

Web Archives Metadata Generation with GPT-4o: Challenges and Insights

Ashwin Nair, Zhen Rong Goh, Tianrui Liu, and Abigail Yongping Huang

ABSTRACT

Current metadata creation for web archives is time consuming and costly due to reliance on human effort. This paper explores the use of GPT-4o for metadata generation within the Web Archive Singapore, focusing on scalability, efficiency, and cost effectiveness. We processed 112 Web ARChive (WARC) files using data reduction techniques, achieving a notable 99.9% reduction in metadata generation costs. By prompt engineering, we generated titles and abstracts, which were evaluated both intrinsically using Levenshtein distance and BERTScore, and extrinsically with human cataloguers using McNemar's test. Results indicate that while our method offers significant cost savings and efficiency gains, human curated metadata maintains an edge in quality. The study identifies key challenges including content inaccuracies, hallucinations, and translation issues, suggesting that large language models (LLMs) should serve as complements rather than replacements for human cataloguers. Future work will focus on refining prompts, improving content filtering, and addressing privacy concerns through experimentation with smaller models. This research advances the integration of LLMs in web archiving, offering valuable insights into their current capabilities and outlining directions for future enhancements. The code is available at <https://github.com/masamune-prog/warc2summary> for further development and use by institutions facing similar challenges.

INTRODUCTION

The digital landscape is constantly evolving, and the need to preserve our online heritage has become increasingly urgent. The Resource Discovery (RD) department of the National Library Board Singapore (NLB) provides cataloguing services for the collections of the National Library (NL) and public libraries in Singapore. NL, like many other archives and libraries worldwide, collects and archives websites in its effort to preserve the mercurial history of the web.¹ This collection is known as Web Archive Singapore.² Each website is manually reviewed and catalogued according to the local application profile based on Dublin Core.³ NL has been crawling websites on a curated basis since 2006, with each website requiring individual website owner consent. In 2019, the National Library Board Act was amended to impart NL with the authority to harvest all websites in the .sg domain without explicit consent from each individual website owner.⁴ The legislative change resulted in the explosive growth of the web collection.⁵ As of June 2024, there are over 98,000 unique websites and webpages from both the domain crawl and

About the Authors

Ashwin Nair (ashwin_nair_madhavan@nlb.gov.sg) is Manager/Librarian (Systems), Resource Discovery Management, National Library Board, Singapore. **Zhen Rong Goh** (zgoh033@e.ntu.edu.sg) is a student in Business and Computing, Nanyang Technological University, Singapore. **Tianrui Liu** (tianrui.liu@u.nus.edu) is a Master of Science student in Computer Engineering, National University of Singapore. **Abigail Yongping Huang** (abigail_huang@nlb.gov.sg) is Assistant Director/Senior Librarian, Resource Discovery Management, National Library Board, Singapore. © 2025.

All authors contributed equally to this work.

Submitted: 25 November 2024. Accepted for Publication: 19 February 2025. Published: 16 June 2025.

curated collection. With the large collection numbers, RD sought to explore whether technology such as generative AI would be viable to facilitate the cataloguing of individual sites harvested in the yearly domain crawl. RD wished to explore the generation of titles, abstracts, and subjects for each website, as these metadata fields were deemed the most crucial and used for display and discovery.

Background and Motivation

This paper addresses the critical need for an efficient and accurate method of generating metadata for web archive collections. We focus on the challenge of managing large-scale web archives, where manual metadata curation is no longer practical due to resource constraints and the sheer volume of data.

Problem Statement

This paper addresses two primary challenges:

1. **Efficiency:** How can we develop an automated system that can process and generate metadata for a large number of websites, significantly reducing the time and resources required?
2. **Accuracy:** How can we ensure that the automatically generated metadata maintains a decent level of quality, accurately representing the content of the archived web pages?

These challenges are nontrivial for collections such as WAS, which aims to preserve the digital heritage of Singaporean life, culture, and history in the 21st century. With an estimated 20,000 Web ARChive (WARC) files created annually that comply with quality control standards, there is a critical need for a scalable, reliable, and cost-effective solution for metadata generation.

RELATED WORK

WARC files are a standardized format crucial for web crawling, archiving, and digital preservation.⁶ They play a significant role in historical web content preservation and the development of large language models (LLMs).⁷ For instance, Common Crawl, a major source of training data for LLMs, provides monthly crawls of billions of web pages in WARC format, significantly contributing to models such as GPT-3.⁸ Additionally, WARC files ensure that historical web data remains accessible and support research into web trends and digital culture by offering detailed snapshots of web pages. This capability is essential for tracking changes and innovations, while the format's design ensures effective accessibility and interoperability for both research and archiving.⁹

Metadata Strategies for Web Archives

This work focuses on the provision of descriptive metadata to support search and access of individual sites in a web archive. There can be varying metadata approaches, such as:

- harvesting the metadata as-is;
- reviewing only basic metadata fields such as Title and Language, or
- providing fuller metadata for specially curated web collections, including assigning subject headings.¹⁰

In many cases, a combination of approaches is used. For institutions conducting large-scale automated crawling (e.g., by domain), this would be:

- for websites that have been curated into themes, metadata is individually reviewed and enhanced;
- for the remainder of the full crawl, metadata is used as-is;
- for both, full-text search is available.

In addition, there are also available metadata guidelines for web archiving such as those developed by OCLC Research's Web Archiving Metadata Working Group.¹¹ In NLB, it was deemed viable for each site (but not subpages) in the domain crawl to be catalogued individually. This is implemented via a local application profile based on Dublin Core. Many fields are populated at the collection level, but cataloguers individually review the following fields, which are used for searching and browsing (alongside full-text search):

- Title
- Abstract
- Subject (Library of Congress Subject Headings and a local categorization vocabulary of 17 top level terms)

Cataloguing every single website in the domain crawl is a time-consuming task and given the explosive progress and availability of artificial intelligence/machine learning (AI/ML) tools in recent years, especially with large language models (LLMs), we were keen to explore integrating this technology into our workflow.

Libraries have naturally explored myriad use cases for AI/ML. In the case of cataloguing, there have been many attempts at using generative AI for cataloguing and metadata enhancement.¹² The IFLA WLIC 2023 programme featured an entire session on "Utopia, Threat or Opportunity First? Artificial Intelligence and Machine Learning for Cataloguing," with speakers from four different countries sharing on the use of AI/ML in cataloguing applications, such as generating catalogue records, detecting metadata from scanned resources, and deduplication.¹³

While results have been promising, most projects embody what was said by Allen in a Library of Congress blog post:

With all the possibilities, why aren't we already embedding machine learning into everything we do? The answer is complicated, but at its simplest, machine learning (ML) and artificial intelligence (AI) tools haven't demonstrated that they're able to meet our very high standards for responsible stewardship of information in most cases, without significant human intervention.¹⁴

Interest in AI/ML has naturally extended to the web archive community. This was embodied in the program for the 2024 IIPC Web Archiving Conference, which featured a dedicated session themed "AI & Machine Learning," in addition to having other lightning and drop-in talks that featured AI/ML use.¹⁵ We know of many projects utilizing AI/ML for web archives such as the Harvard Law School Library Innovation Lab's WARC-GPT, launched in February 2024, which is a custom chatbot for users to explore WARC collections via natural language queries.¹⁶ To date, however, we have not come across specific work on the application of AI/ML for cataloguing by directly processing WARC files.

Large Language Models

The advent of LLMs has revolutionized the field of natural language processing (NLP). These sophisticated models, capable of processing and generating human-like text, have surpassed traditional methods in numerous NLP tasks. What was once considered a complex challenge, such as text summarization, is now routinely managed with impressive accuracy and efficiency by LLMs. Leading the charge in the commercialization of large language models are models such as GPT-4o by OpenAI, Claude by Anthropic, and Gemini by Google.¹⁷ Built upon the groundbreaking transformer architecture introduced in "Attention Is All You Need," these models leverage the attention mechanism to outperform earlier methods for NLP tasks that relied heavily on linguistic techniques like stemming and lemmatization or other neural networks like recurrent neural networks (RNN) and gated recurrent units (GRU).¹⁸ The integration of reinforcement learning from human feedback¹⁹ and direct preference optimization²⁰ further enhances the quality of LLM-generated output, making them more human-like and thus suitable for generation of high-quality content. Fine-tuning LLMs with datasets from a particular domain not only boosts their contextual understanding and expertise in domain specific tasks but also enhances the accuracy of their output.²¹

In-Context Learning

However, given the limitations of time and computational resources, the process of fine-tuning language models for specific downstream tasks often incurs significant costs. As a result of this challenge, in-context learning (ICL) has gained significant popularity. The ICL technique enables LLMs to acquire new tasks without modifying the model parameters, by relying on task-specific examples provided in the input context. ICL has demonstrated strong performance in few-shot and zero-shot learning. Zero-shot learning leverages the model's existing knowledge to deduce task requirements. As illustrated by Brown et al., GPT-3 was utilized to showcase the system's versatility in the absence of task-specific fine-tuning. With few-shot learning, the model is provided with a restricted number of examples to showcase the task, resulting in enhanced comprehension and performance. Radford et al. have demonstrated that few-shot prompting can significantly enhance the model's accuracy across a diverse range of tasks.²²

Multi-query reasoning techniques and single-query reasoning techniques like chain of thought (CoT) prompting are examples of advanced ICL techniques. CoT prompting directs the model to produce intermediate reasoning steps that resemble how people solve problems. This approach is efficient and effective because it guarantees acceptable outcomes in addition to lowering the token count. The use of CoT significantly improves performance on arithmetic and logic-based activities.²³ Limited in number of instances, prompting is a supplementary method that involves providing task-specific examples to assist in generating answers.

Multi-query reasoning techniques, like graph of thought (GoT), tree of thought (ToT), and least-to-most, use multiple LLM queries to extract various tenable reasoning paths.²⁴ By breaking down complicated issues into smaller, more manageable subproblems, these techniques improve the reasoning power of the model. In conclusion, in-context learning strategies are essential for fully utilizing LLMs. The rationale behind us selecting CoT is its capacity to minimize token consumption while yielding superior outcomes.

METHODOLOGY

Data Collection and Preparation

To obtain web data for our inquiry, we collected a total of 112 WARC (Web ARChive) files from the Web Archives of Singapore. The HTML content from these files was extracted using the WARCIO and fastwarc python libraries.²⁵ BeautifulSoup facilitated the extraction of relevant metadata, such as titles and primary text content, from the HTML. We removed unnecessary tags and scripts during this procedure to ensure that the main content is highlighted. We standardized the URLs and conducted a quality assurance assessment to eliminate any substandard or irrelevant data, ensuring uniformity. This involved identifying common indicators of nonfunctional material, such as “404” errors or placeholder text like “lorem ipsum.” In addition, we implemented a deduplication technique to consolidate individual records obtained from duplicate URLs. This ensured the preservation of the information’s originality and relevance.

We employed multithreaded processing using a ThreadPoolExecutor to efficiently handle the vast amount of data, resulting in optimized resource utilization and significantly reduced processing time. The final step was combining the processed data into DataFrames, which were then ready for further analysis and model development. An indispensable requirement for the reliability and robustness of our investigation was a high caliber dataset ensured by this approach.

Heuristics for Data Reduction

To reduce the number of input tokens and thereby achieve cost savings, our methodology involved developing and evaluating various heuristic methods for efficient content extraction and summarization. The three heuristics were:

1. About Page Priority: We prioritized extracting content from the “About” page. If no “About” page existed, we used content from the shortest URL.
2. Shortest URL: We extracted content exclusively from the web page with the shortest URL.
3. Shortest URL with Regex Filtering: We extracted content from the shortest URL and applied regular expression (regex) filters to reduce the token count, thereby optimizing the input for our model.

These heuristics were inspired by observations of professional cataloguers, who typically require only a few pages to make accurate judgments about content categorization. This step reduces the token count across all heuristic methods. By adopting this methodology, we systematically compared various strategies for content extraction and summarization, striking a balance between computational efficiency and accuracy in content representation.

Prompt Engineering for Title and Abstract Generation

After some initial testing with in-context learning in LLMs, we settled on two meticulously designed prompts with contrasting characteristics for generating titles and abstracts to enhance metadata accuracy and website classification. These CoT prompts guide the generation process by providing clear instructions for cataloguing and summarizing website content. The output was then subjected to both automated and manual evaluation processes. This prompt, which builds upon the previous one, provides specific rules for the summarization of a variety of websites, such as corporate websites, personal blogs, and property listings. It guarantees that the summary is appropriate for the website’s type, thereby enhancing the relevance and precision of the abstract that is generated.

Prompt 1:

You are a diligent cataloguer working to create metadata for websites. Let's think step by step to ensure accurate and comprehensive metadata creation:

- 1) Determine the title of the organization or company on the main web page, ensuring it reflects the primary focus or name without additional descriptors. It should match the root domain of the web page.*
- 2) Create an abstract: Summarize the main content of the website in a brief and informative abstract.*
- 3) Format the result: Return the result in JSON format as {'title': [inferred_title], 'abstract': [created_abstract]}.*

This prompt initiates the process by identifying the main title and summarizing the website's content. The emphasis is on creating a concise and informative abstract and ensuring the title aligns with the website's root domain. The result is formatted in JSON to facilitate structured data use.

Prompt 2:

Use the previous prompt and add:

Summarize the content of the website following these rules:

- For company websites: This is the website of (company's name) which offers (services). The website contains information of (contact, operating hours, location, its services, customers' testimonials).*
- For websites selling properties: (Name of project) is a private residential development by (name of company). The project is located at xxx. This website contains information on (the condominium, location, floor plans, developer, and contact details).*
- For personal websites/blogs: This is a website of (Name of person), (role). This website contains information on (work experience, profile, education, research works, projects, publications, professional development, skills, portfolio).*
- For others, create a summary.*

This prompt, which builds upon the previous one, provides specific guidelines for the compilation of cataloguing rules for the summarization of a variety of websites, such as corporate websites, personal blogs, and property listings. It guarantees that the summary is appropriate for the website's type, thereby enhancing the relevance and precision of the abstract that is generated.

These prompts are intended to improve metadata generation by directing the LLM to perform comprehensive, rule-based actions in a methodical manner. In the subsequent segment, we also demonstrate the contrast between the rule-based prompt result and the rule-free outcome. This ensures the development of abstracts and titles that are contextually pertinent, concise, and precise, which are critical for enhancing the searchability and classification of websites.

Evaluation Methods*Automated Evaluation*

We employed an aggregation of two metrics to assess the quality of approaches used.

- Levenshtein distance for title comparison²⁶
- BERTScore (using bert-based-cased to produce the embedding) for title and content similarity²⁷

Our ranking algorithm evaluates heuristic variants based on their ability to generate abstracts and titles through an aggregated scoring system. This method integrates three primary criteria: the minimal median of Levenshtein distance, the maximum median of BERTScore, and the minimum standard deviation. The semantic similarity between reference and produced texts is assessed by BERTScore through contextual embeddings from BERT; a larger median indicates increased quality. The exact match accuracy is measured by the Levenshtein distance; a lower median indicates a closer alignment with the reference. The standard deviation of BERTScores is a metric that indicates consistency; values that are lower indicate more consistent performance. A total score is generated by combining the rankings of heuristics for each criterion. This comprehensive method ensures a fair evaluation of accuracy and consistency, thereby influencing our selection of heuristics for text production. The 2D and 3D visualizations of BERTScore of our generated data are demonstrated in figure 1 and figure 3. In the 2D BERT embedding vector space, it is evident that certain abstract pairs and title pairs are identical. The majority of the distances between abstract pairs is negligible.

Figure 1. Visualization of BERTScore embeddings for titles and abstracts on a 2D plot.




Figure 2. Visualization of BERTScore embeddings for abstracts on a 3D plot.




Figure 3. Bar chart of ranked aggregated scores for the six different combinations of prompts and heuristics. The scores are based on criteria including maximum BERTScore median, minimum Levenshtein median, and minimum standard deviation.

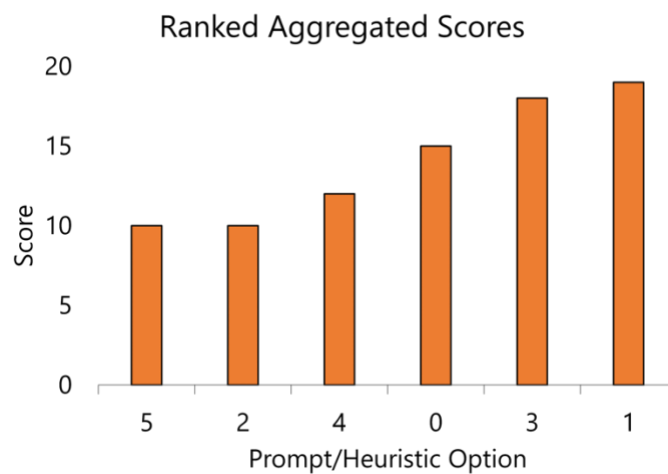


Table 1. Comparison of ranked aggregated scores for different prompt and heuristic combinations. The highlighted rows, heuristic 2 with rules and heuristic 3 without rules, achieved the best overall performance, each with a top-ranked score of 1. These combinations were selected for manual evaluation based on their superior results in automated scoring metrics.

Prompt	Heuristic	Combination	Ranked aggregated scores
With rules	1	0	4
Without rules	1	1	6
With rules	2	2	1
Without rules	2	3	5
With rules	3	4	3
Without rules	3	5	1

The automated evaluation showed that the following approaches had the best scores:

- Heuristic 2 with 2nd prompt with rules
- Heuristic 3 with 2nd prompt without rules

Manual Evaluation

Based on the automated evaluation results, we shortlisted the two approaches for manual evaluation. These titles and abstracts, in combination with a set generated by a human, were then subjected to manual grading by a team of eight trained cataloguers. The results were evaluated using the statistical tests. Cochran's Q test and McNemar's test are statistical tools that are employed to assess categorical data, particularly in situations involving related samples or repeated measurements.²⁸ These tests are employed to compare the performance of human evaluation and various heuristics in our research.

McNemar's test is often used to analyze nominal data that is paired, particularly in before-and-after studies or when comparing two related samples. It is particularly beneficial for binary results in matched pairs of subjects. In this study, we employ McNemar's test to evaluate combination 2 with human assessment. This pairwise comparison enables us to determine whether the automated heuristic and human judgment exhibit statistically significant differences. Cochran's Q test is an expanded version of McNemar's test that is applicable to more than two related groups. It is implemented when three or more conditions that are related generate binary outcomes, such as "pass" or "fail." The test assesses whether the proportions of these binary outcomes vary statistically significantly across the conditions. We apply Cochran's Q test to determine whether the pass/fail rates of the three evaluation approaches (combination 2, combination 3, and human) display any significant variations. This test elucidates whether the approaches are consistently producing distinct results or if the variations could be the result of random chance.

In our dataset, there are numerous evaluations of the same set of objects, including both human and heuristic evaluations. Both tests are qualified by our pass/failure data. The purpose of these tests is to conduct a precise statistical comparison of the performance of various evaluation techniques. McNemar's test enables us to focus on the comparison between two combinations and

human evaluation, which may be helpful in assessing the automated approach’s performance against human judgment.

Software and Hardware Setup for Metadata Generation and Evaluation

To verify if heuristics can effectively reduce the number of tokens needed for metadata generation, we wrote a python package warc2summary which provides a pipeline for metadata creation and evaluation. In short, this package uses WARCIO and fastwarc as the WARC file processor, and for title and abstract generation we use OpenAI API (GPT-4o) as the main model and instructor to constrain the output.²⁹

We ran the software on a Lenovo ThinkPad T14 with an 11th Gen Intel® Core™ i5-1145G7 CPU and 16GB of DDR4 RAM. The pipeline took close to two hours to process the WARC files and another two hours to generate the synthetic data from API calls. We expect faster performance in WARC file processing with a better CPU. In addition, the API call pipelines were called sequentially in order to not breach OpenAI API rate limits. With higher rate limits, it is possible to parallelise the calls for better performance. The WARC files are warc.gz files and are taken from Web Archive Singapore.

This provided us with a DataFrame containing synthetic titles and abstracts.


We identified six different combinations of prompts and heuristics as stated above. We then used a ranked aggregation scoring method, incorporating Levenshtein distance for the title and BERTScore (using bert-based-cased embeddings), with criteria of maximum BERTScore median, minimum Levenshtein median, and minimum standard deviation for both, to shortlist two specific combinations of prompts and heuristics. Finally, a team of eight trained cataloguers performed manual grading.

RESULTS

Data Reduction

From the DataFrame, each heuristic finished within 15 seconds on 112 files. The heuristics were consolidated. Based on our heuristics, the final DataFrame should have an identical number of rows as the number of WARC files. Based on the collection of 112 WARC files we have, we obtained a 99.9% reduction in total token count, resulting in significant cost savings compared to letting OpenAI parse the entirety of the WARC files.

Figure 4. Visualizing cost differences before and after applying heuristics to count tokens with tiktoken, a Byte-Pair Encoding tokenizer for GPT-4o.³⁰



Statistical Test

We performed Cochran’s Q test by having trained cataloguers rate the titles and abstracts, without knowledge of their provenance. Based on a 5% confidence level, we found that the synthetic titles and abstracts are statistically distinguishable from human-generated titles and abstracts. Since the p-value (0.02) is less than our significance level ($\alpha = 0.05$), we reject the null hypothesis, indicating a significant difference between the GPT-4o generated metadata and human crafted metadata.

Figure 5. Results of Cochran’s Q test.


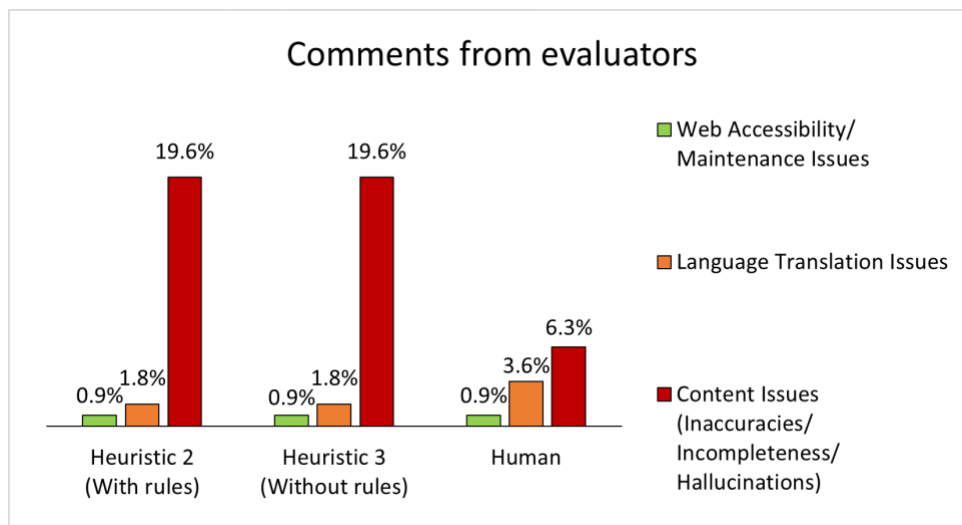


Figure 6. Issues faced from evaluator’s comments.



DISCUSSION

Implications of Results

Our study demonstrates both the potential and limitations of using large language models like GPT-4o for automated metadata generation in web archives:

1. **Scalability and efficiency:** Our approach efficiently processes large volumes of web archive data, addressing a critical challenge in digital preservation. The method achieves a 99.9% reduction in associated costs compared to full WARC processing.
2. **Cost-effectiveness:** By reducing the need for manual cataloguing, our approach significantly lowers the costs associated with metadata creation, allowing institutions to reallocate resources to other critical tasks.
3. **Quality comparison:** Our adaptation of the Turing test, involving eight cataloguers evaluating metadata from different sources, provided valuable insights into the current capabilities of LLMs, specifically GPT-4o in this domain. The synthetic titles and abstracts were statistically distinguishable from human curated metadata, with a p-value of 0.02, suggesting that human generated metadata still maintains an edge in quality. Furthermore, there is no significant difference between LLM based approaches, with rules and without rules.

Limitations and Challenges

Our study revealed several important limitations and challenges, both from our experimental results and broader considerations in the field.

1. **Content accuracy and hallucinations:** 19.6% of LLM generated titles and abstracts had content issues such as inaccuracies, incompleteness, or hallucinations, higher than the 6.3% observed in human-curated metadata. This highlights the known challenge of LLM hallucinations, where generated content can be factually incorrect or not grounded in the input data. Mitigating these hallucinations remains an active area of research.
2. **Language translation:** Both LLM and human generated metadata faced challenges with language translation, highlighting the complexity of handling multilingual content in web archives.
3. **Legal considerations:** The generation and use of metadata from copyrighted web content raises complex legal questions. Navigating fair use and copyright in the context of web archiving and LLM generated metadata requires careful consideration to avoid potential legal complications.³¹
4. **Privacy and dependency concerns:** The automated nature of LLMs for metadata generation poses challenges in honoring individual requests for content removal or anonymization, which are crucial aspects of privacy rights. Additionally, our reliance on proprietary, closed-source LLMs like GPT-4o introduces dependencies that may limit flexibility and control. LLMs can potentially leak training data, raising concerns about privacy and data protection.³² The potential for these models to use web content for further training underscores the need for transparency and alternatives.
5. **Quality of input data:** The adage “garbage in, garbage out” is especially relevant in our context.³³ The quality of LLM generated metadata is directly dependent on the quality of the input data from web archives. Ensuring that our web archive data is accurate, diverse, well labelled, and free from noise is crucial for generating reliable and unbiased metadata.

6. **Evaluation criteria:** Ensuring that LLM generated content accurately represents the website, remains relevant to the topic, and avoids subjective interpretations or biased language remains a significant challenge, as evidenced by our experimental results.

CONCLUSIONS

This study introduces a novel approach to generating metadata for web archives using GPT-4o, combining data reduction heuristics with evaluation methods inspired by the Turing test. Our findings highlight the efficiency and scalability of LLM generated metadata, though it still falls short in quality and accuracy compared to human-curated metadata. GPT-4o produced more inaccuracies and hallucinations but is cost effective and scalable for WARC file metadata creation. This suggests LLMs are best as assistive tools for human cataloguers rather than replacements. Future efforts could focus on changing the prompts to fit the cataloguers' standards, more aggressive heuristics to filter out promotional website content, developing strategies to reduce and identify hallucinations and using smaller language models to circumvent privacy concerns.

Our study marks a significant step towards leveraging AI in web archiving, offering valuable insights into its current capabilities and limitations. Addressing the identified challenges will help us work towards a future where AI enhances the preservation and accessibility of digital heritage, maintaining high standards crucial for the utility of web archives. This research sets a roadmap for future improvements, aiming to bridge the quality gap between AI and human curated metadata.

ACKNOWLEDGEMENTS

We would like to extend our heartfelt gratitude to the following cataloguers from the National Library Board for their invaluable contributions to this project:

- Ann Cheah
- Chang Siew Fen
- Chor Swee Chin
- Munifah Shaik Mohsin Bamadhaj
- Ng Wee Lay
- Rohaya Bte Yacob
- Tan Minli Mindy

ENDNOTES

- ¹ National Library Board, accessed August 10, 2024, <https://www.nlb.gov.sg/main/home>.
- ² "Frequently Asked Questions," National Library Board, Web Archive Singapore, accessed August 22, 2024, <https://eresources.nlb.gov.sg/webarchives/faq>.
- ³ DCMI-Libraries Working Group, *DC-Libraries - Library Application Profile - Draft* (technical report), Dublin Core Metadata Initiative, September 2004, accessed August 10, 2024, <https://dublincore.org/specifications/dublin-core/library-application-profile/>.
- ⁴ National Library Board Act 1995, *Singapore Statutes Online*, accessed August 11, 2024, <https://sso.agc.gov.sg/Act/NLBA1995>.

-
- ⁵ Shereen Tay, “An Archive of Singapore Websites: Preserving the Digital,” *BiblioAsia* 16 no. 3 (October–December 2020), <https://biblioasia.nlb.gov.sg/vol-16/issue-3/oct-dec-2020/website/>.
- ⁶ “WARC (Web ARChive) File Format,” Library of Congress, <https://www.loc.gov/preservation/digital/formats/fdd/fdd000236.shtml>.
- ⁷ Common Crawl, <https://commoncrawl.org/>.
- ⁸ Tom Brown et al., “Language Models Are Few-Shot Learners,” in *Advances in Neural Information Processing Systems* 33, ed. H. Larochelle et al. (Neural Information Processing Systems Foundation, Inc., 2020), 1877–1901, https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf.
- ⁹ Emily Maemura, “All WARC and No Playback: The Materialities of Data-Centered Web Archives Research.” *Big Data & Society* 10, no. 1 (2023), <https://journals.sagepub.com/doi/10.1177/20539517231163172>.
- ¹⁰ A. K. Wong and D. K. W. Chiu, “Digital Curation Practices on Web and Social Media Archiving in Libraries and Archives,” *Journal of Librarianship and Information Science* (2024), <https://doi.org/10.1177/09610006241252661>.
- ¹¹ Jackie Dooley and Kate Bowers, “Descriptive Metadata for Web Archiving: Recommendations of the OCLC Research Library Partnership Web Archiving Metadata Working Group,” *OCLC Research* (February 17, 2018), <https://www.oclc.org/research/publications/2018/oclcresearch-descriptive-metadata/recommendations.html>.
- ¹² R. Brzustowicz, “From ChatGPT to CatGPT: The Implications of Artificial Intelligence on Library Cataloguing,” *Information Technology and Libraries* 42, no. 3 (2023), <https://doi.org/10.5860/ital.v42i3.16295>; E. H. C. Chow, T.G. Kao, and X. Li, “An Experiment with the Use of ChatGPT for LCSH Subject Assignment on Electronic Theses and Dissertations,” accessed August 23, 2024, <http://arxiv.org/abs/2403.16424>.
- ¹³ IFLA, “Session 157: Utopia, Threat or Opportunity First? Artificial Intelligence and Machine Learning for Cataloguing”, 2023, Full Programme, 88th IFLA General Conference and Assembly, accessed August 22, 2024, <https://iflawlic2023.abstractserver.com/program/#/details/sessions/276>.
- ¹⁴ Laurie Allen, “Why Experiment: Machine Learning at the Library of Congress,” *The Signal: Digital Happenings at the Library of Congress*, Library of Congress Blogs, November 13, 2023, <https://blogs.loc.gov/thesignal/2023/11/why-experiment-machine-learning-at-the-library-of-congress/>.
- ¹⁵ IIPC, “WAC 2024 Program”, 2024, accessed August 22, 2024, <https://netpreserve.org/ga2024/programme/wac/>.
- ¹⁶ M. Cargnelutti, K. Mukk, and C. Stanton, “WARC-GPT: An Open-Source Tool for Exploring Web Archives Using AI,” Harvard Law School Library Innovation Lab blog, 2024, accessed August

- 23, 2024, <https://lil.law.harvard.edu/blog/2024/02/12/warc-gpt-an-open-source-tool-for-exploring-web-archives-with-ai/>.
- ¹⁷ OpenAI, *Hello GPT-4o*, May 13, 2024, <https://openai.com/index/hello-GPT-4o/>; Anthropic, *Claude 3.5 Sonnet*, June 21, 2024, <https://www.anthropic.com/news/claude-3-5-sonnet>; Sundar Pichai and Demis Hassabis, "Introducing Gemini: Our Largest and Most Capable AI Model," *Google Blog*, December 6, 2023, <https://blog.google/technology/ai/google-gemini-ai/#sundar-note>.
- ¹⁸ Ashish Vaswani et al., "Attention Is All You Need," in *Advances in Neural Information Processing Systems* 30, ed. I. Guyon et al. (2017), https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- ¹⁹ Long Ouyang et al., "Training Language Models to Follow Instructions with Human Feedback," in *Advances in Neural Information Processing Systems* 36, ed. S. Koyejo et al. (Curran Associates, Inc., 2022), 27730–44, https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf.
- ²⁰ Rafael Rafailov et al., "Direct Preference Optimization: Your Language Model Is Secretly a Reward Model," in *Advances in Neural Information Processing Systems*, vol. 36, ed. A. Oh et al. (Curran Associates, Inc., 2023), 53728–41, https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf.
- ²¹ Chaoyi Wu et al., "PMC-LLAMA: Towards Building Open-Source Language Models for Medicine," 2023, accessed August 10, 2024, <https://arxiv.org/abs/2304.14454>; Xiao-Yang Liu et al., "FinGPT: Democratizing Internet-Scale Data for Financial Large Language Models," 2023, accessed August 10, 2024, <https://arxiv.org/abs/2307.10485>.
- ²² Alec Radford et al., "Learning Transferable Visual Models from Natural Language Supervision," in *International Conference on Machine Learning* (PMLR, 2021), 8748–63.
- ²³ Jason Wei et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," in *Advances in Neural Information Processing Systems*, vol. 35, ed. S. Koyejo et al. (Curran Associates, Inc., 2022), 24824–37, https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf.
- ²⁴ Maciej Besta et al., "Graph of Thoughts: Solving Elaborate Problems with Large Language Models," *Proceedings of the AAAI Conference on Artificial Intelligence* 38, no. 16 (March 2024): 17682–90, <https://doi.org/10.1609/aaai.v38i16.29720>; Shunyu Yao et al., "Tree of Thoughts: Deliberate Problem Solving with Large Language Models," in *Advances in Neural Information Processing Systems*, vol. 36, ed. A. Oh et al. (Curran Associates, Inc., 2023), 11809–22, https://proceedings.neurips.cc/paper_files/paper/2023/file/271db9922b8d1f4dd7aaef84ed5ac703-Paper-Conference.pdf; Denny Zhou et al., "Least-to-Most Prompting Enables Complex

- Reasoning in Large Language Models,” 2023, accessed August 10, 2024, <https://arxiv.org/abs/2205.10625>.
- ²⁵ Ilya Kreymer, *warcio*, 2020, GitHub repository, <https://github.com/webrecorder/warcio>; Janek Bevendorff et al., “FastWARC: Optimizing Large-Scale Web Archive Analytics,” 2021, <https://arxiv.org/abs/2112.03103>.
- ²⁶ Vladimir I. Levenshtein, “Binary Codes Capable of Correcting Deletions, Insertions, and Reversals,” *Soviet Physics Doklady* 10 (1966): 707–710.
- ²⁷ Tianyi Zhang et al., “BERTScore: Evaluating Text Generation with BERT,” 2020, <https://arxiv.org/abs/1904.09675>.
- ²⁸ W. G. Cochran, “The Comparison of Percentages in Matched Samples,” *Biometrika* 37, no. 3/4 (1950): 256–66, <https://doi.org/10.2307/2332378>; Quinn McNemar, “Note on the Sampling Error of the Difference Between Correlated Proportions or Percentages,” *Psychometrika* 12, no. 2 (1947): 153–57, <https://doi.org/10.1007/BF02295996>.
- ²⁹ Jason Liu, *Instructor: Structured LLM Outputs*, 2024, GitHub repository, accessed August 13, 2024, <https://github.com/jxnl/instructor>.
- ³⁰ *tiktoken*, OpenAI, 2023, accessed August 10, 2024, <https://github.com/openai/tiktoken>.
- ³¹ Muhammad Zakir et al., “Navigating the Legal Labyrinth: Establishing Copyright Frameworks for AI-Generated Content,” *Remittances Review* 9 (January 2024): 2515–32, <https://remittancesreview.com/article-detail/?id=1467>
- ³² Nicholas Carlini et al., “Extracting Training Data from Large Language Models,” in *30th USENIX Security Symposium (USENIX Security 21)*, USENIX Association, August 2021, 2633–50. <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>.
- ³³ William D. Mellin, “Work with New Electronic ‘Brains’ Opens Field for Army Math Experts,” *Hammond Times* 10, no. 66 (1957); Charles Babbage, *Passages from the Life of a Philosopher*, (Longman, Green, Longman, Roberts, & Green, 1864).