

Book Indexing and Artificial Intelligence

Elizabeth Bartmess

AI Committee of the American Society for Indexing

ai.indexing@asindexing.org

GW Ethics in Publishing Conference 2025



The American Society for Indexing's AI Committee

- Formed Summer 2025
- Our work:
 - [White paper on AI and indexing](#): “LLM-Generated Book Indexes: Can They Replace Professionally Created Indexes?”
 - [Statement on AI and Book Indexing](#) with recommendations for indexers, authors/editors, and publishers
 - Reviews of AI-based indexing software
 - Outreach/collaboration with publishers, authors, and indexers

Book indexes

- Empower readers to quickly access information and explore thematic connections
- Provide an “exploded” map of the book’s themes, concepts, citations, etc.
- Are not autogenerated but are produced by a human indexer, who
 - Reads the entire book
 - Analyzes/extracts substantive terms/concepts
 - Orders them alphabetically
 - Adds cross-references to express relationships
- Are complex, cognitively intense, and time-consuming to produce

Book index expectations

- Readers expect indexes to accurately reflect the book's contents
- Authors expect indexes to respect their terminology and arguments
- Indexes fail to meet expectations when they:
 - Omit relevant substantive terms or concepts, thus denying readers access to the full content of the book
 - Misrepresent the author's intent, thus violating authorial integrity
 - Contain inaccurate or incomplete page references, thus violating the reader's trust
- Such indexes are consequently unethical

Book index characteristics

- A good index must meet many criteria, some of which may differ based on subject matter and audience
- Three minimum criteria are
 - Completeness: index must guide reader to all relevant info in the book
 - Navigability: index must guide reader to subtopics and related topics
 - Accuracy: should not contain false or inaccurate information, and should reflect the author's terminology

Our experiments

- Asked three AIs to create indexes for three public-domain books
- Compared the AI's output to the books' existing indexes on the three necessary criteria listed above
- Tests done during first half of 2025
- Currently analyzing data from newer models

Als and books used

- Full-length, public-domain books with existing indexes; in PDF format
- ChatGPT 4.5:
 - *The Behavior of Organisms*, by B.F. Skinner
 - *How to Stop Worrying and Start Living*, by Dale Carnegie
- Perplexity.AI with Deep Research Mode:
 - *The Psychology of Day-Dreams*, by J. Varendonck
- Claude 3 Opus:
 - *The Psychology of Day-Dreams*, by J. Varendonck, minus the last ~60 pages

Prompt used:

“Please create a professional-quality index for the attached book, using the same techniques that a professional indexer would use (e.g., including cross-references, subheadings, etc). Ignore the book's existing index; instead, create your own.”

For Perplexity and Claude

- Included a copy of the Criteria for the ASI Excellence in Indexing Award
- Prompt: The same as previously, plus “Please ensure that the completed index adheres to the guidelines in the attached file.”
- Due to file size limitations, for Claude we deleted the last ~60 pages, which included the existing index; hence we omitted the instruction to ignore the existing index

Findings:

- AI generated indexes had issues with technical aspects including:
 - Formatting
 - Alphabetization
 - Failure to detect page numbers
- These are relatively easy (though still necessary) fixes
- More troubling, AI generated indexes also failed the three minimum criteria

Criterion 1: Completeness

- An index should provide access to all indexable information in the book.

Completeness: AI-generated indexes

- Compared to the books' original indexes:
 - AI-generated indexes contained fewer than 1/3 as many main entries
 - Omitted at least one-third (more typically, nearly all) of the indexable names
 - Entirely failed to index titles of works
- AIs omitted material by other measures as well:
 - They provided < 2 entries per page on average (expected range for an index: 2–10 entries per page)
 - ChatGPT failed to index >90% of the pages of one book and 21% of another

Criterion 2: Navigability

- Indexes should guide readers to subtopics and related topics via cross-references (*see* and *see also*)

Navigability: AI-generated indexes

- Als failed to use cross-references effectively to point readers to subtopics and related topics
- AI cross-reference use was incorrect and idiosyncratic

Criterion 3: Accuracy

- Indexes should be accurate: they should contain no false or invented information
 - Page references should be correct
 - Terminology should reflect the author's use
- We used the “copyeditor's rule of thumb” (Mulvany, 2005):
 - Check 10% of page references
 - If only 1–2 errors are discovered, check another 10%
 - If >2 errors are discovered, the entire index must be checked

Accuracy: AI-generated indexes

- All AI-generated indexes failed the copyeditor's rule of thumb
- Significant numbers of pages omitted entirely from the index
- A majority of entries in all indexes had one or more additional problems
 - Imposed terminology instead of the author's terminology
 - Passing mentions (term referenced but not substantively discussed, and should not have been indexed)
 - Excessively long page ranges (typical page range is ~8 pages; AI indexes had ranges up to 99 pages)
 - Mischaracterized the nature of a discussion
- Only 13–26% of the total entries checked were adequate

Summary

- AI-generated indexes
 - Are incomplete, lack navigability, and are inaccurate, thus failing the three necessary criteria
 - Suffer from a host of lesser technical issues which would require extensive human remediation
 - Mislead readers and fail to serve their needs
 - Mischaracterize and misrepresent the author's work
- The fine-grained professional judgment of a trained and skillful human indexer is still necessary to produce an adequate index, let alone an excellent one

Ethical considerations

- Quality: AI-generated indexes are of poor/inadequate quality, thus do a disservice to both author and reader
- Transparency: AI tools are not always transparent about how they use uploaded content, thus pose potential copyright, privacy, and authorial integrity issues
- Honesty: Authors should always be informed about AI usage in production of an index to their book
- Integrity: Financial pressures encourage cheap and fast over costly and time-consuming, sacrificing quality for speed

Recommendations for authors and editors

- Don't use AI-generated indexes
- See ASI's [Advice for Editors & Authors](#) for guidance on hiring a professional indexer

Recommendations for Publishers

- If you work with authors who provide their own indexes:
 - Make them aware of the shortcomings of AI-generated indexes
- If you are considering adopting an AI product for indexing:
 - Be aware that issues with AI-generated indexes are not evident on a surface read
 - Check ASI's [AI News](#) category to see if the product has been reviewed
 - Request a review, or hire an indexer to evaluate the product
- If you use AI for indexing despite its flaws:
 - Disclose your use to authors ahead of time, including the specific AI engine and what it may or may not do with the uploaded files
 - Offer your authors the option to provide a professionally created index instead

Resources

- ASI's AI Committee: ai.indexing@asindexing.org
- [White paper on AI and indexing](#): “LLM-Generated Book Indexes: Can They Replace Professionally Created Indexes?”
- [Statement on AI and Book Indexing](#) with recommendations for indexers, authors/editors, and publishers
- [Advice for Editors & Authors](#)
- ASI's [AI News](#) category