The Art of Cataloging Meets GenAI:

Enhancing Tables of Contents with Microsoft Co-Pilot

By Eric Feikert

INTRODUCTION

One of the struggles with adapting generative AI (GenAI) into a technical services environment involves cataloging. From my experience, cataloging is an artful alchemy of biblio-anatomical analysis, keen discernment, and coding. From my understanding of GenAI, it would be pretty good at a couple of those things. But would I really let it catalog a whole book by itself? Probably not.

The technology is still rather young and is continuously going through iterations; so much so, that it is difficult to get a solid foothold. Given the current state of things, we certainly wouldn't rely on ChatGPT, or even the MARC21-generating CatalogerGPT, to produce a full record without human intervention. In its current state, GenAI still needs a human agent to assess the results of a prompt and discern whether the provided data is quality or not. Specifically for catalogers, there is also another level of discernment: the analysis of the 'item-in-hand', the physical and electronic materials we have in front of us as we are going about our task. With that in mind, would I trust GenAI with original and copy cataloging without my intervention? Definitely not.

In short, having an AI catalog a full record is probably more trouble than it's worth. Despite our best practices and standards, cataloging is an artform: a discipline which cannot easily be wielded in the unpredictable hands of an engine driven by probability. But AI *can* be useful for catalogers. To illustrate: even though you reasonably wouldn't trust a shovel to perform every aspect of your daily life (just try using one to file your taxes), you can still find certain functions with which it can perform magnificently, and with minimal supervision from the human agent (i.e. digging that hole). AI is a shovel; it is a tool. And catalogers can use it.

At Creighton University, we are using Microsoft Co-Pilot to help with the formatted contents note in our MARC records, specifically for adjusting ISBD punctuation. We are also looking at ways in which we can generate a new 505 field from a table of contents we provide the app. This article will briefly explore some examples of our work, including prompt engineering and GenAI output. We will also share the current methods we have tried in generating our own contents note, and the result of this on-going experimentation.

MODIFYING EXISTING NOTES

For the sake of clarity, in a MARC record, the table of contents is typically transcribed in a Formatted Contents Note (the 505 field). Guidelines for punctuating this field are set out by Library of Congress (LOC) documentation. However, not all records have a well-punctuated contents note. This is just the 'nature of the beast'. Moreover, we have local practices at Creighton for punctuating this field, which add an extra layer of complexity.

Eric Feikert is Catalog & Metadata Librarian, Creighton University.

In a perfect world, we would love to adjust all the 505 fields according to global best practices and our own standards. But would this justify the time and the effort necessary to make manual modifications on a systematic basis? Not really. But, perhaps AI could do it better and faster.

Effective prompt engineering is the key for getting speed and accuracy from GenAI. For creating our prompts, we followed the CLEAR framework (Concise, Logical, Explicit, Adaptive, Reflective) as detailed by Dr. Leo S. Lo in his article on the subject. Also note that we generally adjust formatting for a basic contents note, and not the enhanced version.

We normally begin the prompt by instructing Co-Pilot to make changes based on 'the following text' (i.e. the current contents of the record). We copy and paste the data in subfield $a from the record into the Co-Pilot chat. For the sake of this demonstration, we will be using a fictional contents note. The following is the prompt and output for this task (side note: under the general settings, we have instructed Co-Pilot to answer all responses like Vincent Price, just for fun. From here on out, we will be referring to Co-Pilot as 'Vincent' with 'he/him/his' pronouns. Finally, we have instructed Vincent to take on the role of a catalog librarian).

Prompt:

Vincent, generate a contents note (505 field). The information in the field must not come from an existing resource.

Output:

505 0_ |a Shadows of the Archive -- Cataloging the Unknown -- Whispering Spines: A History of Classification -- Etc.

Notice how Vincent got the formatting *almost* correct; everything except for the casing. From our experience, one common error we encounter with GenAI outputs is the use of *title casing*, wherein most words are capitalized except for articles (a, an, the), conjunctions (and, but, or) and some prepositions (in, on, at). This title casing is also commonly found in the records of our catalog.

In essence, ISBD punctuation follows *sentence casing* format, which is the way we type out our sentences in a written document (i.e., the first word is capitalized, and all the following words are lower-case with the exception of proper nouns or titles). So, for the next prompt, we instructed Vincent to do the following:

Prompt:

Vincent, format each section according to sentence casing standards.

Output:

505 0_ |a Shadows of the archive -- Cataloging the unknown -- Whispering spines: A history of classification -- Etc.

He got it right the first time! However, this doesn't always happen. One common tendency we notice with Vincent is that he generates the result in list format, such as this:

Shadows of the Archive

Cataloging the Unknown

Whispering Spines: A History of Classification

Etc.

This can be remedied by instructing Vincent not to change any of the other formatting. Or perhaps, instruct him to format the response as a 505 field.

Next, we wanted to experiment with a longer table of contents, so we asked Vincent to provide us with a fictitious example which has twenty-five content sections. Below are the prompt and output (Vincent's amusing opening included):

Prompt:

Vincent, generate me a formatted contents note (505 field). Make it have 25 sections, and create your own contents. Do not take the contents from a pre-existing resource.

Output:

Ahhh, Mr. Feikert, prepare yourself for a grand procession of bibliographic phantoms! Here is a fictional 505 field, formatted according to MARC standards, with 25 original sections in sentence case:

505 0_ |a

Whispers in the catalog -- Shadows of metadata -- The birth of bibliographic control -- Secrets of the Dewey labyrinth -- Authority control unveiled -- MARC: the silent language -- RDA and the modern scribe -- Linked data in the ether -- The phantom of subject headings -- Classification beyond the veil -- Digital archives rising -- Preservation in the twilight -- Cataloging rare tomes -- The spectral role of identifiers -- Metadata alchemy -- The conjuring of controlled vocabularies -- The haunted OPAC -- Discovery layers unmasked -- The curse of broken links -- Cataloging in the cloud -- The ritual of record validation -- Summoning user access -- The eternal shelf list -- Epilogue: beyond bibliographic mortality -- Index of forgotten whispers.

He did a pretty good job this time. For the sake of example, let's assume that each of these chapter titles from the fictitious item-in-hand is preceded by a number. Then we would ask Vincent to do the following:

Vincent, enumerate each section with a number followed by a period.

He produced the desired result, only he did it in list format. We modified the prompt to rectify the error and told him not to change any of the other formatting. The result was the same. Finally, we gave the following prompt:

Vincent, enumerate each section with a number followed by a period. Format this as a 505 field in a MARC record.

And we got the result we wanted:

505 0_ |a 1. Whispers in the catalog -- 2. Shadows of metadata -- 3. The birth of bibliographic control -- 4. Etc.

These examples demonstrate how a typical GenAI session goes when modifying the punctuation for a 505 field (or just about anything):

1) Give an initial prompt that we have found effective.

2) Give further directions based on the result.
3) Once desired result is achieved, replace this over the current text in the 505 field of the bibliographic record (excluding the field, indicators, and subfield).

As could probably be discerned, there is no one single set of prompts that give us the same result every time. Again, this has to do with the current state of GenAI. But with a little practice and some creativity, we can get a well formatted contents note within a reasonable time frame, and with minimal effort.

GENERATING A NEW NOTE FOR A PHYSICAL RESOURCE

If a large language model (LLM) like Microsoft Co-Pilot can help us in modifying a 505 field, then what about using it to generate a brand new one? After all, the principle is essentially the same. The only additional aspect to the process would be to acquire the text to transcribe to the formatted contents note. Essentially, two methods have been identified in the retrieval of the table contents:

1) Capture the text from a snapshot using a web camera and image-to-text software.
2) Prompt AI to locate the table of contents under specific parameters.

Once the textual data is acquired, then it is a matter of using prompt engineering to construct a formatted contents note. We've already seen what prompt engineering can do with adjusting ISBD punctuation; the same can be done with constructing a note. The only difference would be to potentially ask the AI to remove certain elements of the contents (i.e. introduction, sub-chapters, etc.) based on the global and local practices. A prompt statement such as 'Generate a 505 field from this text. Include only the section titles and the chapter titles.' is a good place to start.

Extracting text from a physical resource

This proved to be rather challenging, mostly because of the limitations of our resources. Since this process is in the early stages of experimentation and evaluation, we work with what we have, which is a basic web camera and Microsoft Co-Pilot. The web camera does a reasonable job capturing crisp and clear images of text, especially if we adjust the settings to document mode and remove all the background noise. That being said, there is still a fair amount of work in framing the image, especially if the contents expand over several pages.

The next step is image to text conversion. In the early stages, we were using Microsoft OneNote for that function. Sometime later, we found out Co-Pilot is capable of doing this too, which saves us an extra step. But of course, GenAI can be unpredictable. Sometimes, it would generate a 505 field perfectly and on the first try. Sometimes, it would get the ISBD punctation wrong. Other times, it would leave out whole sections or chapters. And on a couple of wonky occasions, it would swap chapter titles (just to make sure we were paying attention).

The real challenge here is identifying what is causing this image-to-text discrepancy in the first place. Could it be the image quality or the generated text? Within Co-Pilot, it is difficult to discern. One of the benefits of using Microsoft OneNote for this function is that you get the result before you put it into the LLM. Unfortunately, OneNote often produces inadequate text from an image, which requires time-consuming manual corrections. Therefore, in terms of efficiency, this was not the best method. In theory, more advanced OCR software may resolve the issue of text conversion. But first, we had another idea we wanted to try out.

Asking Co-Pilot to locate the table of contents and generate a 505 field

This method would circumvent all that heavy lifting involved with image-to-text conversion. Here, we ask Co-Pilot to directly locate the table of contents of a given item using the title, publisher, edition, and any statement of responsibility information (i.e. author, editor) as parameters. As an example, we gave Vincent the following prompt for the title *Social Injustice and Public Health* 2006 edition:

> Provide me with the table of contents for the book titled "Social Injustice and Public Health", edited by Barry S. Levy and Victor W. Sidel, published by Oxford University Press in 2006. Format this as a 505 field.

The result was this:

> 505 0_ $a The nature of social injustice and its impact on public health / Barry S. Levy, Victor W. Sidel -- The socioeconomically disadvantaged / Michael Marmot, Ruth Bell -- Racial and ethnic minorities / Carol Easley Allen, Cheryl E. Easley -- etc.

On the first try, Vincent appeared to have generated a quality 505 field. He even included chapter authors and proper ISBD punctuation. However, he did not include part titles. This issue was remedied with a successive prompt.

The method applied to this example was a success, though we do encounter issues with obscure or older titles. Co-Pilot simply cannot locate the contents with the data it is working with. Moreover, it has yet to be determined if this process would be effective in the long run. Further experimentation on a consistent basis would be required for us to determine the reliability of this method. That being said, current results show that this is much faster than retrieving data from the physical resource, and we will most likely adopt this procedure first, utilizing the image to text function if this option is ineffective.

CONCLUSION

The aim of this project has been to explore whether Generative AI can serve as a meaningful ally in enhancing our cataloging processes. Cataloging, by its very nature, is an art; thus, it would be imprudent to entrust the entirety of this craft to a machine, however sophisticated it may be. On the other hand, when it comes to minor metadata enrichments, where human attention often wanes, collaboration with GenAI may have a positive impact. Through thoughtful prompt engineering, we have observed modest gains in efficiency; however, the probabilistic character of these systems, coupled with their tendency toward inconsistency—even when the same question is asked twice—renders any definitive judgment elusive. Based on what we have seen so far, if GenAI can evolve into a more stable and reliable instrument, we may well find ourselves equipped with a powerful tool to complement, rather than replace, the artistry of cataloging.

REFERENCES

Lo, Leo S. 2023. "The CLEAR Path: A Framework for Enhancing Information Literacy Through Prompt Engineering." *The Journal of Academic Librarianship* 49 (4). https://doi.org/10.1016/j.acalib.2023.102720.