worked.
2. The conclusion is larger than the experiment. One dataset does not support a claim about models in general.
3. Missing hyperparameters, so nobody could rerun it.
4. Figures without axis labels or units.
5. A literature review that summarises papers one after another without ever comparing them.
6. No statement of limitations.

Questions about any of this go to sairc.support@gmail.com. A person reads them.