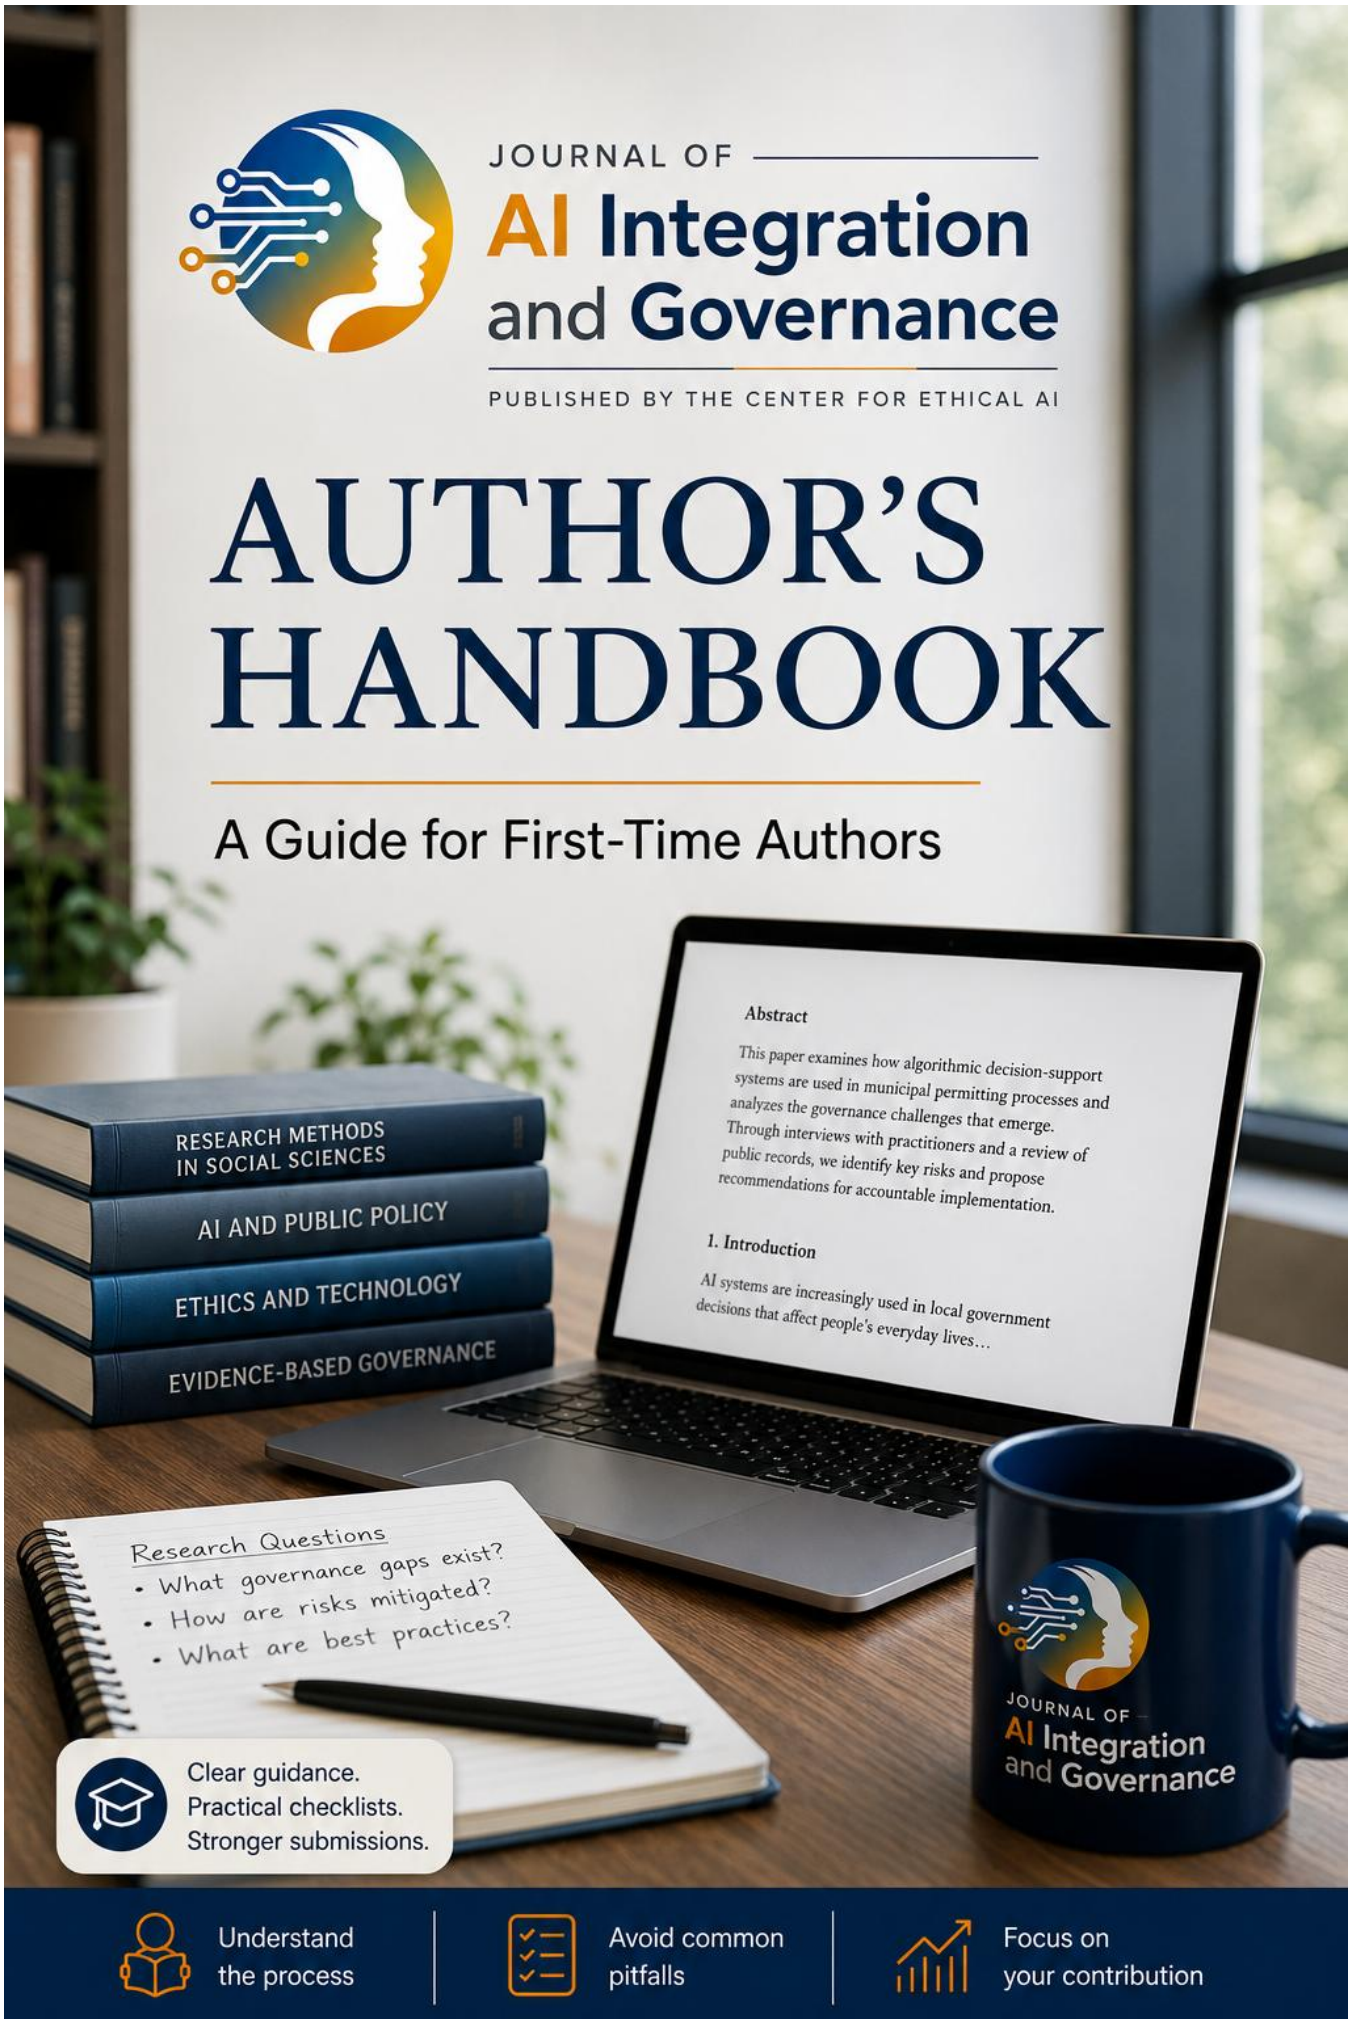
JOURNAL OF
AI Integration and Governance
PUBLISHED BY THE CENTER FOR ETHICAL AI
AUTHOR'S HANDBOOK
A Guide for First-Time Authors
RESEARCH METHODS IN SOCIAL SCIENCES
AI AND PUBLIC POLICY
ETHICS AND TECHNOLOGY
EVIDENCE-BASED GOVERNANCE
Abstract
This paper examines how algorithmic decision-support systems are used in municipal permitting processes and analyzes the governance challenges that emerge. Through interviews with practitioners and a review of public records, we identify key risks and propose recommendations for accountable implementation.
1. Introduction
AI systems are increasingly used in local government decisions that affect people's everyday lives…
Research Questions
• What governance gaps exist?
• How are risks mitigated?
• What are best practices?
JOURNAL OF
AI Integration and Governance
Clear guidance.
Practical checklists.
Stronger submissions.
Understand the process
Avoid common pitfalls
Focus on your contribution

JAIIG

Author's Handbook

Version 1.2 | 2026

SUGGESTED CITATION

The Center for Ethical AI. (2026). *Journal of AI Integration and Governance Author's Handbook* (Version 1.2). The Center for Ethical AI. [Journal URL forthcoming]

PUBLISHER

The Center for Ethical AI

1507 Shumate Dr, Little Rock, AR 72212

Contact: admin@thecenterforethicalai.com

DOCUMENT IDENTIFIERS

Version	1.2
Document Type	Author Guidance Companion
Status	Non-authoritative companion
ISBN	Pending

COMPANION DOCUMENTS

This Handbook is a non-authoritative companion to the Manuscript Format Specification v1.2 (authoritative source for format rules) and the Author Disclosure & Verification Policy v1.1 (canonical author-facing AI disclosure policy). The authoritative documents prevail in case of any conflict. For the canonical map of all JAIIG documents, including versions, status, and cross-references, see Document Set Index v1.0.

DISCLAIMER

This Handbook guides authors preparing manuscripts for submission to JAIIG. It does not constitute legal, professional, or editorial advice for any specific manuscript. Where this Handbook differs from the authoritative governance documents listed above, the authoritative documents prevail. The Center for Ethical AI does not warrant that following this Handbook will result in any specific editorial outcome.

ACKNOWLEDGMENTS

This Handbook was developed by the Editor-in-Chief of the Journal of AI Integration and Governance with reference to peer venues and to scholarly publishing best practices. The journal currently operates under the Editor-in-Chief's unified editorial authority; once an editorial board is constituted, future revisions of this Handbook will reflect input from that board.

REVISION HISTORY

Version 1.0 — Initial publication, 2026.

Version 1.1 — 2026-05. Aligned with Author Disclosure & Verification Policy v1.0 (four-element disclosure standard); updated AI tool examples to current model designations; updated editorial structure references to reflect current no-board operational reality; consolidated contact information.

Version 1.2 — 2026-05. Added AI text detection policy disclosure (Part 7); expanded Reference Integrity Module description with four-layer probe detail (Part 7); added submission portal interaction description (Part 6); aligned Peer Review Rubric criterion 6 name (Part 8); replaced companion document list with Document Set Index cross-reference.

Welcome

If you are reading this, you are likely preparing to submit a manuscript to JAIIG for the first time. This handbook is for you.

Academic publishing is a craft that nobody is born knowing. The conventions, the unwritten expectations, the things experienced authors do without thinking — these have to be learned, and most of us learn them through trial and error, often expensive error. This handbook tries to compress some of that learning so you do not have to repeat it.

JAIIG is committed to serving as a platform for emerging scholars, particularly graduate students and early-career researchers. The journal does not assume that first-time authors will write like senior ones. The pre-screen and review processes are designed to surface fixable issues early, not to gatekeep against unfamiliarity with conventions. Your first manuscript will not be perfect. That is expected. Your job is to make it as good as you can, submit it, and engage constructively with the feedback you receive. The journal's job is to help you do that.

How to use this handbook

Read it once, end to end, before you start writing. The chronological organization is intentional: the parts follow the order in which you will encounter the issues they address. Skipping ahead means you will face decisions without context.

After your first reading, return to specific parts as needed. The table of contents is your navigation. Part 10 (Frequently Asked Questions) is the part you are most likely to revisit.

What this handbook is not

This handbook is not authoritative. The Manuscript Format Specification v1.0 governs format rules, the AI Pre-Screen Specification v1.0 governs pre-screen behavior, and the Peer Reviewer Scoring Rubric governs peer review. Where this handbook differs from those documents, those documents win. The handbook explains that the specifications govern.

A note about the running example

Throughout this handbook, you will encounter references to a fictional manuscript on "Algorithmic Decision-Support in Municipal Permitting." This is an invented topic used for illustration only. It is not a real study, and the findings, citations, and disclosures shown for it are not real. We use a single fictional example so you can see how a coherent manuscript is built across sections, citations, and disclosures, rather than encountering disconnected fragments.

Contents

Part 1. Understanding JAIIG

Before you start writing, take fifteen minutes to understand what kind of journal you are submitting to. This determines what you should write, how you should frame it, and which conventions matter.

What the journal publishes

JAIIG publishes scholarly work on the oversight, regulation, ethical design, and institutional management of artificial intelligence systems. The journal sits at the intersection of policy, technology, and ethics. Its audience includes academics, policymakers, practitioners building AI systems, and standards bodies. This is a broader audience than a purely computer science venue or a purely ethics venue. Your writing needs to be intelligible to readers from any of these communities, even if your work draws most heavily from one.

The journal accepts several types of contributions:

- Empirical research that examines how AI systems are governed in practice or how they should be.
- Conceptual or theoretical work that develops frameworks for understanding or addressing governance problems.
- Policy analyses that examine specific regulations, standards, or institutional arrangements.
- Case studies of governance approaches, deployments, failures, or successes.
- Methodological work on how to evaluate, audit, or oversee AI systems.

Each type has different conventions about what counts as a contribution. An empirical paper contributes new evidence about how things are; a conceptual paper contributes new ways of understanding which questions are relevant. Be honest with yourself about which type you are writing, because reviewers will evaluate your work against the conventions of that type.

What does “contribution” mean here

Reviewers will repeatedly ask what your contribution is. This is the single most important question in academic publishing, and the one first-time authors most often answer poorly.

A contribution is not the same as a topic. “This paper is about algorithmic bias” is a topic. “This paper shows that algorithmic bias in municipal permitting reproduces existing zoning inequities through a feedback loop specific to administrative AI” is a contribution. The first describes what the paper covers; the second describes what the paper adds.

A contribution is also not the same as effort. The amount of work you put into a manuscript does not, by itself, make it a contribution. Reviewers respect effort, but they evaluate work. If you wrote 8,000 words restating what others have already established, the manuscript will not be accepted, regardless of how long the writing took.

Useful test: if your paper were absent from the literature, what would readers be missing that they cannot get elsewhere? If you cannot answer this question in one sentence, your contribution is not yet clear, and you need to clarify it before the manuscript is ready to submit.

RUNNING EXAMPLE — Stating the contribution

"This paper provides the first empirical examination of bias in AI-assisted municipal permit approval, drawing on three years of decision data from a hypothetical mid-sized US city. We show that the system reproduces historical zoning bias not through training data alone but through a feedback loop in which appeals processes systematically disadvantage applicants from minority neighborhoods. Our findings have implications for how municipalities procure and audit administrative AI systems." — This sentence states a topic (municipal AI bias), a contribution (the feedback-loop mechanism), and an audience (procurement and audit decision-makers).

How JAIIG differs from adjacent journals

If you have read papers in adjacent venues — Big Data & Society, AI and Ethics, AI & Society, Minds and Machines, Journal of Responsible Innovation — you have a sense of the broader space. JAIIG differs from these in three ways worth noting:

First, governance is the centerpiece, not a corollary.

Many adjacent journals publish work in which governance implications appear in the final paragraph. Here, governance is the focus of the paper. If your paper's governance discussion is confined to the conclusion, the work is probably not the right fit, regardless of how interesting the technical or ethical content is.

Second, the journal values applied work alongside theoretical work.

Reviewers will not penalize practical, deployment-oriented contributions for a lack of theoretical novelty. They will penalize practical work for lacking useful insight. Conversely, they will not penalize theoretical work for lacking immediate application. They will penalize theoretical work that rediscovers established conclusions in new vocabulary.

Third, the journal expects standards-aligned reasoning.

When relevant, work should engage with established governance frameworks (ISO/IEC 42001, NIST AI RMF, sectoral regulations, professional standards). This does not mean every paper must cite every standard; it means that work claiming governance relevance should engage with the institutional mechanisms through which governance actually happens.

Cross-reference

For the journal's mission and editorial philosophy in their authoritative form, see "Why the JAIIG Manuscript Template Is Structured This Way" v1.0.

Part 2. Planning Your Manuscript

Most manuscript problems are planning problems in disguise. The work that happens before you write the first sentence determines whether the writing goes well or badly. This part covers the planning decisions worth making explicitly.

Decision 1: What is your contribution?

We covered this in Part 1, but it bears repeating because it is the decision from which everything else flows. Before you start writing, write a single sentence stating what your manuscript adds to the literature. If you cannot, you are not ready to write yet. Read more, think more, talk to your advisor, and return when the sentence is clear.

Write the sentence on a sticky note. Put it where you can see it while you write. Every paragraph you draft should serve that sentence. If you find yourself writing material that does not, the material does not belong in this manuscript.

Decision 2: Who is your audience?

JAIIG readers come from policy, technical, and ethics backgrounds. You cannot write equally well for all three. Choose one as your primary audience and write for them, while keeping the other two as secondary readers who should be able to follow.

If your primary audience is policymakers, your technical detail should be calibrated to support a policy claim rather than to demonstrate technical mastery. If your primary audience is technical, your policy framing should be responsible but need not engage every nuance of regulatory theory. If your primary audience is ethics scholars, your case grounding should be concrete enough to engage with, but the philosophical analysis should carry the weight.

Your choice of primary audience should match your contribution. An empirical study of permit-approval bias has a primary policy or technical audience; the same work framed as an exploration of fairness ethics has a primary ethics audience.

Decision 3: What type of paper are you writing?

The manuscript template is intentionally agnostic about type. Empirical, conceptual, policy, case study, and methodological papers all use the same section structure. But the conventions within each section differ by type, and you should know which type you are writing before you start.

If you are writing empirical work:

- Methodology must specify data sources, sample sizes, analytical methods, and limitations.
- Results must be reproducible from the description.
- Discussion must distinguish what the data show from what you infer from what the data show.

If you are writing conceptual or theoretical work:

- Methodology must explain your analytical approach (literature synthesis, framework development, philosophical analysis).
- Results take the form of the framework or argument itself, not data.
- Discussion must engage with how the framework would be applied or how it differs from existing frameworks.

If you are writing policy analysis:

- Methodology must specify the policy or regulatory artifacts being analyzed and the analytical lens applied.
- Results take the form of policy claims supported by the analysis.
- Discussion must engage with implementation, second-order effects, and alternatives considered.

RUNNING EXAMPLE — Choosing a paper type

Our hypothetical municipal permitting study could be written as empirical work (we collected data and found a feedback loop), as a case study (this is what happened in this hypothetical city, with lessons), or as a policy analysis (this is what procurement frameworks should require). Each version emphasizes different evidence, has different reviewer expectations, and produces a different manuscript. Choose one before you start writing.

Decision 4: How long should your manuscript be?

JAIIG does not enforce a strict word limit, but most published articles fall in the range of 6,000 to 9,000 words, excluding references. Conceptual and policy papers can run shorter; empirical papers with substantial methodology can run longer. Manuscripts under 4,000 words usually feel underdeveloped; manuscripts over 12,000 words usually contain material that should have been cut.

Length is a derivative of contribution, not a target. A 5,000-word paper that makes a clean argument is better than a 9,000-word paper that pads it with unnecessary literature review.

Decision 5: Have you done what you need to do?

Before you start writing, confirm three things:

1. You have read enough of the relevant literature to know what your contribution adds to it. Reviewers will identify any gaps in your reading.
2. You have completed your analysis. Manuscripts that promise analyses they have not yet performed (“future work will examine...") are weaker than they need to be. Do the analysis first.
3. You have someone, ideally your advisor or a peer, who will read a draft and give you honest feedback before you submit. First-draft submissions almost always have problems that an internal review would catch.

Common mistake

Treating manuscript writing as a way to figure out what you think. Some writers do work this way successfully, but it is risky for first-time authors. The first draft of a manuscript should be written when you already know what you are trying to say. Writing-as-thinking belongs in your notes, not in your submission.

Part 3. Writing the Manuscript

This part walks through each section of the template in the order you should write them. The order is not arbitrary. Most experienced authors do not write the introduction first, despite its position in the manuscript.

Recommended writing order

4. Methodology (you know what you did; this is the easiest section to draft)
5. Results (you know what you found, write it down before you forget the details)
6. Discussion (now interpret what the results mean)
7. Literature Review (now you know what you are positioning against)
8. Introduction (now you know what the paper is, you can introduce it)
9. Conclusion (write last, when the paper exists in full)
10. Abstract (write last, when you can summarize what is already there)

Writing the introduction first is one of the most common mistakes for first-time authors. You cannot introduce a paper that does not yet exist. Drafts that start with a comprehensive introduction often include framing claims that the rest of the paper cannot support, and these need to be rewritten.

Title

Your title should describe what your paper does, specifically. Vague titles ("Some Considerations on AI Ethics") signal that the author has not committed to a contribution. Titles with colons are common in academic publishing because they let you state both the topic and the specific claim: "Algorithmic Decision-Support in Municipal Permitting: A Feedback-Loop Analysis of Bias Reproduction."

Avoid titles that describe the form of the paper rather than its content. "A Study of..." or "An Analysis of..." wastes the first words on information the reader will infer anyway.

Abstract

The abstract is the most-read part of your manuscript. Many readers will read only the abstract before deciding whether to read further. It needs to do real work in 250 words.

A useful structure, not the only one:

- One sentence on the problem or context.
- One sentence on the gap in existing work or the question you address.
- One to two sentences on what you did (methodology in brief).
- Two to three sentences on what you found (the substantive contribution).
- One sentence on implications.

Avoid filler. Phrases like "in this paper, we discuss..." waste your word budget. The abstract is not a meta-description of the paper; it is a compressed version of the paper itself.

RUNNING EXAMPLE — Abstract draft

Municipal permit approval is increasingly automated, but the governance implications of administrative AI in zoning contexts have not been examined empirically. This paper presents a hypothetical three-year study of permit-approval data from a mid-sized US city deploying an AI decision-support system. Using regression analysis and stratified appeals-process data, we identify a feedback loop in which the system's bias against applicants from minority neighborhoods is amplified by unequal access to appeals. The mechanism is not captured by training-data bias analysis alone. We discuss implications for procurement standards, audit requirements, and the design of administrative AI oversight.

Keywords

Choose three to eight keywords that a reader searching for work like yours would actually use. Resist the temptation to include every possible term. Keywords are for discoverability, not for staking claims. "Algorithmic bias" is a useful keyword; "the deontological implications of administrative AI in permit contexts" is not.

Introduction

The introduction has three jobs: establish that the problem matters, establish that existing work has not adequately addressed it, and announce what you will do about it. These map roughly to the first three subsections in many introductions.

Common first-time author mistakes in the introduction:

- Starting too broad. "AI is transforming society" is too far from your specific contribution. Start one or two steps closer to your actual topic.
- Reviewing the literature exhaustively. The introduction motivates; the literature review reviews. Save the comprehensive treatment.
- Hedging the contribution. "This paper attempts to consider some potential implications of..." is weaker than "This paper shows that...". If you cannot make a strong claim, you may not have a contribution yet.
- Promising what the paper will not deliver. Do not say the paper will address something it does not actually address.

Literature Review

The literature review should establish two things: that you know the relevant scholarship, and that there is space in it for your contribution. It should not be a comprehensive survey of everything anyone has written on related topics. That is a thesis chapter, not a journal section.

Organize thematically, not chronologically. "Smith (2018) argued X. Then Jones (2019) argued Y. Then Brown (2020) argued Z" is a list. "Three threads of work have addressed this question. The first, which we will call X, ..." is a review.

Engage critically, not just descriptively. Reviewers want to see that you can evaluate the work you cite, not just summarize it. “Smith showed that...” is weaker than “Smith showed that..., though the analysis assumes...”.

Cite work you have actually read. Citing a paper you have only seen referenced by others is a recognized form of weak scholarship. It produces the most common form of citation error: claims attributed to a source that the source does not actually make.

Methodology

The methodology section is where reviewers most often identify problems. Your job is to make it impossible for a reviewer to claim they cannot tell what you did.

For empirical work:

- Specify your data: what it is, where it came from, how much of it you have, and what its limitations are.
- Specify your analytical methods: what you ran, why you ran it, and what assumptions it requires.
- Specify your sample: who or what is included, who or what is excluded, and on what basis.
- Acknowledge limitations honestly. Reviewers will identify them anyway; you may as well do it first.

For conceptual work:

- Specify the analytical approach: literature synthesis, framework construction, normative analysis, comparative analysis.
- Specify the scope: what you are claiming about, what you are not.
- Specify the foundations: what frameworks, traditions, or assumptions you are building on.

Limitations are not weaknesses to hide.

First-time authors often write defensive limitations sections that minimize the limitations of their work. Reviewers see this immediately and treat it as a credibility signal in the wrong direction. A paper that honestly acknowledges what it cannot show is more credible than one that hedges. Limitations belong in the methodology and discussion sections; they should be specific rather than generic.

Results

Results sections describe what you found in sufficient detail that a reader could evaluate whether your interpretation is justified. They are not the place for interpretation itself; that goes into discussion.

Common first-time author mistakes:

- Writing results as narrative rather than as findings. “We then examined X, and we found Y, and then we considered Z,” reads like a research diary. Results sections should be organized around the findings, not the chronology of the research process.

- Mixing results with discussion. If you are explaining why a finding matters, you have moved from results into discussion.
- Burying the main finding. Lead with the most important result. Reviewers should not have to hunt for it.

Discussion

The discussion is where you interpret your results and connect them to the broader scholarly and practical conversation. The template requires three subsections: Policy Implications, Practice Implications, and Limitations.

Policy Implications.

What should regulators, lawmakers, or governance bodies take from your work? Be specific. "Policymakers should consider these findings" is not a policy implication. "Procurement standards for administrative AI should require documentation of feedback-loop testing."

Practice Implications.

What should practitioners, developers, or organizations actually do differently? Same standard of specificity. The audience for this subsection is people who are building or deploying AI systems, not the abstract "the field."

Limitations.

What can your work not show? Be specific and honest. "Our findings are based on data from one hypothetical city, and may not generalize to municipalities with different procurement structures or population profiles" is a useful limitation. "More research is needed" is filler.

Conclusion

The conclusion summarizes your contribution and gestures toward what comes next. It should be substantively distinct from the abstract: the abstract previews; the conclusion reflects.

Avoid the conclusion that restates the abstract. If a reader who has read the whole paper finds nothing new in your conclusion, the conclusion has failed.

Part 4. Citations and References

JAIIG (*Journal of AI Integration and Governance*) uses a hybrid citation-and-reference convention called Option C-2. This is unusual. Most journals use either pure APA (author-year in-text) or pure IEEE (numeric in-text). The journal pairs IEEE-style numeric in-text citations with an APA 7-style reference list, alphabetized in the manner of a standard APA bibliography. This part explains the convention, why it exists, and how to follow it correctly.

The convention in summary

- In-text citations use bracketed numerics: [1], [1, 3, 7], [1-4].
- References use APA 7 element formatting (authors, year, title, venue) wrapped in [N] numbering.
- The reference list is sorted alphabetically by the surname of the first-listed author, or by organization name for institutional and corporate sources.
- Bracket numbers [1], [2], [3]… are assigned according to that alphabetical position, not according to the order in which sources are first cited.
- The same source reuses the same number on every citation. You do not renumber mid-draft, but you must confirm each source's number against the alphabetized reference list before submission.

Why this convention

AI governance research draws from policy, technical, and ethics traditions. Pure APA pairs naturally with social science writing but interrupts technical prose with parenthetical author lists. Pure IEEE pairs naturally with technical writing but provides less bibliographic detail than humanities and policy readers expect. This convention provides technical readers with compact in-text citations while giving every reader a reference list organized as APA bibliographies: alphabetically by author.

Alphabetical numbering keeps the reference list self-describing. A reader looking for the Mitchell et al. paper, or for the ISO standard, can find it by name in its alphabetical position, the same way they would in any APA bibliography, without first searching the body text to discover which bracket number was assigned to it. This matters in AI governance scholarship, where standards and regulatory sources (ISO, NIST, the EU AI Act) are frequently consulted as reference material in their own right, independent of how a given article uses them.

The trade-off is that bracket numbers no longer track citation order. [7] may be cited before [3] in the body. Readers should not infer sequence from the numbers, and authors cannot assign numbers simply by drafting in order. A source's number is fixed only once the reference list has been alphabetized, which in practice means numbering is finalized in a late formatting pass rather than during drafting.

How to use citations in text

A note on drafting order.

Because bracket numbers depend on the alphabetized reference list rather than the order in which sources are introduced, do not assign [N] numbers while drafting. Two workable approaches:

- Draft with author-surname placeholders (e.g., [Chen], [Mitchell et al.]) and convert to [N] form in a single pass once the reference list is complete and alphabetized.
- Build the alphabetized reference list first, then write the body against the resulting [N] assignments.

Either approach is acceptable during drafting. What the pre-screen checks is the final manuscript: every in-text bracket number must match its source's alphabetical position in the References list, and no author-surname placeholders may remain in the submitted version.

Acceptable forms:

```
Recent work on algorithmic accountability [1] has shown...
Several studies [1, 3, 7] support this finding.
Multiple authors [1-4] have examined...
Smith's analysis [1, p. 47] specifies...
The argument runs across the chapter [1, ch. 3].
```

Forms that are NOT acceptable:

```
(Smith, 2020) argued that...          ← Author-year, not permitted.
Smith [1] argued that...              ← Surname before bracket, redundant
[Smith, 2020] argued that...          ← Bracketed author-year, invalid
(Smith [1], 2020) argued that...      ← Mixed form, invalid
```

Mixing styles fails the pre-screen

If your manuscript contains both [N] form and (Author, Year) form anywhere in the body, the pre-screen will FAIL §5.4.2 and the manuscript will be returned for correction before peer review begins. If you migrate from a previous draft that used author-year, search and replace systematically before submitting.

How to format reference entries

Each reference entry begins with [N] and follows APA 7 element conventions. The general pattern:

GENERAL PATTERN

```
[N] Author, A. A., & Author, B. B. (Year). Title in sentence case.
    Italicized Venue, Volume(Issue), Pages. https://doi.org/10.xxxx/xxxxx
```

The worked examples below are presented in the alphabetical order they would occupy in a References list, so that source types appear intermixed rather than grouped by type, the way they would in an actual submission.

JOURNAL ARTICLE

```
[1] Chen, L., & Martinez, R. (2024). Procurement Frameworks for
    Administrative AI: A comparative analysis. Journal of Public
    Administration, 12(3), 245–271. https://doi.org/10.1234/jpa.2024.0312
```

BOOK CHAPTER

```
[2] Crawford, K. (2021). Classification. In Atlas of AI: Power, politics,
    and the planetary costs of artificial intelligence (pp. 123–149).
    Yale University Press.
```

REGULATION

[3] European Commission. (2024). Regulation (EU) 2024/1689 of the European Parliament and the Council on artificial intelligence (AI Act). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

BOOK

[4] Floridi, L. (2023). The ethics of artificial intelligence: Principles, challenges, and opportunities. Oxford University Press. https://doi.org/10.1093/oso/9780198883098.001.0001

ONLINE RESOURCE

[5] Hardt, M., & Recht, B. (2024). Patterns, predictions, and actions: Foundations of machine learning [Online textbook]. https://mlstory.org

DATASET

[6] Hugging Face. (2024). The Stack: A 6.4 TB code dataset [Dataset]. Hugging Face Hub. https://huggingface.co/datasets/bigcode/the-stack

STANDARDS DOCUMENT

[7] International Organization for Standardization. (2023). ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system. ISO. https://www.iso.org/standard/81230.html

CONFERENCE PROCEEDING

[8] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model Cards for Model Reporting. Proceedings of the Conference on Fairness, Accountability and Transparency, 220–229. https://doi.org/10.1145/3287560.3287596

GOVERNMENT DOCUMENT

[9] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

PREPRINT (arXiv)

[10] Park, P. S., Goldstein, S., O'Gara, A., Chen, M., & Hendrycks, D. (2024). AI deception: A survey of examples, risks, and potential solutions. arXiv. https://arxiv.org/abs/2308.14752

Citation integrity: what NOT to do

JAIIG runs an automated reference integrity check on every submission. The check verifies that DOIs resolve, that arXiv IDs exist, and that the metadata in your reference list matches the metadata in the underlying source. This is not a hostile check; it catches both honest errors and deliberate fabrication. But it does catch fabrication, and the consequences are serious.

The most common citation integrity problems:

Citing without reading.

If you cite a paper you have not actually read, you are likely to misrepresent what it argues. This produces citations that the underlying source does not actually support. Reviewers catch these. The pre-screen catches them. Your credibility as an author depends on not doing this.

Citation padding.

Adding citations to make a passage look more scholarly than it is. If a claim comes from one source, cite that source. If you have not read three other sources you are citing alongside, do not cite them.

AI-generated citations.

Generative AI tools fabricate plausible-looking citations that do not exist. Authors using AI assistance for literature search must verify every citation against the actual source. The pre-screen will catch fabricated citations, and the consequences include investigation under publication ethics policies. This is the single most important reason the AI Assistance Disclosure exists.

Citing retracted papers.

A paper that was once cited may have been retracted since. Before submitting, check the current status of any citation older than five years through Retraction Watch or the journal's website. Citing retracted work without noting the retraction is a recognized integrity issue.

What happens if the pre-screen flags fabrication

Manuscripts with confirmed fabricated citations are not simply returned for correction. They are routed to the integrity editor for investigation in accordance with the journal's publication ethics policy. Honest errors are handled differently from patterns of fabrication; the system distinguishes between them. But the cost of honest carelessness is still high. Verify your citations before submission.

Part 5. The AI Assistance Disclosure

Every manuscript submitted to JAIIG must include an AI Assistance Disclosure section. This section affirms or denies the use of AI tools in manuscript preparation. The requirement is not optional, and it is not pro forma.

Why this section exists

Generative AI is now widely used in academic writing. Authors use it for literature search, drafting, editing, translation, summary, and analysis. Some uses are unproblematic; others raise integrity concerns. The disclosure enables readers, reviewers, and editors to evaluate AI-assisted manuscripts with full information on how the tools were used.

JAIIG specifically requires this disclosure because the journal publishes work on AI governance. Failing to disclose AI use in a paper about AI governance would be inconsistent with the journal's stated values. The disclosure is a small act of governance practice that the journal expects of authors who write about governance.

What requires disclosure

You must disclose AI use if:

- You used a generative AI tool to draft any portion of the manuscript text.
- You used AI for literature search and incorporated results into your reading.
- You used AI to summarize sources you then cited.
- You used AI to edit prose, restructure sections, or translate from another language.
- You used AI to perform any portion of your data analysis.
- You used AI to generate figures, tables, or other visual content.

You do not need to disclose:

- Standard spelling and grammar checking by the word processor features.
- Standard search engine use that does not involve AI summarization.
- Use of AI in your underlying research that does not bear on the manuscript itself (e.g., AI in the system you are studying).

How to write a compliant disclosure

A compliant disclosure has four elements. Listing all four is required for affirmative disclosures, per Author Disclosure & Verification Policy §B.1.

Element 1: The tools.

Name each AI system, vendor, and version. "Claude (Anthropic, Sonnet 4.6)" or "ChatGPT (OpenAI, GPT-5)" are specific. "AI tools" or "generative AI" is not.

Element 2: The stages.

Where in the manuscript pipeline did you use each tool? Ideation, literature search, drafting, editing, code generation, data analysis, citation gathering, figure preparation, summarization, translation, or visualization. Be specific about which sections or activities each tool was applied to.

Element 3: The human verification.

How did you confirm that AI-generated content is accurate, that AI-suggested citations exist, and that AI-generated analysis is correct? This element is the most important. AI tools can fabricate plausible content; your verification is what protects the manuscript's integrity.

Element 4: The limitations.

What material limitations of the AI tools affect a reader's interpretation of the work? If no material limitations apply, state so explicitly. This element supports the reader's evaluation of AI-assisted scholarship.

An acceptable affirmative disclosure:

ACCEPTABLE AFFIRMATIVE DISCLOSURE

```
Generative AI (Claude, Anthropic, Sonnet 4.6) was used during the
drafting and editing stages to assist with prose refinement and
Literature synthesis for Sections 1 and 2. All AI-generated content
was reviewed and verified against primary sources by the authors.
The tool was not used to generate analytical results, fabricate
citations, or produce content not subsequently verified by the
authors. No material AI-tool limitations affected the manuscript's substantive content.
```

An acceptable negative disclosure:

ACCEPTABLE NEGATIVE DISCLOSURE

```
No AI tools were used in the preparation of this manuscript.
```

Phrasings that fail the pre-screen

Ambiguous disclosures are treated as non-disclosures and fail §4.10 of the Pre-Screen Specification. Specifically, the pre-screen will return a manuscript if the disclosure includes any of the following phrasings:

- "AI may have been used." Either it was, or it was not.
- "AI tools were used in some capacity." Specify the capacity.
- "Standard AI assistance." There is no such thing.
- "Used AI as needed throughout." Specify where and how.
- Disclosure that omits the tool name, the stage, or the verification.

If you are uncertain whether to disclose

Disclose. Over-disclosure is not a problem. Under-disclosure is. If you cannot decide whether something rises to the level of disclosure, disclose it and let the editor decide whether the disclosure is necessary.

Part 6. Before You Submit

Before you submit, run through this checklist. The pre-screen will catch most of these issues, but catching them yourself first saves you a round trip.

Content checklist

- Your contribution sentence is clear, and the abstract reflects it.
- Every claim in the manuscript is supported by either evidence (for empirical claims) or argument (for conceptual claims).
- Every citation in the body has a matching reference entry, and vice versa.
- Every figure and table is referenced in the body text. Figures and tables that are not referenced should be removed.
- The discussion connects the results back to the introduction's framing. If the introduction promises X, the discussion delivers X.

Format checklist

The Author Quick-Start Guide v1.0 contains a complete pre-submission checklist mapped to the §5 rules of the Format Specification. The most common pre-screen failures, in order:

11. Title and headings using manual font formatting instead of Word styles.
12. Missing AI Assistance Disclosure section.
13. Missing Conflicts of Interest section.
14. References list ordered alphabetically rather than by appearance.
15. DOIs in bare form (10.1234/...) rather than URL form (https://doi.org/...).
16. Mixed citation styles in the body (some [N], some (Author, Year)).

Integrity checklist

- Every citation maps to a source you have actually read.
- Every quoted or paraphrased claim accurately represents the source.
- No citations are AI-generated without verification.
- Co-authors have all read the final manuscript and agreed to submission.
- If your work involves human subjects, IRB approval is documented.
- If your work uses datasets, the data availability statement is specific.

Before submission read

Before you submit, do one final read of the entire manuscript, end to end. Read it as if you are a reviewer who has not seen it before. Three things to watch for during this read:

First, internal consistency. Does the abstract describe the same paper that the body delivers? Does the introduction promise what the conclusion claims? Are terms used consistently throughout?

Second, places where the writing is unclear. If you find yourself rereading a passage, the passage probably needs editing. Reviewers will not reread. They will conclude that the passage is unclear and treat it as a flaw.

Third, places where the contribution claim is weak. The introduction, discussion, and conclusion should all gesture at the contribution. If any of them is vague, the others will not save you.

What to expect at the submission portal

Submission proceeds through the journal's submission portal at [portal URL]. The portal accepts only .docx manuscripts conforming to the Manuscript Template v1.1.

When you submit, the portal will:

- Validate that your file is in .docx format and conforms to basic file constraints.
- Present a structured AI Assistance Disclosure field with a minimum 150-character requirement. You must enter your four-element disclosure (Tools, Stages, Verification, Limitations) directly into this field, even though it also appears in your manuscript.
- Present five required confirmation checkboxes covering: ownership of the manuscript content, accuracy of authorship attribution, accurate representation of references, compliance with the Governed Data Boundary Policy, and authorization to submit on behalf of all co-authors.
- Provide FAQ anchor links for questions that arise during submission.

The portal interaction is specified in the Author Disclosure & Verification Policy §G. Authors who encounter issues with the portal should contact admin@thecenterforethicalai.com.

Part 7. After You Submit: The Pre-Screen

When you submit your manuscript, it enters the AI-Augmented Pre-Screen process before peer review begins. This part explains what the pre-screen does, what it does not do, and how to interpret its outcomes.

What the pre-screen evaluates

The pre-screen runs three modules:

Template Compliance Module.

Checks that all required sections are present, that the section structure is correct, and that disclosures are complete. This is the largest module and produces most of the findings that the authors see.

Format Fidelity Module.

Checks that the manuscript uses the correct paragraph styles, citation format, and reference formatting. This is where Option C-1 is enforced and where most format-related issues are caught.

Reference Integrity Module.

Verifies that your references actually exist and accurately represent their source. Within the module, an automated four-layer probe runs against each reference: Layer 1 (identifier presence) checks whether the reference has a DOI or equivalent identifier; Layer 2 (identifier resolution) verifies the identifier resolves; Layer 3 (metadata coherence) confirms the resolved record's metadata matches the citation; Layer 4 (citation–claim alignment) inspects whether the cited source actually supports the claim made in the manuscript. The module produces a four-band classification: PASS, WARN, FLAG, and FLAG with editorial escalation. The bands are calibrated against published baselines in accordance with the Reference Integrity Threshold Calibration Specification v1.0.

A note on AI text detection

The journal does not use stylometric AI text detection tools (such as GPTZero, Originality.ai, Turnitin AI, or comparable tools) to evaluate manuscripts. The rationale is documented in Author Disclosure & Verification Policy §C.2: such tools exhibit bias against non-native English writers and produce unreliable distinctions between AI- and human-written text. The journal relies on author disclosure (per ADVP §B), reference integrity verification (per ADVP §C and the pre-screen Reference Integrity Module), and human editorial judgment.

What the pre-screen does NOT evaluate

The pre-screen does not evaluate:

- Whether your work is novel.
- Whether your argument is persuasive.
- Whether your methodology is appropriate for your research question.
- Whether your conclusions follow from your evidence.
- Whether your work contributes to the field.

Those evaluations are what peer review is for. The pre-screen is a structural check that ensures the manuscript is ready for substantive review, nothing more.

Interpreting pre-screen outcomes

Each pre-screen rule produces one of four outcomes. You may receive any combination of these in your pre-screen report.

PASS.

The rule is satisfied. No action required.

WARN.

Revision is recommended. WARN findings do not, on their own, block peer review, but the editor may bundle them into a correction request before peer review begins.

FAIL.

An objective omission has been detected. Peer review is blocked until the issue is corrected. Common FAIL outcomes: missing required section, missing AI Assistance Disclosure, and missing Conflicts of Interest statement.

FLAG.

Human editorial review is required. FLAG outcomes are confidence-graded findings that the system cannot resolve deterministically. The editor will adjudicate.

Reference integrity bands

The Reference Integrity Module produces a four-band classification specific to citation health.

Band 1 (PASS).

The reference list's integrity is within expected variance for legitimate authorship. Manuscript advances.

Band 2 (WARN).

Elevated error rates detected. The manuscript is returned with an itemized correction request. This is the most common non-PASS outcome and almost always reflects fixable formatting issues rather than integrity concerns.

Band 3 (FLAG).

Statistically improbable error rates detected. Manuscript is routed to integrity review for editor adjudication. This does not mean fabrication has been confirmed; it means the rate exceeds what is expected for legitimate manuscripts and warrants closer examination.

Band 4 (FLAG with escalation).

Clear pattern of fabrication detected. The manuscript is routed for a COPE-aligned investigation in accordance with the publication ethics policy. This is rare, and the consequences are serious.

If you receive a Band 2 outcome

This is the most common reference integrity finding for first-time authors and almost always reflects fixable issues: bare DOIs that should be URLs, missing italics on journal names, and occasional metadata mismatches. Read the Editor's Report carefully, address the listed items, and resubmit. Most manuscripts that receive Band 2 outcomes are passed on for resubmission.

What to do if your manuscript is returned

If the pre-screen returns your manuscript, you will receive an Editor's Report listing the specific findings. Each finding cites the rule clause that triggered it (§5.1.1, §5.6.9, etc.) and explains what to fix.

Your response should:

17. Address every finding in the report. The editor expects a complete revision, not a partial one.
18. Make the requested changes without arguing about the rules. The rules are codified in the Format Specification and are not negotiable at the pre-screen stage.
19. Include a brief response letter explaining what was changed. The letter is for the editor; it does not need to be elaborate.
20. Resubmit promptly. Long delays between submission and resubmission make the manuscript stale.

If you believe a finding is incorrect

Pre-screen findings can occasionally be incorrect. If you believe a finding does not apply to your manuscript (e.g., a section detected as missing is actually present under a different heading), include the explanation in your response letter, and the editor will adjudicate. Do not silently leave the issue unaddressed; the resubmission will fail the same check again.

Part 8. Peer Review

Once your manuscript clears the pre-screen, it enters peer review. This is the substantive evaluation by external reviewers. Peer review is the part of academic publishing that most surprises first-time authors, often in uncomfortable ways.

What peer reviewers do

Peer reviewers are scholars or practitioners with relevant expertise who evaluate your manuscript on behalf of the journal. They read your work, write a structured review, and recommend an editorial outcome. Reviewers are independent of the editorial board; they do not know each other or coordinate.

JAIIG uses a structured rubric. Reviewers score your manuscript on seven criteria: relevance to AI governance, originality and contribution, methodological soundness, evidence and citation integrity, practical and policy relevance, clarity, structure, and coherence, and ethical and governance awareness. Each criterion is scored 1 to 5. Reviewers also write narrative comments and make a recommendation: accept, accept with minor revisions, revise and resubmit, or reject.

Cross-reference

The complete rubric, with score anchors and reviewer guidance, is in the Peer Reviewer Scoring Rubric document. Reading the rubric before you submit helps you understand what reviewers will look at.

What reviewer reports look like

A reviewer report typically contains: a summary of what your paper is about (this confirms the reviewer understood it), an assessment of strengths, an assessment of weaknesses, specific comments on sections or claims, and a recommendation. The narrative is usually 500 to 2,000 words.

Reviewers write candidly. They will identify problems they consider serious. Some of their comments will be useful; some will be wrong; some will be abrasive. This is normal. The reviewer's job is to help the editor make a decision, not to be polite.

What to do when reviews arrive

When you receive your reviews, three things help:

Wait before responding.

First-time authors often have strong emotional reactions to negative reviews. Wait at least 24 hours before responding. The first reading is almost always worse than the second.

Read with the goal of understanding.

Even comments that seem unfair often contain a grain of truth about how the manuscript is being read. Try to identify what the reviewer was responding to, even if you disagree with their conclusion.

Talk to your advisor or a peer.

Showing the reviews to someone who has been through the process before helps calibrate your response. Experienced authors recognize patterns ("this is a normal Reviewer 2 complaint") that first-time authors miss.

Responding to revisions

If your manuscript receives a "revise and resubmit" or "accept with minor revisions" decision, you will write a response document along with the revised manuscript. The response document does three things:

21. Lists every reviewer comment.
22. Describes how you addressed each comment in the revision.
23. Where you disagree with a comment, explain why and what you did instead.

Editors and reviewers read the response document carefully. Skipping comments or providing vague responses is the surest way to a rejection on resubmission.

When you address a comment, point to the specific section and paragraph in the revised manuscript where the change was made. "We have added a paragraph to section 3.2 (page 11, lines 14-22) addressing this concern" is much more useful than "We have addressed this concern."

When you disagree with a comment, do so respectfully and substantively. Reviewers are sometimes wrong. A polite, well-reasoned disagreement is acceptable. Defensive or dismissive responses are not.

On disagreeing with reviewers

First-time authors are often reluctant to disagree with reviewers, even when the reviewer is mistaken. This deference is misplaced. Editors expect you to defend your work where you have grounds to do so. Capitulating to every reviewer comment, especially incorrect ones, weakens the manuscript and signals that you do not understand your own work. Disagree when you have reason to. Just do it well.

Part 9. After Acceptance

When your manuscript is accepted, you have done the hard work. The remaining steps are administrative.

Production

Your accepted manuscript enters production. The journal copyedits the prose for grammar, clarity, and style consistency. You will receive a proof — the typeset version of your article — for review. Read the proof carefully. Errors at the proof stage are easier to fix than errors after publication.

Common issues to watch for in proofs:

- Author names spelled correctly. Affiliations are correctly formatted.
- Equations and special characters rendered correctly.
- Figures and tables in the right place, captions correct.
- References intact, especially DOI URLs.
- Hyperlinks resolve to the correct destinations.

DOI assignment and online publication

Once the proof is finalized, the journal assigns a DOI to your article and publishes it online. The DOI is permanent. Once assigned and published, the article cannot be retracted casually; corrections and retractions follow established publication ethics procedures.

After publication

Your responsibilities continue after publication. You should:

- Cite your published article correctly when others ask for the citation.
- Respond to reasonable correspondence from readers about your work.
- Notify the journal if you discover an error in the published version.
- Comply with any data sharing commitments you made in the Data Availability Statement.

Cross-reference

For copyright, licensing, and author rights, see the journal's Copyright and Privacy Statement. JAIIG publishes under a non-commercial Creative Commons license; you retain copyright in your work.

Part 10. Frequently Asked Questions

These are the questions first-time authors most often ask. The answers here are concise; for authoritative detail, follow the cross-references to the governance documents.

Q: Can I submit work I previously presented at a conference?

A: Yes, provided the manuscript is substantively expanded beyond the conference version and the conference paper is cited. Work-in-progress presentations and poster abstracts are not considered prior publication for these purposes.

Q: Can I submit work that is also under review at another journal?

A: No. Concurrent submission to multiple journals is a recognized integrity violation and grounds for rejection.

Q: Do I need IRB approval for a conceptual paper?

A: No. IRB approval is required for work involving human subjects or personal data. Conceptual, theoretical, and policy analyses do not require IRB approval. Write “Not applicable” in the IRB section.

Q: How long does the review process take?

A: Pre-screen typically completes within five business days. Peer review typically takes 8 to 12 weeks from the clearance of the pre-screen to the first decision. Revise-and-resubmit cycles add another 4 to 8 weeks each. The total time from initial submission to acceptance typically ranges from 4 to 9 months.

Q: What if I disagree with the editor’s decision?

A: If the decision is rejected and you believe the rejection was based on a clear misunderstanding of your work, you may write to the editor-in-chief to request reconsideration. Such requests succeed only rarely. Most rejected manuscripts are better served by revision and submission elsewhere.

Q: Can my advisor be a co-author if they did not write any text?

A: Authorship requires substantial intellectual contribution to the work, not just writing. If your advisor contributed to the conception, methodology, analysis, or framing of the work, they may legitimately be a co-author. If they only supervised your work in a general sense, they belong in Acknowledgments, not authorship. The CRediT taxonomy used in the Author Contributions section helps clarify this.

Q: How do I list authors when there are many?

A: List all authors who contributed substantially. The Author Contributions section should specify each author’s role using the CRediT taxonomy. For 21 or more authors, list the first 19, an ellipsis, and the final author in the reference list (this rule applies only when others cite your work, not when you list authors on your own paper).

Q: Can I cite my own prior work?

A: Yes, when relevant. Self-citation is normal in academic writing. Excessive self-citation (citing every previous paper you have written, regardless of relevance) is a recognized integrity issue. Cite your prior work where it actually supports a claim in the current manuscript, not to inflate your visible citation count.

Q: Do I need to share my data?

A: It depends on the data type. The Data Availability Statement specifies what is available and how. For empirical work using public data, the data should be accessible to readers. For sensitive data (proprietary, personal, regulated), the statement should explain access conditions. Vague phrases such as “available upon request” do not specify the pre-screening conditions.

Q: What if my data has confidentiality restrictions?

A: Specify the restrictions in the Data Availability Statement. Acceptable: “The data are not publicly available because they contain personal information protected under [specific regulation]. Aggregate statistics underlying the analyses are available from the corresponding author.” The restriction is acceptable; the vagueness is not.

Q: Can I use AI tools to help with my manuscript?

A: Yes, with disclosure. See Part 5 of this handbook. AI use is not prohibited, but it must be disclosed in the AI Assistance Disclosure section, and the authors must verify AI-generated content.

Q: What if I find an error in my published article?

A: Notify the journal promptly. Minor errors (typographical, citation formatting) can usually be corrected through an erratum. Substantive errors (incorrect data, incorrect claims) may require a correction notice or, in extreme cases, retraction. Acting promptly when errors are discovered is itself a sign of integrity.

Q: Can I share preprints of my submitted manuscript?

A: Yes. JAIIG permits preprint posting before, during, and after submission. Cite the preprint version in your submission’s reference list if you cite earlier versions of work that became part of your submission.

Q: What if my manuscript is rejected?

A: Most rejected manuscripts are eventually published, often after revision and submission to another venue. Read the rejection letter and reviews carefully, identify what can be improved, and submit elsewhere. Rejection is not a verdict on your work or your career; it is a single editorial decision in a particular venue at a particular time.

Q: How do I find a journal if JAIIG is not the right fit?

A: Adjacent venues include AI and Ethics, AI & Society, Big Data & Society, Minds and Machines, Journal of Responsible Innovation, Journal of Information Policy, and Policy & Internet. Each has a different scope and conventions. Read recent issues of any journal you consider before submitting.

Closing

Academic publishing is a long apprenticeship. Your first manuscript is the first project; it is not the only one. Many of the things this handbook covers will become second nature with practice, and you will develop your own sense of what works in your subfield, with your collaborators, for your audience.

JAIIG exists to publish good work on important questions. The journal's commitment to first-time authors is genuine. The pre-screen, the rubric, the structured documentation, the layered governance — all of it is designed to make the process more transparent and fairer, particularly for authors entering academic publishing without the institutional advantages some have.

Your work matters. The questions you are working on matter. Submit it, engage with the feedback, and keep going.