Digital Humanities in the India Rim

Contemporary Scholarship in Australia and India

Hart Cohen Ujjwal Jana Myra Gurney

https://www.openbookpublishers.com

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at https://archive.org/web

Updated digital material and resources associated with this volume are available at https://www.openbookpublishers.com/product/0423#resources

ISBN Paperback: 978-1-80511-387-4
ISBN Hardback: 978-1-80511-388-1
ISBN Digital (PDF): 978-1-80511-297-6
ISBN Digital eBook (EPUB): 978-1-80511-389-8
ISBN HTML: 978-1-80511-390-4
DOI: https://doi.org/10.11647/OBP.0423

Cover concept by Thalan-Harry Cowlishaw, designed using original image provided by Freepic, https://www.freepik.com, CC BY-NC

Cover design by Jeevanjot Kaur Nagpal

14. Digital Humanities for a different purpose

Julian Walker, Miyuki Hughes, Madeleine Leehy and Peter Mauch

Abstract

This chapter explores the recent application of Digital Humanities (DH) methods to a database for Australia-Japan Research and Industrial Collaboration (DAJRIC) and assesses their potential utility to a project outside of traditional scholarly purposes. The primary objective of this chapter is twofold: (i) to evaluate the effectiveness of this approach in serving our purposes, and (ii) to consider the scalability of the pilot database to accommodate numerous yet unfunded Japanese-Australian research projects. By developing a Heurist database, the project can harness the intuitive design principles that make DH methods so effective and appealing for scholarly purposes to users unfamiliar with these research fields. Throughout this process, the project team has discovered the challenges of raw data, ontology development, and bilingual functionality that face a project of this scale, whilst also realising the potential of Digital Humanities' techniques in providing improved user interactivity and search functionality through record types and their connections with each other, and through data visualisation. These techniques, when applied with knowledge organisation techniques, enable a scalable database that can organically grow with the hundreds of projects to be entered in the future. As such, this project provides a valid example of how scholarly techniques within the Digital Humanities can

 https://doi.org/10.11647/OBP.0423.14

be applied successfully to projects that act as a gateway between academia and other sectors.

Keywords

Digital Humanities; Heurist databases; databases; ontology development; translation; data visualization; knowledge organisation.

Introduction

This chapter introduces readers to a recent attempt at applying Digital Humanities techniques for reasons other than the scholarly: we are developing a database that will showcase a large range of Japanese-Australian collaborative technological-industrial-scientific-academic ventures. Funded by the Australia-Japan Foundation, the database aims to provide a gateway for investors and to facilitate corporate funding for Japanese-Australian translational research.

The project seeks to serve an emerging community between collaborative researchers and industry funding in the bilateral Australia-Japan relationship. Functioning as a repository for this information, a significant amount can be learned from digital libraries and their definitions of sustainable preservation of digital metadata. That said, there are varied definitions of what sustainability entails, especially when dealing with a database that seeks to connect academia and industry in such a way. Therefore, sustainability in this case is best defined as the ongoing maintenance of the platform to ensure continued access for both researchers and funders (Edmond & Morselli, 2020, p. 1022).

Digital Humanities techniques are applicable because the search functions of the database must necessarily be intuitive and all users, regardless of their level of familiarity with the Japanese-Australian research landscape, need to be able to find projects of interest with ease. This includes the ability to simplify technical jargon across various research fields. The discussion in this chapter centres on a pilot project, which involved the construction of a database based on ten research projects funded by the Foundation for Australia-Japan Studies. The success (or otherwise) of the pilot database would determine whether

(i) the DH techniques under discussion in this chapter would serve our purposes, and (ii) the pilot database could be scaled up so that it included several hundred as yet unfunded Japanese-Australian research projects.

This process has uncovered various challenges and obstacles that stem from dealing with dense raw data specifically tailored for specialist audiences in their original form, and the objective of bilingual functionality. Furthermore, the significant desire to create a database that develops organically in a self-sustaining manner raises concerns about scalability in both key terms and search functionality. Yet for each challenge, the nature of DH methodologies provided alternative solutions to our problems, such as creative search functionality through Heurist record types and their connections with one another; data visualisation through a map to demonstrate the transnational research collaboration simply; and the usage of Heurist's database structures to ensure scalability for such projects.

Raw data

The raw data for the pilot project were the research proposals that research teams submitted to the Foundation for Australia-Japan Studies. These proposals fell into one of the following priority areas: (i) medical science and technology; (ii) materials and energy technologies and systems; (iii) computing, digital and AI applications and developments; (iv) food, marine and agriculture science and technologies; and (v) social and physical infrastructure.

Each proposal included (i) project title, (ii) project summary, (iii) project salience, (iv) innovation, (v) project team members' track record, (vi) networks, (vii) sustainability, (viii) diversity, and (ix) financial contributions from sources other than the Foundation for Australia-Japan Studies.

The raw data is dense, full of discipline-specific phrasing and terms, and not necessarily easy for the layperson to digest. The principal issue for the database involved making this data digestible for a wider audience, so that, for example, an employee at a capital-rich corporation looking for projects to support financially can quickly and easily locate projects of potential interest.

Database construction

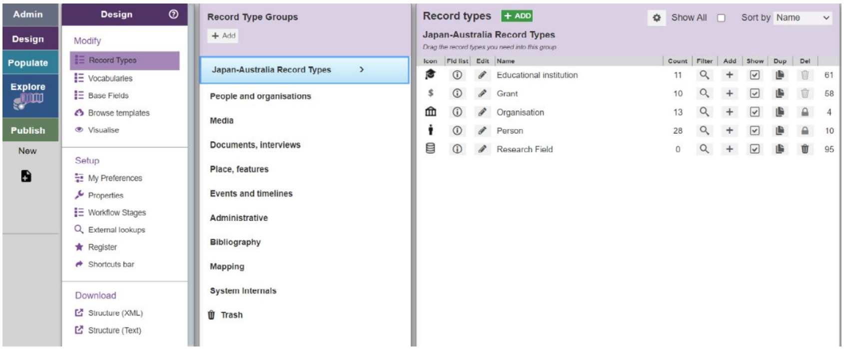

Fig. 14.1 Database interface.

The project team used Heurist, which is an open-source online database builder and CMS publisher designed for Humanities research data and collections. In Heurist, we created our own "record types," creating records for "educational institutions" and "organisations", "persons" (meaning researchers), and the "research field" of a project (see Figure 14.1). By creating these categories, database users can search projects by research field, keywords, institution, organisation, keywords, etc.

While our initial steps looked simple, we later realised these "record types" actually affect the database's user-friendliness. For example, we initially grouped universities and companies in one category called "organisations." But if a third party is looking in this database for a project to fund, they will probably be interested in, e.g., industrial partners on any project. The database therefore needs to filter companies from universities. We had to differentiate and separate these record types so that the database now includes separate entries for "educational institutions" and "organisations".

Heurist includes DH functionality that is beneficial to the database. We have been able to make connections between items so that, e.g., the database connects "researchers" to "educational institutions", "organisations" to "projects", or "projects" to "researchers". This allows users to also search by a researcher's name to see what project they are working on, or to search by the university to see what other projects are based there. This functionality contributes directly to the user-friendliness of the database.

Ontology building

As part of producing content for the database regarding the project proposals, we were tasked with building an ontology. This involved identifying keywords and phrases while considering the specific users of the database, such as academics and individuals from the corporate world (Noy & McGuiness, 2001, p. 6). We concerned ourselves with likely search items that would enable database users to find proposals and key information (Jones et al., 1998). Ultimately, we were concerning ourselves with how users would interact with the database.

The process of ontology building began with reading through each proposal a couple of times to understand both the intent of the author and the material being presented. While reading through the proposal, speciality terms such as 'fluorescence in situ hybridisation' were researched, defined and noted within their context.

When building the list of keywords, the aim was not to duplicate the terms within the proposal document as these were to be uploaded into the database to enable a free search of terms. For example, if you search 'fluorescence', even though this was not included within the list of keywords, the proposal would still emerge as it was included in the raw data.

There were a range of things to consider in building an ontology as a distinct entity from the data present in the proposal document. For example, we had to consider pluralisation; various terms within the proposal document were treated as singular, but what happens if the database user, e.g., attempts to search for "blood cancers" rather than "blood cancer"?

We also had to consider multiple expressions for the same term. The use of synonyms proved the most expansive element of the list. There are so many different terms used for the phrase "blood cancer research", what happens if a database user searches for "blood cancer interventions", or "blood cancer imaging", or "blood cancer monitoring", or "blood cancer study", or "blood cancer diagnostics"? The list just keeps going and continues to grow if we move beyond the more generalised terms to include higher-order concepts. For example, instead of "blood cancer research" database users may search for more abstract terms such as "haematology", "biomedical research" and "phenotype research".

As mentioned, the material is dense and not easily digested by the layperson, and so keywords and phrases had to be simplified and redefined to increase accessibility for database users. For example, users interested in interventions within blood cancer research may not be searching for the specific term "in situ hybridisation" but may use layperson terms such as blood cancer treatments and blood disorder initiatives.

The last thing we considered was spelling localisation, for example using an 's' or 'z' in 'hybridisation' as per British English or American English. This localisation was especially important as generally American English is used within Japan. However, British English is used in Australia, leading to differences in search terms as well as the free search of the proposal depending on what type of English the author used.

The ten projects in our pilot database gave rise to a total of 1353 keywords. This is a mean number of keywords of 140.5 per project. As mentioned, by using Excel, data was produced on a line-by-line basis, thus it could be easily transposed to the database after checking the lists for clarity and spell-checking. Thus, keywords can be easily updated, which is essential as the number of keywords is always growing. There will always be new terms to add and consider and more terms will develop over time. This is an essential aspect of ensuring that the database is organic and growing constantly as it incorporates new data and new projects.

Bilingual issues

Target users of the database will mostly be based in Australia and Japan. If the database is to be user-friendly, it needs to be a Japanese-English bilingual database. This has created various issues. For reasons of brevity, we have restricted our discussion to several illuminative examples (Figure 14.2).

		Taro Tanaka
Kanji	Chinese characters	田中太郎
Hiragana	The basic Japanese phonetic alphabet. (It is primarily used for function words and inflexions.)	たなか　たろう
Katakana	The basic Japanese phonetic alphabet. It is often used for the transcription of loan words from foreign languages.	タナカ　タロウ

Fig. 14.2 Japanese has three spelling systems.

As this table demonstrates, the Japanese language has three types of characters. Should we enter every Japanese word in all three characters? In the context of personal names, as a rule, personal names are written in kanji (Chinese characters). But if a database user wants to search by researcher name, even though they don't know the exact kanji that the researcher uses for their name, the database user is likely to use either hiragana or katakana—they are essentially searching by phonetics. In short, a search for a particular researcher in the database might take one of the following three forms: "田中太郎"; "たなかたろう"; "タナカタロウ".

One problem is that the raw data includes no Japanese text (the proposals were written only in English), so the team constructing the database cannot always be certain of the kanji that a researcher uses to write their name. At the very least, it involves a lot of cross-checking, by consulting, for example, at the university where the researcher is based, to see whether the raw data includes the researcher's name written in kanji.

Circumstances	**Issues**	**Examples**
When we know how to pronounce the name.	It would have some variations depending on how you hear it.	Madeleine [ˈmadəlɪn, -lèɪn] マデライン;マドレイン; マドリーン;マドレーン; マドレーヌ
When we do not know how to pronounce the name.	We may transcribe it completely wrong.	Peter Mauch (How do you pronounce it?) ピーター・マウチ (Mauch?) ピーター・マーク (Mark?) ピーター・モーク (Mork?)

Fig. 14.3 Non-Japanese names.

The second question was, if we enter Japanese researchers' names in both English and Japanese (Fig. 14.3), should we do the same for non-Japanese researchers? Would it affect the search results if only Japanese researchers had entries in. both languages? To answer this, I thought about how we could add non-Japanese researchers' names in Japanese. Two problems arose.

The first problem arises when we know how to pronounce the name. We can write a non-Japanese person's name in Japanese according to the phonetic sounds. However, we are just sorting the sounds that do not exist in Japanese into the closest Japanese sounds. The result would have some variations depending on how you hear it. This may cause problems later. For example, Figure 14.3 indicates the possible spellings of Madeleine in Japanese, and they all sound similar but are spelled slightly differently.

The second problem is when we do not know how to pronounce the name. In this case, we could record it completely incorrectly. Let us use the surname 'Mauch' as an example. If we don't know how to pronounce this name, the way we write it phonetically in Japanese becomes nothing more than guesswork. Possible renderings in Japanese are マウチ, マーク, and モーク. They are quite different.

Cancer
癌 Cancer in epithelial tissue e.g., Skin cancer, stomach cancer
がん All types of cancer e.g., Insurance for cancer, cancer research centre The visual impression is softer than Katakana
ガン All types of cancer

Fig. 14.4 Keywords.

Similarly, a database user might conduct a keyword search for 'cancer' in one of three different ways: 癌、がん、or ガン. These variables become more difficult and more complex when we consider variations in English-language terms. Sometimes each English-language variation offers multiple possible Japanese-language translations, so that the process of database construction becomes almost endless. Figure 14.5 below includes a few examples.

Possible variations	**Examples**
Use of hyphen	shelf-life / shelf life
Singular/ plural	tumour / tumours
Multiple expression	medical implant / biological implant / biomedical implant chromosomal abnormalities / abnormal chromosomes
Acronyms	Fluorescence in situ hybridisation / FISH Magnesium / Mg
Spelling (US/UK/AUS)	haemato- (US) / haemato- (UK) Tumor (US) / Tumour (UK)
Technical term / layman's term	Multiple myeloma / Bone marrow cancer

Fig. 14.5 English variations and multiple translations.

Visualising data

We have also considered how users might choose to visualise the data. A map seemed like a logical choice as it demonstrates not only the transnational Japanese-Australian connections between the proposals, but it also increases user interactivity by making searching through the database functional, rather than just relying on the free search/ categorical search function. Geotagging was used on both the proposals and educational institutions, which can be searched and identified on the map, ultimately linking to the proposal as the final function. The map widget also has a built-in timeline function that demonstrates when the proposed research began and finished, in other words, the lifespan of the proposal.

The use of a map tool to present data within our database gives rise to additional issues. As mentioned, one of the initial goals of the database was for it to be bilingual. As the map is primarily visual, it allows people with limited language capability to access information, which is highly important in its current state as a largely English-language-dominant site (while the more difficult elements of producing a bilingual database are resolved). Using the map also allows the layperson to engage with it, as it requires a minimal level of skill—just clicking around and exploring a map, rather than having to click through various menus or attempt to freely search if a user has no idea what they are trying to look for or is attempting to engage with the database for various purposes.

As mentioned previously, using the map as a visual representation of data demonstrates the global connections between proposals. As one of the major goals of this database was to produce international collaborations between both Japan and Australia, and to allow engagement with various stakeholders external to close-knit academic circles, the map draws various connections and reinforces the close links between Australia and Japan, especially regarding research. Further, mapping international data points draws attention to foreign facilities that may not be recognised by someone overseas but are generally well known in the country of origin, as demonstrated by the mass connections to this data point.

In another vein, it may also draw attention to smaller facilities that have great potential to produce future work and demonstrate unexpected connections. For example, one of the proposals regarding autonomous robots generated contributions from researchers in Jamaica, outside of the realm of an exclusive Australian and Japanese relationship. This demonstrates unexpected connections that would not have been picked up on the initial reading of the proposal or the database entry. In contrast, you can also examine which proposals involve prestigious organisations—for example, the National Institute for Quantum Science and Technology in Japan was involved in various proposals. Users may be interested in what these large organisations are focusing on, as well as the connections between pieces of work, which is demonstrated clearly through the visual map function.

We have begun to explore how filters, such as searching for organisations, countries, and whether the data points are linked to an industry partner impact the effectiveness of establishing connections and increase the discoverability of proposals. For example, the field "industry partner" enables users to examine organisations outside of close-knit academic circles for possible collaboration. The timeline function supports the aim for the database to be a living thing, and also helps to build the reputation the grant proposal program as a whole, as it demonstrates the history of providing funding and the length of the body of work.

Although a map appears to be a clear benefit to the database and a valuable tool for representing the data, we also attempted other forms of modelling, such as network maps to demonstrate connections between different identifiers including the grants, organisations, researchers, institutions, etcetera. But it became eclectic very quickly as there are too many variables to consider at this early stage and we have not yet identified a key area of focus for connections and how to filter these. This might be something to discuss as we move into presenting the database to focus groups—what are the most important factors that individuals who use the database want to search for, and do visualisations such as network maps aid in the discoverability of proposals and in people's understanding of the database as a whole?

Scalability and usability

One of the challenges with large-scale projects such as this database is scalability and usability for end users. This requires good knowledge organisation (i.e., the organisation of the data itself), and streamlined ontologies that are easily understood by the end user. Such a requirement often emerges in Digital Humanities (DH) tools where the intuitive design makes substantial projects possible for broad audiences (Golub & Liu, 2022, p. 2). The Heurist platform enables humanities scholars to enter data and create front-end websites, just as easily as an end user can search and use that database for their purposes. But this hurdle can result in overwhelming ontologies both in size and in complex jargon. While terminology is beneficial for those familiar with the applicable fields, it is fundamentally confusing to lay users. Therefore, this project in its second year has devoted time to streamlining the vast ontology to become more direct for point A to point B searches, whilst also laying the foundations for large-scale knowledge organisation that further acts as a filter for the expansive categorisation of projects.

Knowledge organisation and databases

How a database organises its data is an important task, and thus it is the focus of much scholarly inquiry, especially in the Digital Humanities. DH researchers commonly face the realities of ad hoc solutions for data collection management and find themselves suffering from scalability issues (Golub & Liu, 2022, p. 1). Part of this stems from the sheer scale of digitalised material Digital Humanities researchers collect, annotate, and analyse (Robinson, 2014, p. 247).

For this database, however, we focus on the closely related field of information retrieval. While the database, and how data is organised, is crucial, the objective of the database is to be quickly and easily searchable for the layperson outside of science, technology, engineering, and medical (STEM) fields. This requires data to be organised for the fundamental purpose of retrieval rather than analysis. It is critical to note that, while users may be competent in searching during their day-to-day lives, they will not necessarily be experts nor familiar with these cutting-edge projects (Hjørland, 2015, p. 1570).

Nested vocabularies

Yet searchability is vital for the success of the database's objective, which is to connect funding to collaborative research projects. All projects within this database are located within priority areas as defined by the Foundation of Australia and Japanese Studies, and as such, these categories become the foundation of the nested vocabularies. This top-down approach takes broad conceptions, the priority areas, and narrows them down into specialisations or specific areas of interest (Noy & McGuiness, 2001, p. 6). For example, the generalised priority area of "Food, Marine, and Agriculture" has split into areas of specific interest such as "food", "marine", or "agriculture". In a similar sense, it also required us to further define the meaning of the priority area of "Social and Physical Infrastructure" into categories of "renewable energy", and "resource security" to enable users to immediately understand what the priority area holds.

Further examples include "Medical Science and Technology" being broken down into smaller categories familiar to broader audiences, such as "cancer", "genetics", and "surgeries". This also applies to the priority area of "Materials, Energy Technologies, and Systems", where sensor and measurement devices are bundled together with "systems" and "energy". This approach allows the database, for those managing it, and for the end users, to take broad categories and begin to narrow them down for their respective purposes. This is a core principle for nested vocabularies, whose emphasis on filtering down vast categories into smaller sections is one solution to the scalability problem faced by a database that seeks to be as organic as possible.

Streamlined ontologies

The key to knowledge organisation within large-scale projects aimed at non-specialist users lies in an easily understood, usable, and streamlined ontology. This can be achieved in multiple ways, involving a "standard and systematic procedure to use, and categorise information" that can be searched and obtained from a database (Ezhilarasi & Kalavathy, 2023, p. 54). This project has approached this standard and systematic process through the construction of an ontology during the pilot year. This ontology was applied to both searchable tags and simple searches. The streamlined ontology, however, uses the currently existing ontology as a foundation to create a set of tags that improves the user's ability to find information easily (Hepp, 2007, p. 90). This new ontology is designed to improve scalability by being specific yet not overly technical or redundant. Tags used in the search fields aim to be clear, precise, and concise in their usage. For example, the various ways of searching for implants have been reduced to a single keyword. Specific terms, such as types of cancers, have remained for scalability purposes. This enables the user to traverse the projects in the database in a fast and direct manner in three ways. They can either: 1) use the nested vocabularies of priority areas to find a certain selection of projects; 2) use these streamlined tags to find a certain project; or 3) search the database more broadly using the vast ontology previously developed.

Another vital aspect of streamlined ontologies is its maintenance and application. A dedicated ontology can be additionally purposed for user self-entry into the database and would require little maintenance, whereas an overly expansive one can quickly run rampant as more users enter data into the database. This allows the database to become more of a repository that other sectors such as agriculture, information and technology, or scientific fields already utilise significantly. The database would evolve from one that primarily stores data for the sole purpose of connecting funder to researcher, into a self-sustaining repository for these projects that grows over time.

Future enhancements

The scalability changes have been successfully implemented in the pilot database. Furthermore, these enhancements for knowledge organisation and optimised ontology have been extended to the Japanese sister site, aligning with the database's ultimate goal of bilingual functionality. Looking ahead, the most significant challenge the database will encounter is archiving completed projects from the Search Projects page.

Currently, this issue is addressed through a filtering process in search fields. However, as the database grows, a maintenance process will be necessary to ensure that completed projects are appropriately archived in a separate location. This measure is crucial to prevent finished projects from overshadowing unfunded projects and maintaining the integrity and relevance of the database's content.

Discussion

In this section, we discuss how the evolution of the database over the pilot process has many implications for the use of Digital Humanities technologies and techniques in connecting academic approaches to external opportunities. Does it justify the use of Digital Humanities techniques in non-academic spaces, and is this approach scalable to accommodate large-scale and continuous data entry?

The primary obstacle to applying DH approaches to databases like this is the acquisition of raw data. However, our research has shown that not only can this obstacle be overcome, but it can also be implemented without necessitating knowledge from the end user. DH functionalities provided by platforms like Heurist, such as establishing "relationships" between different types of records (e.g., "researchers" and "organisations"), enhance user-friendliness by enabling users to search for specific researchers, universities, or industries of interest directly.

Furthermore, our study has highlighted how DH approaches can address the challenges of bilingual databases in Australian-Japanese collaborations. For instance, we explored complexities related to the Japanese writing system, which complicates English-to-Japanese translations for keywords developed in the ontology, as well as translating English names into Japanese. Additionally, managing English variations and multiple translations poses ongoing challenges as the database progresses beyond the pilot stage. However, DH techniques offer innovative solutions to overcome these challenges.

One such solution provided by Heurist is the ability to leverage data visualisation tools such as maps, which not only showcase transnational collaborations between researchers, organisations, and industries but also enhance usability by allowing users to interact with the database intuitively, without needing to know specific search terms. This demonstrates the versatility and effectiveness of DH techniques in improving accessibility and usability for diverse audiences. Yet, the exploration has also revealed that this process requires strict streamlined ontologies and knowledge organisation to ensure that projects remain scalable and useable for the end user. This challenge is unique to non-scholarly applications of DH techniques, as the database's primary purpose is to become a self-sustaining repository for research projects, as opposed to its traditional purpose as a tool to analyse vast quantities of data. By using Heurist's intuitive vocabulary functions, we can streamline the ontology that can grow alongside the vast number of projects in the future.

The limitations of this chapter are primarily attributable to the project being in its pilot phase. However, this limitation also serves as an opportunity for future research and a re-evaluation of how DH techniques can be further applied to the database once it becomes self-sustaining. Despite being in the early stages, these findings hold significance as they demonstrate the feasibility and success of DH approaches for non-academic purposes. With careful refinement, these techniques can be scaled to accommodate large and ongoing data entry.

The insights gleaned from this study should encourage others in similar or related fields to explore the potential of DH techniques beyond academic research. By showcasing the adaptability and effectiveness of DH methodologies on a broader non-scholarly scale, this chapter provides a compelling case for the wider adoption of such approaches in various contexts.

Conclusion

This chapter presents a novel application of Digital Humanities techniques aimed at lay audiences in the creation of a database highlighting a diverse array of Japanese-Australian collaborative research endeavours spanning technological, industrial, scientific, and academic domains. The primary objective is to establish connections between corporate or industry funding and these collaborative projects in a manner that is transparent and easily accessible by all stakeholders.

The database contains a vast amount of complex raw data, posing a significant challenge in simplifying specialised terminology for lay audiences. However, the Heurist platform offers innovative and user-friendly solutions to address this issue by facilitating simple and filtered searches, thus enhancing accessibility.

Furthermore, the database's aim to support Japanese-Australian research projects necessitates bilingual capabilities, which pose distinct challenges. Translating English keywords into Japanese and selecting the appropriate Japanese spelling system present initial hurdles. Moreover, the intricate task of translating Japanese text back into English adds another layer of complexity.

For a project of this magnitude, managing hundreds of projects awaiting data entry, and anticipating future expansions necessitates a scalable ontology and search functionalities. To tackle these challenges, nested vocabularies have been employed to enhance knowledge organisation and establish a flexible structure capable of accommodating the organic growth of the database. This approach significantly enhances the ability of the database to achieve bilingual functionality.

This project's application of Digital Humanities methods is not only justified, but it also serves as a model for collaborations between academia and industry within the DH space. By leveraging techniques commonly employed in traditional DH projects dealing with large datasets, this project demonstrates the potential to effectively overcome knowledge and language barriers. Consequently, it showcases a viable and efficient pathway for future projects using DH approaches outside the realm of academia.

Works Cited

Department of Foreign Affairs and Trade. (2024). Australia-Japan Foundation. https://www.dfat.gov.au/people-to-people/foundations-councils-institutes/australia-japan-foundation

Ezhilarasi, K., & Kalavathy, G. (2023). Development of Contextual Crop Ontology for Effective Information. In *Proceedings of 3rd International Conference on Recent Trends in Machine Learning, IoT, Smart Cities and Applications: ICMISC 2022* (pp. 53–65). Springer Nature Singapore. https://doi.org/10.1007/978-981-19-6088-8_6

Foundation for Australia-Japan Studies (n.d.). https://www.fajs.org/

Golub, K., & Liu, Y.H. (Eds). (2021). *Information and Knowledge Organisation in Digital Humanities: Global Perspectives*. Taylor & Francis. https://www.doi.org/10.4324/9781003131816

Gunjan, V.K., & Zurada, J.M. (2023). Development of contextual crop ontology or effective information retrieval. In *Proceedings of 3rd International Conference on Recent Trends in Machine Learning, IoT, Smart Cities and Applications* (Vol. 540, pp. 53–65). Springer. https://doi.org/10.1007/978-981-19-6088-8_6

Hepp, M. (2007). Possible ontologies: How reality constrains the development of relevant ontologies. *IEEE Internet Computing 11*(1), 90–96. https://doi.org/10.1109/MIC.2007.20

Hjørland, B. (2015). Classical databases and knowledge organization: A case for Boolean retrieval and human decision-making during searches: Classical databases and knowledge organization. *Journal of the Association for Information Science and Technology 66*(8), 1559–1575. https://doi.org/10.1002/asi.23250

Jones, D., Bench-Capon, T., & Visser, P. (1998). *Methodologies for Ontology Development*. University of Liverpool. https://cgi.csc.liv.ac.uk/~tbc/publications/itknows.pdf

Noy, N.F., & McGuinness, D.L. (2001). *Ontology development 101: A guide to creating your first ontology*. https://corais.org/sites/default/files/ontology_development_101_aguide_to_creating_your_first_ontology.pdf

Robinson, P. (2014). Digital Humanities: Is bigger, better? In P.L. Arthur & K. Bode (Eds). *Advancing Digital Humanities.* (pp. 243–257). Palgrave Macmillan UK. https://doi.org/10.1057/9781137337016_16