Mary Aycock

Prompting Generative AI to Catalog

The Promise and the Reality

As libraries shift their budgets toward investing in digital resources and content, employees must also streamline work processes to accommodate thousands, if not millions, of titles. Missing or substandard metadata can hinder discoverability, impacting the library's return of investment in these resources, not to mention the opportunity costs that result for our users.[1]

Thus, our library faced a quandary when notified in fall 2024 about nearly one hundred ebook conference titles lacking associated MARC bibliographic records. Not only was the vendor unable to supply the records, but they didn't even exist in WorldCat, our customary bibliographic to-go database of MARC records. How could we fulfill our commitment to provide access to these MARC records with the limited resources at our disposal?

Enter the ever-present hope that technological advancements can save time for the cataloger. Because the recent advancements of large language models (LLMs) pose disruptive ramifications for those working in knowledge industries, many library leaders have recommended a proactive approach in experimenting with these tools.[2] Some cataloging and metadata departments have responded to this call with experimentation and skepticism.[3]

Perhaps this need to catalog one hundred ebook conference titles could supply an opportunity for our own test case using a specialized generative artificial intelligence (AI) called CatalogerGPT.[4] This plugin for OpenAI will generate MARC records or fields based on prompts and uploaded files. The output is generated in an easy-to-view and easy-to-edit format, a mnemonic text file familiar to most who use MarcEdit.[5] These files can be easily copied and pasted into a blank MarcEditor file.

Although open source models exist, they often require technical expertise and time to implement, which may propel many catalogers to turn to CatalogerGPT or other commercial models instead. Our department's own brief experimentation with this tool demonstrated that this model has the capability to draft descriptive cataloging and supply access points. Particularly impressive was the ability to generate a table of contents from an uploaded file. Perhaps such AI-generated records might prove better than brief or skimpy machine-generated records. Whether this tool would prove to be an oracle that could spit out a good enough cataloging record remained to be seen.

Mary Aycock, formerly of Texas State University, is systems and cataloging librarian at Del Mar College, email: maycock@delmar.edu.

Prompting the Oracle, or a Journey, of Many, Many Prompts

An employee downloaded the front matter for these ebooks. The frontmatter PDF files consisted of a title page, a title page verso, introductory material, and a table of contents. I thus embarked on the journey of multiple prompts in December 2024 to discover the best words and approach to obtain the desired record output. The chat transcript is available for viewing as well as through selected screenshots.[6] Unfortunately, CatalogerGPT limited output to three MARC records a day using the uploaded files (unless we subscribed to a paid version per a pop-up box).

Experimentation over a few weeks yielded some observations. The first prompt requested "Create a MARC record from the attached content." This one-shot prompt yielded a subpar record, not only needing extensive editing but also missing critical fields, such as a conference heading access point, genre headings, etc.

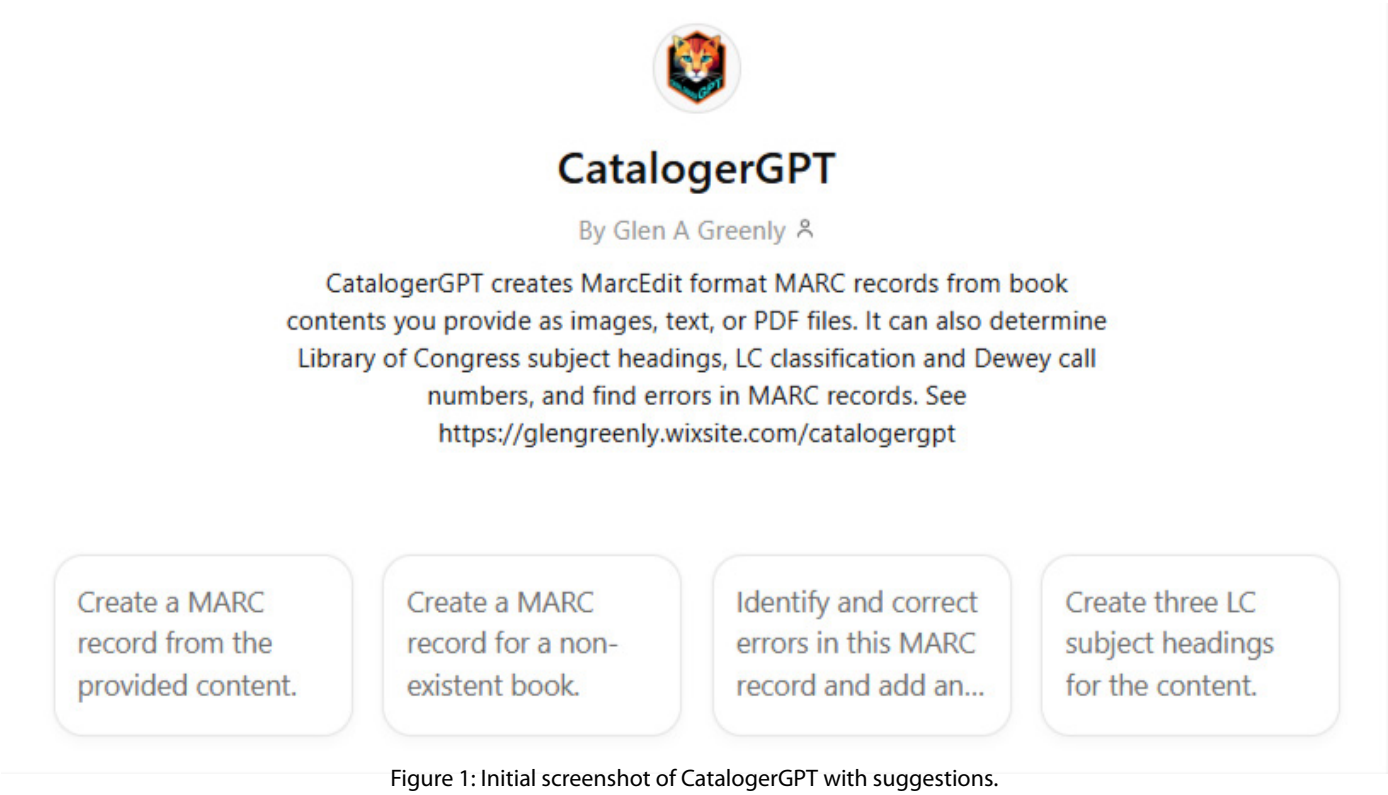

Figure 1: Initial screenshot of CatalogerGPT with suggestions.

ISBNs and titles were not transcribed correctly for some of the initial titles. Having an accurate title and ISBN comprised a bare minimum requirement because the records would be submitted to WorldCat.

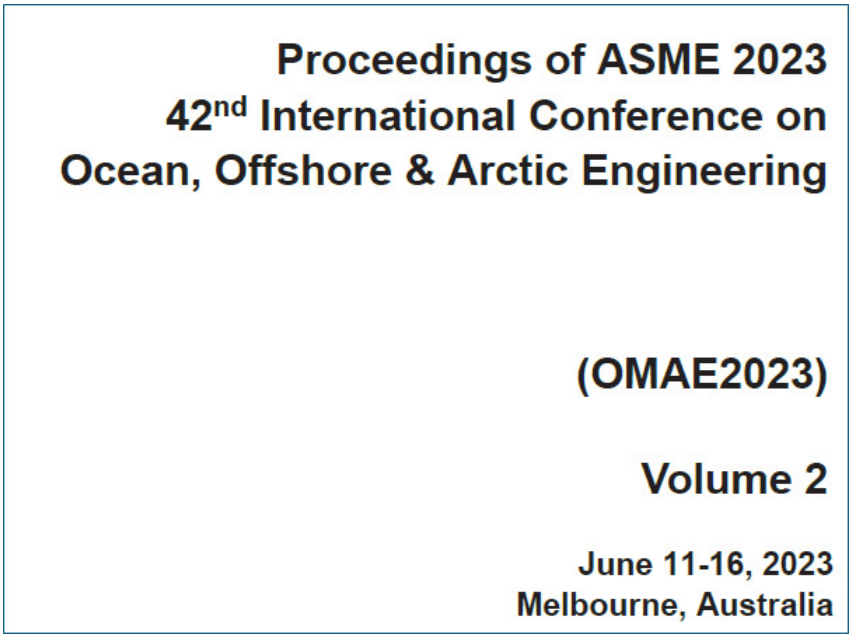

Figure 2: Title page of example volume.

CatalogerGPT provided the following erroneous output for title field shown in Figure 2:

=245 10$aProceedings of the 2023 Ocean, Offshore, and Arctic Engineering Conference (OMAE2023)$nVolume 2 :$bJune 4-9, 2023, Melbourne, Australia.

The ISBN listed on title page verso for the example if 978-0-7918-8684-7. But CatalogerGPT generated the following erroneous output for the ISBN field:

=020 \\$a978079188

Revising a previously generated record to include in the prompt required time and cataloging knowledge (Figure 3). However, this extra effort reaped rewards by producing more reliable MARC output.

```
=LDR  04543nam a2200409li 4500
=008  241209t20232023nyua\\\\ob\\\\10\\0\eng\d
=040  \\$aTXI$beng$erda$cTXI
=020  \\$a9780791888687
=050  \4$aTC1505$b.I58 2023 v.2
=111  2\$aInternational Conference on Ocean, Offshore, and Arctic Engineering Conference$n(42nd :$d2023 :$cMelbourne, Australia)
=245  10$aProceedings of the ASME 2023 42nd International Conference on Ocean, Offshore & Arctic Engineering (OMAE2023)$nVol
=264  \1$aNew York, N. Y. :$bThe American Society of Mechanical Engineers,$c[2023].
=264  \4$c@2023.
=300  \\$a1 online resource$billustrations.
=336  \\$atext$btxt$2rdacontent
=337  \\$acomputer$bc$2rdamedia
=338  \\$aonline resource$bcr$2rdacarrier
=500  \\$aProceedings include peer-reviewed papers from the 2023 Ocean, Offshore, and Arctic Engineering Conference, held in Melb
=504  \\$aIncludes bibliographical references.
=505  0\$aCollision and Crashworthiness: Numerical Simulation of Ship Collision Scenarios -- Damage Assessment of FPSO Structur
=520  \\$aThe proceedings of the 2023 Ocean, Offshore, and Arctic Engineering Conference (OMAE2023) Volume 2 focus on structure
=588  0\$aSome metadata was created with AI assistance on December 9, 2024.
=650  \0$aOffshore structures$xReliability$vCongresses.
=650  \0$aStructural analysis (Engineering)$vCongresses.
=650  \0$aMooring of ships$xReliability$vCongresses.
=650  \0$aRenewable energy sources$xReliability$vCongresses.
=650  \0$aRisk management$vCongresses.
=650  \0$aHydrodynamics$vCongresses.
=650  \0$aShips$xHydrodynamics$vCongresses.
=655  \7$aConference papers and proceedings.$2lcgft
=710  2\$aAmerican Society of Mechanical Engineers.$bOcean, Offshore, and Arctic Engineering Division.
```

Figure 3: Example MARC record.

Not surprisingly, extra care needed to be taken with the model record because any inadvertent errors, such as with subfields, would be faithfully copied. In the 245 field below, $r should be subfield $c.

=245 10$aProceedings of the ASME 2023 42nd International Conference on Ocean, Offshore & Arctic Engineering (OMAE2023)$nVolume 2 :$bJune 11-16, 2023, Melbourne, Australia /$rConference sponsor: Ocean, Offshore and Arctic Engineering Division.

Here the generated field faithfully followed the model record to transcribe an erroneous $r in the 245 field.

=245 10$aProceedings of the ASME 2023 42nd International Conference on Ocean, Offshore & Arctic Engineering (OMAE2023)$nVolume 1 :$bJune 11-16, 2023, Melbourne, Australia /$rConference sponsor: Ocean, Offshore and Arctic Engineering Division.

Apart from human introduced mistakes in the model, wrong subfields were sometimes generated, such as in this conference heading:

=111 2\$aDesign of Medical Devices Conference$n(2024 :$dMinneapolis, Minn.)

The subfields in the heading should be:

=111 2\$aDesign of Medical Devices Conference **$d**(2024 ***:$c***Minneapolis, Minn.)

Specifying very narrow instructions seemed to limit the model to generating a minimal record, lacking fields that were not explicitly mentioned, as if micromanaging the model constrained it (see Figure 4).

Please catalog records using the attached PDFs and the following criteria:
1. Do not include 001 field
2. Do not include 245 $h or 245 $c
3. Do not include 082
4. Do not include 490/830 (ASME Digital Collection)
5. Do not include 090
6. Include 588 field: Some metadata was created with AI assistance on {today's date}.
7. Create 505 fields using all the titles, but not the section heading or the authors, such as: $a UMBRELAA: Design of a variable-sized left atrial appendage occlusion device for stroke prevention -- Blood pressure prediction from photoplethysmogram signal using artificial intelligence –
8. Include 655 Conference papers and proceedings.$2lcgft
9. Include 300 $billustrations
10. Include a 111 field with the appropriate subfields, such as: $aDesign of Medical Devices Conference $d(2023 : $cMinneapolis, Minn.),
11. 264 should be New York, N. Y.
12. 040 should be $aTXI$bengl$erda$cTXI

Figure 4: Prompt with narrow instructions.

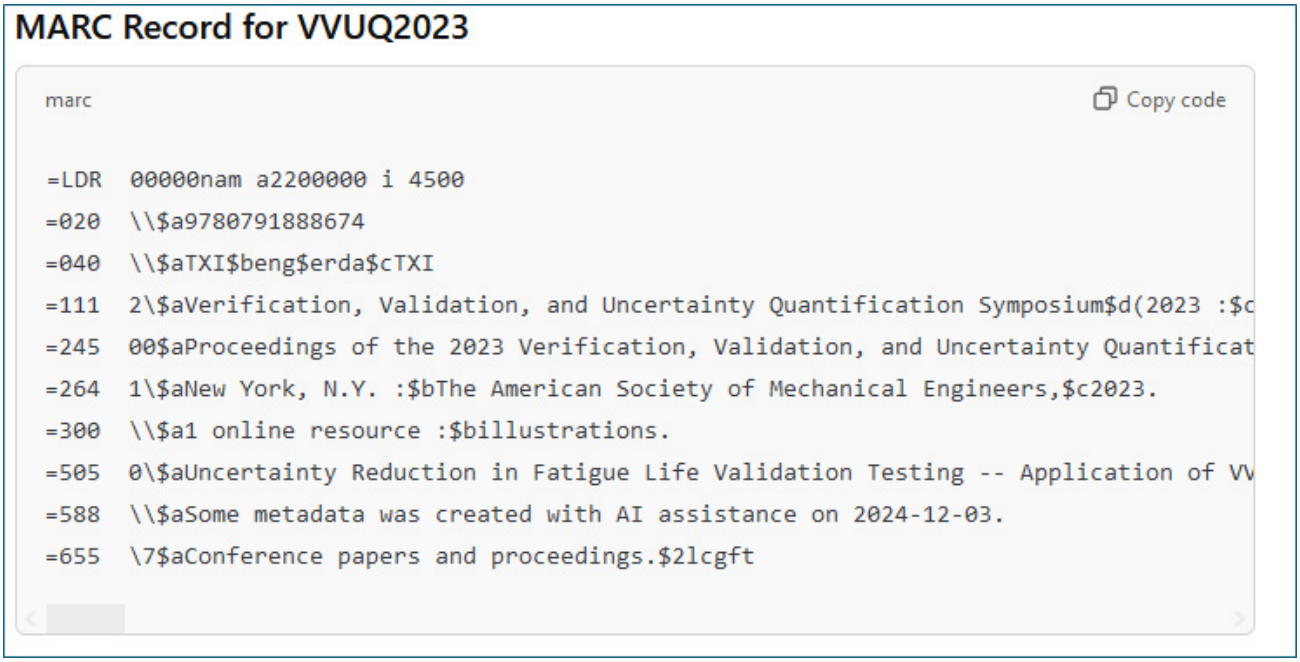

Figure 5: Generated minimal record.

The fields missing from the generated MARC record (Figure 5) included:

- Call number (050 field)
- 33x fields that are standard in current records (336, 337, 338 fields)
- Notes about bibliography (504 field)
- Summary (520 field)
- Subject headings (6xx fields)
- Sponsoring organizations (710 field)

We had high hopes about the possibility of obtaining a granular table of contents incorporating titles and authors of individual conference papers (metadata that normally would be too time-consuming for catalogers), but that exploration proved problematic:

- The model would transcribe the first page of the table of contents and would need continuous prompting for subsequent pages.
- It would often hallucinate titles of papers, requiring too much work reviewing and editing the records.

Even requesting *valid* Library of Subject Cataloging Headings did not necessarily result in authorized ones. Just when the prompt seemed refined enough to declare "finished," another error would pop up in the record. At least pointing out an error yielded a gratifying response (Figure 6).

What was the best prompt? It turns out that prompt engineering can also be delegated to CatalogerGPT—as shown, a question asking about what prompt to use and the resulting long response (Figure 7 and supplemental document).

In contrast to the nineteen words generated by the directive prompt (Figure 4), the open-ended question from Figure 7 produced a verbose response of 827 words excluding the MARC records. It also seemed a bit repetitive to include both an "ideal prompt" and

Figure 6: Pointing out error and CatalogerGPT response.

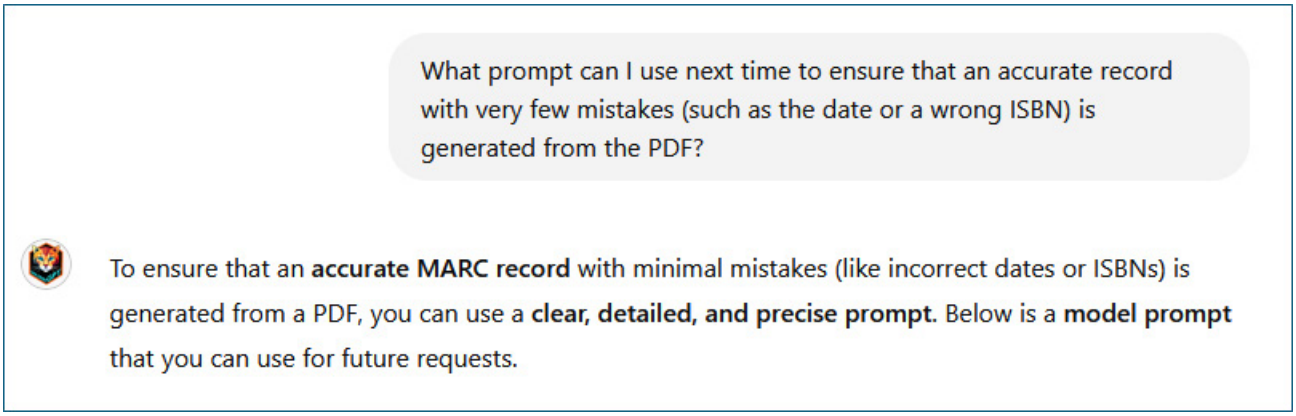

Figure 7: Open-ended question about better prompt.

an "example prompt" in its answer. It's also clear that CatalogerGPT did not generate the featured MARC records from a vacuum but incorporated much of our previous prompts and example records in them. Catalogers will still need use their cataloging expertise in prework regardless.

Was this the prompt to end all prompts? Nope. Subsequent prompts still required continual tweaking, including emphasizing that the ISBN in particular should be accurate.

Evaluation of the Output (Oracle Answer)

Did the cataloging oracle live up to expectations and generate a good enough catalog record that saved time? It depends on your expectations.

Drafting a record via generative AI proved helpful but required constant vigilance to ensure the accuracy of transcription fields (title, ISBN) as well as the relevance of the access points. Due to the limit of three a day, the routine for this project included generating MARC records each day, importing them into Connexion cataloging software, and revising them. The review included checking the following fields:

- ISBN and title field (critical for identification)
- Call number
- Conference heading
- Table of contents

- Subject headings (via OCLC, controlling the headings revealed which ones were valid at a glance)
- Access point for the organization as well as the sponsoring committee
- Date of the metadata note was often wrong: "Some metadata was created with AI assistance on 2024-12-20"

The jagged technological frontier visualizes the boundary at which AI can be an asset versus a detriment for the user.[7] Due to the inaccuracies of the generated MARC records in this project, MARC record generation barely landed on the favorable side of this frontier *if* using an effective prompt. However, even these records could not be trusted without a cataloger in the loop. Admittedly there is an art to cataloging, but it must be grounded in reality-based adherence to standards and norms, not one of creative writing.

There can be a steep learning curve to gaining the expertise of a cataloger, which by one recent estimate requires three to five years of experience to obtain.[8] An experienced cataloger who uses all the tools of their trades (deriving records, macros, quick editing) can often catalog accurately and rapidly, particularly if the records are uniform enough. I do not routinely catalog in my current position, but in fifteen minutes, I was able to draft nine ebook conference volumes—three times more than CatalogerGPT would allow per day. By the end of the project, some of the records had been generated via generative AI and some through manual processes. All records needed further enhancement and review.

Yet framing expert catalogers against untrustworthy AI is an oversimplification and overlooks the advantages of a beneficial partnership. While the project proved more time consuming than anticipated, it still provided an enlightening exploration of the capabilities and limitations of a specialized generative AI at this current time. These models demonstrate clear potential to assist catalogers in their work but only under close supervision. These results also agree with several other research articles that concluded that LLMs could be useful in drafting records but still required human oversight (preferably with enough cataloging knowledge to efficiently evaluate the output).[9,10]

Conclusion

These AI tools could be especially helpful for generating potential subject headings and summaries in those areas that the cataloger lacks subject matter expertise with the caveat that any generated fields would still need to be validated.

Based on my experience, I have the following suggestions for any metadata worker considering such a project with LLMs.

1. Decide what fields are important and emphasize these in the prompt.
2. Create a model record (either from scratch or revising an initial draft generated by the LLM).
3. Use a prompt similar to the one CatalogerGPT suggested (supplemental document) or ask for suggestions on an effective prompt.
4. Prepare to review output records, particularly if they will be submitted to a cooperative cataloging database.

Hopefully, such suggestions can assist other catalogers and metadata workers in reaching an effective prompt in fewer attempts than this project (more than twenty-five prompts), thus compensating for some of the energy consumption expended.

Notes

1. Heather Moulaison-Sandy, Hyerim Cho, and Felicity Dykas, "Approaches to Conceptualizing the Cost of Academic Library Cataloging: Discourses on Metadata Creation Cost, Value, and Worth," *Library Trends* 70, no. 3 (January 2022): 387–408, https://doi.org/10.1353/lib.2022.0001.

2. Jeehyun Davis, "Artificial Intelligence (AI) and Academic Libraries: A Leadership Perspective," *College & Research Libraries News* 85, no. 8 (September 11, 2024): 347, https://doi.org/10.5860/crln.85.8.347.

3. Program for Cooperative Cataloging Task Group on Strategic Planning and AI, *Final Report of the PCC Task Group on Strategic Planning and Artificial Intelligence*, February 6, 2023, accessed December 26, 2024, https://www.loc.gov/aba/pcc/taskgroup/TG-Strategic-Planning-AI-final-report.pdf.

4. Glen Greenly, CatalogerGPT, "Home," accessed January 7, 2025, https://glengreenly.wixsite.com/catalogergpt.

5. MarcEdit Development, "HowTo: Editing Your First Recordset," November 30, 2022, https://marcedit.reeset.net/howto-editing-your-first-recordset.

6. ChatGPT, "ChatGPT - MARC Record Creation," accessed January 22, 2025, https://chatgpt.com/share/67915ba6-73b0-8006-ba4c-b45d69d917a3.

7. Fabrizio Dell'Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," *SSRN Electronic Journal* (2023), https://doi.org/10.2139/ssrn.4573321.

8. Heather Moulaison-Sandy and Zach Coble, "Leveraging AI in Cataloging: What Works, and Why?" *Technical Services Quarterly* 41, no. 4 (October 1, 2024): 375-83, https://doi.org/10.1080/07317131.2024.2394912.

9. Eric H. C. Chow, T. J. Kao, and Xiaoli Li, "An Experiment with the Use of ChatGPT for LCSH Subject Assignment on Electronic Theses and Dissertations," arXiv, July 10, 2024, https://doi.org/10.48550/arXiv.2403.16424.

10. Shoichi Taniguchi, "Creating and Evaluating MARC 21 Bibliographic Records Using ChatGPT," *Cataloging & Classification Quarterly* 62, no. 5 (July 3, 2024): 527–46, https://doi.org/10.1080/01639374.2024.2394513