

Connecting the Dots

A Semantic Web Solution for Enhanced Library Resource Discovery

Min Hoon Ee, Ashwin Nair, and Robin Dresel

ABSTRACT

The National Library Board (NLB) of Singapore has made significant strides in leveraging data to enhance public access to its extensive collection of physical and digital resources. This paper explores the development and implementation of the Singapore Infopedia Widget, a recommendation engine designed to guide users to related resources by utilizing metadata and a Linked Data Knowledge Graph. By consolidating diverse datasets from various source systems and employing semantic web technologies such as Resource Description Framework (RDF) and Schema.org, NLB has created a robust knowledge graph that enriches user experience and facilitates seamless exploration.

The widget, integrated into Infopedia, the Singapore Encyclopedia, surfaces data through a user-friendly interface, presenting relevant resources categorized by format. The paper details the architecture of the widget, the ranking algorithm used to prioritize resources, and the challenges faced in its development. Future directions include integrating user feedback, enhancing semantic analysis, and scaling the service to other web platforms within NLB's ecosystem. This initiative underscores NLB's commitment to fostering innovation, knowledge sharing, and the continuous improvement of public data access.

BACKGROUND

NLB's Open Data Initiative

The National Library Board (NLB) of Singapore advocates the use of open data as they help eliminate the barriers caused by proprietary formats and access standards. This allows for the broader utilization of NLB's metadata and catalogue records and expands on its continuing efforts to enhance public data access and foster innovation and knowledge sharing through collaborative efforts. Some of our application programming interfaces (APIs) and data services are available on NLB Labs (<https://www.nlb.gov.sg/main/partner-us/contribute-and-create-with-us/NLBLabs>), a platform that facilitates collaboration, allowing the public to utilize them for a common purpose and spotlight new services.¹

Linked Data in Libraries: Successes and Potential

Linked data, which involves assigning identifiers to items on the web and allowing them to be referenced and connected through statements in a standardized format known as Resource Description Framework (RDF), is viewed as a progressive method to consolidate datasets from various sources and across different domains, enabling them to be interconnected.² Institutions

About the Authors

Min Hoon Ee (ee_min_hoon@nlb.gov.sg; corresponding author) is Senior Librarian, National Library Board of Singapore. **Ashwin Nair** (ashwin_nair_madhavan@nlb.gov.sg) is Librarian, National Library Board of Singapore. **Robin Dresel** (robin_dresel@nlb.gov.sg) is Senior Librarian, National Library Board of Singapore. © 2026.

Submitted: 2 June 2025. Accepted for Publication: 22 October 2025. Published: 16 March 2026.

that have successfully embarked on this trajectory include the Library of Congress and the Online Computer Library Center (OCLC).

The Library of Congress's Bibliographic Framework (BIBFRAME) model seeks to transform bibliographic descriptions from literals to entities, thus allowing such information to be used by the library community and facilitating the broader use of curated library vocabularies.³ Under the Linked Open Initiative (developed by the World Wide Web Consortium), major search engines such as Google will be able to understand library data and make these authoritative records available to web users. OCLC's CONTENTdm Linked Data Pilot project sought to understand how the libraries, archives, and museums (LAM) community can transition and benefit from linked data.⁴ Five representative institutions participated in the pilot project: The Huntington Library, Art Museum, and Botanical Gardens; the Cleveland Public Library; the Minnesota Digital Library; Temple University Libraries; and the University of Miami Libraries. The focus was on "developing efficient workflows for transforming metadata, evaluating existing interfaces to leverage linked data, and testing applications built in the Wikibase environment for managing the newly created linked data." Through the pilot, it was apparent that to strongly enhance discovery and data management of cultural materials, having a "shared and extensible data model," transforming text headings into linked data entities with relationships and allowing these entities and relationships to be searched in a unified platform are key. Assigning a persistent uniform resource identifier (URI) to entities will aid in authority control and lead to improved precision and recall when it comes to search experiences. Information from external knowledge sources enriches entity information.

Implementation Challenges and Limitations

Linked data implementation is not without its challenges. The sheer volume of the metadata involved, in machine-readable cataloguing (MARC) or another format, makes the transition to BIBFRAME or other data models a complex endeavor.⁵

The OCLC Linked Data Pilot project has shown that for a project of such scale to succeed, there is a need to set up a community of experts who are equipped with the necessary tools and workflows to transform traditional metadata into linked data entities and are able to commit to making the transition a reality.

Having different types of resources compounds the issue further, as in the case of NLB. NLB manages an extensive collection of physical and digital resources, including those coming from the National Archives. The metadata of these resources come from different source systems and are presented in different formats. BIBFRAME was created with bibliographic resources in mind and does not extend well to cover other resource types.⁶

NLB's Linked Data Journey

Since 2013, in its effort to use semantic web technologies, NLB has worked with data management vendors to create bespoke data models for describing bibliographic resources and managing authority and archival data.

In 2021, NLB awarded a contract to a vendor consortium made up of Metaphacts GmbH (Germany), Data Liberate (UK), and Kewmann (Singapore) for delivering the most recent version of the Linked Data Management System (LDMS).⁷ Each member of this three-party consortium brings specialized capabilities: Metaphacts provides graph data technology expertise, Data Liberate contributes library data system knowledge, and Kewmann functions as the business integrator, offering local presence and implementation services.

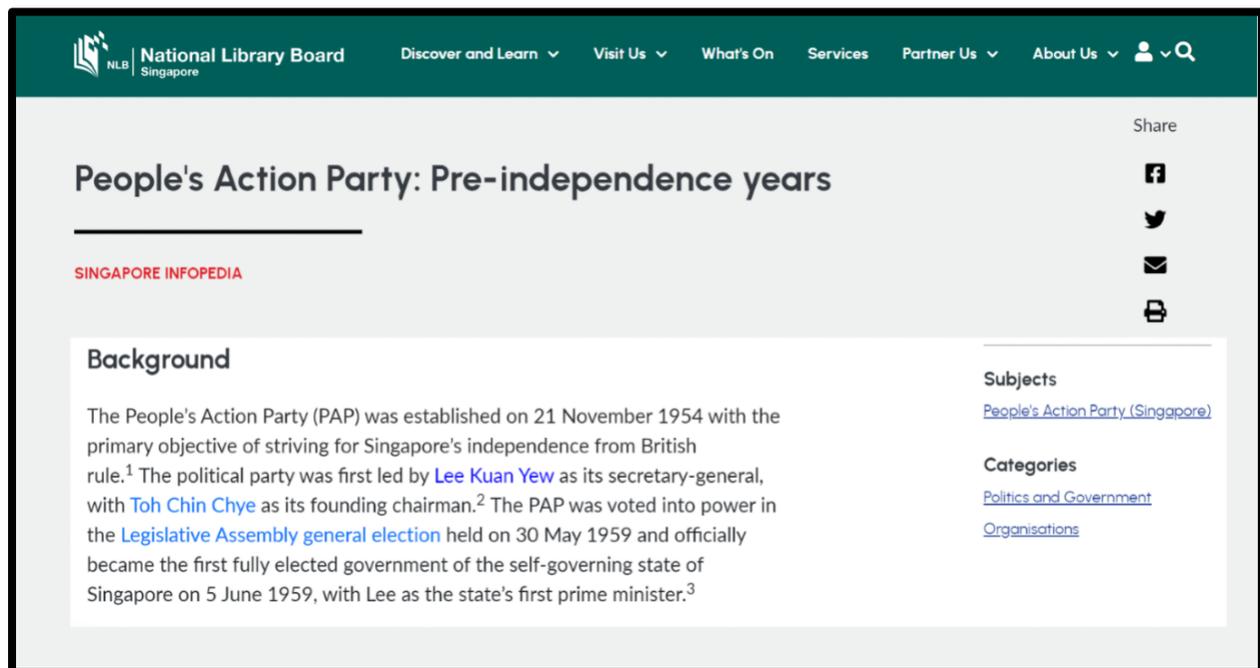
PROBLEM STATEMENT AND OBJECTIVES

Current Limitations

The Singapore Infopedia is an online encyclopedia featuring articles on the nation's history, culture, notable figures, and major events.⁸ It can be accessed through the National Library Online platform, providing detailed insights into diverse topics with references for further exploration.⁹

Figure 1 shows an example of an Infopedia article on the People's Action Party. Prior to the linked data implementation, there were no links to related content (physical or digital) for the user to explore further if they were interested in finding out more about entities mentioned in this article, for example, the late Minister Mentor Lee Kuan Yew.

Figure 1. Infopedia article on "People's Action Party."



Proposed Solution

To bridge this gap, NLB developed a service called the Singapore Infopedia Widget. This widget functions as a recommendation engine, guiding users to resources related to the article they are viewing. Utilizing metadata from each article, the widget suggests entities (Work, Person, Organization, Place, and Topic) based on connections within our Linked Data Knowledge Graph. This will aid students and researchers in their research efforts, as they are the primary users of Infopedia.

Objectives

The primary objectives of the Singapore Infopedia Widget are the following:

1. Enhance user discovery of related resources within NLB's extensive collection
2. Leverage the Linked Data Knowledge Graph for improved content recommendations
3. Provide a seamless and intuitive user experience for exploring interconnected information
4. Support researchers and students in their information-seeking processes

IMPLEMENTATION OF THE SINGAPORE INFOPEDIA WIDGET

Data Consolidation and Standardization

The LDMS serves to consolidate resources, including physical and digital materials, from the various systems under the National Library, the National Archives, and the Public Libraries. As data from these systems are presented in different formats, conversion is required to create a cohesive graph centered around entities.

These source systems, as indicated by Dresel, include the following:

- The Integrated Library System (ILS) that contains the MARC records for bibliographic content
- The Content Management System (CMS) that aggregates digital objects from various source systems (including those from the National Archives of Singapore in the General International Standard Archival Description [ISAD-G] archival standard) and their metadata in the Dublin Core format
- The Taxonomy and Thesaurus Editor (TTE), which serves as a vocabulary management system for local authority records that include descriptions and identifiers for People, Places, Organizations, and so on, in CSV format¹⁰

To standardize the different vocabularies, Schema.org was selected as the common framework for mapping all datasets. This ensures that once the data are made accessible on the web, they are presented in a format that search engines can readily understand.

MARC data coming from the ILS are converted into BIBFRAME using the open-source program provided by the Library of Congress.¹¹ BIBFRAME was selected due to its support for Resource Description and Access (RDA) and RDF, both of which are standards adhered to by NLB.¹² Also, BIBFRAME acknowledges “entities, attributes, and relationships between entities.”¹³

Records in the CMS are captured first in DC XML format before conversion into DC RDF entities for enrichment with Schema.org structured data. Data in TTE are mapped directly to Schema.org.

NLB's Knowledge Graph

NLB has consistently been creating and maintaining “user-centric, Singapore-centric and Southeast Asian-centric name authorities.”¹⁴ The authority records include standardized terms for entities in Singapore and Southeast Asia, highlighting local contextual information such as dialect names for ethnic Chinese. Moreover, entries for entities not encompassed in international authorities, such as the Library of Congress Name Authorities (LCNA), are incorporated to meet the requirements of digital cataloguing.¹⁵

The comprehensive description of entities enhances the interconnections within the Knowledge Graph, which facilitates more diverse pathways for discovery.

Figure 2 shows the authority record of Lee Kuan Yew. It includes links to other Person entities such as his parents, spouse, and children, as well as Works of which he is author, contributor, or subject.

Figure 2. Authority record of Lee Kuan Yew.

Lee, Kuan Yew, 1923-2015
 Person | Prime ministers

The first Prime Minister of Singapore (1959-1990), he is recognized as the nation's founding father. He was Senior Minister (1990-2004); Minister Mentor (2004-2011); Member of Parliament for Tanjong Pagar (1955-1991) and Tanjong Pagar GRC (1991-2015). He co-founded the People's Action Party (PAP) in 1954 and was its Secretary-General between 1954 and 1992. He passed away at 3.18 am on 23 Mar 2015. National mourning from 23 Mar to 29 Mar 2015 was declared, culminating in a State Funeral Service on 29 Mar 2015.

URI: <https://eresources.nlb.gov.sg/linkeddata/primary-entity/person/ffc65509-37df-49be-8dbd-51e40c6db488>

Also known as: Li Guangyao, LKY, 李光耀, Harry Lee Kuan Yew, Li, Guangyao, 1923-2015, Li, Kuan Yu, 1923-2015, Yi, Kwan Yu, 1923-2015, Yew, Lee Kuan, 1923-2015, Li Kuanŷg hŷ, 1923-2015, Lee, Harry, 1923-2015, 李光耀, 1923-2015, 1923-2015, 李光耀, Li, Kuan-yao, 1923-2015, Yi, Li Kuan, 1923-2015, Li, Quang Diu, 1923-2015, 李光耀, 1923-2015

Birth place: Singapore

Birth date: 1923-09-16

Death date: 2015-03-23

Spouse: Kwa, Geok Choo, 1920-2010

Children: Lee, Hsien Loong, Lee, Hsien Yang, Lee, Wei Ling

Parent: Lee, Chin Koon, Lee, Chin Koon, Mrs.

Related: Lee, Freddy Thiam Yew, Lee, Hoon Leong, Lee, Suet Fern

Nationality: Singaporean

Awards: Order of the Companions of Honour (1970), Bintang Republik Indonesia Adi Pradana, First Class Order of the Rising Sun, Japan (1967), The Most Esteemed Family Order, DK, Brunei [The Darja Kerabat Laila Utama], Freedom of the City of London, G.C.M.G. (1972), The Most Honourable Order of the Crown of Johore, First Class, Grand Cordon of the Order of the Paulownia Flowers, Japan (2015)

Alumni: Fitzwilliam College, Raffles College, Choon Guan School, Telok Kurau English School, Raffles Institution (Singapore)

Affiliation: Malayan Forum (Great Britain), People's Action Party (Singapore)

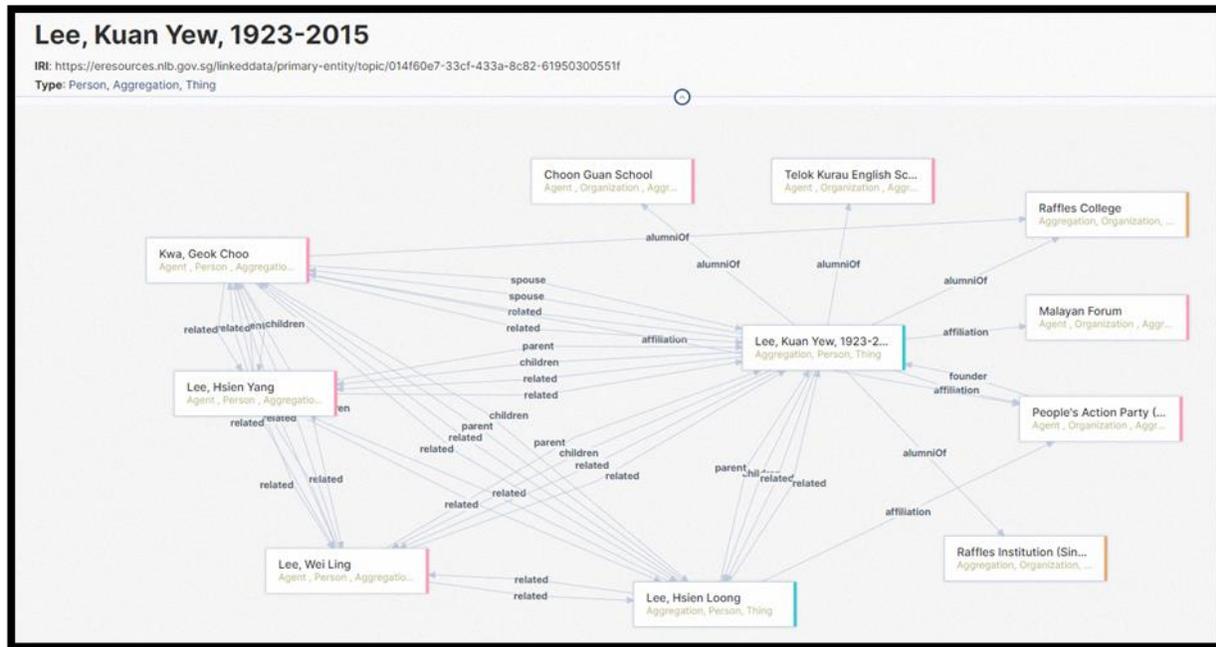
Works by Lee, Kuan Yew, 1923-2015: 114 items

Contributions by Lee, Kuan Yew, 1923-2015: 1257 items

Works About Lee, Kuan Yew, 1923-2015: 444 items

Figure 3 is a graph representation of Person and Organization entities that are linked to him. The Singapore Infopedia Widget aims to make visible the relationships among these entities by capitalizing on the available metadata. Going back to Figure 1, the article on the People’s Action Party references Lee Kuan Yew in its content. If the user wants to find out more about him, they may do so by selecting his name in the widget. Doing so will bring up the authority record of Lee Kuan Yew, detailing vital information about him. From there, the user is also presented with relevant resources about him in multiple formats (text, audio, and video). The user will be able to traverse not just among related entities within the article but also discover other related content that are part of the wider NLB network of resources. These include physical items on our Online Patron Access Catalog (OPAC) and digital content on Archives Online and National Library Online, thereby enriching the user’s discovery experience.

Figure 3. Person entities related to Lee Kuan Yew.



To identify named entities for each article in the Infopedia collection, we used Named Entity Recognition (NER). This allowed us to extract named entities in context, effectively utilizing unstructured data to establish connections to additional resources within the extensive NLB collection.¹⁶

NLB has also leveraged Generative Artificial Intelligence (GenAI), particularly through a large language model (LLM), to incorporate mentioned entities into the metadata of the articles.¹⁷ This aims to enhance the number of relationships within the Knowledge Graph and consequently expand related resources. Specifically, we have utilized Generative Pre-trained Transformer 4 (GPT-4), known for its robustness and superior performance compared to OpenAI’s previous LLMs.¹⁸

Widget Architecture

To achieve these goals, the idea of a widget was developed to surface the data and integrate them into existing web services to aid customers in finding items related to the resources they were currently browsing. A decision was made to trial this concept with NLB’s Singapore-centric encyclopedia, with the option to scale the widget to other services.

Generic API for Adaptability

We decided to develop the service to be generic to allow for potential scaling. The design idea was to fetch all available data from the Knowledge Graph and let the user interface (UI) decide on which fields and values to display. This approach would achieve two objectives:

- 5. Allowing the service to be integrated into another web service, with unique needs, thus requiring different data to be displayed
- 6. Being flexible enough to allow for customization of the UI to adapt to the respective web design of the service it would be integrated into

Context Switch

The widget allows users to select from the subjects assigned to the Infopedia articles during the cataloguing process. These subjects will determine the widget’s behavior, the content displayed for the Knowledge Card (KC), and the related resources.

To add further value and display additional resources related to the respective articles, a separate NER process was used to identify Singapore-related People, Places, and Organizations mentioned in the articles. The information was stored in the Knowledge Graph alongside the subjects to be used to populate the context switch.

Knowledge Card

The KC is the location of the entity description. Depending on the entity (Person, Place, Organization, Topic), different values will be displayed, as in the following examples:

- A person will have their birth date and occupation displayed (see Figure 4)
- An organization will have its founding date (see Figure 5)
- A place will show its address
- A topic will have its broader term shown

Figure 4. Anatomy of the widget.

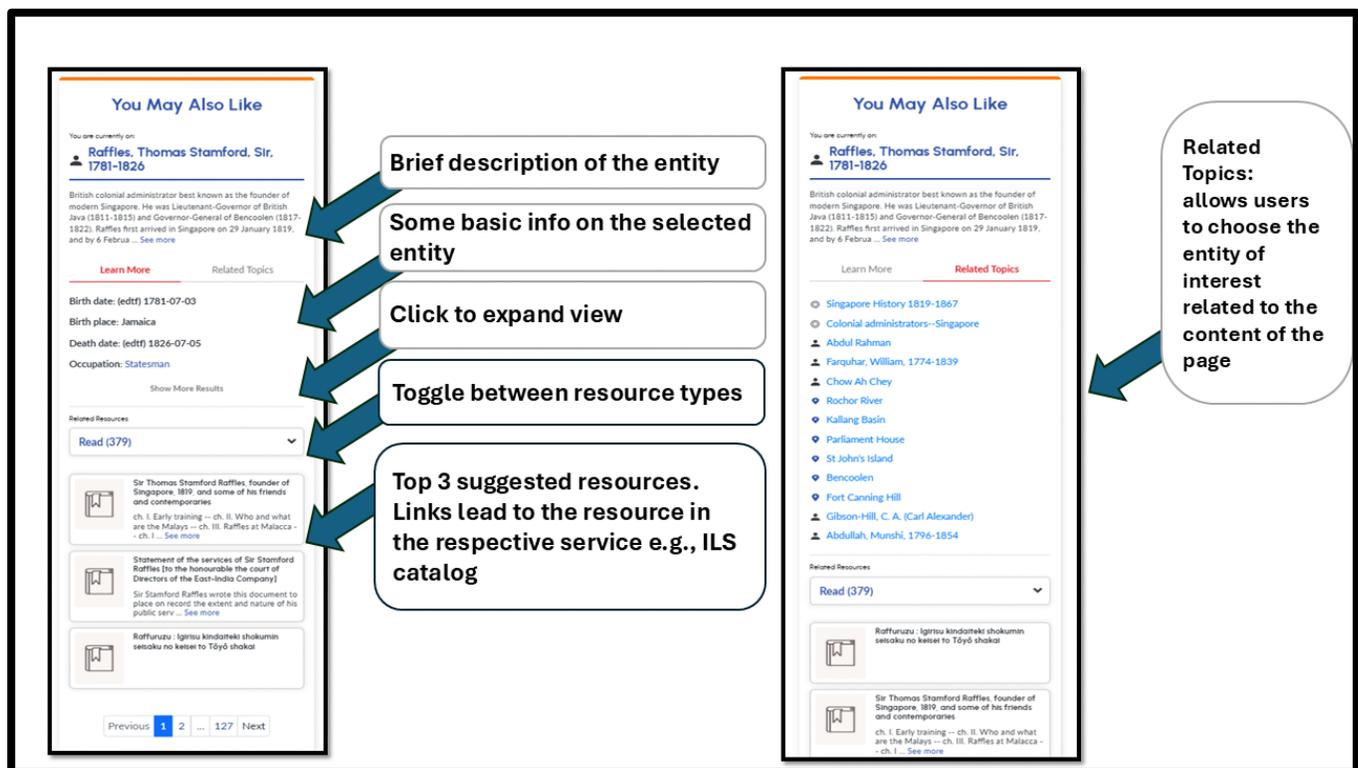


Figure 5. Knowledge Card of the National Library Board of Singapore.

Singapore. National Library Board

Organisation

Statutory board established to oversee the National Library, the National Archives of Singapore, a network of public libraries, and libraries belonging to government agencies, schools and institutions in Singapore. The aim of the Board is to promote reading and learning, as well as to improve information literacy among Singaporeans.

URI: <https://eresources.nlb.gov.sg/linkeddata/primary-entity/organization/68bec5ce-9ab6-4760-9ce6-6a970dc3f495>

Also known as: NLB
National Library Board, Singapore
Xinjiapo guo jia tu shu guan
新加坡國家圖書館
Perpustakaan Negara Singapura
Singapore. Lembaga Perpustakaan Negara

Parent Organisation: Ministry of Digital Development and Information

Sub Organisation: Lee Kong Chian Reference Library
National Archives (Singapore)

Founded: 1995-09-01

Former Name: National Library (Singapore)

Street Address: 100 Victoria Street, #14-01

Postal Code: 188064

Country: Singapore

External Links 2 results

- IAIF - Virtual International Authority File
- Wikidata

Resource Type Choice

The resources linked to the respective topics are grouped into four categories based on their format type:

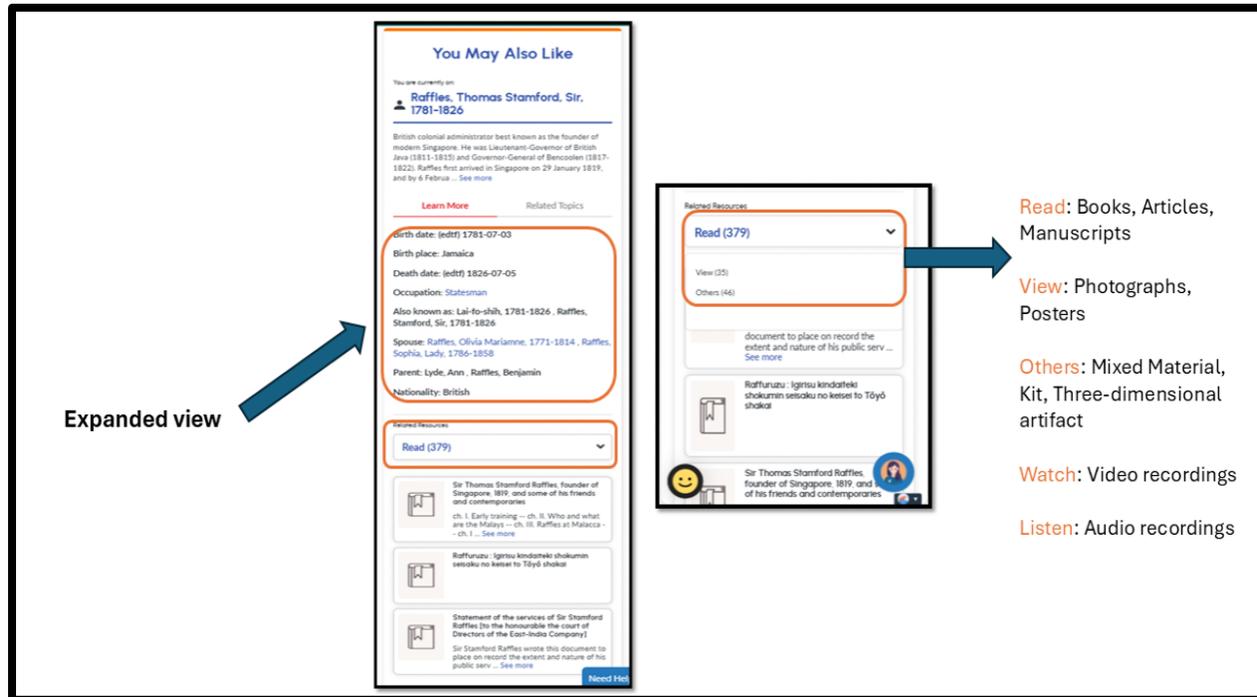
- Text-based materials are clustered into “read”
- Still images are clustered into “view”
- Moving images are clustered into “watch”
- Audio materials are clustered into a “listen” category

For usability, it was decided to omit items that could not be classified clearly into these four categories based on source data. Having a category of “others” was considered confusing to users if there was a potential overlap with existing categories.

Resource Listing

The resources themselves are displayed in order of relevance, where relevance is determined through an algorithm. The display also includes a thumbnail from the collection as well as a direct link to the resource in its source system, via the LDMS. As the LDMS stores the information about the entities, as well as the source systems of the resources, it becomes the invisible link between the source systems and the widget displayed on the Infopedia web service. Users are unaware of the underlying technological solution that provides this convenience.

Figure 6. Resource Type Choice.



CHALLENGES AND FUTURE PLANS

Ranking Algorithm

The widget API’s ranking algorithm plays a crucial role in the recommendation and display of resources within the LDMS, significantly affecting user engagement and satisfaction. This section explores the current operational mechanics of the widget API’s ranking algorithm, identifies the principal challenges it faces, and outlines proposals for its ongoing refinement to ensure future robustness and adaptability.

The ranking algorithm uses relational type scores and local entity scores to determine the relevance and presentation of resources.

Relational Type Score

- **schema:mentions:** Indicates entities mentioned within the article. This property helps recognize indirect relationships, which, although not central, may provide valuable context or additional exploration avenues.
- **schema:about:** Defines the main themes or entities that the content directly addresses. In the widget API, resources tagged with “schema:about” receive higher scores, reflecting their direct relevance to the core subject matter of the article.

Primary Scoring Metrics

- **Mentions Only:** Specifically awarded for only the presence of schema:mentions, indicating indirect relevance.
- **About Only:** Specifically awarded for only the presence of schema:about, indicating direct relevance.
- **Mentions and About:** Assigned when both schema:about and schema:mentions are detected, indicating optimal contextual relevance.

Secondary Scoring Metrics

- **Non-Local Entity:** No additional points for topics without local entity connections, maintaining a neutral position on globally relevant topics.
- **Local Entity:** Additional points are awarded for topics linked with local source entities, boosting cultural and regional significance.

The schema:mentions triple is retrieved using Entity Extraction and Entity Matching, particularly for enhancing the metadata of digital encyclopedia articles in Singapore Infopedia.¹⁹ This metadata facilitates the dynamic linking of mentioned resources within the articles. The integration introduces complex challenges, necessitating thoughtful, strategic planning. The process involves the following:

1. **Entity Extraction:** Leveraging GPT-4 to accurately identify entities such as people, places, and organizations within texts, guided by sophisticated prompt engineering techniques.
2. **Entity Matching:** Post-extraction, entities are aligned with the NLB Named Authorities Dataset through precise string matching, ensuring accurate linkage within the Knowledge Graph and enhancing article metadata.

Challenges in Current Ranking Methodology

The current ranking methodology faces several challenges that impact the effectiveness of resource recommendation and user engagement:

1. **Over-Reliance on Fixed Scores:** The static nature of the scores may not adequately reflect the evolving interests of users or the dynamic content landscape within the Knowledge Graph.
2. **Limited Contextual Adaptation:** The current scoring system inadequately adjusts for varying levels of user engagement or historical interaction patterns, which could allow for deeper personalization and user-specific content curation.
3. **Complexity in Balancing Scores:** Managing the balance between “Mentions” and “About” poses significant challenges, particularly when striving to maintain relevance across broader or more nuanced topic connections.

Future Enhancements and Strategic Directions

To address these challenges and enhance the widget API’s functionality, the following strategies are proposed:

1. **User Feedback Integration:** Develop a feedback aggregation system where user interactions and preferences are periodically analyzed to inform updates to the SPARQL queries. This system would operate on a batch processing model, where feedback is accumulated over a period and then analyzed to determine trends and preferences, which are subsequently reflected in the graph queries and adjust the ranking algorithm in the SPARQL query.
2. **Semantic Analysis Enhancements:** Expand the schema to include additional properties or create a more layered approach to semantic analysis. Investigate other schema properties that could be relevant and might capture different dimensions of content relationships.
3. **Dynamic Scoring Adjustments:** Implement real-time content analytics to dynamically adjust the scoring of resources based on evolving content trends and relevance. This approach allows the algorithm to adapt to changing content landscape and user engagement patterns in real time, ensuring that the scoring reflects the current state of the

Knowledge Graph and its relevance to the overall user base, while maintaining user anonymity.

These extensive strategies and insights into the widget API's ranking algorithm underscore the ongoing efforts to refine and enhance how information is curated and presented to users of the widget in Singapore Infopedia, ensuring an enriched and engaging user experience.

CONCLUSION

The Singapore Infopedia Widget represents a real-life application of a semantic discovery service in the context of library resources. By leveraging semantic web technologies and artificial intelligence (AI), these tools demonstrate the potential for libraries to evolve in the digital age.

This project showcases the power of linked data in enhancing resource discovery, illustrates the potential of AI integration in library systems, and provides a model for creating interconnected knowledge networks applicable to other institutions.

This project not only enhances NLB's services but also contributes to the broader discourse on the future of libraries and information management. It exemplifies how traditional knowledge institutions can adapt to meet the evolving needs of users in an increasingly digital world, potentially influencing future developments in library science and information management practices globally.

As the field continues to evolve, the principles and technologies demonstrated in this project may serve as a foundation for further innovations in how libraries organize, present, and make accessible their vast collections of knowledge.

ENDNOTES

- ¹ Min Hoon Ee, "Mining Text, Linking Entities – NLB's Journey," paper presented at the IFLA World Library and Information Congress, Athens, Greece, August 24–30, 2019, <https://repository.ifla.org/bitstreams/e166ef01-bcf5-4727-8252-0bff0992a8d8/download>.
- ² "Linked Data," World Wide Web Consortium (W3C), last modified September 24, 2023, <https://www.w3.org/wiki/LinkedData>.
- ³ "BIBFRAME Model, Vocabulary, Guidelines, Examples, Notes, Analyses," Library of Congress, accessed October 6, 2025, <https://www.loc.gov/bibframe/docs/index.html>.
- ⁴ Greta Bahnemann et al., "Transforming Metadata into Linked Data to Improve Digital Collection Discoverability: A CONTENTdm Pilot Project," OCLC Research, January 2021, <https://doi.org/10.25333/fzcv-0851>; Eben Carle, "Ask a Techspert: What Is Generative AI?," Google – The Keyword, last modified April 11, 2023, <https://blog.google/inside-google/googlers/ask-a-techspert/what-is-generative-ai/>.
- ⁵ Irfan Ullah, Shah Khusro, Asim Ullah, and Muhammad Naeem, "An Overview of the Current State of Linked and Open Data in Cataloging," *Information Technology and Libraries* 37, no 4 (2018), 47–80, <https://doi.org/10.6017/ital.v37i4.10432>.
- ⁶ Richard Wallis, "BIBFRAME Dilemmas for Libraries: Challenges and Opportunities," *Data Liberate* (blog), November 6, 2024, <https://dataliberate.com/2024/11/06/bibframe-dilemmas-for-libraries-challenges-and-opportunities/>.
- ⁷ Robin Dresel, "Designing a Linked Data Service across Borders and Timezones: The National Library Board's Experience," paper presented at the International Conference on Dublin Core and Metadata Applications, Daegu, South Korea, November 6–9, 2023, <https://doi.org/10.23106/dcmi.953335608>.
- ⁸ "Infopedia," National Library Board Singapore, accessed June 3, 2025, <https://www.nlb.gov.sg/main/onesearch/result?type=infopedia&nlonline=true>.

-
- ⁹ “National Library Online,” National Library Board Singapore, accessed June 3, 2025, <https://www.nlb.gov.sg/main/onesearch/result?nlonline=true>; Ashwin Nair, Min Hoon Ee, and Robin Dresel, “Metadata Enrichment with Named Entity Recognition Using GPT-4,” paper presented at the International Conference on Dublin Core and Metadata Applications, Ontario, Canada, October 20–23, 2024, <https://doi.org/10.23106/DCMI.952445840>.
- ¹⁰ Dresel, “Designing a Linked Data Service.”
- ¹¹ “lcnetsdev/marc2bibframe2,” Library of Congress, GitHub repository, last modified June 13, 2017, <https://github.com/lcnetsdev/marc2bibframe2>.
- ¹² Hanna Hussein, “Linked Data @NLB,” *Singapore Journal of Library and Information Management* 44 (2015): 20–34, <https://www.las.org.sg/wp/sjlim/files/SJLIM2015-Hussein-2.pdf>.
- ¹³ “Bibliographic Framework Initiative,” Library of Congress, last modified May 19, 2019, <https://www.loc.gov/bibframe>.
- ¹⁴ Puay Eng Tang, Glenn Hong, and Haliza Jailani, “Authoritative Content to Build Trust in an Age of Information Overload: The National Library Board of Singapore’s Experience,” paper presented at the IFLA World Library and Information Congress, Kuala Lumpur, Malaysia, August 24–30, 2018, <https://repository.ifla.org/handle/20.500.14598/6347>.
- ¹⁵ Ee, “Mining Text, Linking Entities.”
- ¹⁶ Rachael Goh, “Using Named Entity Recognition for Automatic Indexing,” paper presented at the IFLA World Library and Information Congress, Kuala Lumpur, Malaysia, August 24–30, 2018, <https://repository.ifla.org/items/f25d31af-cba5-4238-a898-de82e6f661d9/full>.
- ¹⁷ Nair, Ee, and Dresel, “Metadata Enrichment.”
- ¹⁸ Carle, “Ask a Techspert”; OpenAI et al., *GPT-4 Technical Report* (arXiv, March 15, 2023), last updated March 15, 2024, <https://arxiv.org/abs/2303.08774>.
- ¹⁹ Nair, Ee, and Dresel, “Metadata Enrichment.”