Custodians & Midwives

The Library of the Future

Australian National University

NATIONAL LIBRARY OF AUSTRALIA

2/m

My dear Jeannette

We got an account of your wedd[ing]
week ago - Mamma Tante Mimine and Mary all
Tom to David - it appears that Mr Fell blundered
observes - you are not the less securely
will be all the same a hundred years hence
Cards - had not Emil an unofficial one a
them doing the wedding twice over at all -
same breakfast at Broadlands and only the ~~same~~
it would have been much more sensible to h
at the same time that we were - and so saved al
it is all right - and Ambleside has been edifie
wedding's in a week - I am now wishing to hear
& how you enjoyed London - whether you manag
enjoyed your journey up the Rhine - whether you
Rackwork - and whether old nurse said "Por mein, we
did not think you would get married so soon" ! I
managed Emil's arriving without notice - and "bei
that you looked more like a Lamm opfer the
of Wilhelm - no one ever mentions him
he said of his two cousins proceedings - your M
to me about you and your prospects - and I
no trouble by not agreeing for this last year. I
say about their being placed at different scho
as that had not been done and it could ha
about the boys knowing what it was for - it
do it - and give that eclat to the matter - So
it cannot be done and I dare say Dick
Charles & yours even if the idea of their hav
not mollify them the last six months - I
that Dummy had made some queer observati
no one mentions it - how soon am I to have a
Do you know I actually could not read Aug
had not taken the precaution to send me an
should have been no wiser for it. I suppose you
will remind you of Jane Raven - whom by
you met in Kew Gardens - She has not written
you will send me a long account of the Berlin
return for which you shall be personally introduced
and ducks that we keep at Prahran - We have
already - and he is going to our house tomorrow
the day after ... will be quite ready - and the

Custodians & Midwives
The Library of the Future

School of Cybernetics

The Australian National University
Canberra ACT 2601 Australia
www.anu.edu.au

CRICOS Provider No. 00120C

Custodians & Midwives:
The Library of the Future

Research and analysis on artificial intelligence and emerging technologies at the National Library of Australia by the School of Cybernetics, Australian National University, June 2021

Authors

Genevieve Bell, Alex Zafiroglu, Zena Assaad, Charlotte Bradley, Ned Cooper, Ellen O'Brien, Kathy Reid, and Lorenn Ruster.

Supported by Sarah Copley, Katherine Daniell, Mikaela Jade, Charlotte Jones, James Liu, Brenda Martin, Johan Michalove, Juliette Parsons, James Taylor and other members of the School of Cybernetics. The future scenario illustrations in this report were created by Leib Leventhal (Monrotype) in collaboration with Charlotte Bradley and Charlotte Jones.

Project Management

Paul Wong, Zena Assaad, and Ellen O'Brien.

Acknowlegement Of Country

The authors and the ANU School of Cybernetics acknowledge the Ngunnawal and Ngambri people, who are the Traditional Owners of the land upon which this report was prepared. We pay our respects to their elders, past and present. This Country, from which we benefit, has an ancient history that is both rich and sacred and will continue to hold a space for future generations to come together, learn from Country and one another. We recognise Aboriginal and Torres Strait Islanders are the first innovators, technologists, and engineers of this continent, and through custodianship have sustained knowledge systems that have informed and inspired generations.

Contents

"There are two functions in which I can see the library to serve in the social fabric of the future, one as being a custodian for books, the other one as being a midwife for those who wish to give birth to new insights and ideas … The question now arises, what are the inner workings of such a system in which you can act the double role of a custodian of books and of a midwife for new ideas and insights, thus maintaining the concept of a library as being a place where knowledge can be acquired?"

Heinz von Foerster
(1984, pp. 213–14)

Executive Summary

In July 1970, Heinz von Foerster, a leading Austrian American scientist, delivered an address to what is now the University of Wisconsin's I-School, and what was then the Library Institute. He had been asked to reflect on the ways in which emerging computational technologies would impact the future of the library and library staff. Rather than directly addressing this brief, von Foerster instead argued that library staff had considerable agency in deciding not only what technologies might be successful in libraries, but what technologies might be successful writ large. In so doing, he made clear his belief in the power of libraries and their centrality to what he poetically described as the ‘'social fabric of the future''. Von Foerster articulated two important and complementary roles for libraries and library staff – as custodians of collections and midwives of ideas. More than half a century later, his words still make considerable sense.

Custodians and Midwives is the outcome of a research collaboration between the National Library of Australia (NLA) and new School of Cybernetics (SoCy) at the Australian National University (ANU). This collaboration was undertaken in the first half of 2021. During the collaboration, SoCy researchers investigated and evaluated the potential new dynamics that will emerge if, and when, the NLA integrates tools and processes enabled by promising artificial intelligence (AI) technology capabilities into core work processes of the organisation's mission to "collect, preserve and make accessible" library materials (*National Library Act 1960* (Cth)).

The new School of Cybernetics sits within the College of Engineering and Computer Science at the ANU. Established in January 2021, the School of Cybernetics builds on the pioneering work done in the 3Ai Institute (2017-2020), which led the world in designing education on an approach to safely, sustainably, and responsibly scale AI.

Under the leadership of Distinguished Professor Genevieve Bell AO FSTE FAHA, the new school is updating cybernetics for the 21st century. The team is applying an approach to thinking about systems of technology, people, and environment that responds to, and modifies, the technocratic approach that has typified the development and evaluation of the increasingly complex digital-enabled systems of the last half century. As AI systems are increasingly embedded in everyday processes and spaces, the cybernetic approach presents a compelling way of thinking about how to design and manage complex systems that are dynamic, interactive, and relationship-dependent.

This approach seems particularly relevant to the NLA as it seeks to fulfil its role to "collect, preserve and make accessible" resources of national significance, as defined by the *National Library Act 1960* (Cth). In *Custodians and Midwives*, we take a cybernetic systems approach to the questions of how new and emerging technology capabilities might shape and be shaped by the future of the NLA.

In this report, we provide a library collections-centric analysis of clusters of AI capabilities to support the NLA in making decisions regarding the integration of AI-enabled technology systems into the organisation's processes now and in the future. We prioritise clusters of *capabilities* defined as (1) the services a class of technologies provide, and as (2) knowledge that can be applied to specific technology products and processes. Focusing on AI capabilities, rather than specific products and processes in the market or deployable now, enables us to outline a pathway for the NLA to evaluate opportunities, risks, pitfalls, and issues (ORPI) with AI over the timescales and rhythms that characterise the library as a cybernetic system.

The report is divided into five discrete sections:

- In Section 1, we detail the foundations for an approach to evaluating the suitability and fit of emerging AI capabilities for the NLA. These emerging AI capabilities must be considered not only in isolation but also within the broader system of the whole of the NLA – its people, its systems and processes, and its data, collections, and extant infrastructure. This systems approach is critical to successful technology deployments, and given the complexities of the library, we propose a cybernetic systems approach. We explore the NLA as a cybernetic system and the dynamics of this system from a collection-centric perspective.
- In Section 2, we turn to the technological innovation in and of libraries, and the pathways that have led to new AI capabilities in these systems, focusing on automation, digitisation, and AI. There is an overview of the landscape, emerging trends, and approaches taken by others in the library sector, including peer national libraries. We detail the landscape of AI and machine learning (ML) capabilities that are most relevant to library services, including a discussion of dominant threads in the Responsible AI conversations.
- In Section 3, we use science fiction prototyping to explore the future state of the library, focusing specifically on four sets of AI capabilities identified as priorities by the NLA collections branch: Optical Character Recognition, Machine Transcription, Machine-Actionable Collections and Transparent AI. We provide visions for the longer-term outcomes of the incorporation of these four clusters of current AI capabilities into the NLA collection workflows, and a detailed review of each cluster, including attendant opportunities, risks, pitfalls, and issues. We complement this section with a further overview of the science fiction prototyping exercises (see Appendix B).
- In Section 4, we provide pathways forward for decision makers at the NLA as they encounter specific solutions or approaches that may be incorporated into library practice. Whereas in Section Three, we provide an evaluation of current emerging *capabilities*, in this section we offer a technology reference guide for the NLA to evaluate options for incorporating AI into library practice, attentive to the role of the NLA now and in the future.
- The "Cybernetic Star Guide" – our technology reference guide – is a cybernetically-grounded way to evaluate prospective AI applications. It comprises a conceptual perspective on the library as a cybernetic system that sits within many other systems outlined in Section 1 (political, economic, regulatory), as well as a set of tools for supporting confident decision making regarding the direction of AI-enabled collections solutions. The Cybernetic Star Guide enables a comprehensive examination of the opportunities, risks, pitfalls, and issues with a proposed technology application from the perspective of the internal work processes of the library (implementation issues, knowable risks, automation, and workforce) through to more audience-centric perspectives (ethical concerns, unintended consequences, and impacts on users). We complement this section with additional reference materials and worksheets in Appendix C.
- In Section 5, we provide a glimpse into the emerging socio-technical, business, and regulatory trends that might be relevant to the National Library in coming years, and reflect on the current and future state of the library.

Taken in its totality, we believe this report represents a valuable contribution to the NLA in its dual roles as custodians and midwives for all Australians, today and into the future.

"I feel you should not wait and see what technology will do with you – pardon me, technologists will say 'for you' – but rather you should tell technologists what you want to be done."

Heinz von Foerster
1984, p. 213

01

Cybernetic Systems and the Library

The National Library of Australia (NLA) is a statutory authority within the Communications and Arts Portfolio of the Australian Government Department of Infrastructure, Transport, Regional Development and Communications. It was established under the *National Library Act 1960* (Cth), and as such, has particular legal requirements and obligations, which shape its collection, preservation, and engagement strategies and activities.

There are a range of new and emergent AI technological capabilities that seem especially relevant to the NLA's charter and obligations. These forces impact how the NLA fulfils its role to collect, preserve and make accessible resources of national significance, as defined by the *National Library Act 1960* (Cth).

Following von Foerster and others, we are interested in understanding the NLA contexts into which new and emerging AI capabilities will be applied, and how these new capabilities will necessitate different ways of thinking and doing for the NLA. We are especially interested in understanding the NLA as a system, one that extends beyond its obvious technological infrastructure to include the contexts, data, infrastructure, processes, and people that make up the whole-of-library. Whilst it might have been possible to hold separate some of the broader components of the library in different organisations and functional groups, the future state of the Library requires much greater dialogue about the dynamics produced by and between these components. It is less a case of adding AI-enabled capabilities to the existing library, and more a matter of fundamentally reconceptualising the whole-of-library as a complex system.

In fact, a systems approach is critical to many successful technology deployments, and given the complexities of the Library, we propose applying a cybernetic systems approach both to the technological capabilities themselves and the ways in which they will be deployed. This conceptual approach, which accounts for the interplay between the technological, human, and ecological, owes a great deal to the field of cybernetics and the theories of cybernetic systems (Figure 1).

Specifically, we analyse the NLA as a cybernetic system with a set of dynamic relationships between the technological, the human, and the ecological. And whilst there are many systems in the NLA, for the purposes of this report, we have explored the dynamics of the NLA cybernetic system from a collection-centric perspective.

Figure 1: The NLA is a cybernetic system.

1.1 A Cybernetic Systems Approach

"Practitioners of cybernetics use models of organizations, feedback, goals, and conversation to understand the capacity and limits of any system (technological, biological, or social); they consider powerful descriptions as the most important result."

Paul Pangaro
(2006, n.p.)

The need to design, build, regulate, and adapt increasingly complex systems emerged as a critical need in the 1940s: in the US, it was at places like Bell Labs, and the National Aeronautics and Space Administration USA (NASA) shortly thereafter (Hall 1962, Brill 1999); in Japan, and elsewhere, it was associated with post-war reconstruction and large-scale factory automation (Liker *et al.* 1995, Duke 2014). What these sites all have in common is large, diverse organisations attempting to create complex machinery with increasingly sophisticated computation power (Schlager 1956, Goode and Machol 1957, Hall 1962, Fagen *et al.* 1975). These systems were dynamic and expansive, more open than had previously existed and thus some of their properties would be emergent rather than established in advance.

To tackle the challenges of these complex systems that combined increasingly powerful computers with human processes and a great deal of large-scale infrastructure, systems engineers looked to cybernetics for its theoretical and methodological approaches. As a result, the field of systems engineering is deeply intertwined with cybernetics, especially around theories of control and communications, feedback loops and dynamics, and in the insistence that a system is more than just its technological parts.

In turn, cybernetics owes its own twentieth century intellectual roots to Norbert Wiener (1948, 1950) and to a remarkable series of interdisciplinary conversations in the 1940s and 1950s (von Foerster 1984, Heims 1993, Kline 2015, Pias 2016, Rid 2016). As defined by Norbert Wiener, cybernetics sought "to develop a language and techniques that will enable us indeed to attack the problem of control and communication" (Wiener 1948, p. 9). These problems, he believed, were uniquely raised by the development of computing in the 1940s – how we might control these new technologies, and communicate with and through them, built on Wiener's early work around feedback systems and control theory.

Cybernetics argued persuasively that one had to think about the relationships between humans, increasingly smart computing, and the broader ecological world as a holistic system with dynamic feedback loops. Cybernetics highlighted a wholly new kind of feedback loop that held together biological, technological, and human systems. It was a way of thinking about large, complex systems made up of parts that shape each other as they come together. Put another way, cybernetics is an "[a]ssembly of parts into a system – including human, environmental, and mechanical components – [which] generated new, emergent properties" (Bell and McLennan 2020, p. 200). For us, cybernetics is an approach which privileges the system as the core unit of analysis (rather than, for instance, a piece of technology), as well as a framework by which to encounter and re/structure such systems. Cybernetics is, of course, not without its critics and critiques especially around its totalising narratives and ambit claims as a general theory (Bowker 1993, Hayles 1999, Keller 2004, Malapi-Nelson 2017, Swann 2018) and we do not propose to resurrect mid-century cybernetics for this project.

However, a cybernetically informed approach to analysing and adapting complex systems, one which argues for human, technological, and ecological dynamics, seems a perspective that is especially pertinent to the challenges of the twenty-first century, and those of the NLA.

1.2 The NLA as a Cybernetic System

We are not the first to suggest that a cybernetic approach to libraries might be valuable. In 1967, two American archivists – Barbara Fisher from the Library of Congress and Frank Evans from the Office of the National Archives – created a bibliographic review of material related to digitally-powered automation in the context of archives and manuscript collections (Fisher and Evans 1967). They foresaw changes in the ways libraries collected material and made sense of them. In framing those changes, which included impact on retrieval and documentation techniques, they enjoined the archivists of the late 1960s to consider the importance of theories of control, data processing and system studies, and to acquaint themselves with cybernetics and cybernetic theory as a way to manage the future of technology in their collections and collection practices (Fisher and Evans 1967, p. 340).

It does not seem like a stretch to assert that libraries have always been cybernetic systems. They have produced and been produced by relationships between technological, ecological, and cultural forces; they are complex dynamic systems; and they require feedback loops and distinctive models of control and communication. The NLA is no exception.

To the casual outside observer, the NLA might rank among the most stable of national institutions. Housed in its current iconic building on the shores of Lake Burley Griffin for over fifty years – spanning the twentieth and twenty-first centuries – it is the archetype of stability, authority, and constancy. Of course, the NLA is more than just its impressive architecture and statutory framing. It is also a place, a collection, and set of practices and processes. Through interviews with Library employees and analysis of NLA materials, we came to realise that the NLA may be better understood as a constantly regenerating and emergent cybernetic system (Figure 2).

In this report, we posit that the NLA is a cybernetic system characterised by feedback loops and a set of dynamic relationships between the technological, the human, and the ecological. We believe that approaching the NLA this way is helpful as it lets us see the NLA as more than a site for technology adoption, and more than just a test bed for next generation AI capabilities. Comprehending and acknowledging this complexity allows us to offer a more holistic set of ways for the Collections Branch of the NLA to approach new and emerging technology capabilities.

Figure 2: Components of the NLA cybernetic system.

To facilitate this approach to the NLA, we have focused on five components of this cybernetic system: *people*, *data*, *process*, *infrastructure*, and *context* and their interrelationships (Figure 3). We represent these relationships by arranging these components into a star shape we have named the Cybernetic Star. The advantage of the star shape is that no one point is necessarily primary, and it can be rotated and rearranged to reflect different orientations.

From our research, we anticipate that these five components are key to understanding how a technology or capability might shape and be shaped by the dynamics between the components of the Library as a system. Exploring these interdependencies helps to not only think through the implications of technology transitions at the library, but also to interpret signals that can alert decision makers and teams at the library to shifts in the system, indicating possible opportunities, risks, pitfalls, and issues for the NLA to navigate.

After all, some of these different components might currently function discretely, and there might be challenges in aligning them or bringing them into dialogue. Furthermore, some of the emerging dynamics in this cybernetic system, especially as it expands to accommodate new AI-enabled capabilities, might render tacit or implicit knowledge more visible and explicit, and raise (new) questions about how to strike the right kind of balance.

Although these five components represent only one lens through which to view the NLA as a cybernetic system, our intent in using them is to draw focus to a network of relationships that will most clearly explicate the opportunities, risks, pitfalls, and issues in deploying AI-enabled technologies in the Library, as summarised in Figure 4.

Figure 3: A Cybernetic Star is one way to represent components of the NLA Collections cybernetic system.

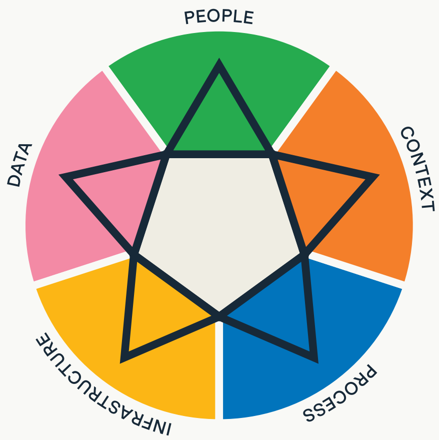

People: NLA audiences, workforce and third parties references in library materials.

Context: The external and intersecting systems, which may include global information ecosystems, that cross commercial markets, regulation, funding and more, which impact the NLA's operation.

Process: The chains of activity that happen in the NLA to deliver services. For example, the process that occurs when a collection item is accessioned and eventually accessed by NLA audiences.

Infrastructure: The information technology services in place to support NLA activities and services. The capability embedded in the NLA's infrastructure underpins the links between technologies and services.

Data: The information collected and stored by the NLA, including the records of the collection as well as information on NLA audiences.

Opportunities	Risks	Pitfalls	Issues
Advantageous possibilities afforded by the technology capability	Potential negative consequences of the *intended* application of the technology capability	Potential negative consequences of the *unintended* application of the technology capability	Known challenges with implementing the technology capability

Figure 4: Opportunities, risks, pitfalls, and issues (ORPI) outlined in this report.

1.3 General Dynamics of the Collections System

Viewed from the perspective of collections, the NLA is a cybernetic system that continuously ingests new items. Its employees grapple with making confident, informed decisions about what to collect and how various types of collection items proceed through idealised processes for acquisition through to discoverability and accessibility by audiences in the context of multiple temporal rhythms.

Because of these multiple temporal rhythms, the NLA occupies a curious temporal space. To deliver on its mission, it must be fully active in the present, while always accountable to the past and to the future. Indeed, one cannot escape the primacy of time and its passage when analysing the collections of the NLA; “We collect today what will be important tomorrow” reads the mission of the NLA’s *Collecting Strategy 2020–21 – 2023–24* (National Library of Australia 2020a).

The collections decisions made by the Library decades ago have dramatic repercussions and unexpected political effects today. The implications of activities and decisions made today may not be fully understood for decades. Because the Library is continuously ingesting new items, the collections of the NLA are like the river Heraclitus wrote of so long ago (as described by Barnes 1982): just as one never steps into the same river twice, one never twice encounters the same collection, for not only are other items continually flowing into the Library, but existing items are edited, made more accessible, and new connections made among them.

Examining some of the relationships among processes, people, infrastructure, data, and context in various “todays” and “tomorrows” in this mission statement leads us to a few of the most significant dynamics at work in this cybernetic system. It is important to remember that these general dynamics represent different ways to understand the work and functioning of the NLA – they are not the only dynamics, but they are the ones we believe are most relevant to collections.

Here we highlight five general dynamics at work in the Collections (cybernetic) system at the NLA. These dynamics are one way of expressing the kinds of opposing forces that any organisation needs to navigate and balance as they scale and grow. New AI-enabled capabilities will add pressure to these dynamics, and potentially necessitate new kinds of decisions, decision-making, and even decision-making forums. We have described these general dynamics by reference to the forces they attempt to balance: *Quality & Quantity*, *Nodes & Networks*, *Formats & Fluidity*, *Agents & Audiences* and *Serendipity & Indexing*. Each of these general dynamics involves, to varying extents, each system component. We have chosen to highlight a different primary component at the top of the Cybernetic Star for each as an entry point for considering the forces that animate changes in the system.

> Five general dynamics at work in the Collections system at the NLA express the kinds of opposing forces that the organisation needs to balance as it scales and grows. New AI-enabled capabilities will add pressure to these dynamics, and potentially necessitate new kinds of decisions, decision-making, and even decision-making forums.

Quality & Quantity (Data). There has been a remarkable increase in the volume of material being created, collected, and curated by the NLA; it has thrown into stark relief the complexities that come with trying to balance quality and quantity. As the NLA's collection has scaled in the past decades, regulating and upholding the standards of cataloguing and access from earlier eras has become unsustainable, and choices regarding updating processes (e.g. what metadata to prioritise), infrastructure (e.g. what new automation is available), people (e.g. how to prioritise work based on likely audiences) and context (e.g. new e-deposit agreements with partner libraries) have ensued. Inaccuracy in describing items is a recurring concern. Meta-data created generations ago by earlier staff members, or meta-data generated more recently through automated systems can inaccurately describe an item. As the collection grows, there is limited capacity to catch and fix these errors among staff and through crowdsourcing. Meanwhile, the increasing amounts of digital data being created and increasing amounts of storage capacity becoming available suggest that the scale of the collection could grow exponentially. The challenges of sustaining quality at scale feel acute.

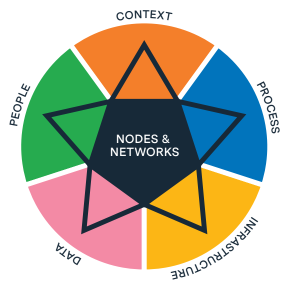

Nodes & Networks (Context). The NLA – and in particular Collections – is not a closed cybernetic system, in part because of the constant influx of new collection items, and in part because the context in which it operates is relatively open with multiple dependencies between it and other systems and organisations. Going forward the NLA needs to continue to find ways to balance its own needs against those of other organisations with whom it has relationships and the nature of the networks these organisations form. The NLA's pioneering work with Trove captures the collections of one thousand institutions into infrastructure run by the NLA. The NLA sits in a broad network of information institutions, characterised by standardised information and data sharing processes. Yet an ongoing challenge for the NLA is ensuring control over, and values alignment with, organisations (beyond other libraries) and technology systems that connect to the NLA's collection. For example, search engines like Google have the potential to provide increased ease of access to the NLA's collection by Australian and international audiences. Yet, over time, Google has limited the discoverability of Trove's data because the collection does not generate revenue for the company. Looking forward, the potential use of proprietary systems to perform increasingly sophisticated collections capabilities raises similar issues of values alignment. The NLA's capacity to create its own systems with advanced capabilities is frequently constrained by funding and resourcing capacity. Being connected to more organisations and tools that support the NLA's goals and collection strategies requires navigating the terms of those relationships to the Library's benefit.

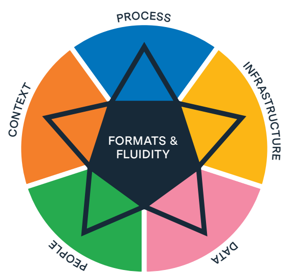

Formats & Fluidity (Process). Even as the digital side of the NLA collections grows in complexity, reach, and networked relationships, the physical content of the collection is also continually changing. This is not as simple as digital versus physical formats; formats change between physical and digital over time, and digital formats are intermittently updated along an item's journey in the collection, requiring new infrastructure and data, even sometimes a physical manifestation. Items are collected, archived, and made accessible in cycles as the ways the Library processes the collection transition over time. Connections between items in the collection emerge, re-emerge and multiply as people (employees and audiences) and infrastructure (new automated systems) add and edit over time. The public interprets and reinterprets the collection as they access it in different ways and in different contexts. Every facet of the collection will change over time.

This flux means that to make decisions, the Library must anticipate how the collection and the infrastructure that underpins it will be used in future cases; it must also maintain enough flexibility to manage through the ways that new tools might transform old tools and old data, and vice versa. Processes, data structures, and infrastructure need to be adaptive to multiple scenarios. NLA leaders recognise that capabilities such as data modelling are evaluated for their flexibility to address use cases that are unanticipated or not yet possible, and with an eye towards longevity of compatibility with tools and systems to access digital content (see also Owens 2018).

Agents & Audiences (People). At the heart of collections, and of the mission of the Library, is preserving materials created by Australians and making them accessible to audiences. As new processes and infrastructure make these materials – this data – more accessible, faster, and with more precision, issues regarding privacy and ownership of data are raised. The library already has an established contract with its audience guaranteeing that the data it makes available is a matter of public record. This suggests that the public, rather than individual members of the public, have rights to this data. As access to the NLA's data is amplified through increasingly accessible databases and search engines, this contract has become harder to navigate. The context when this contract was set did not anticipate the broad and deep accessibility to collection items to the extent that is possible now, nor the engaged audiences who may request (or act themselves through new means of involvement with collections) to obfuscate access to personally sensitive information. Creating digital forums in which the public engages with and acts upon the collection and with each other has led to unanticipated behaviours that have implications for the collection and the community of people around the NLA. People who were formerly strictly audiences are now at times agents who, rather than receiving collection items as described, indexed, and made discoverable by NLA staff and processes, now act – have agency – in these processes. The highly-engaged community of Newspaper Archive correctors on Trove is an early example of the types of agency in shaping the NLA collections that will characterise the next 20 years. Such agency will also extend to computing systems that can act upon and in conjunction with the NLA machine readable digital collections. These new contexts suggest that the NLA requires a renewed set of contracts with its audience as emerging technologies transform the ways people interact with each other and the collection.

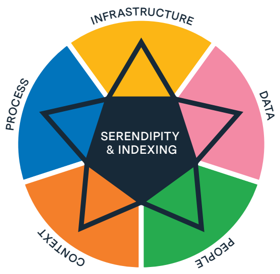

Serendipity & Indexing (Infrastructure). The NLA's goal is not only to collect and catalogue information, but also to create opportunities for people to connect information together into new strings of meaning. Returning to von Foerster's description of a library as a system in which employees "act the double role of a custodian of books (sic) and of a midwife for new ideas and insights" (1984, p. 212), libraries do not just hold information; they foster knowledge through the connections they enable among collection items. NLA employees speak of the delicate balance they must strike between 'mess' and 'control' in describing, organising, and creating tools to access the collections. We gloss these as fostering serendipity whilst still providing enough sense-making to ensure items are discoverable. Fostering discovery and new ideas and insights will require the design of new kinds of search tools. Today's tools emphasise wayfinding in a sea of data, which requires anchoring data to search paths that make it easy to find. This often means that each collection item has multiple layers of data attached to it. For example, adding geographic designators to existing articles helps connect items to a place and time. A data layer, such as this, helps people discover items using digital search capabilities. It also curates the items into networks that have more rigid connections between items than the connections created by indexing systems for physical books. If search tools increasingly curate the collection, this influences the ways that people who access the collection find their own narratives among the collection items. More rigid networks can mean fewer possible ways of stringing together items through an individual lens. The ways that emerging technologies continue to shape search will influence the ways that the collection is curated, and how porous the curation is to serendipitous discovery and emergent meaning making.

1.4 Conclusions: The Systems Scene is Set

In identifying these five general dynamics, we hope to make visible some of the landscape into which new AI capabilities might be enabled, and to explicate some of the conversations that such new capabilities might necessitate. We are not recommending the NLA attempt to resolve or remove these general dynamics. Rather, by explicitly articulating the forces which must be balanced and navigated – *Quality & Quantity*, *Nodes & Networks*, *Formats & Fluidity*, *Agents & Audiences* and *Serendipity & Indexing* – we hope to inform debate and considered decision making.

We set the scene with these general dynamics in this section and will return to them in more detail in our analysis of clusters of AI capabilities in Section 3. In Section 4, they inform the foundation of our technology reference guide. Prior to that, in the following section, we explore some of the ways these general dynamics came into being and some of the new pressures they might face.

Straddling the dual role of custodian of documentary resources relating to Australia and the Australian people, and of midwife to the discovery, learning, and creation of new knowledge by the Australian Community, the National Library of Australia is a constantly regenerating and emergent cybernetic system. It continuously ingests new resources and grapples with making confident, informed decisions about collections in the context of multiple time frames and feedback loops.

"Librarians are now engineers…"

Verner Clapp
(1951, p. 305)

"Today one can no longer think of just the computer. One must think in the more comprehensive terms of information technology or information systems … the computer is emerging from its glassed-in throne room, and as it becomes increasingly accessible to those needing its services, the links between it and society proliferate both in number and in complexity."

John Diebold
(1966, p. 130)

02

Automation, Digitisation, and AI

In Section 1, we introduced and elucidated the analytical framework informing this report – a cybernetic systems approach which we believe will help make sense of the Library as both place and practice. Having established that the NLA is usefully conceived of as a cybernetic system and having explored five general dynamics that animate change from a collections-centric perspective, here we turn to the technological innovation in and of libraries, and the pathways that have led to new AI capabilities in these systems.

In this section, we offer an analysis of the broader landscape of the history of innovation in the library sector, including early automation, digitisation, and contemporary AI. These histories are relevant to the ways in which the current general dynamics have emerged, and the additional pressures these dynamics might encounter in the future. We have included summaries from several recent reports on the topic of AI-enabled technology in libraries, drawing out some of the recurrent themes visible in the sector, and their consequences for the NLA.

The NLA has been on its own multi-decade journey of innovation and technological transformation, one that continues today (Figure 5). Its early history with mainframes and punch cards, and the initiatives to create various biographic reference standards, shared resources, scientific and technological information services and networks, plus its early embrace of the internet and digitisation have all helped propel the NLA into a global leadership position regarding library innovations and library futures (Bryan 1991, Fielding 1991, Webb 1991, Middleton 2006, Cathro and Collier 2010, Hider 2019).

The continual application of new technologies has shaped how libraries store, retrieve, and navigate information (Too 2010, Lyons 2011, Lee 2012). Libraries have also been sites of curating knowledge: deciding what to preserve and discard; deciding who does and does not have access to knowledge; deciding how access is granted; and deciding how such material might be classified and indexed (Dewey 1876, Chan 2007, Krajewski 2011). The technology capabilities that have been applied to storing, retrieving, and navigating information have also changed the way that libraries perform curation, in time and in space. Given the focus of many libraries on collection, preservation, and accessibility, it is hardly surprising that libraries have, historically, adopted and adapted a range of automation and digitisation technologies into their workflows, practices, and spaces.

Libraries have been built on, and benefited from, advances in techniques and technologies for storing, preserving, curating, circulating, and accessing material, as well as advances in techniques for producing the material itself. Innovations including the printing press, letterpress printing, typing and carbon paper, mimeograph, microfilm, offset printing, photocopying, transcription, optical character recognition, and scanning have had enormous impacts on the library's capacity to store and recall information – the complex interplay between how format influences retrieval mechanisms and vice versa. In deciding what would be collected and preserved, and how it might be accessed, when, and by whom, libraries and their library staff were indeed, to borrow from von Foerster, weaving the social fabric of the future (von Foerster 1984, p. 213). That weaving pulls together everything from the materiality of the collections, the spaces that are required to keep them, as well as the people, processes, and contexts.

Figure 5: A summary of the NLA's Digital Library Infrastructure Replacement (DLIR) Program (National Library of Australia 2017, n.p.).

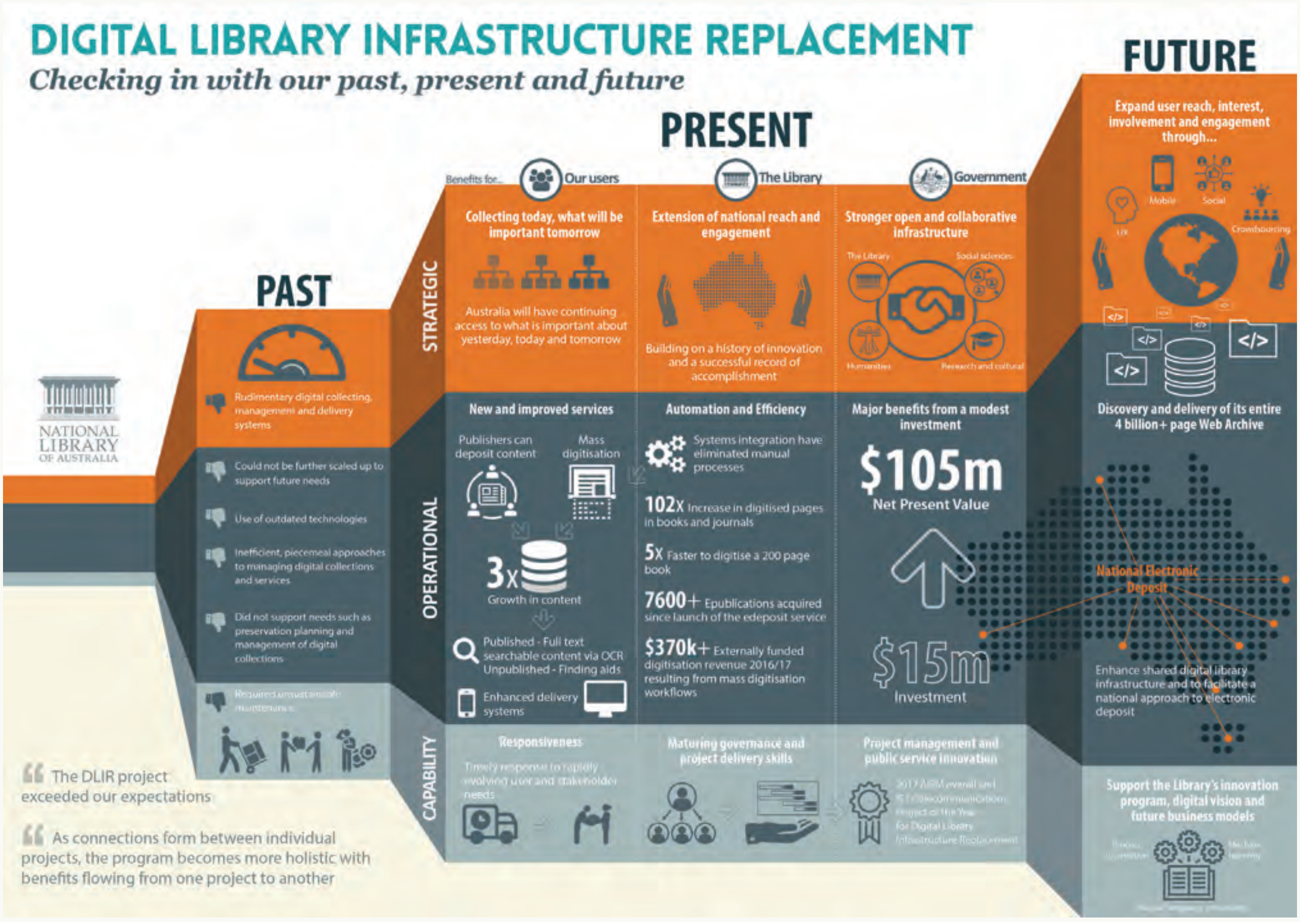

2.1 Digitisation and Automation

> "…automation is fundamentally a mechanical or electronic extension of the traditional logic – or deductive and inductive reasoning – to problem solving and the performance of work. Since archivists and manuscript curators have always used logic in establishing physical and intellectual control over the research materials in their custody, they should have an active interest in the application of techniques of automated logic to these specialised information sources."
>
> **Barbara Fisher and Frank Evans**
> (1967a, p. 334)

In 1957, a computer named EMERAC made history. It was, by all accounts, the first computer to star in a major motion picture (Morris 2017). EMERAC, or Emmy, was styled by IBM, which is listed as a consultant on the movie. The movie was, of course, *Desk Set* and the computer was intended to replace the library staff in a large media company's reference department (*Desk Set* [film] 1957). It was already a familiar trope – the automation of human labour and endeavour and its accompanying threat of job loss and economic peril. In *Desk Set* the work of converting knowledge to data is made surprisingly visible; and the notion that computing was about information, not insight, is one of the film's recurring themes. That Emmy is oriented to facts rather than sense-making is another persistent theme – hello curfew, not Corfu – and one that is both funnier and sadder now, as we have much greater insight into the perils of a data-driven world.

At the time, critics were not convinced that computers would have much impact on libraries or library staff, and many scoffed at the premise of the movie; let alone that any computer could best Katherine Hepburn. However, the arc of the second half of the twentieth century is one in which computers, and computing more broadly cast, have profoundly altered the practice and place of libraries (Hilton 1963, 1973, Kilgour 1970, Ceruzzi 1998). The information context has vastly scaled and automating storage, retrieval, and navigation has been the focus of innovating the library using technology. Perhaps the key dynamic that libraries have focused on with the advent of automation and computerisation has been a dependency between capturing increasingly large stores of relevant information whilst ensuring that the collection is still accessible to the public (see the NLA's general dynamics in Section 1.3 – *Quality & Quantity*; *Serendipity & Indexing*). To manage this dependency, libraries have become highly attentive to developing indexing systems at scale as a way to control large collections for ease of item retrieval.

Here, as Lang had suggested on the big screen in *Desk Set*, insights and collections and sense-making became flattened into ever-larger information sets and flows.

The embrace of computer-driven automation and digitisation is not without its challenges; whilst the work of computing and the work of libraries have often shared overlapping pre-occupations and concerns, there are also significant departures. To paraphrase Kelly (1992), computers want data and to perform complex calculations; libraries want collections and to curate and make meaning. The distance between these desires has expanded and contracted over the arc of the last sixty years, and libraries might presently find themselves less aligned to the desires of computers and computing rather than more.

For much of the twentieth century though, libraries certainly benefited from the ways in which computing, and computers, could rationalise, expand, and deepen various forms of indexing and storage. Of course, the idea of the index was not without its critiques. In his profoundly important essay "As We May Think", Vannevar Bush (1945) describes the "artificiality" of indexing as a mechanism to find meaningful insights in large stores of information. This essay, which ultimately shaped the work of everyone from J.C.R. Licklider to Doug Engelbart to Tim Berners-Lee (Waldrop 2001) and helped give rise to the internet, the personal computer, and the world-wide-web, provided a new way to think about indexing. For Bush, early twentieth century indexing worked on a linear spectrum to find more and more specificity, ordering items by strings of subclasses, while in comparison, the human mind works by making associations, which is more like drawing connections within a network of different ideas. Bush suggests an alternate approach to indexing would be to create "associate trails" of indexing (1945, p. 35), which represents a possible next innovation for libraries to make. Of course, in many ways, this "associative trails" notion captured the work that library staff have often done in navigating a collection – the associative trails were a sense-making practice, not an informational one.

Building on Bush's insights, but pushing them further than Bush might have intended, librarian Verner W. Clapp described a transition away from thinking about libraries as storehouses of knowledge and library staff as navigators of the collection who used their great knowledge of books and their contents to help people find their way through troves of knowledge, and towards libraries as organisations seeking control over flows of information (Clapp 1951). In his formulations, library staff were no longer experts on literature; they had become experts on the mechanics of how items of the collection flowed through processes of storage, preservation, and retrieval. From this perspective, Clapp argued, a library collection is best understood as data points and library staff as engineers of data that are flowing in and out of the library (Clapp 1951). This point of view was considerably strengthened with the ongoing commercial development of computing and the launch of more organisationally-friendly computing, like the IBM 360 in 1964 (Kilgour 1970, Ceruzzi 1998).

This imagining of the library was continually reinforced by emerging capabilities of mechanising data management. In the 1960s, as computers becoming more ubiquitous, mechanisation and automation became hugely influential trends for library staff to cast forward what the future of libraries might be (Egan and Shera 1949, Clapp 1964, Fisher and Evans 1967). Groups such as Mechanization and Automation in American Libraries within the USA Chapter of the International Federation of Library Associations (IFLA) were established, in which thinkers such as Clapp influenced conversations to unpack the possibilities of emerging technology on libraries (Fisher and

…computers want data and to perform complex calculations; libraries want collections and to curate and make meaning. The distance between these desires has expanded and contracted over the arc of the last sixty years, and libraries might presently find themselves less aligned to the desires of computers and computing rather than more…

Evans 1967b, McCallum 2003). Discussions at the time centred on technological capabilities, including: more efficient storage of information; less mechanical work to maintain a manual storage system; formats such as microfilm that enabled large sets of information to be compressed; faster retrieval achieved through search filters; and faster, cheaper facsimile production (Clapp 1964).

For Clapp and many others, the value that lay in these capabilities was the ability to circulate more information at cheaper cost, which was the basis for achieving easy access to information for everyone – a democratic ideal. Of course, what is lost in these techno-deterministic narratives is the idea of sense-making, and certainly of serendipitous discovery (see the NLA's general dynamics in Section 1.3 – *Serendipity & Indexing*). The underlying metric in these narratives is about ease of retrieval and access, not necessarily the depth or possibility of the insight. It is a deeply reductive understanding of library functions, and of the value of its collections.

Throughout the latter half of the twentieth century and early twenty-first century, the way that libraries have considered applying new technologies follows this emphasis on information being enabled to flow beyond the library to people who would use it (see the NLA's general dynamics in Section 1.3 above – *Agents & Audiences*). The key parameter around introducing new capabilities, therefore, has been to enhance access to information more so than gaining other outcomes. Along this trend, libraries have broadly focused on transitioning to new data formats, considering how to organise and mechanise their data, and responding to user reactions to automation (McCallum 2003).

Post-1970s, this approach expanded to thinking about libraries within networks of other libraries. As countries built national networks, libraries began to concentrate on forming networks both locally and across the globe. This led to new communities working to develop standard data processes to enable trans-border flows of information between libraries. In the 1990s, this approach expanded to encompass the internet, particularly with the development of web application; this, in turn, allowed libraries to become increasingly networked and porous in terms of the flows of information between them (see the NLA's general dynamics in Section 1.3 above – *Nodes & Networks*). Trove is an excellent example of what this capacity has ultimately led libraries to achieve (Holley 2009, 2010, Neilson 2010, Ayres 2013).

In the late 1990s, "Digital Libraries" were the next formulation of the fusing of information and communications technologies with the library; the overlay of the internet and the continued expansion of networked consumers led to a new conception of libraries as being fully integrated into massive flows of information (Cleveland 1998, Borgman 1999, Schwartz 2000, Lynch 2002, McCallum 2003, Calhoun 2014). The "Digital Library" represented an important shift in how libraries made sense of new information and communication technologies: it was not entirely a discontinuous break with the past automation practices. Collections – or at least their digital instantiations – were no longer quite as located in time and in place as they once had been; plus, they were differently accessible. In 2001, the Australian National Archives launched a digitisation-on-demand service which built on its early *PhotoSearch* online digital image service. While the public immediately embraced it (Ling and Mclean 2004), issues of data privacy, cultural sensitivity, and ease and timeliness of access all surfaced quickly. Beyond the new capacities for access, there was also recognition that digital formats were not only a method of preservation and communication of large volumes of information; collection items were increasingly born-digital (see the NLA's general dynamics in Section 1.3 above – *Formats & Fluidity*). This transition meant libraries needed to start thinking about collecting, storing, and providing access to born-digital material, which gave rise to new sets of capabilities, such as web-scraping, that characterise the mechanics of contemporary libraries. It is against this backdrop of changes in both the material and practice of collections, and sense-making within these, that AI arrives in the library.

2.2 The Arrival of AI

> "The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."
>
> **John McCarthy et al.**
> (1955, p. 2)

Computing in the 1950s was torn between the necessary but banal application space of increasing calculations, speed and storage, and the fantastical possibilities of replicating or mirroring human intelligence. For the men who gathered at Dartmouth College for the world's first conference on AI in 1956, they looked to the latter as their vision of how computing might unfold. For them, it was about a kind of automation and replication of the human intellect. This vision of AI, as unlikely as it may seem, was built on much of the work and insights of the contemporary cybernetics conversations (Rid 2016, Bell and McLennan 2020, Bell 2021a, 2021b). Indeed, many of the attendees at the 1956 event had also been part of the cybernetics debates in the preceding decade. However, their AI departed from those conversations in two critical ways: it had not been conceptualised as a system per se, and it was silent to issues of people and their environments; or put another way, it focused on the technological and ignored the cultural and the ecological.

It took more than fifty years for the AI described back in 1956 to seem viable, and now over the last decade, AI-enabled technologies have started to significantly change the way people, civic and civil institutions, and commercial enterprises access and process information, among other things. Various AI deployments are already powering the growth of several large multinational companies, as well as forming the backbone of various governmental agendas and real-world, state-initiated deployments. However, the residual effects of a technology-centric regime developed without a clear reference to cultural and ecological frames have been clearly felt. As a result, in the past few years there have been a range of critical commentaries exploring those who build AI systems, the places and corporations in which that is transpiring, and the kinds of algorithmic regimes that are being created therein (Ananny 2016, Buolamwini 2016a, 2016b, 2018, Finn 2017, Ananny and Crawford 2018, Buolamwini and Gebru 2018, Noble 2018, Besteman and Gusterson 2019).

Much of the current debate is characterised by a persistent slippage between AI, machine learning, and algorithms. Furthermore, there is a tendency to describe AI as though it were singular and monolithic. It is important then to pause and redefine terms. In a recent Australian governmental report, AI was defined as "a collection of interrelated technologies used to solve problems that would otherwise require human cognition. Artificial intelligence encompasses a number of methods, including machine learning (ML), natural language processing (NLP), speech recognition, computer vision and automated reasoning" (Walsh *et al.* 2019). Here it is important to see that ML is defined as a subset of AI. Elsewhere Kate Crawford, Meredith Whittaker and their colleagues have defined AI as "a constellation of technologies, including machine learning, perception, reasoning, and natural language processing" (Crawford and Whittaker 2016, p. 2). In both instances, AI is a constellation of technologies that entail sensing, sense-making, and actuation, but it is not necessarily seen as a system or treated as such.

The appeal of AI-enabled technologies to the library sector is intense, and it has been on the radar for several years (International Federation of Library Associations 2013). Fortunately, it appears to go beyond the imagination of AI as simply next generation automation and digitisation technology to something far more generative and rich. The new wave of AI-enabled capabilities suggests a different endpoint might be possible for automation in libraries than what was previously possible through the twentieth century. Rather than focusing on the efficiency of storage, retrieval, and preservation in the context of scaling digital collections, new AI-enabled capabilities suggest an opportunity for managing the complexity of large digital collections, particularly regarding enhancing the capabilities of discoverability. These possibilities are also couched in challenges. Some of these challenges are overtly familiar to information scientists, such as managing bias and assuring accuracy in managing data.

Many libraries already have experience with limited AI-enabled technologies. This includes experience using AI to produce metadata, manage collections, conduct discovery and search, and to extract knowledge (EuropeanaTech AI for *et al.* 2020). A recent survey of European libraries indicated that of these areas, libraries were most interested in further using AI to enable knowledge extraction, or "for facilitating the exploitation (and to some extent, the production) of their digitized collections" (EuropeanaTech AI for *et al.* 2020, p. 8). Whilst there has been a lot of interest, to date, the majority of applications of AI to library collections have been characterised as experimental (Jakeway 2020). A survey of leading research on university libraries in the United States and Canada in 2019 found that engagement in projects or labs related to AI has been "startlingly" low (Wheatley and Hervieux 2020, p. 353). That said, there are a surprising number of encouraging examples. It is interesting to contemplate what might happen here as an increasing array of ordering and ontological schemas are layered over the collection: How will they be rationalised? Do they need to be rationalised? And what new kinds of knowledge, tools and experiences might be produced as a result of these new kinds of practices?

Several national libraries have begun experimenting with AI-enabled technologies applied to their collections. Many of these experiments have been led by national library labs. For example, the Library of Congress Labs (LC Labs) sponsored a "Season of Machine Learning" in 2019, which included partnering with the University of Lincoln-Nebraska to conduct explorations of ML and collections management (Lorang *et al.* 2020). LC Labs also hosted an "innovator in residence" to explore the visual and textual content of its newspaper datasets using ML (Lee 2020). The British Library Labs have also partnered with universities to experiment with technological infrastructure and digital content to benefit humanities researchers (Baker 2015). In the Netherlands, the Cultural AI Lab, a partnership between the Koninklijke Bibliotheek (the National Library of the Netherlands) and other institutions was recently established to "harness the potential of AI for cultural research and make the technology aware of cultural context" (Innovation Center for Artificial Intelligence 2021, n.p.). Other national libraries have embedded AI-enabled technologies in core functions of their collections branches: Kansalliskirjasto (the National Library of Finland) developed "Annif" using a combination of existing natural language processing and machine learning tools, while Nasjonalbiblioteket (the National Library of Norway) has experimented with applying ML to automate Dewey Decimal classification and uses AI to assist with analysing the collection (identifying, for example, similar books and similar images) (Wetjen and Brygfjeld 2018, Suominen 2019).

It is also worth noting that AI capabilities are being adopted in contexts well beyond that of collections. Some libraries are also using these new capabilities for aspects of management and engagement. For example, Singapore's National Library Board uses predictive modelling to forecast demand for its collection, guiding the distribution of stock to each of its 25 public libraries (SAS Denmark n.d.). The Cambridge Public Library in Massachusetts, Massachusetts Institute of Technology Libraries and Harvard metaLAB have hosted an installation entitled "Laughing Room," in which patrons can engage with AI in a room that plays a laugh track whenever something is said that the algorithm in the room deems funny (The Laughing Room 2018). It is likely that we will continue to see more AI capabilities at play and at work in libraries more broadly in years to come.

Alongside changes to services and processes of the traditional library, AI-enabled technologies will change the capabilities required of library staff. Here too, a cybernetic lens might be helpful (Beer 1959, 1972, 1985, Dubberly *et al.* 2014). This will have workforce implications for libraries, with some authors suggesting that library technicians, library assistants, and library staff more broadly could be replaced with robots within twenty years (Arlitsch and Newell 2017) – it is *Desk Set* (1957) all over again.

Some have criticised the library sector for "complacency" with regard to the potentially transformative and disruptive effects of AI, suggesting this reflects a tendency in the sector to wait until technology reaches market saturation before reacting (Wood and Evans 2018, p. 9). This hesitancy has largely been attributed to a lack of staff expertise or experience with AI and ML (Cordell 2020, Wheatley and Hervieux 2020, p. 353). Analysts have noted few library or information science programs offer training in AI and ML methods, trained ML specialists command substantial salaries elsewhere, and training within the library requires significant time and money (Cordell 2020). As a result, while many library staff do not have an active understanding of what constitutes AI or how to identify potential AI applications (Hervieux and Wheatley 2021), there is a significant opportunity to engage more deeply and there is optimism regarding the possibilities of AI technologies for the sector.

The new wave of AI-enabled capabilities suggests a different endpoint might be possible for automation in libraries than what was previously possible through the twentieth century. Rather than focusing on the efficiency of storage, retrieval, and preservation in the context of scaling digital collections, new AI-enabled capabilities suggest an opportunity for managing the complexity of large digital collections, particularly regarding enhancing the capabilities of discoverability.

2.3 Three Key AI Reports (2019–2020)

It is clear from the broader context of public and private sectors that AI best practices are an evolving area of concern and opportunity. It is also clear that the regulatory frameworks in Australia and globally are in a state of flux, in response to changing community and societal expectations and experiences.

That said, some practices could be established, even as the broader context continues to evolve. Three reports delivered in 2019-2020 – one by the International Federation of Library Associations and Institutions (IFLA) (2020), one for the Library of Congress (US) (Cordell 2020), and one for OCLC, Inc. (Padilla 2019) – considered the roles of AI and ML in the context of National Libraries. Each points the way to various activities and actions libraries might take, including statements of value regarding ML and AI, increasing data, ML and AI awareness and literacy among library staff, and developing more crisp articulations of the ways in which libraries and library staff understand data and its various discontents. These reports represent significant bodies of work and bear further analysis. We have provided a brief contextual summary and identification of the key threads most relevant to the NLA in Sections 2.3.1-3.

Embedding AI responsibly is always contingent on understanding the context of where AI is embedded, and of how an AI system is shaped by and shapes this context. Generalised approaches to understanding the implications of AI are always limited.

2.3.1 IFLA Statement on Libraries and Artificial Intelligence

The IFLA Governing Board released a *Statement on Libraries and Artificial Intelligence* in September 2020 (International Federation of Library Associations 2020). The Statement outlines considerations for the use of AI-enabled technologies in the sector and provides guidance to libraries on the attendant risks of AI applications for libraries. This includes challenges to standards of intellectual freedom for researchers, data privacy, and data bias. The Statement suggests providing AI literacy to patrons as one way to manage these risks.

This report emerges in line with a discourse on AI and ethics that is converging on a set of principles for 'ethical AI' (see Fjeld *et al.* 2020 and Section 2.4 of this report). These principles emerge from technology-centric contexts and while they provide a relevant pulse-check on the types of concerns and issues emerging in building and deploying AI-enabled technologies, they are not a holistic guide to responsibly, safely, and sustainably taking AI to scale. Embedding AI responsibly is always contingent on understanding the context of where AI is embedded, and of how an AI system is shaped by and shapes this context. Generalised approaches to understanding the implications of AI are always limited and this is true of how the Statement thinks about the risks of AI.

While each of the concerns raised are important for libraries to delve into, they are more complex and situational than the Statement indicates. For instance, thinking about libraries in the same way one might think about social media is only useful to a point, because the intents and drivers that inform these systems and their users are different. Overall, each of these risks requires a greater level of specificity in defining what the implications are for the library setting.

Regarding data privacy, for example, it is important to think about *whose* privacy and in *what* context. As Cordell makes clear in his report on ML and libraries (discussed below), capabilities that monitor, evaluate, and deploy machine learning on user information, such as search behaviour on the collection, are places where libraries need think about managing privacy, as well as the transparency of these operations. This includes thinking about ways for users to have agency while using library systems, for example through opt-in/opt-out functionalities. From speaking with the NLA, we also recognise that in certain situations, privacy of individual users is a lesser concern than public record. For example, a person featuring in a collection item may not want the information findable on Google but depending on the content and its significance for public record, the patron's privacy might not always be privileged. Discerning and addressing the complexity of general risks inherent in AI as they relate to the library context will be key to managing them.

2.3.2 Machine Learning and Libraries: A Report on the State of the Field

In 2020, the Library of Congress (US) commissioned Ryan Cordell, an Associate Professor of English at Northeastern University, to provide a "wide-ranging view into the current applications and practices of applying machine learning in libraries and other cultural heritage organizations" (Cordell 2020, p. v). The Cordell Report, published in July 2020, provides a survey of the field, covering cautions for ML in libraries, promising applications of ML in libraries, and challenges for implementation. The report concludes with a set of practical recommendations for libraries and other cultural heritage institutions seeking to undertake ML experiments.

While not the main argument of the work, Cordell's report outlines an inflection point in library automation: a shift in emphasis toward enabling serendipity in searching large digital collections rather than a focus on enhancing only retrieval capability. Cordell demonstrates that this shift in emphasis is a direct outcome of the possibilities that ML and Machine-Actionable collections offer (see Section 3.2.3). He provides several examples of where these capabilities enable audio, visual, tabulated, and textual formats to be searched for correlations between like items, which allows users new ways to navigate digital collections and new ways to make sense of digital collections. Like Vannevar Bush's imaginings of the future of automation, this suggests a function of digital search that is more akin to associative connections, and therefore more in line with the intent behind ancient libraries as places that help people to navigate stories in the collection.

Cordell's thorough exploration of the opportunities and challenges of machine learning and machine-accessible collections raise further lines of enquiry that will be important for libraries to consider regarding these capabilities. One opportunity for further enquiry relates to control: How much control should library staff have over facilitating sense-making of the collection? How much autonomy should ML have over facilitating this sense-making? How will audiences at the library assess the outputs of ML and perceive levels of accuracy? What bias is present and how will this be signalled? At what points can library staff intervene in these systems to manage the collection with oversight?

Beyond search, Cordell also points to other applications of ML that have potential in the library. One is to link collections using metadata. Another is to apply ML to collection management processes, which could include using ML to help 'weed' the collection (take out duplicates etc.) or using ML to monitor and predict the deterioation of physical objects using computer vision and ML models. These examples indicate the possibilities of ML that are currently at the periphery of concern among library staff compared to the potential applications of ML to aid a more sophisticated approach to discoverability. However, applications of ML to automate laborious tasks relating to collections management may become more prevalent as more technology applications are developed and become available.

Regarding the immediate challenges of deploying ML in libraries, Cordell's report outlines a range of necessary considerations. One is producing Machine-Actionable data, which is the basis for enhancing serendipity in search through Machine-Actionable collections. Cordell notes that standard training datasets typically focus on contemporary materials and, often, dominant linguistic or cultural groups. For the library context, where it is important to navigate collections that contain data across time periods and beyond the dominant culture, ML systems need to be trained on diverse data sets that include content from multiple time periods and contexts if search capabilities are to find more accurate, meaningful narratives in the collections.

Cordell's report also raises an interesting challenge for incorporating ML in libraries, which relates to how we perceive the role of libraries compared to other information archives. In particular, we expect libraries to have data and metadata that is high quality, and we accept that both are created slowly. ML generates useful data and metadata quickly and at a larger scale, but with flaws; therefore, while ML offers efficiencies that library staff have hoped for and anticipated in managing digital collections through automation, the readiness of ML in library contexts still needs further development before it can meet library staff expectations of the quality of data that libraries give access to. This will most likely come through continued experimentation with ML in libraries, led by library staff.

In regard to the readiness of libraries for ML, Cordell suggests that there two challenging aspects to consider before deploying ML today. First is the cost of infrastructure, specifically ML-capable hardware. This is currently a barrier to ML being deployed in libraries, though Cordell notes that costs are decreasing rapidly. The second is developing staff expertise in ML. While most library staff interviewed for the report had a general understanding of ML applications, and a desire to explore ML methods with their collections, many had less understanding of the full ML workflow. As a result, Cordell points to a need for library staff to develop further expertise on ML workflows.

2.3.3 Responsible Operations: Data Science, Machine Learning and AI in Libraries

Published in 2019, Thomas Padilla's "'Responsible Operations: Data Science, Machine Learning and AI in Libraries" was commissioned by the OCLC (formerly the Online Computer Library Center) to form the blueprint for a research agenda "to help chart community engagement with data science, machine learning and artificial intelligence" (2019, p. 7). The Padilla Report is particularly focussed on the growth of digital collections as a proportion of the whole collection and the ways that emerging technologies amplify, and in some cases transform, what these new automation capabilities can achieve for libraries as well as addressing the associated risks. In thinking about the risks and opportunities, the Padilla Report insists that libraries cannot separate the technological details of implementing AI from discussion of the effect of AI on the responsibility of the library.

Padilla suggests that library staff should orient their focus towards three broad goals: engaging with challenges and opportunities attendant to scaling digital collections; enacting diversity as a broad principle; and, managing the unequal distribution of emerging technology capabilities across libraries. In addressing these goals, Padilla participates in a broader movement to think about ideas like scale, bias, and the unequal distribution of new technologies more explicitly.

Padilla highlights several considerations that relate to scaling digital collections in the context of increasing the use of data science, machine learning, and AI in libraries. For example, he discusses enhancing description at scale by automating semantic metadata using a series of capabilities such as computer vision to describe video materials, full-text summarisation using machine learning, genre determination in addition to text material description, and speech-to-text transcription for audio. These capabilities enable the possibility to create links between data in the collection, from which unanticipated research questions can emerge. He also discusses governing collection use at scale through aggregating collections from different libraries and making more content available. Here, complexity has emerged regarding legal rights, where one access portal can encompass competing legal paradigms that govern the use of content.

To address the issue of the inherent biases that emerge through treating digital collections as data, Padilla acknowledges the role libraries have always played in navigating biases in the collection. As bias is enacted through capabilities such as machine learning in libraries, Padilla highlights that driving a culture that privileges diversity is necessary to continually manage bias. The activities that fall under this banner include things like ensuring community representation, enhancing the capabilities of technology applications so that they work effectively beyond the dominant culture (for example, natural language processing that works for languages other than English), and building a workforce that has diverse perspectives as well as the capacity to work across disciplines.

Padilla is also attentive to the implications of emerging technologies being embedded in libraries at varying rates, levels of sophistication, and with various applications. He suggests that centring capabilities such as data science, ML and AI will inevitability occur, but navigating a context where these capabilities are peripheral requires sharing experience across libraries, and defining data standards and best practices that can lead broader practices in information contexts, even beyond libraries.

2.4 Responsible AI Systems: An Emerging Conversation

Globally, there have been significant conversations regarding the emerging AI regulatory, legislative and policy frameworks and framers (see for instance Dignum 2019, Turner 2019, Wischmeyer and Rademacher 2020). These conversations extend beyond conversations about AI and ethics to include notions regarding safety, responsibility, and sustainability, among other issues. It seems clear that these broader conversations are increasingly preoccupying regulators, governments, and industry and they allow for a more nuanced and generative framework that addresses AI-driven systems as well as specific AI applications.

In the Australian context, there has been a great deal of activity over the last three years regarding approaches to ensuring responsible, rights-respecting, and socially beneficial use of AI capabilities. In November 2019, the Australian Government in collaboration with CSIRO's Data61 published the *Artificial Intelligence (AI) Roadmap*, identifying strategies to help develop a national AI capability (Hajkowicz *et al.* 2019). The Roadmap spans a range of issues associated with AI, including ethics and trust in AI systems; its focus is primarily on driving productivity and research innovation in AI, using AI to solve problems at home, and exporting these solutions to the world. The roadmap focuses on 'natural resources and environment', 'cities and towns', and 'health, ageing and disability' as three areas of specialisation for Australian researchers and industry.

The Australian government also released *Australia's Artificial Intelligence Ethics Framework* in 2020 which offered a voluntary set of eight principles to help guide the design, development, integration, and use of AI to ensure better outcomes, including the reduction of negative effects and the practice of the highest standards of ethical business and good governance. Several large Australian and international companies are committed to trialling these principles, including National Australia Bank, Commonwealth Bank, Telstra, Microsoft, and Flamingo AI.

In addition to these documents, the Australian Council of Learned Academies (ACOLA) horizon scanning report on *The Effective and Ethical Development of AI* (Walsh *et al.* 2019) was published, as was the Australian Human Rights Commission's *Human Rights and Technology Discussion Paper* (Farthing *et al.* 2019) and final report (Farthing *et al.* 2019). The latter report calls for an AI Commissioner, as well as a moratorium on the use of facial recognition in high-stakes situations until adequate regulation is in place. It also calls for legislation to ensure accountability around AI-informed decision-making. Here, the focus is on tangible actions and accountability frameworks, as well as a clear articulation of the perils of particular and specific technological components of an AI system in real-world applications.

Conversations about responsible applications of AI have been amplified by the actions and inactions of various commercial entities and trans-national companies. Emerging from this discourse is a set of sectoral norms that cluster around eight principles for Responsible AI systems (from Fjeld *et al.* 2020):

1. *Fairness and non-discrimination*: AI systems should be designed and used to maximise fairness and promote inclusivity.
2. *Privacy*: AI systems should respect individuals' privacy.
3. *Accountability*: Accountability for the impact of AI systems should be appropriately distributed, and adequate remedies provided.
4. *Transparency and explainability*: AI systems should be designed and implemented to allow for oversight.

5. *Safety and security*: AI systems should be safe, perform as intended, and be secure.
6. *Professional responsibility*: Calls on the professionalism and integrity of individuals involved in developing and deploying AI systems to ensure appropriate stakeholders are consulted and long-term effects are planned for.
7. *Human control of technology*: Important decisions should remain subject to human review.
8. *Promotion of human values*: The ends to which AI systems are devoted, and how they are implemented, should correspond with core human values and promote humanity's well-being.

These converging principles suggest a recurring series of dilemmas emerging in building and deploying AI systems, relating to common technical constraints, organisational motives, and cultural values that are mostly linked to the places where current AI systems are being built and deployed. They offer some useful insights for thinking about managing some of the emergent challenges of deploying AI systems, but they are not exhaustive.

Early generations of automation and digitisation have helped support the library as a complex system; now the promise of AI-enabled technology capabilities is to amplify the Library's capability to manage this increasingly complex, open system within a world of similar systems. To do this successfully, AI technology capabilities must still allow for the collections to maintain a state of flux. Planning for this outcome requires thinking about the Library as a whole system that has transient boundaries to the world beyond, and that is shaped by relationships between context, people, and technology. Transforming the ways we can collect and organise information requires grappling with how people and context interact with, and are changed by, technology. It requires a cybernetic approach.

Nevertheless, these principles could provide a valuable starting point for the NLA to frame its own conversation about, and its approach to, responsible AI.

Of course, these principles will need to be considered within the context of the library sector, which is already alive to concerns around the ways in which AI and ML could perpetuate or amplify existing bias (Lee 2020). NLA staff similarly shared this concern with us, reflecting an understanding that the decisions of libraries around what materials to digitise often reflects the biases and priorities of researchers and audiences from dominant groups over others (Cordell 2020).

In this broader formulation of responsible AI, there are also debates about individual technology capabilities, for instance, facial recognition (Crawford *et al.* 2019, Whittaker *et al.* 2019, Partnership on AI 2020) and the appropriateness of continuing to support its development and use (Keyes 2018, *All Things Considered* 2018, Crawford 2019, American Civil Liberties Union 2020, Conley and York 2020, Greene 2020, Farthing *et al.* 2021). In 2020, major American technology companies, including IBM, Amazon, and Microsoft decided to cease supporting facial recognition systems to American police departments until federal laws were put in place to regulate their use (Greene 2020). Beyond the particulars of specific AI-enabling technologies, the COVID-19 pandemic has tested regulatory frames regarding data and data privacy, and the trade-offs between data privacy and public health (Kemp and Greenleaf 2020, Bell 2021c, 2021b). These new dynamics, whilst still nascent, could have implications for collections processes in the library sector, as well as some of the temporal dimensions of those processes.

There are also other emergent regulatory frames. In July 2020, the Aotearoa New Zealand Government launched a first-of-its kind charter which is "a commitment by government agencies to carefully manage how algorithms will be used to strike the right balance between privacy and transparency, prevent unintended bias and reflect the principles of the Treaty of Waitangi" (StatsNZ 2020, p. 1). In so doing, the charter brings together various legislation

and principles in New Zealand regarding data, privacy, and AI but also highlights the importance of First Nations.

There is also a small but growing body of work regarding Indigenous AI, including research being conducted concerning the role of decolonisation of technology and data sovereignty in Australia, New Zealand, America, Canada and parts of Africa, as well as proposing different kinds of conceptual and epistemological ways to approach ML and AI (Kukutai and Taylor 2016, for example Abdilla 2018, Global Indigenous Data Alliance 2018, Kesserwan 2018, Lewis 2019, Mohamed *et al.* 2020, Goodchild 2021). This work includes new Indigenous protocols arising for emergent AI capabilities (see for instance Jade 2021), including for augmented, virtual and mixed reality, that build on earlier protocols regarding photography and other imagery. Again, this is an area where further attention and engagement will be required.

As the NLA grapples with all these new emergent AI capabilities, and how they might think about their own ethical, moral, and responsible approaches, the specific context of the library remains a critical frame; after all, libraries are not the same things as large multinational corporations or other commercial enterprises. As Cordell (2020, p. 1) reminds us, it is important that libraries implement AI-enabled technologies “through existing commitments to responsibility and care” and statutory obligation.

2.5 The Future(s) of the Library

Libraries have served to shape broader forms of collective sense-making, such as shared imaginings of history, science, technology, politics, and culture, not to mention those regarding the nation-state, countries and even empires. As a result, the library, as an idea and a place, has always been an investment in, and a vision about, the future.

Perhaps it is unsurprising then that those libraries have also come to represent sites of anxiety, conflict and competing views of that self-same future. Indeed, libraries have formed part of very many imagined futures, some of which centred on the disappearance of books, or library staff, or even the buildings themselves. Others have focused on the future of the twin roles of the library and imagined new futures for custodianship and midwifery, with familiar socio-technical imaginaries of distributed knowledge, robotic systems, and democratised access. Lang's *Desk Set* ([film] 1957) is an early example of the "robots will take our jobs" storyline in the context of a reference library, for instance. Those various kinds of speculative futures found their way into government reports, academic publications, movies, and even the Sunday comic sections of daily newspapers where, in the 1950s at least, consumers were alive to the idea that libraries might change (see Griffen 1987, as well as Figure 6).

Figure 6: The Electronic Home Library as imagined by Arthur Radebaugh in the Chicago Sunday Tribune in 1959 (Radebaugh 1959).

Throughout their history, libraries have been shaped by, and have in turn shaped, their broader social, political, and cultural contexts. As centres of knowledge, libraries have sometimes been complicated extensions of public intellectuals and their capacity to weave associations between items of knowledge and tell stories of history and science (Ong 1982, Blair 2011). They have been impacted by various forms of regulation, social transformation, and upheaval. Libraries have found themselves transformed and remade by the changing economics of publishing and the confluence of emerging middle classes and leisure time, which gave rise to subscription and lending library business models and gendered models of engagement with libraries (Scott 1986, Watson 1994, Raven 2006, Bowman 2017). In Australia, the remarkable expansion of the Mechanics' Institutes and their attendant libraries spoke to a very particular set of social impulses (see Kelly 1952, Rochester 1990, Lowden 2006) and shaped the ways Australian citizens understood and imagined library collections, in sharp contradistinction to the Carnegie model of public libraries in the United States and the United Kingdom (Prizeman 2012). Of course, the history of libraries in Australia is also a history of the federating states built on earlier colonial arrangements, where the diversity of library systems and processes recall the complexities of other state systems, for instance the railways (Biskup and Goodman 1982, Rochester 1990, Biskup 1994). The state libraries were modelled on the British Museum Library and saw themselves as 'national' institutions, with a mandate to collect 'books of all languages and descriptions'. Until the 1950s they remained the backbone of the Australian library system. By 1962, with the expansion of university education, the holdings of the university libraries for the first time equalled the combined resources of the state libraries and the NLA, which was not without its own particular complications and history (Biskup 1983, Biskup and Henty 1991, Bryan 1991).

Laws and regulations changed the frame too. Those deliberately targeting libraries, such as the various public library laws like the *British Public Libraries Act 1850* (Minto 1932, Murison 1971, Hoare and Black 2006) or similar laws in Australia, for instance the *NSW Library Act 1939* (Remington and Metcalfe 1945, Jones 2005, Maguire 2016), sought to position libraries as centres of class mobility and regulation. Others shaped the world in which libraries existed (i.e. mandatory schooling laws, and changing labour laws and opportunities). More recently, regulations and laws regarding privacy, data, security versus national security, and even algorithms and facial recognition have all impacted the ways libraries can and do function.

Libraries have also changed daily life and shaped the way we design technologies to capture, store, organise, search, and represent information. Libraries are the driving metaphor for how we view "indexing" on the internet and are therefore a kind of blueprint for some of the most transformational technologies we use today. And here is the thought experiment: if library indexing has informed the structure of search on the internet, how might the library inform new ways of connecting these large caches of information to generate meaning? The processes, structures, networks, and relationships that make a library have a lot to recommend to the way that we might build AI-enabled systems to help generate meaning out of vast datasets.

It is important to remember that even the idea of the library is, in and of itself, a remarkable innovation. Although early conceptions of the library were tied to metaphors of human memory, and reach, and the ability to store and recall information, the sheer scale and enduring capacity of the concept of a library is noteworthy. In its early instantiations, the library concept was born through both the physicality of such libraries as Alexandria with its expansive collections of books from around the world (Too 2010), as well the assertion of a sense-making schema and techniques such as those offered in *Etymologie* by Isadore of Sevilla, which proffered alphabetisation as a novel ontology

(Blom 2005). In this way, the library as both place and practice transformed the scale and conception of information that could be both preserved and accessed, and upon and through which meaning can be made.

Early generations of automation and digitisation have helped support the library as a complex system; now the promise of AI-enabled technology capabilities is to amplify the library's capability to manage this increasingly complex, open system within a world of similar systems. To do this successfully, AI technology capabilities must still allow for the collections to maintain a state of flux. Planning for this outcome requires thinking about the library as a whole system that has transient boundaries to the world beyond, and that is shaped by relationships between context, people, and technology. Transforming the ways we can collect and organise information requires grappling with how people and context interact with, and are changed by, the technology. It requires a cybernetic approach.

Future steps for the NLA may be to pivot from thinking about AI only as something to be implemented at the Library towards considering the role of the NLA in shaping the future of AI. This suggests that focusing on the ways emerging technologies will shape libraries is only one side of the coin; the other would be to focus on how libraries will shape emerging technologies – in particular, how we think about them and how they are built. As von Foerster (1984) helpfully reminded the librarians-to-be at the University of Wisconsin all those years ago, libraries need not just be consumers of other people's solutions, they can lead the conversation on how these technologies might evolve too.

> Libraries are the driving metaphor for how we view "indexing" on the internet and are therefore a kind of blueprint for some of the most transformational technologies we use today. The processes, structures, networks, and relationships that make a library have a lot to recommend to the way that we might build AI-enabled systems to help generate meaning out of vast datasets.

"The role of the archivist and manuscript curator, thus transformed by electronic recording and transmission equipment of millisecond speed and virtually unlimited storage capacity, becomes as a consequence an increasingly dynamic one."

Barbara Fisher and Frank Evans
(1967, p. 335)

"The future is already here – it's just not very evenly distributed."

William Gibson
(2003)

"We collect today what will be important tomorrow."

National Library of Australia
(2020b, p. 1)

03

Building the Future Library

In the next half century, the AI capability clusters explored in this section will potentially radically transform the NLA, enabling a more open, more complex, and more inclusive library than has been possible to date. The scope and edges of the cybernetic library will be more diffuse, and its publics, as agents and audiences, will engage with the NLA in settings far removed from the shores of Lake Burley Griffith, and far from the computers, tablets and smartphones that mediate their digital library experience today. The collection will engage Australians in streets, in courts, in bush and more.

In the previous sections, we established the analytical perspective of the NLA as a cybernetic system and explored the history and dynamics of technological innovation in libraries, ending with the adoption of recent AI capabilities in the global library sector. In Section 3, we explore the most likely dynamics that will emerge when the NLA integrates solutions involving four very different clusters of current AI capabilities into their collections system. For each, we also provide a summary of the attendant opportunities, risks, pitfalls, and issues these new dynamics will generate.

We focus specifically on four very different clusters of AI capabilities:

- Optical Character Recognition (OCR)
- Machine Transcription
- Machine-Actionable Collections (MAC)
- Transparent AI

By AI capabilities, we refer to both the services that a class of technologies provide, as well as knowledge that can be applied to specific technology products and processes.

Certainly, these are not the only capabilities that may be relevant to the NLA. These specific areas were prioritised by the NLA Director of Collections, in consultation with our researchers, as most relevant to the branch's remit to collect, preserve, and make accessible library materials. Moreover, whilst these are all AI capabilities, they are a diverse constellation of technologies, approaches, and processes, so exploring them has broad applicability.

In exploring these four clusters of capabilities, we draw on a matrixed analysis of materials drawn primarily through several research methods. First, through desktop research on AI capabilities relevant to the creation and management of collections. Second, through a review of existing NLA materials related to its mission, strategies, organisation, and workflow processes. Third, through expert interviews with NLA employees. Fourth, and finally, through the creation of speculative future scenarios using the method of science fiction prototyping. For more information on these methods and the research and analysis activities associated with each, please see Appendix A.

In our analysis, we have prioritised clusters of AI capabilities rather than specific products and processes in the market or that are deployable now. Doing so enables us to outline a pathway for the NLA to evaluate opportunities, risks, pitfalls, and issues with AI over the time scales and rhythms that characterise the library as a cybernetic system.

3.1 Current AI Capability Clusters and the NLA Collections

We start with Optical Character Recognition (OCR), a set of capabilities the NLA has over a decade of experience deploying at scale with the newspaper digitisation project. OCR is thus both a familiar starting point and a useful test case of the fundamental challenges and reconfigurations of the collections dynamics that will emerge when the NLA integrates Machine Transcription, Machine-Actionable Collections, and Transparent AI capabilities. A cybernetic analysis of OCR and identification of the opportunities, risks, pitfalls, and issues it introduces is a model for preparing for the next onslaught of capabilities that will take the library from its current dual incarnation as analogue/digital, brick and mortar/online to a much more complex, open, and porous cybernetic system.

We will move from the specificity of OCR to the superset of media-to-machine readable data and metadata covered by Machine Transcription capabilities. The examples provided in this portfolio of capabilities can be imagined as younger siblings in the same family of media analysis and transformation capabilities. As the NLA moves to render other data legible to machines, some of the same ORPI will arise as with OCR, but others will be different – much as siblings may be similar but are ultimately unique.

To return to – and play a bit with – von Foerster's imaginings of the role of the library in the future social fabric of society, for OCR and Machine Transcription, the NLA collections system will act primarily as *custodian*. Capability pathways are relatively well formed and easily projected into the near and mid future; care in thoughtfully ushering and integrating them into the existing systems dynamics is called for. Together, these two clusters point towards how the NLA can integrate AI capabilities to transform analogue collections and born-digital collections into a single, digitised collection that can act and is actionable.

The last two clusters oblige the NLA collections system to act more as *midwife* to relatively less well-formed *approaches* and *processes* for managing the data, meta-data, and associated digital processes with the fully digitised and machine-readable collections promised by the application of the first two clusters. Considered as a whole and imagined in a general continuum from now and near, to emerging and further, the AI capability clusters we analyse below suggest the possibility of radical transformation of the NLA cybernetic system into a new system configuration that meets a broader, larger, and more inclusive audience than has been possible to date (Figure 7).

Figure 7: Evolving the library. A summary of pathways to the future through four clusters of AI capability. NLA collections system will act as custodian and midwife depending on the technological capability being imagined and implemented.

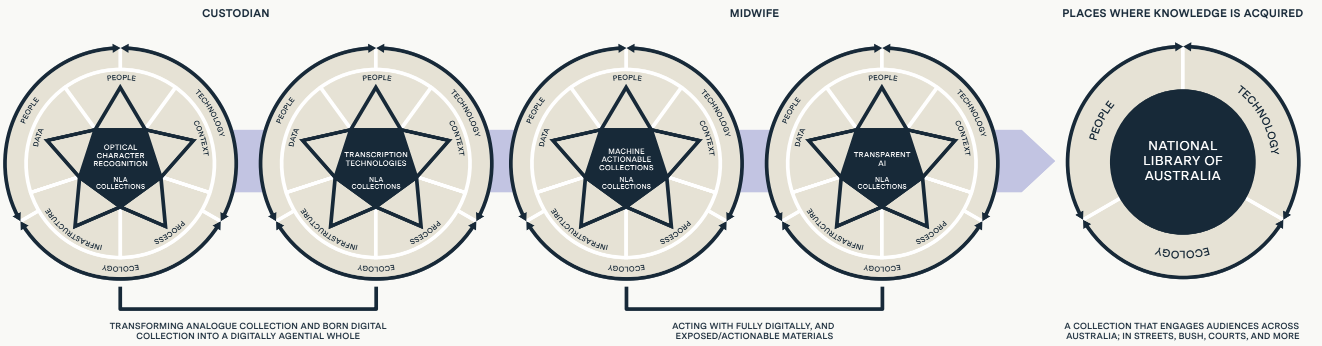

3.2 A Cybernetic Systems Analysis of Capability Clusters

For each cluster, we begin with a speculative glimpse into the future, generated through science fiction prototyping methodologies. We start with these visions of the longer-term outcomes of the incorporation of capabilities into the NLA collection workflows to illustrate the opportunities for new services and experiences with NLA collections. These new experiences and opportunities will be enabled by increased content, and new ways of discovering, accessing, making connections among, and interacting with, content in the Library's collections. A more detailed unfolding of the creation of each of the long-term speculative visions, including inputs for developing each, and a template for the NLA to create additional visions, is contained in Appendix B.

After defining the cluster and providing a brief overview of the history of the capabilities and their applications in the library sector, we turn to an analysis of how the capabilities may be integrated into the NLA's cybernetic system. We start with describing the most salient aspects of *context*, *people*, *data*, *processes*, and *infrastructure* in relation to the capability, and note which of the general dynamics of the collections system are most relevant to an analysis of the capability cluster (see Figure 8).

We then turn to specific applications of the capability in the library sector or in the NLA and evaluate the potential new dynamics among people, processes, infrastructure, data, and context that will emerge if, and when, the NLA integrates these capabilities into approaches to collect, preserve and make accessible Library materials. These new dynamics inform a summary of the risks, pitfalls, and issues with incorporating current and future capabilities into the cybernetic system of the NLA.

Figure 8: General dynamics of the collections system reviewed in this report.

Scenario 1:
Deakin Speaks, World Listens

A lawyer representing the experimental neural implant company Affect Link in the High Court of Australia is trending worldwide today after calling a digital avatar of Alfred Deakin, the second Prime Minister of Australia, as an expert witness in a constitutional law case against the Commonwealth.

The lawyer claims that the avatar, trained on notes from the constitutional convention debates and Deakin's personal records held by the National Library of Australia, can represent the original intention of the founders of the Australian Federation.

The lawyer for the Commonwealth argues that there is no way of assuring either the accuracy of the avatar or the underlying data. "We have no way of knowing if that machine dotted every 'i' and crossed every 't'," the lawyer told the court. "A stray apostrophe could render unto Caesar that which is due elsewhere."

This news follows unverified claims leaked from Parliament earlier this week that the current Prime Minister, Karen Wong, has consulted with the National Library's Robert Menzies avatar on policy issues.

Affect Link is opposing the Commonwealth's attempt to legislate against communication via neural implant. The case turns on whether the Affect Link device, which allows one user to share a physical sensation with another, falls under the Commonwealth's constitutional power to regulate "postal, telegraphic, telephonic, and other like services".

In the year 2035, artificial agents trained on the personal data of individuals are a compelling proxy for interacting with the person themself. Of course, these artificial agents are only ever as good as their data sets, and tiny errors in transcription could have significant ripple effects. And yet we still feed them everything: old-fashioned text- and image-based data like emails and social media posts, voice recordings from smart speakers, gesture tracking from wearables and even dreamscapes captured by experimental neural implants are among the data sources used to train digital avatars that stand in for a real person.

This trend began with small, personal projects. In the late 2010s, a handful of technologists were reported to have trained text-based ‘chatbots’ on the personal data of deceased friends and family members. In 2020, Microsoft was granted a patent for a similar method, though they expanded the range of targets to include any “past or present entity ... such as a friend, a relative, an acquaintance, a celebrity, a historical figure,” who would be rendered in 2D or 3D animation. Squeamishness over the ethics of such an undertaking, though, meant that it was not Microsoft but a start-up, Gemini, who offered the first commercial avatar building service in 2025. Gemini’s offering was extremely expensive and restricted to relatively young people for whom a vast corpus of born-digital data was available. But it launched a market that, over the next 10 years, drove significant improvements in avatar technology. This, combined with improvements in machine learning and synthetic data techniques, means that developers now claim to be able to build realistic avatars from pre-digital data like printed newspapers and handwritten notes.

By the late 2020s, the most common form of avatar was the so-called ‘grief bot,’ an interactive memorial for the living to converse with the dead via social media or an electronic messaging service. There are now a range of platforms offering this service. Grief bots are relatively affordable compared to other forms of avatar because they require less data and less training to be made fit for their private purpose. At the higher end of the market, high-fidelity avatars are built to represent their living subjects publicly. For example, in 2028 the CEO of Amazon began offering executive coaching services through her avatar. Her advice to young entrepreneurs? Invest early in your own avatar so you can be in two places at once. In 2031, a tech billionaire combined the two concepts of grief bot and digital representative, appointing his avatar as an executor on his will. The legality of the appointment is yet to be tested.

The NLA was among the first libraries to experiment with offering digital avatars as a novel way of interacting with their collections, partnering with a successful Australian high fidelity avatar provider, Perpetua. This caught the attention of an aging Australian novelist, who offered his personal data to the library along with an exclusive license to his identity by way of bequest, with the stipulation that it only be made publicly available after his own passing and the passing of his six children. Further, the current prime minister has made it known that they plan to donate their identity to the library for this purpose. Already, the NLA and Perpetua have created digital avatars of five deceased Australian prime ministers. Indeed, exclusive identity licenses have become a promising source of alternative revenue for the NLA, which has also sought and established similar agreements with other public figures.

3.3 From Future to Present with Optical Character Recognition

> The creation of machine-readable text from manuscripts, government records, newspapers and typewritten speeches is the first step in radical new experiences with the library's collection. Optical Character Recognition, a set of algorithms which detect and extract letterforms and other written markers from digital images and convert them into machine-readable text, is the starting point of this journey.

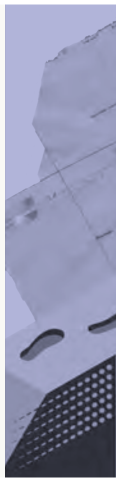

The nucleus of this speculative glimpse into the future of the NLA is today's AI-enabled OCR capabilities, some of which are already in use at the library. In the *Deakin Speaks, World Listens* scenario, information created in the final decade of the nineteenth century during the constitutional convention debates, preserved first on paper and later 'freed' from digital image to machine readable text, forms the basis for what we now may consider extraordinary applications.

Yet the technology alone is insufficient: It is the considered forethought of NLA collections staff to "collect today what will be important tomorrow" that makes this scenario feasible. The creation of machine-readable text from manuscripts, government records, newspapers and typewritten speeches is the first step in radical new experiences with the library's collection. These experiences are built on new syntheses of existing materials, and user interfaces for connecting the past with the future.

3.3.1 What is Optical Character Recognition?

Optical Character Recognition (OCR) refers to a set of algorithms which detect and extract letterforms and other written markers from digital images and convert them into machine-readable text. In this section, we focus on OCR as applied to typewritten text and not handwritten text (sometimes referred to as intelligent character recognition), which is discussed in the next section, the broader superset of Transcription Technologies (see Section 3.4.1).

The early history of OCR is equally enmeshed with the automation of business processes and with broadening access to written texts for the visually impaired. OCR emerged from telegraphy in the early nineteenth century and found more stable form in the twentieth century (Schantz 1982). In 1914, Emanuel Goldberg developed a machine that read characters and converted them into standard telegraph code – it was a form of automating and standardising information creation and input (*ibid.*).

Over the last century, OCR systems have reached wider audiences and helped to not only automate and standardise but also to make information more accessible. In the 1970s, Ray Kurzweil's invention of omni-font OCR systems was motivated by a particular focus on creating computing systems that could read digital texts and, in combination with a text-to-speech synthesizer, read them out loud to visually impaired users (Kleiner and Kurzweil 1977). Since the early 2000s, OCR has become more widely available to consumers through cloud-based services, accessed through desktop or mobile applications, although it may not always be obvious. For example, Adobe Acrobat offers support for OCR on any PDF file, and Google offers real-time translation of foreign-language signs, built on OCR.

OCR was first used by libraries for newspaper digitisation in the early 1990s. This included projects at the British Library with the Burney collection, and at the NLA with the Ferguson collection; both were unsuccessful as a result of early challenges with OCR and historic newspapers (Holley 2009). In the early 2000s, the NLA commenced the Australian Newspaper Digitisation Program, working with OCR contractors to undertake mass digitisation of historic Australian newspapers (National Library of Australia 2021). The NLA was a pioneer in crowdsourcing of corrections to post-OCR text through Trove, recognising a desire in the Australian public to contribute and participate in the work of its national library (Ayres 2013).

OCR is a subsection of the broader AI field of computer vision, which is concerned with creating computing systems that can sense and understand information contained in digital images (both picture and video) and actuate based on these understandings. Current OCR solutions are predominantly based on supervised machine learning; an OCR engine is trained on images in labelled documents and then used to detect text in images contained in unlabelled documents (Cordell 2020).

The general workflow for OCR proceeds across four stages: (1) Pre-processing of a scanned image (i.e. de-skewing, de-noising, binarization), (2) Layout analysis (i.e. region, line, and word segmentation), (3) Application of OCR (i.e. feature detection and extraction, and classification), and (4) Post-processing (i.e. automated and manual error detection). While this idealised workflow takes a scanned image through to exporting machine-readable text, there are a range of other systems and processes which OCR interacts with in any application context. For example, at the NLA, the OCR workflow interacts with Trove and its various component systems during the post-processing stage – through crowdsourcing of manual corrections. OCR also impacts other systems that may rely on the output of this workflow, for example making collections Machine-Actionable for NLP tasks (van Strien *et al.* 2020).

There are a number of critical questions the NLA should explore with regard to OCR. These include:

- What are the benefits and limitations of OCR (see: *Quality & Quantity* dynamic)?
- Can changes within established workflows improve accuracy of transcriptions for digitised newspapers and journals in the collection (see: *Formats & Fluidity* dynamic)?
- How does digitising information change audience expectations (see: *Agents & Audiences* dynamic)?
- How does OCR influence sense-making of collections (see: *Serendipity* & *Indexing* dynamic)?

3.3.2 Integrating OCR into the NLA's cybernetic system

How might we make sense of OCR in the context of the NLA? Here we turn to a cybernetic analysis of these capabilities in the system of the library. To start, we outline the aspects of each system component – context, people, data, infrastructure, and process – and the general dynamics among them that we consider important to understand in relation to OCR at the NLA. We then consider the emerging dynamics among system components through three examples of the application of OCR capabilities in the library sector. We consider potential changes to dynamics along a continuum from the *internal-facing* (implementation issues, knowable risks, automation, and workforce) through to more *audience-facing* (ethical concerns, unintended consequences, impact on users). We also consider potential changes to dynamics and feedback loops along the multiple time rhythms that rule collections work at the NLA.

CONTEXT	• OCR has been available to the public via cloud-based services since early 2000s. • The most significant use of OCR at the NLA to date is the Australian Newspaper Digitisation Project. • Globally, many other libraries also use OCR to digitise historic newspapers, and other historic materials. • Increasing focus in the NLA on promoting access by the public to the Library's collection, not only researchers. • NLA is well known in the global library sector, and among Australians, for its crowdsourcing of error corrections in newspapers.
PEOPLE	• NLA's audience: newspaper digitisation and search on Trove using OCR accounts for 80+% of search. • Existing Library staff manage quality assurance of OCR. • NLA has a dedicated and engaged audience who correct OCR errors on Trove.
DATA	• Predominantly based on supervised machine learning – OCR engine is trained on images and labelled documents to then detect text in images contained in unlabelled documents. • Analogue items can be challenging for audiences to search, especially audiences beyond trained researchers. • The quality of data collected by the NLA on historic materials is affected by scanning (resolution, brightness, straightness, discoloration etc.) and textual issues (layout, font type and size, contrast etc.). • Pre-processing, layout analysis, and the application of OCR engines can address some of the quality of data issues, though many errors remain.
PROCESSES	• Inaccuracies generated by OCR affect the application of other techniques or capabilities further downstream e.g. named entity recognition is increasingly difficult with inaccurate text produced through OCR. • Post-processing includes some automated error detections, but most OCR workflows still require manual detection and correction. • Platformisation of historical newspapers and other items through OCR changes the reading experience and the materials that the audience most engages with.
INFRASTRUCTURE	• Building domain-specific AI requires building one's own models or fine-tuning existing models. • OCR interacts with Trove during the post-processing stage. • OCR impacts systems that rely on its output downstream, e.g., making collections Machine-Actionable for NLP.

3.3.3 Salient General Dynamics for OCR Capabilities

Collections staff at the NLA now have over a decade of experience identifying, processing, and assuring newspaper materials that have been that have been run generated through OCR workflows. While OCR has enabled the NLA to digitise and make accessible a larger corpus of newspapers, and recently journals, than would have been possible manually, the transcriptions that result from OCR are less accurate than would have been achieved manually. In this way, the *Quality & Quantity* dynamic manifests in the tension between the priority of staff for slow but accurate processes for data and metadata tagging, and the fast but flawed processes of OCR for data and metadata tagging. Navigating this dynamic for the application of OCR to materials others than newspapers and journals requires the NLA to distinguish between those that require manual recording for the highest quality, and those materials for which that discoverability would be enhanced through mass processing by OCR.

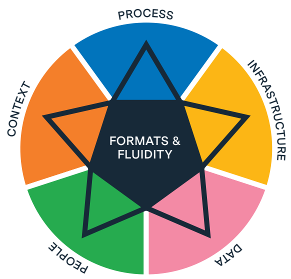

The workflows established by the NLA for OCR of newspapers and journals – from scanning materials through to exporting text – provide guidance for the implementation of workflows for other transcription technologies (see Section 3.2.2). Equally the links between OCR and other processes within NLA, and in the broad ecosystem, demonstrate how digitisation may be tied to publication, for example through Trove. However, improving the accuracy of transcriptions for digitised newspapers and journals in the collection will require changes to steps within established workflows and the reapplication of OCR to those materials. Further changes to workflows are likely required for materials not yet digitised. The fluidity of the workflow for OCR will continue as other techniques and methods are applied downstream to the outputs of OCR, requiring additional or different data and metadata to be added to materials produced through OCR workflows. These ongoing updates and improvements, to OCR workflows and the formats they generate, demonstrate that the growing breadth and depth of the collection demands continual review of the application of any AI-enabled technology within the library.

OCR is the first step in the Library's process of digitising and publishing a grand collection of newspapers and journals, ubiquitously accessible through a single platform. For OCR, the *Agents & Audiences* dynamic manifests primarily in two ways. First, digitising newspapers and journals drives the expectation of audiences for access to other materials in the collection. While newspapers and journals are generally intended for a broad, public audience when published, other materials – typewritten, handwritten and in other formats – may not. Where applying OCR to new materials that may have been intended for a private audience, the NLA will need to consider additional steps to negotiate public access. Second, digitisation fundamentally changes how an audience engages with the collection through the platformisation of once-physical materials.

The digitisation of newspapers and journals has resulted in an explosion in interest in those materials, from researchers and the public alike, in Australia and internationally. The platform has increased demand for digitised materials, in turn increasing demand for OCR of physical materials, which will further increase engagement with digitised materials over their original physical sources – a reinforcing feedback loop. This prioritisation of digital over physical resources – if unacknowledged – has the potential to change how researchers, and the public alike, treats materials in the collection. This is exemplified in the *Deakin Talks*, *World Listens* scenario, as the typewritten notes from the constitutional convention debates – secondary sources of Alfred Deakin's contributions to the debates – and Deakin's own personal records – primary sources – are treated equally as training data for the avatar that represents Alfred Deakin in the High Court.

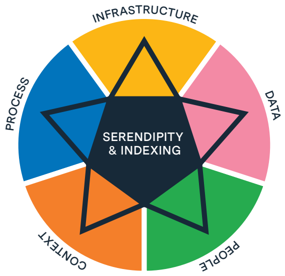

OCR has enhanced discoverability in the library by enabling keyword and full text search of newspapers and journals. This increases sense-making, as the NLA expands the corpus of materials in its collection for its audience to connect. Each item is ascribed multiple layers of data and metadata that relate to its creation, processing, and publication, enabling multiple possible layers of connection between a growing number of items in the collection. These layers of data, when combined with the search capabilities of Trove, form a rigid network of newspapers and journals in the collection, directing the audience swiftly and clearly to items of interest. OCR does not, however, promise to support the audience to form connections between the items. Forming new and/or unexpected connections requires support from staff, or the expertise of a user themselves. The challenge of connecting items in the collection to create new meaning is not addressed by OCR: indeed, it may be exacerbated as OCR creates rigid networks from the layers of data ascribed to items. In short, while OCR may provide more materials for the audience to connect, it does not offer capabilities for the audience to make new connections between items.

3.3.4 Current and Projected Applications of OCR Capabilities

OCR aids discoverability across many library collections around the world. The capability is widely deployed across the sector, overcoming much of the cost and time that would be required to manually transcribe typewritten text from large collections of images. OCR is also an example of how libraries have balanced the desire of staff for highly accurate data with the drive from audiences for enhanced discoverability. OCR demonstrates that technological capabilities in the library sector are dynamic – they require ongoing updates and improvement, and consideration for application to domains broader than their initial implementations. These updates and improvements are particularly important for OCR, as the text it outputs is the basis for other capabilities further downstream. As such, the quality of OCR that may have been acceptable at one time for certain purposes may not be acceptable at another time as those purposes change or grow.

Using two current examples, and projecting forward to mid- and far-future applications, we consider the potential new dynamics that will emerge if the capability is further developed or applied to other materials and the attendant possibilities, risk, pitfalls, and issues each example presents.

Example 1: Improved accuracy for historical materials

Historic materials often have low paper quality, use various historical fonts or require the recognition of complex page layouts to separate text from images and tables. As a result, applying OCR to materials published prior to the early twentieth century results in lower accuracy than can be achieved for materials from the mid-twentieth century onwards (measured, generally, by the character error rate). To date, the NLA has focused on post-processing steps to correct errors for historic materials. This has included outsourcing to low-cost, offshore services to manual correct titles, subtitles, and the first four lines of article text for over twenty-one million newspaper pages (as of July 2016) (Bode 2016). It has also included crowdsourcing corrections through Trove; already by July 2013 “more than 100 million lines of text had been corrected through crowdsourcing” (Ayres 2013, p. 1).

Many library digitisation projects use a single OCR system, with few pre-trained font and layout models that may not adapt to the range of character forms and complex page layouts found in historical data. Since 2015, however, the German OCR-d project has focused on gathering, combining and improving various OCR software components to integrate them into a single, open-source OCR workflow (Neudecker *et al.* 2019). The project is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) and led by a consortium of partners across the library and university sector: The Berlin-Brandenburg Academy of Sciences and Humanities, the Herzog-August Library Wolfenbüttel, the Berlin State Library and the Karlsruhe Institute of Technology. The project focuses on historical prints, with workflows that can be fine-tuned for recognition of specific corpora. This contrasts with many commercial offerings that, while easy to use, are less customisable and focus on modern prints. Integrating multi-source layout and recognition tools into scalable OCR workflows focuses more attention on the stages *prior* to post-processing stages to improve accuracy for historic materials.

In Scenario 1, *Deakin Talks, World Listens*, artificial agents trained on text from historic materials that has been detected and extracted by OCR are held out as accurate representations of those materials. The reliability of these future applications depends on very high accuracy of OCR, though present systems cannot produce such accuracy from historic materials. Additionally, manual post-processing of materials run through an OCR workflow, while applied to a larger set of materials at the NLA relative to many other libraries around the world, has resulted in the review and correction of only a small proportion of the materials generated through OCR at the Library. Existing, singular systems and the focus on manual post-processing, then, will not be sufficient to enable the kinds of applications envisaged in the *Deakin Talks, World Listens* scenario.

Drawing on the OCR-d and related projects to combine software components may improve the accuracy of OCR for historic materials earlier in the OCR workflow, reducing the reliance on post-processing for corrections. The NLA may also consider contributing to the OCR-d and related projects using the large corpus of historical materials the NLA has already run through OCR workflows and corrected through manual post-processing (including crowdsourcing). This could assist the creation of datasets for training and evaluation of new OCR workflows across the sector. This would support the consolidation of the NLA's position as a leader in the Galleries, Libraries, Archives and Museums (GLAM) sector with respect to sector-wide innovation to improve the accuracy of OCR.

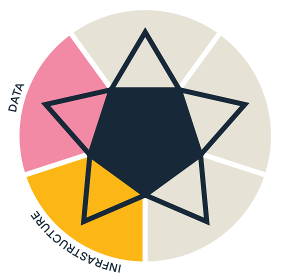

Shifting the emphasis from manual to automated post-processing of text generated through OCR may also increase accuracy on a greater set of materials than is presently possible through existing workflows. This has been a focus of recent library and information science related scholarship, for example the International Conference on Document Analysis and Recognition (ICDAR). The accuracy of error detection achieved in winning entries to the competition at the ICDAR 2019 Competition on Post-OCR Text Correction was above 95%, while error correction remained low (<50%) (Rigaud *et al.* 2019). This indicates an opportunity to incorporate new techniques at the post-processing stage of the OCR workflow to automatically detect errors in materials formerly digitised with older OCR workflows, as well as newer, challenging documents. The NLA could also, then, engage the crowdsourcing community through semi-automated systems. For example, a semi-automated system might identify materials with significant error rates and present them to a person for correction. Alternatively, if automated error correction capabilities were included in a system, the system might allow a person to select the right correction within a list of system-generated candidate corrections.

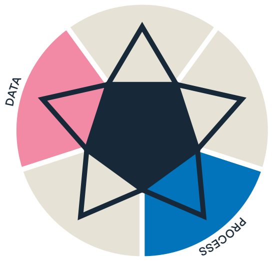

Example 2: Training domain-specific OCR

OCR was originally developed to process English-language documents. Large training datasets of the English language have developed over time, as well as language models for post-processing, resulting in high accuracy of OCR for documents of the same type. However, for documents in languages other than English and multilingual documents, particularly those that do not use Latin script, error rates are higher. The accuracy of OCR for any document depends on the training data available, and therefore the development of more robust and diverse collections of training data is required to reduce error rates for many languages other than English.

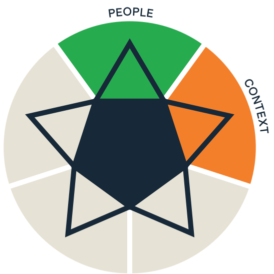

Some national libraries are seeking to improve discoverability for documents in their own languages and for multilingual documents by focusing on domain-specific training and modelling of OCR engines. For example, since the late 1990s. the National Library of Finland (NLF) has digitised historical newspapers, journals and ephemera published in Finland using OCR (Kettunen *et al.* 2020). Recent statistical analysis of the resulting corpus suggests that approximately 70–75% of the words were likely right and recognisable; of approximately 2.4 billion words in the collection, 600–800 million words were incorrect (Kettunen *et al.* 2018). Materials published in Finland are produced in the country's two official language: Finnish and Swedish. Thus about half of the materials digitised by the library are in Swedish, particularly those created in the nineteenth century when Swedish was the main publication language of newspapers and journals (Kettunen *et al.* 2018). Discoverability of materials in Finnish and Swedish was reduced because of erroneously identified text in these languages, particularly relative to English words. The NLF therefore started compiling its own sets of training data on Finnish language and Swedish language words to improve OCR for Finnish and Swedish language newspapers and journals (Kettunen *et al.* 2020). Over the last five or so years, the NLF has used these training datasets to re-OCR the data in the collection by fine-tuning an open-source OCR engine called Tesseract (Kettunen *et al.* 2020). Ultimately, the open-source engine was selected as the cost of re-OCRing the collection with a proprietary engines such as ABBYY FineReader was deemed too high, and the open-source engine was more customisable (Kettunen *et al.* 2020).

In Scenario 1, *Deakin Talks, World Listens*, future applications rely on domain-specific training of artificial agents using text generated from OCR. Under the supervised machine learning approach to OCR, the artificial agents discussed in the scenario would require training data sets specific to any individual that is intended to be represented, in their language and reflecting their orthographic style. This may be understood as an extension of a trend towards domain-specific training of OCR engines using training datasets specific to a language or group. However, where the current approach to computational linguistics says that "more data is better data" (Kettunen *et al.* 2020), the future applications envisaged in the scenario may require representation from training datasets decreasing in size (i.e. from one person, at one time of their lives).

The NLA has some experience with OCR for languages other than English, for example the digitisation of the Bulletin of the Indonesia Resources and Information Programme, named "*Inside Indonesia*", which required an Indonesian language plug-in at the OCR processing stage of digitisation (Berthon and Hickie 2018). However, the NLA may consider developing domain-specific training datasets to fine-tune OCR engines for languages other than English, in areas that align with its strategic goals. This may provide a complementary opportunity for the Library to interrogate the process for training and deploying an OCR engine, to better understand the stages of an OCR workflow. The NLA could, for example, support broader efforts in Australia to revive First Nations languages by developing training datasets and building systems from any typewritten texts that are available in First Nations languages. This may include accessing materials in its own collection, or in others' collections e.g., the Living Archive of Aboriginal Languages maintained by Charles Darwin University library, and all materials maintained by the Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS).

Note, however, that many First Nations languages are retained predominately through oral histories, so typewritten documentary history may be limited to early settler accounts of the languages. Caution is required when using documents such as these as training datasets, as the colonial history of the documents is embedded in, and transformed by, OCR engines. This caution has been encouraged by the *Primeros Libros de las Américas* project, which worked to automatically transcribe sixteenth-century books in a mixture of Latin, Spanish, Nahuatl, and other languages (Alpert-Abrams 2016). As Alpert-Abrams (2016) notes, cultural interactions embedded in early colonial transcriptions remind us that copies are never neutral. Similarly, the training data used in OCR embeds cultural biases in the output (Alpert-Abrams 2016). This suggests a need to accompany text run through an OCR workflow from historical accounts of any language – particularly colonial accounts – with the context of the text's production, providing a mode for analysing digital texts as historical texts.

3.3.5 Summary of Opportunities, Risks, Pitfalls, and Issues with OCR Capabilities

The opportunities, risks, pitfalls, and issues (ORPI) identified through this Cybernetic Systems analysis of OCR include:

OPPORTUNITIES

- Improve accuracy on a large corpus of historic newspapers by combining open-source OCR software components and increasing automation of post-processing.
- Contribute to sector-wide innovation through training datasets for historic newspapers.
- Continue to engage the crowdsourcing community while improving accuracy for a larger corpus of materials through semi-automated systems.
- Support broader efforts in Australia to revive First Nations languages by developing training datasets and building systems from any typewritten texts available in First Nations languages.

RISKS

- Embed cultural biases in the output of OCR for languages other than English where colonial transcriptions are used as training data for OCR engines.
- Emphasising typewritten text for OCR may lead the NLA to focus on languages for which typewritten text is available, rather than focusing on other techniques for different data types. For example, audio-to-text annotation based on oral history collections of First Nations languages.

PITFALLS

- Existing, singular systems and manual post-processing may limit the kinds of future applications envisaged in Scenario 1, *Deakin Talks, World Listens.*

ISSUES

- Known limitations – textual and scanning issues affect OCR accuracy, particularly for historic materials, and need to be accounted for in any of the examples listed.
- Effect of OCR inaccuracies on employing other techniques downstream e.g. named entity recognition is increasingly difficult.

Scenario 2:
Magnifica Emerges from the Web

Amateur botanists across Australia are celebrating after one of their own, Jarli Grange, discovered a small grove of supposedly extinct Australian orchids growing in the wild, aided by a series of clues left in historical maps, oral histories and other resources held online by the National Library of Australia.

The caladenia magnifica, or magnificent spider orchid, so named for the striking size and colour of its long spidery petals, is endemic to Central Victoria, but had not been seen there since 1979. That is, until Grange's discovery earlier this month.

"I couldn't believe what I was seeing," Grange told In Media Res *of his experience discovering the lost orchid. "I've been bush-bashing my way across Clydesdale for years, but it turns out I was looking in the wrong place the whole time. I should have been looking online!"*

Grange's trail of discovery began when he came across a reference to orchids and a sloped area of stony soil – the ideal habitat for the caladenia magnifica – on an early pastoral map of Victoria held by the National Library, in an area not previous associated with the flower.

While the map had been available online for years, a library spokesperson says it was a more recent technological innovation that led to the historical find. The map, and many others like it, has been carefully read and annotated by a new kind of librarian: an artificial intelligence (AI) system that finds information 'buried' in archival documents and brings it to the surface, transcribing handwritten, audio and visual information into catalogue records, which can then be searched by people like Grange.

Having found the map, Grange picked his way through a series of additional clues scattered through oral histories, handwritten letters, and photographs from the area, also made accessible by AI. One notable find was a 1923 letter from Ms Dymphna Marston to her fiancée in Cairns, which contained a fragile pressed flower – none other than the caladenia magnifica – and a detailed description of the pleasant afternoon she spent walking near Avoca, where she found the bud nestled under a tree on a hill. Ms Marston's letter was accompanied by a small, grainy Kodak Brownie photo taken of her and her friends with the tree in the distance. The same tree was identified in the background of several photos taken by the Avoca Walking Club in 1982, along with a recurring flash of red annotated by the AI as a 'native flower'.

"It notices things and connections we don't," said Grange. "It found that same damn tree again in an oral history with a local done after the 1985 bushfires. Poor bloke. Lost most of his farm, but all he wanted to talk about was how Avoca's one tree hill kept standing. I'm pleased it did, 'cos I followed it all the way."

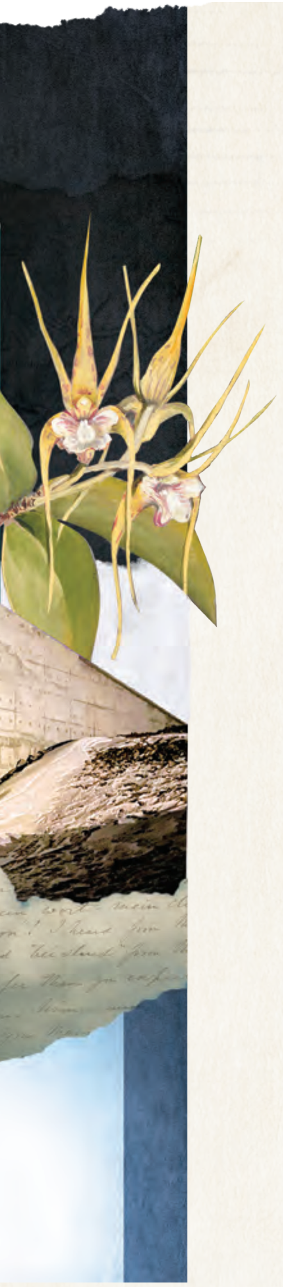

For all his enthusiasm about the new technology, though, Grange, a Wurundjeri man, is ambivalent about the map that set him on his journey. "When I saw those words [orchid, stony soil], I got so excited" he said. "But then I zoomed out, and all I could see was someone cutting up our country, taking it for themselves."

Grange's discovery marks a small bright spot in an otherwise dark year for Australian biodiversity. Three endemic Australian species that were living before colonisation in 1788 have been listed as extinct in 2032, taking the total tally to 122. "I reckon it's only a matter of time before there are more flowers in the library than the bush," said Grange.

In the year 2032, the National Library of Australia has established an international reputation as a leader in applying Machine Transcription technologies to their collections. After several tentative years of experimenting with speech-to-text translation of oral histories in the mid-2020s, the library branched out to object recognition in photographs, and text extraction from maps, among other AI services. In each case, the extracted information was attached to items first as metadata, and gradually made more public as the library gained confidence in the technology and relevant governance processes.

The case of pastoral and other nineteenth century maps is a particular source of pride to the NLA. Inspired by the success of its AUSTLANG code-a-thon, which saw more than 8,000 records in the Australian Bibliographic Database coded with a unique Indigenous Australian language identifier, the library took a similar approach to de-colonising the records for its colonial era maps. After using Machine Transcription processes to pull text out of maps into a searchable database, the NLA recruited a pool of expert volunteers to apply new subject headings to the map records, which had been developed by the Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS) to better serve the needs of Indigenous Australians.

It was one of these human-applied subject headings – not the AI-enabled system alone – that led Grange to his map. In fact, without these AIATSIS subject headings, Grange may not have found the information that led to his discovery, as the NLA struggled to adjust its catalogue search algorithms to cope with the influx of new metadata and information generated by its Machine Transcription efforts.

While its initial experiments were achieved with small digital transformation grants, the NLA required a more sustainable business model to fund its expanded Machine Transcription activities. After contemplating an exclusive partnership with a major machine learning platform provider, the NLA eventually opted to develop a shared Machine Transcription roadmap with the state libraries, who were excited by the results of the NLA's experiments. Seed funding for that project was provided by a philanthropic campaign targeting successful tech entrepreneurs who also formed an advisory council for the roadmap, with additional ongoing funding provided by the states. Together, the libraries developed a suite of Machine Transcription tools, built on top of a range of commercial and open-source platforms and data models, uniquely suited to the broader Australian collection.

Unfortunately, Grange's remarks on the future of Australian biodiversity proved all too prescient. Extinction rates continued to increase with global warming throughout the twenty-first century. By 2056, Synthetic Biology labs have started mining the library's collections for reference material, to help re-create lost plant- and eventually animal—life. Drawing inspiration from the Svalbard Global Seed Vault, the NLA and Te Puna Mātauranga o Aotearoa, the National Library of New Zealand, co-found an Oceanic Life Vault (OLV) in 2078 to store images and sound recordings relating to endemic flora and fauna to aid in their eventual reconstruction. One unexpected but significant source of data for the OLV was animal and bird calls caught in the background of other audio and video items held by the library, identified through audio event detection.

3.4 From Future to Present with Machine Transcription Capabilities

The seed of this speculative glimpse into the future of the NLA are today's AI-enabled Machine Transcription capabilities. In Scenario 2, information previously analogue in physical maps, photographs, and other graphics, later digitised or born digital in non-visual formats such as oral histories are made newly accessible through AI-enabled Machine Transcription. The new accessibility of information enables increased discoverability, multiplying research possibilities and thereby new pathways for NLA's audiences to create new knowledge.

The power of Machine Transcription is that it fundamentally shifts the primacy of human cognition in sense-making of visual and audio materials. As Paglen (2016) argues, Machine Transcription moves images from a representation requiring human viewers, to machine readable information, which no longer requires a "human in the loop". That is to say, a computing system does not need to turn an image into a form that a person can read in order to act upon it. Much the same can be said for audio sense and sense-making capabilities (Natural Language Processing, Audio Event Detection). Machine Transcription ultimately creates machine-readable text, whether that be a transcription of handwriting in manuscripts, of voice from oral histories and other audio material, a visual image from a photograph, painting, or drawing, as well as annotations about each. This has profound implications for the management of data in many applications, including libraries.

3.4.1 What are Machine Transcription Technologies?

The Machine Transcription cluster is a suite of tools, processes, and algorithms that convert multi-media data into machine accessible text, such as video-to-text, speech-to-text and image-to-text. Some of these are already in use by the NLA. They are a critical component in an array of applications to increase content availability, discoverability, and to improve search and user experience for library collections. Libraries often hold substantial collections that include media in formats other than type-written text, including manuscripts, diaries, letters and other handwritten genres, musical scores, as well as oral histories, other audio recordings, photography, maps, paintings, drawings, and images. For an institution such as the NLA, these holdings are substantial, and have been largely inaccessible to computational analysis. The promise of transcription technologies is to automate the transformation of media into text-based meta-data, and thus make these materials accessible not only to computational analysis but in turn to a broader audience, with speed, and at scale. As increasingly sophisticated techniques are developed for analysing images, video, and non-standardised text data, libraries can change how existing and new forms of media are described, discovered, and used by both the library and by its audience, in expected and in new ways.

Some critical questions the NLA should consider for transcription technologies include:

- How is meaning impacted with Machine Transcription (see: *Quality & Quantity* dynamic)?
- What are the impacts on time rhythms that may arise with the incorporation of Machine Transcription (see: *Formats & Fluidity* dynamic)?
- Does processing materials more quickly directly correlate to increased accessibility (see: *Agents & Audiences* dynamic)?
- Does transcription technology improve quality of meta-data as measured by human sense-making (see: *Serendipity & Indexing* dynamic)?

Three current Machine Transcription capabilities that informed the *Magnifica Emerges from the Web* are summarised here.

Handwriting Recognition Technology

Like OCR that transforms images of typewritten text into a machine-readable format, handwriting recognition technology uses computer vision or optical pattern recognition and transcribes handwriting into machine-readable text. Handwritten text, for all the careful instruction of generations of teachers, is highly variable. Fonts standardise typewritten texts in ways that regional, national, and time period handwriting conventions do not override individual handwriting styles. Thus, handwritten text presents a different type of transcription challenge than typed text and is a challenge that is particularly relevant to archives and library collections which include large volumes of handwritten notes, memos, letters, diaries, and other document types.

Whilst OCR capabilities can be traced back to innovation in telegraphy, handwriting recognition activities can be seen part of a suite of advances associated with putting thoughts into visible words and resulted in the first working typewriter made in 1808 (Weiner Grotta 2012). Handwriting recognition has been explored since the 1950s and can be broadly divided into two categories – firstly, handwriting analysis for authentication purposes (such as signature verification), and secondly handwriting recognition for Machine Transcription (Leedham 1994). Handwriting can be processed on-line (using a pen or stylus that interacts with a writing surface that can digitise and process the handwriting in real time) or off-line (at any point after it has been written). Off-line recognition enables processing to occur at many times the usual speed of handwriting and it is this efficiency and scale that is transformative compared to current practices of manual transcription.

Today, Handwritten Text Recognition (HTR) methods based on deep learning – a category of ML that attempts to replicate human decision making – are used with wide success in the library context to provide discoverability and access to manuscript collections. The applicability of HTR methods, however, is limited to certain types of manuscripts, where there are known symbol sets that can be used for supervised HTR models. This approach is not so relevant for more unusual manuscripts of which libraries are proud custodians, such as historical ciphers – a type of manuscript which has encrypted information using a variety of unknown symbols – where unsupervised machine learning models may provide a way forward (Baró *et al.* 2019).

The power of Machine Transcription is a computing system never needs to turn an image into a form that a person can read in order to act upon it. Machine Transcription is a suite of tools, processes, and algorithms that convert multi-media data into machine accessible text, such as video-to-text, speech-to-text and image-to-text. They are a critical component in an array of applications to increase content availability, discoverability, and to improve search and user experience for library collections.

Audio Machine Transcription: Speech to Text, Audio Event Detection, Music/Audio Data annotation

Audio Machine Transcription technologies can take multiple forms. Speech-to-text technologies transcribe audio recordings to text automatically, enabling increased accessibility by people and machines to search the audio material. Audio event detection involves the detection of different kinds of acoustic events such as gun shots, laughter or clapping. Audio data annotation labels audio data for increased searchability.

Audio Machine Transcription and speech recognition technologies have been developed in some form since the 1930s and popularized through advances in the 1980s and 1990s (Rabiner 2004). They are connected to the emergence of Natural Language Processing (NLP), an area of research that explores how computer systems can understand and manipulate natural language in similar ways to humans (Chowdhury 2003). NLP systems operate at several levels, including analysis at a word level (considering pronunciation, suffixes, and prefixes etc.), at a semantic level (dealing with the meaning of words and sentences), at the discourse level (focused on structure) and at a level that takes into account external context (Liddy 1998). Analysis at each level enables different accessibility features. For example processing at the semantic level can enable disambiguation of words with multiple senses and/or add synonymous equivalents of the query terms, whereas processing at the discourse level can be used to understand whether a piece of information is a conclusion, an opinion or a fact (Liddy 1998).

Emerging machine learning techniques have rapidly moved the technology closer to achieving "human parity" with error rates as low as 5.9 percent (Hang 2017). Accuracy rates are a function of a variety of factors such as audio quality and use of dialects, jargon, acronyms, and so on (Draxler *et al.* 2019). Audio transcription needs to take into consideration theoretical and methodological approaches to transcription. For example a naturalized approach to transcription would capture all the sounds that are heard including length of pauses and intonation, whereas a psychosocial method may include notes on emotional reaction or body language if audio transcription is produced from a video (Bucholtz 2000). The type of meta-data that will be attributed to the recording is also of concern, with various metadata schemas in place as guidance (Draxler *et al.* 2019).

There are many examples of audio Machine Transcription technologies in use today by library staff, archivists and researchers, such as otter.ai and Konch, the latter of which works with university libraries such as New York University to automatically transcribe parts of their media content through an API (Konch 2020). In a commercial context there are a range of off-the-shelf speech-to-text services available including Google Cloud Speech API, IBM Watson Speech to Text and Microsoft Custom Recognition Intelligent Service (CRIS) (Hunt *et al.* 2017). Many of the underlying datasets associated with the machine learning models have been developed in an English language context. The data lab (KBLab) at the National Library of Sweden is working to ensure a Swedish counterpart to sentiment analysis of texts by developing a Swedish-trained model based on Google BERT architecture and utilising access to the Vega EUROHPC JU system (an EU supercomputer) (EuroCC National Competence Centre Sweden 2021).

Audio Machine Transcription is currently used in libraries of today as a way of digitising oral history collections in particular and improving access. Its maturity has been spurred by the Web Content Accessibility Guidelines (WCAG) (Web Accessibility Initiative 2018), which mandate libraries, as government-funded institutions, to provide text-based alternatives to multimedia content as part of its responsibilities under the *Disability Discrimination Act 1992* (Cth). The State Library of NSW has created the Amplify Machine Transcription and crowdsourcing platform where NSW public libraries can create machine generated transcripts for oral history collections using the Amplify machine learning service and then engage with the local community to review and correct the transcripts (State Library New South Wales 2020).

Audio event detection endeavours to recognise sound events like clapping or gunshots within audio and video recordings, generally by employing machine learning algorithms trained on annotated datasets (Elizalde *et al.* 2017). Compared to image-based datasets like ImageNet for example, datasets are limited for different acoustic soundsets. Consequently, researchers are looking for ways to circumnavigate the need for annotated data, trialing methods using weakly labeled data (Kumar and Raj 2016) for example. This technology is commonly employed in surveillance style applications, such as the Amazon Alexa which has a "Guard Mode" trained to listen for broken glass or alarms which then trigger a mobile notification and other security actions (McCue 2019). In a library setting, audio event detection can be helpful in searching audio and video material for important events or similar types of events across different sources.

Audio Data Annotation seeks to label audio data with granularity and precision. Recent developments in machine learning have also enabled new forms of cultural heritage research into audio-based collections, such as the Dig That Lick project, whose approach facilitates "automatic recognition of musical structures and their linkage through metadata to historical and social context". (Dixon *et al.* 2019). Basaran *et al.* (2018) describe the ability to estimate a main melody in an audio recording. This then enables music to be grouped by melody similarity and audiences are able to search for patterns across different collections, time periods, or genres. Similarly, a program called OLIVE was developed to produce an automatic index from television or radio soundtracks, enabling keyword search and access to a preview version of the material before retrieving it from the archive (Chanod 1999).

Object recognition and Visual Data Annotation

Object recognition is a computer vision task that comprises image classification (predicting the type or class of an object in an image), object localisation (identifying the presence of objects in an image) and object detection (combines both classification and localisation to locate the presence and type of object in an image). Modern object recognition research stems back to the 1960s when office automation systems were first created and in 1973 when the world's first machine vision process was applied to the assembly and verification of semiconductors (Andreopoulos and Tsotsos 2013). Biomedical research (Gallus and Regoliosi 1974) and medical imaging (McInerney and Terzopoulos 1996) were also early targets of object recognition models, as were biometric vision-based systems for fingerprint recognition. In libraries, object recognition is particularly important in the context of recognising and classifying non-text (for example seals or diagrams) in manuscripts and also in the context of photographs in order to produce meaningful text that audiences can use to search for these materials.

Object recognition is generally highly dependent on training data, which is difficult to come by in the context of manuscript collections where training data would require considerable manual annotation of seals or drawings, rendering most learning-based methods inappropriate for manuscript research (Mohammed *et al.* 2021). Learning-free pattern detection methods which do not require pre-processing steps have been proposed by Mohammed *et al.* (2021) to make digitised manuscripts searchable for visual patterns in addition to its text.

Visual data annotation includes applications of annotating video and image data beyond type and handwritten texts, making them more tractable for search and research applications. Such technologies draw upon and integrate a wide range of techniques, such as concept classification, person recognition, video OCR and similarity search (Mühling *et al.* 2019). This includes applications such as sorting items by visual similarity, or automatic video content analysis and retrieval to facilitate search of historical video collections (Mühling *et al.* 2019). For example, Pop-Eye extracts time-coded text from video material subtitles and generates multi-lingual indexes, enabling users to access a still image or short clipping from the video for quick judgement as to its relevance (Chanod 1999). Other forms of visual data annotation include image-based word clouds, a way of visualising manuscript collections by arranging the most representative words in a collection in a cloud (Wilkinson and Brun 2015).

3.4.2 Integrating Machine Transcription into the NLA's cybernetic system

How might we make sense of Machine Transcription in the context of the NLA? Here, we turn to a cybernetic analysis of these capabilities in the system of the Library. To start, we outline the aspects of each system component – context, people, data, infrastructure, and process – and the general dynamics among them that we consider important to understand in relation to Machine Transcription at the NLA (see Section 1.3). We then consider the emerging dynamics among system components through three examples of the application of Machine Transcription capabilities at libraries. We consider potential changes to dynamics along a continuum from the *internal-facing* (implementation issues, knowable risks, automation and workforce) through to more *audience-facing* (ethical concerns, unintended consequences, impact on users). We also consider potential changes to dynamics and feedback loops along the multiple time rhythms that rule collections work at the NLA.

CONTEXT	• The statutory nature of the NLA and its legal deposit requirements must not be compromised through the introduction of new technologies. • The NLA caters to audiences now and in 20 years' time (or longer). • The NLA is operating in a time of reduced budget and demands for greater efficiency and sustainability. • The NLA occupies a leadership position in the Australian GLAM sector, and globally, in the adoption of OCR (a subset of Transcription).
PEOPLE	• Collections staff are aware of the potential of Machine Transcription to speed up processing of collection items, while cautious of the potential shortcomings of the accuracy of Machine Transcription solutions. • Collections staff prefer more accurate text though they believe that from the perspective of NLA audiences, there is general acceptance of machine-generated text. • Library staff working alongside current transcription systems see themselves as both beneficiaries and mediators of transcribed material. • IT professionals in the library will need to be trained in ensuring that the various data models used in Machine Transcription are fit-for-purpose and regularly monitored and updated. • For the NLA's audiences, Machine Transcription technologies enable more of the collection to be accessible, but audiences are not yet attuned to error rates, and gaps.

DATA	• Machine Transcription systems need to be trained on data that is relevant to the materials at hand. • For language (speech and handwriting) solutions, given the plethora of script styles and/or accents across different time periods, cultures and languages, there is a wide degree of variation in the applicability of models to the different types of multimedia data. • Underlying models informing a Machine Transcription can be developed in-house, to adopt and modify an existing solution, or an already existing solution can be adopted.
PROCESSES	• Regardless of the transcription solution, human review of the output is needed. In the post-processing of outputs, Library staff review the transcripts to search for defamatory information, create time summaries, and other human labor-intensive activities. • IT professionals will need to be trained in ensuring that the various data models used in Machine Transcription are fit-for-purpose and regularly monitored and updated. • Workflow process documentation will need to be reviewed and updated to include transcription technologies. • Guidelines for interactions between Library staff and audience online will need to be reviewed and updated.
INFRASTRUCTURE	• Existing infrastructure and workflows need to be updated to accommodate machine-derived metadata. • Energy intensive workloads might drive new operating expense costs.

3.4.3 Salient General Dynamics for Machine Transcription Capabilities

Like OCR, for Machine Transcription capabilities, the *Quality & Quantity* dynamic manifests primarily in two ways. The collection staff are uneasy about the inaccuracy of described items, and also the adherence of meaning to the item itself, rather than to the item in context. Whilst Machine Transcription can process more items quickly, it has the potential to atomise meaning and to make mistakes in categorisation that defy "common sense". On the other hand, transcription will increase the quality of library collection items, particularly those that are currently handwritten, as the content will be more legible to audiences.

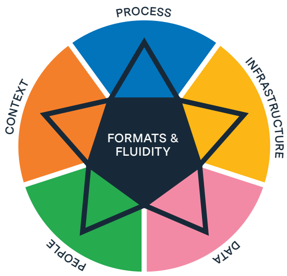

Building on the standards and workflows established by OCR, additional transcription capabilities will become part of the digitisation process at the NLA. This will require developing new pre-processing and post-processing workflows, and tie-ins to existing processes, such as the workflows for Trove. As transcription capabilities are integrated into digitisation workflows, making more content digitally accessible, audience expectations may also change to expect anytime, anywhere access to collections.

Integrating Machine Transcription into digitisation processes may decrease the time interval between acquisition of new analogue materials, and when the materials can be made available in the collection, decreasing the backlog of materials and supporting better discoverability of existing materials.

The ability to speed up processing of collection items may drive new expectations from audiences for access to materials which may not be compatible with the use rights agreed upon by donors of personal papers, manuscripts, and similar collections. In other words: just because materials *could* be processed more quickly does not mean that doing so will speed up accessibility. Faster processing may also create new dynamics in existing relationships among audiences engaged in updating metadata for machine-transcribed collection items (such as with newspaper content in Trove). First, how audience members compete to rank in correcting collection items will change as more machine-transcribed items are available. Second, while agreeing upon correction of text generated by OCR already leads to conflict among audience assessors on Trove, coming to agreement on about object or a scene captured in an image and then rendered into text by a machine will likely generate more audience controversy. A picture may be worth a thousand words, and those thousands of words are always grounded in the biases of the present. Much as collection staff grapple now with updating tags created by earlier staff to describe images, they may increasingly grapple with machine-generated text.

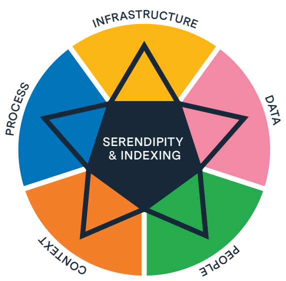

Finally, Machine Transcription capabilities promise to automate and simplify labelling and description of collections items, increasing sense-making and supporting keyword and full text search. The quality of the resulting metadata as measured by human sense-making (is that an image of "a man in room" or "Bushranger in Coombs jail cell"), may be far less impressive than we imagine. As better training data becomes available, re-running Machine Transcription on items to correct metadata descriptions could become a common process. This is crucial for handwriting recognition technologies to successfully integrate into existing archival practices and scholarly research practices. A growing user network has been crucial to the success of Transkribus for example (Muehlberger *et al.* 2019). Individual collections items may be thus more easily searchable and accessible, though the essential capabilities do not promise to increase opportunities to create new connections among collection items, for that we must turn to the set of capabilities, *Machine-Actionable Collections* (see Section 3.5).

3.4.4 Current and Projected Applications of Machine Transcription Capabilities

For libraries, today's Machine Transcription capabilities are potentially as transformative to the capability and service model of the library as Guttenberg's printing press was in the fifteenth century. Machine Transcription capabilities promise to make the currently time-consuming to capture, and labour-intensive to search information present in born-analogue collection materials easier to generate, to discover, and to access. Both what is available and how, and to whom it is available – born-digital versus handwritten, online versus in person, general descriptions versus specific keywords – are ripe for revolution.

Using three current examples and projecting forward to mid- and far- future applications, including those in Scenario 2, *Magnifica Emerges from the Web*, we consider the potential new dynamics that will emerge if the capability is integrated into the current known dynamics of the NLA collections system and the attendant possibilities, risk, pitfalls, and issues each presents.

Application Example 1: Scaling handwriting recognition applications: *Transkribus*

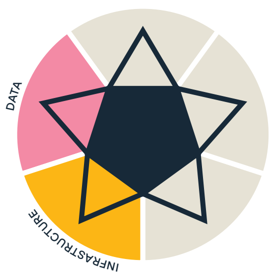

Relevant data models and feedback loops are essential to the success of implementing Machine Transcription in the NLA. Handwriting recognition technology requires the establishment of an error-free reference consisting of a segmentation of document images into individual text lines or words used to test and train a recognition system such as Hidden Markov Models (HMM) or BLSTM neural networks. For modern handwritten documents, training data is more readily available, for example through datasets such as UNIPEN and IRONFF for on-line data, or CEDAR for off-line data. These datasets (and many others) have enabled comparison points for further development of recognition technologies for modern texts. For historical documents, however, there are less dataset sources to draw upon. The University of Innsbruck's *Transkribus* (Muehlberger *et al.* 2019) addresses this issue of dataset sources for historical documents by building a platform that facilitates the initial capture of training data needed to create Handwritten Text Recognition Models. As a result, *Transkribus* can read, transcribe, and search for historical manuscripts of different dates, languages and formats. *Transkribus* has over 50,000 users as of June 2021 (Transkribus n.d.). This growing user network is core to *Transkribus'* success, as the machine learning model underlying the platform is strengthened with each document that is processed through *Transkribus*.

Looking ahead, *Magnifica emerges from the web* explores the potential benefits of collaboratively acquiring data sets. After considering a closed, exclusive system, the NLA decides instead to deliver on Machine Transcription through the creation of a more open collaboration with state libraries, following a series of experiments that are funded philanthropically. The benefit of structuring the collaboration in this way enables further datasets to be created and shared that are relevant to the materials stored in the NLA. The NLA also leverages its audience through user-generated meta-tagging of content, for example through expert volunteers who apply subject headings developed by AIATSIS.

Connecting handwriting recognition to a collection management service as a way of recouping costs may seem like a useful revenue generation activity, however, there is value to researchers and other individuals being able to engage with the development and application of the technology. This is crucial to the ability of handwriting recognition technologies to be successful in integrating into existing archival practice and scholarly research. A growing user network has been crucial to the success of *Transkribus* for example (Muehlberger *et al.* 2019).

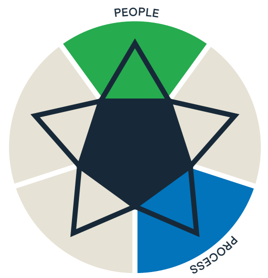

There is power in having a volunteer base involved in checking the accuracy of Machine Transcription technologies, such as crowdsourcing the human check and correct post-processing work flows, particularly for metadata tagging/annotations. This has already been trialed in the NLA as a part of newspaper digitisation processes and is explored further in Scenario 2. As highlighted in an NLA expert interview, it must also be noted that the scale of accuracy checking that can be done by crowdsourcing/volunteers is relatively low compared to the quantum of accuracy checking required. To exclusively use humans to accuracy check the amount of materials produced through Machine Transcription of library collection items could lead to the ethically fraught practice of "Ghost Work". Coined by Gray and Suri (2019), Ghost Work describes an invisible "underclass" of low or no-wage, disposable workforce, focused on menial tasks that ultimately assist with teaching machines to learn.

As Machine Transcription work processes are developed, library collections teams should ensure that user input and bug-reporting feedback loops are in place, and connect into development processes. This will contribute to assuring Machine Transcription workflows are fit-for-purpose.

Application Example 2: Speech-to-text transcription for Norwegian language collections

Machine Transcription technologies are generally reliant on large, generally annotated datasets. Given many Machine Transcription data models are based on English-language training data, applications in other languages need consideration. The National Library of Norway collaborated with software development and project management company Schjønhaug AS, to develop a speech-to-text transcription service (Writbeat n.d.). The National Library of Norway has made the tool available for free to encourage companies to invest in the Norwegian language.

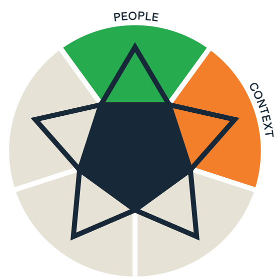

In this example, the Library is thinking about its audiences widely. In addition to audiences interested in library material, they have expanded their remit to also consider the role the Library can play in creating a rare and otherwise expensive-to-develop dataset that can be leveraged by Norwegian businesses. This has flow-on impacts to enabling companies to invest in the Norwegian language, preserving its use for future generations. The inclusion of speech-to-text transcription technologies has expanded the library's responsibility beyond accessibility of their collection, to accessibility and maintenance of the dataset and transcription tools and platforms.

The concept of preserving language resonates deeply for First Nations people. Scenario 2 touches on decolonising records using a unique Indigenous Australian language identifier. While the scenario does not delve deeply into the preservation of Indigenous language, it does illustrate a narrative of collaboration and opportunity with diverse audiences.

In terms of implementing the speech-to-text transcription tool, The National Library of Norway decided to work with an industry provider, Schjønhaug AS, who in turn was working with transcription tool Benevis in order to tailor an appropriate service. The Library leveraged the expertise of an external provider who had experience in Machine Transcription tools, recognising this path as the most efficient and effective way to build an appropriate tool. They also recognised their strengths in integrating the tool to their already existing Språkbanken (the language bank), enabling high visibility and accessibility of the tool in the already-existing library online infrastructure.

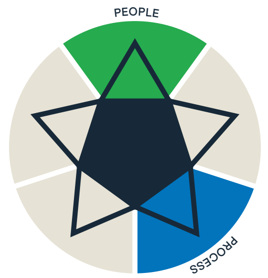

Before the speech-to-text transcription tool, staff at the library manually transcribed audio files. The new tool led to hiring “language technologists” who focused on editing Machine Transcribed text and correcting errors. Other library staff were thus able to focus on more “high value” tasks such as adding contextual cues to records or engaging in more detailed and complex annotation and service provision. Automating transcription shifted work loads and work content, and necessitated new job titles.”

Application Example 3: Object recognition enabling new enagagement with the 1891 Norwegian census

The work of Thorvaldsen *et al.* (2015) describes the experiences of using Machine Transcription in two different historical demographic databases: more than 600,000 records from the Barcelona Historical Marriage Database and forms from the 1891 Norwegian census. Different Machine Transcription technologies are utilised across the two case studies, this example will focus on object recognition technologies applied to the transcription of the 1891 Norwegian census. The initial software created read the scanned images of the census forms in low resolution and determined which answers had been underlined on the form.

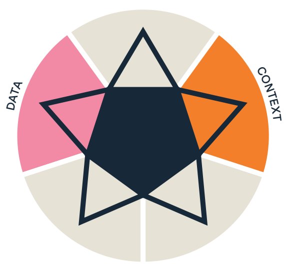

In many cases, viewing a historical database document, such as census data or marriage data, may involve visiting the library, accessing the catalog, making a request to view the material, waiting for its delivery, and then viewing physical materials in the reading room (amongst other steps). There are in-built delays in the system regarding how quickly materials can be viewed and understood. These factors influence how quickly a judgement can be made on the relevance of the item to the audience’s purpose and ultimately how quickly new meaning can be made.

In this example, however, with the application of Machine Transcription technologies to create machine readable text at scale, these systems dynamics change. There is an overall shorter period of time between searching for an item and viewing it to make a judgement on its relevance, because the material is now digitised and likely accessible from anywhere and can be searched via keyword search. In Scenario 2, this speed of finding materials is further augmented using an AI.

The flow-on impacts of machine-readable text could lead to faster meaning-making; it could also lead to a tendency for audiences to find more and more items given the effort to do so is relatively low, which eventually results in slower meaning-making as a larger quantum of materials are digested. Overall, although search processes may be more efficient with Machine Transcription technologies, the time taken to do sense-making on an exploded set of search results may proportionately increase.

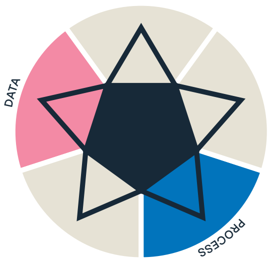

With Machine Transcription of the Norwegian census forms there are several steps involved with detecting whether a particular part of the form is underlined or not: (1) The picture is delimited through the identification of the four corners of the page. (2) The specific fields containing the information to be analysed are identified. (3) A comparison is made between the underlined word and a pre-recorded image of a non-underlined word and a "likelihood underlined" score is generated. (4) A consistency check is applied, for example if male gender is underlined, then family position cannot be "wife" or "daughter". Of course, these kinds of consistency checks embed and embody whole world views that are rarely neutral even when ostensibly culturally accurate. They highlight that bias in AI applications is a constant and formidable issue; using automated processes to analyse historic materials presents challenges to reconcile – or directly engage – different cultural ontologies. Should our automated processes follow century old Norwegian state bureaucratic gender categories, and reproduce a cisgendered, heteronormative vision of families and households, or can the check instead act to uncover possible expressions of gender diversity hidden in plain sight? Are all errors 'true' errors or just misunderstood data?

Consideration as to the potential errors that could occur along each step of the process needs take place. As part of the project, human checks were applied to detect errors in detection and transcription , before creating a more advanced version of the software that transcribes the fields and copies the strings to separate image files. Errors to date have been connected to "noisy documents" where people who had originally completed the forms had crossed out or corrected their responses for example.

In addition to errors within the realm of detection, it is also worthy to consider flow-on errors that could occur in the wider system. For example, with an increase of machine-readable text associated with these census forms, keyword search is now more possible. This method provides a precision not seen previously. Overall, more materials are made accessible, and more results are surfaced, which could create "errors" of "clogging" the search results.

The increase in volume of available, discoverable material can be mediated through search capabilities such as a catalog and its capabilities to do relevance ranking of search results. Our expert interviews revealed a reluctance from the NLA to accept a role of giving Google-like relevance rankings to search results because there are too many questions around how to determine what ranks highest. Making relevance judgements on behalf of unknown NLA audiences of today and the future is counter to the NLA's views on their role as custodian of materials, not meaning-maker. The Library endeavours to remain as neutral as possible so that audiences can ascribe their own meaning, but risks deterring people through the overwhelm of information inadequately "controlled". The introduction of Machine Transcription technologies increases this risk.

3.4.5 Summary of Opportunities, Risks, Pitfalls, and Issues with Transcription Capabilities

The opportunities, risks, pitfalls, and issues (ORPI) identified through this Cybernetic Systems analysis of Machine Transcription include:

OPPORTUNITIES

- New accessibility to manuscripts and to oral history data enables new sense-making and research (clearing of backlogs faster).
- Utilising Machine Transcription instead of analogue transcription by Library staff, frees up library staff to do higher value work drawing on their expertise.
- Build a global user community connected to transcription technologies at NLA.
- Possibility of new attribution of documents with particular authors through computational analysis of handwriting.

RISKS

- Increased ability to find materials leads to more confusion / overwhelm and less meaning-making.
- Unknown limitations of the data models underlying a solution, if sourced externally, leads to suboptimal performance at best and unintended, harmful bias at worst.
- Lack of sustainable model to keep the Machine Transcription technologies operational.
- Search results are "clogged" with too much information that is not meaningfully ranked, leading to audience dissatisfaction (particularly those expecting a Google-like user experience).
- Lower audience satisfaction when interacting with materials online than when engaging directly with Library staff in person as Machine Transcription indexes and labels collection items but is not designed to enable associative trails.
- The value proposition for Machine Transcription technologies is nascent and garnering support for them might be hindered by their newness and lack of successful and established proof points.

PITFALLS

- Higher accessibility to materials previously difficult to access leads to new connections that were not intended to be made, potentially violating the intention of a creator and/or original owner and/or accessibility rights assigned to a particular material. These issues might be especially acute for First Nationals materials and contributors.

ISSUES

- Data models that are underlying Machine Transcription technologies may not be fit for purpose for the breadth of different materials (for example the diversity of handwriting styles).
- Consistency checking technologies can reproduce old, or introduce new, forms of cultural logic and make meaning stable where it perhaps should not be.
- The output of Machine Transcription is often keyword searchable text. Searches may lose some nuance and potential for serendipitous finding of materials through additional precision offered by keyword search.
- Challenges in reading data in "difficult formats" e.g., tables.

Scenario 3:
Great Granddaughters Take Back the Library

A group of young Australians braved the elements in Ngunnawal country this morning to follow in the footsteps of their great-grandmothers, re-enacting the landmark Women's March 4 Justice that occurred there 30 years ago today.

Each participant followed the path taken by an individual protestor, which had been recreated for them from extremely rare location data made accessible to the public by the National Library of Australia.

Marchers lifted replica signs as they followed the directions projected through their smart rigs and listened to a play-by-play account of the march pieced together by organisers from a cache of data from a 'micro-blogging website'. The website comprises short user-generated text content, retrieved from a vintage physical storage device found at the bottom of an Australian National University researcher's closet after they died in 2046.

As the marchers gathered on Capital Hill, a young woman appeared on their rigs. "Change took a long time," march organiser Amelia Chen told the crowd, "But it came."

When asked what inspired her to organise the re-enactment, Ms Chen told In Media Res *that she "wanted people to remember the good things our great-grandparents did for us. I know we have reason to blame them for a lot, but they weren't all bad."*

"I don't have anything physical that belonged to my great-grandmother, just some low-definition Internet-era stuff. And I'm lucky, most people don't even have that. Walking in her footsteps today, it made me feel physically connected to her for the first time."

In the year 2050, smart eye- and ear-wear have replaced smart phones as the ubiquitous personal computing device. 'Smart rigs', as they're called, generally comprise one or two contact lenses for receiving and displaying visuals, bone-conduction pads for audio input and output, a discreet biometric sensor,

and a central computing and networking hub, usually implanted in a molar. Users receive a constant stream of carefully calibrated visual and audio information, projected over and into the real world. Meanwhile, a constant stream of user data is taken in return. This blending of the cyber and physical worlds is referred to as the 'spatial web' or Web 3.0.

Before smart rigs, smart glasses overtook smart phones from around 2030. This transition, combined with the shift from Web 2.0 to 3.0, and balkanisation of the internet by competing global and corporate powers, meant that much of the content produced online in the 2010s and 20s became very difficult to access. Internet-era powerhouses, repositioning themselves in the new AI-everything environment were loath to share their private stores of machine learning training data. Publicly accessible caches like that found in ANU researcher's closet are rare treasures.

Like the smart phone before it, the smart rig generated new ways of interacting with information, and conducting research, including embodied explorations like the March 4 Justice re-enactment. Researchers can literally feel their way through a collection, exploring twisting pathways through real and virtual landscapes of images, sounds, and text. In the early days of the 2030s, the major limit to public participation in such experiences was access to the computational power required to render such detailed information worlds. In 2035 the NLA launched its first public Virtual Reading Room, a cloud computing environment in which users can readily work with the library's Machine-Actionable collections in a virtual space.

The location data used to guide each individual participant in the re-enactment was collected by the COVIDSafe app, a primitive crowd sensing application promoted by the Australian Government during the COVID-19 Pandemic of 2020–2022. The COVIDSafe data was donated to the National Library of Australia in 2032 and placed under indefinite restricted access. In addition to containing personal information, the COVIDSafe dataset was huge, making it both cognitively and computationally inaccessible to the general public. The NLA has only recently begun supporting public use of the dataset by providing access via a data visualisation tool, which offers access to anonymised data in the form of representative, rather than actual, trajectories.

3.5 From Future to Present with Machine-Actionable Collections

The precursors of this speculative glimpse into the future of the NLA are seen in today's efforts to create Machine-Actionable collections. In Scenario 3, information collected from the new born-digital landscape, including internet, social media, Twitter feeds, Internet of Things, check-in systems (such as the COVID app data), are made actionable by audiences. Here, there is a symbiotic relationship; as audiences interact with, integrate, and imbricate elements of the collection, they in turn shape it through the "constant stream of user data" that is gathered in return and ingested by AI-enabled applications. This new ability – to action, to actuate information – enables the physically embodied experiences of the Library. New collection mechanisms furnish the NLA's midwife role with new dimensionality; collaborating with audiences/agents to create new knowledge, and new ways of experiencing knowledge.

3.5.1 What are Machine-Actionable Collections?

Machine-Actionable Collections (MAC) refer to the processes, standards, and infrastructure required for digital collections to be made available, as structured data, to other computing systems that can read and act upon them. MAC orient libraries toward creating large, open, machine-readable datasets for use by researchers. Moreover, these systems provide the necessary data to fuel AI-enabled systems. MAC facilitate new approaches to digital discoverability and accessibility for collections. They present new opportunities for audiences to engage with, and co-create from, the collection. MAC, therefore, incorporate not only collections as data, but the attendant technologies that can read this data and act upon it. As such, MAC represent systems in which collection items become *actuators* upon which machines react with multiple possible intents, outputs, and implications for libraries (Padilla 2019).

Through shifting focus to digital *collections* as the unit of activity and analysis rather than digital collection *items* as was the case in the preceding section on Machine Transcription, MAC present several trajectories for transformation in libraries. Some are outlined in our speculative Scenario 3, *Great Granddaughters Take Back the Library*. One trajectory might be to think about how born-digital and digitised content from multiple sources and collections can relate as part of one system of data and, through computation, shape new approaches to sense-making. Another is to think about how libraries may need to support publics to discover, access, and encounter the collection through a range of emerging technologies and cultural contexts. Having collections Machine-Actionable decouples them from time, space, and platform; they can now be accessed through different technologies, and in different contexts. There is a tension here; as MAC move to the centre from the periphery, the tether with which they are coupled to the Library weakens. The Library has safekeeping of the collection, but not necessarily control, moving away from the role of custodian to the role of midwife of new knowledge.

Machine-Actionable Collections (MAC) are the processes, standards, and infrastructure required for digital collections to be made available, as structured data, to other computing systems that can read and act upon them. They present new opportunities for audiences to engage with, and co-create from, the collection.

For example, if one were to use a Machine-Actionable, open, public collection, and co-create an application from it, this would change the cultural context, and platform through which the collection is encountered (and possibly the publics who encounter it).

Some critical questions the NLA should consider for Machine-Actionable Collections include:

- Can and should the the NLA take on an educational role for audiences to use tools and techniques on MAC? (see: *Agents & Audiences* dynamic)?
- Will supporting applications for MAC shift the role of the NLA too far towards interpreting collections for the public? (see: *Serendipity & Indexing* dynamic)?
- What roles does NLA wish to have in facilitating MAC standards and interoperability across Australia and New Zealand (see: *Nodes & Networks* dynamic)?

The Semantic Web: Foundations for thinking about MAC

MAC sit at the intersection of two trajectories: libraries' long-held role in shaping data practices and the Semantic Web.

First proposed in the late 1990s, the vision of Sir Tim Berners-Lee for Semantic Web was to have a world-wide network in which computers became capable of making meaning from all the data on the web – "the content, links and transactions between people and computers" (Berners-Lee and Fischetti 1999, n.p.). Twenty years later, this aspiration is becoming increasingly fulfilled as platforms generate, exchange, connect and transform structured data.

While MAC represent a current discourse on how libraries are creating and deploying machine-readable data from their collections, the basic principle of generating and collecting machine-readable data has circulated in the library context since the 1970s (Rowe and Ryan 1974, Isaacson 1982). Since then, and in line with the development of technology capabilities, libraries have increasingly oriented toward opening and linking data sets. Concomitantly, new tooling such as ontologies, thesauri and schemas have been created to make data more easily machine readable (Siwecka 2018). As a result, open data principles in libraries have been an engine for a suite of innovations, both in the library and in information science more broadly.

Through the development of an international ecosystem of practices and standards, libraries have created various tools to support Linked Open Data (LOD) of their digital collections and metadata sets. Specifically, libraries globally have developed and adopted library-specific Resource Description Frameworks (RDFs), including MARC, MARCXML, and Dublin Core, making them available using the SRU, OAI-PMH or Z39.50 protocols (Siwecka 2018). RDFs such as these are integral infrastructure for enabling the semantic web; they inform standards for structuring data so that data sets may be interoperable. More recently, libraries have also begun publishing open data stored in a machine-readable format, conducive to the development of the semantic web and linked MAC (e.g., British National Bibliography provided in RD F/XML format).

In addition, libraries have established myriad other Knowledge Organisation Systems (KOS), including ontologies, thesauri, vocabularies and more to help standardise disparate indexing systems. KOS have formed the basis for developing standardised data description models and many national libraries are currently using KOS to enable the semantic web. Linked KOSs are also currently being employed to improve machine-generated library metadata. Some libraries have also made their own KOS available online: for example, the Finnish National Library provides access to nearly forty ontologies, classifications, and thesauri (Siwecka 2018).

As libraries have gradually grown their digitised and born-digital collections, schemas such as these support libraries to make collections available as LOD by structuring them into systems that enable machines to read them. These practices are foundational to creating and deploying MAC for computational analysis. In this legacy of libraries as midwives of the semantic web, MAC represent a gear shift toward thinking about what technology applications are made possible by creating Machine-Actionable data sets into AI-enabled systems.

Computational Analysis using Machine-Actionable Collections in Libraries

MAC in libraries partly describes systems of open and accessible data that are made available to researchers to perform computational analysis on the libraries' collections. MAC is also a movement among libraries toward enabling digital experimentation on their collections of both digitised and born-digital data sets (Candela *et al.* 2020). This movement has manifested a burgeoning network of virtual labs situated within libraries (Padilla 2019), as well as a number of flagship projects such as Always Already Computational, and the OpenGlam movement more broadly, which have advanced libraries toward creating "collections as data'" (Padilla *et al.* 2019).

The NLA already has experience creating, deploying and maintaining Machine-Actionable collections. Trove's application programming interface (API) has enabled digital humanities researchers to process 152 million articles using Named Entity Recognition to identify 27-million person name mentions in 17 million articles (Burrows and Davidson 2015). With further processing and clustering, researchers have been able to identify the frequency of mention of certain people over time. The intent for this project is to eventually provide links between the name data in the Humanities Networked Infrastructure (HuNI) Virtual Laboratory and articles in Trove. Other cultural heritage collections have also been made available in Australia via the Alveo API, which self-describes as a "virtual lab for human communication science" (Alveo 2021), prompted by the National eResearch Collaboration Tools and Resources project.

Broadly speaking, there are three meta-types of MAC, which include: digital or digitised collection items; metadata; and newly created data, which are born from augmenting either of the two aforementioned data types. These augmented data sets are typically created through enriching collection data with other data (like Wikidata and GeoNames data), or by running analysis methods like text mining over collection data to carve out sub-sections of a larger collection. Typical examples of MAC today are often discrete datasets classified by content type such as text, images, maps and metadata (Candela *et al.* 2020). These are increasingly made available openly to researchers through data deposit platforms like the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH), DataCite and Zenodo (ibid.).

Sets of practices around the usability of MAC are emerging in tandem with broader availability of MAC within GLAM institutions. Various organisations are starting to develop fellowship programs with researchers who might wish to use their collections, and more are sharing their prototypes and exemplars of computational analysis and scholarship with the public (Wittmann *et al.* 2019). For example, the Danish National Library has published a series of engaging prototypes based on digital methods such as topic modelling, machine-learning and visualisation tools that are published on their website (KB Labs 2021). To improve ease-of-use of MAC and the transparency of the computational methods behind analyses using MAC, institutions are also experimenting with ways to publish documentation. A good example is the Archives Unleashed Notebooks prototype, which interactively explores and filters the information, extracted full text and visualisation data. Across the sector, Jupyter Notebooks such as these are a popular tool for MAC users, because they provide reproducible, reusable, and transparent data analysis by publishing detailed workflows, narrative text, and visualisation results.

In addition to these usability practices, several frameworks for creating collections as data are maturing. Although they support the end goal of creating a collection from clean, machine-readable data, they exist across the data life cycle. These stages of the life cycle include identification; access and retrieval; cleaning and reformatting; data enrichment; and packaging data sets for computational analysis (Candela *et al.*

2020). Interlinked with these processes are several commonly used standards and tools that are assisting libraries to create MAC, including open-source platforms such as OpenRefine and goodtables.io that help with data cleaning and reformation. The International Image Interoperability Framework (IIIF) is another – a set of APIs that support interoperability between image repositories. To support standardising formats, the World Wide Web Consortium (W3C) provides recommendations for several machine-readable standardised formats that are commonly adopted by the GLAM sector.

As MAC continue to scale across the global library ecosystem, a prescient issue will be how to link and share data more efficiently and effectively across GLAM institutions. Defining these standards could be an important role for the NLA in the Australia-Pacific region.

Applying emergent trends to MAC in the Library

There are rich possibilities for what Machine-Actionable Collections at scale can achieve for the NLA. These possibilities largely map to two capabilities: discoverability and interactive access to digital collections.

As digital collections in libraries have scaled, discovery systems become the "front door" of the library for many audiences. The benefits of having a reliable and effective discovery layer to access the digital collection has become especially pertinent as COVID-19 has demonstrated the benefit of remote access to library collections (Fitzgerald 2020). However, making library collections accessible in line with remote and distributed user patterns more broadly has long-been an agenda for the NLA (National Library of Australia 2013). Libraries have become increasingly embedded with computational tools that support their audiences to do targeted item retrieval, recommended search and to support other global search and discovery tools.

Linked open data provides an opportunity for discovering *knowledge* rather than just *data*. That is, by combining multiple tools, the audience can be furnished with additional temporal, spatial, linguistic, historical, political, and social context. These tools include a series of broad capabilities such as text mining, word embedding, image search, and recommender systems. Each of these capabilities harnesses a range of AI-enabled data analysis methods, such as algorithms, data modelling, computer vision, and NLP (Kangas 2019).

Text-mining allows users to query unstructured or semi-structured data to extract information of possible interest (Cai and Sun 2009). This creates data that is then actionable by humans or other technology systems. As unstructured information grows within digital libraries, text mining has useful applications for discovery and creation of machine-readable data that can be acted upon by other tools for various purposes (Talib *et al.* 2016, Antons *et al.* 2020). That is, text mining helps to turn unstructured data into structured data.

Word embedding enables the relationship between words to be represented mathematically – as being "closer" or "further apart" depending on the frequency with which they appear in the same texts (Brownlee 2019, Gu *et al.* 2020). Word embedding facilitates functionality such as recommender systems – which suggest content items that are "close" to each other – and allows collections to be queried using natural language rather than formalised programming languages, such as SQL.

Recommender systems predominantly follow an algorithm that places audiences into categories based on relationships between items in a collection and how the audience engages with those items, such as following, rating, or liking an item. Relationships between users can also be identified in this way. Key examples occur in entertainment (e.g. Amazon) and content streaming (e.g. Netflix) (Melville and Sindhwani 2010, Fatourechi 2015, Pandey 2019). One possible future of these systems could be increasingly customised environments for audiences, including accounts that sense and respond to real-time factors such as mood, time of day, location, sleep cycle, and energy output.

Machine translation also leverages word embeddings; mapping between words and phrases in one language to provide near-fluent rendition in another (Koehn and Knowles

2017). Similar to speech recognition systems mentioned earlier, machine translation currently works better in languages for which high volumes of written content is available; and more poorly for “low resource” languages. All Australian Indigenous languages are considered “low resource” languages (Foley *et al.* 2018).

Another class of technologies that use linked open data are *generative* tools – such as BERT and GPT-3 (Devlin *et al.* 2019, Brown *et al.* 2020). These tools are trained on huge text datasets to create text in response to a prompt; both BERT and GPT-3 have received significant media coverage due to the realistic nature of the writing they produce (Floridi and Chiriatti 2020). Generative tools currently operate on general or broad contexts; advancements in this space are focused on creating accurate responses in very specific domains such as health, medicine, and industrial applications.

All these capabilities extend the ways in which audiences can encounter and interact with digital collections. Many possibilities exist for re-conceptualising discovery and accessibility. These imaginaries include natural language interaction with collections at a place, space, and time of the audience’s choosing; virtual or mixed reality which visually displays items from a collection, or the ability to interact with a collection in the audience’s preferred language. This highlights the shift in emphasis of the library’s role away from being an arbiter of collecting and mediating access to items, to one of increased curation of collections, informed by open linked data.

Libraries as Cyber-physical Systems

Machine-Actionable data can also be cast as one component within broader cyber-physical systems, in which this type of structured data is increasingly used to connect the physical world with the online world through sensors and actuators. In libraries, CPS enable new affordances that have implications for the collection and use of Machine-Actionable data.

Again, there are several possibilities. One is automating preservation monitoring, where computer vision is employed to evaluate and predict the physical deterioration of collection items, flagging when human intervention is required, reducing the need for manual inspection. A project is underway to do this at the US Library of Congress (Cordell 2020). Other possibilities include radio frequency identification (RFID) tracking of books and other objects. Here, radio frequency signals broadcast the location of an object in a way that is trackable by computers, again reducing the need for manual search (Shahid 2005). Bluetooth Low Energy (BLE) Beacons are another way to track objects, and they can also interact with smart phone. This is an approach of the BeaLib beacon-enabled smart library system, theorised in New Zealand and Singapore, but not yet at prototype stage (Uttarwar *et al.* 2017).

Another avenue that remains under-explored is that of visualising Machine-Actionable data – for both exploratory and explanatory purposes. Although some work has been done in this space (Eaton 2017, Datig and Whiting 2018), there is scope to use data visualisation techniques across multiple channels – digital signage, the web, mobile channels – to both showcase collections, uncover hidden gems, and render visible their cultural value.

3.5.2 Integrating Machine-Actionable Collections into the NLA's cybernetic system

How might we make sense of MAC in the context of the NLA?

Here, a cybernetic analysis is undertaken of these capabilities by outlining each system component – context, people, data, infrastructure, and process – and their salient inter-relationships for the NLA (see Section 1.3). Emerging dynamics are then explored through two examples of the application of MAC capabilities at libraries. We consider these changing dynamics through a continuum from the *internal-facing* (implementation issues, knowable risks, automation, and workforce) through to more *audience-facing* (ethical concerns, unintended consequences, impact on users). Further, we also consider changes to dynamics and feedback loops through the lens of temporal rhythms that govern collections work at the NLA.

CONTEXT	• The NLA caters to audiences now and in 20 years' time (or longer). • The NLA is operating in a time of reduced budget and demands for greater efficiency and sustainability. • The NLA occupies a leadership position in the Australian GLAM sector, and globally, in linking data sets, for example via TROVE. • The NLA's collection has various constraints around usage rights. These usage rights need to be subscribed to by NLA's audience irrespective of context, which requires the NLA to consider what linked data is made available to whom, and for what purpose. • The NLA is making metadata of its digital collection available via the Digital Collections Manager database to the Open Archives Initiative (OAI) – a protocol for harvesting metadata. • The NLA already has already made MAC from Trove data.
PEOPLE	• Library staff prefer developing textured ways of handling the collection that allows for nuanced ways of treating materials, and do not wish for overly prescribed ways of thinking about the materials. • Library staff are largely still navigating their role in mediating the outputs of data models to inform search techniques, interactive interfaces and accessibility to digital collections for the creation of new objects across their digital collections. • Some roles of Library staff, such as curators and custodians of the collection, will need to be reconsidered in relationship with autonomous computing. • IT professionals in the Library will need to be trained in ensuring that the various data models used in MAC technologies are fit-for-purpose and regularly monitored and updated.

PEOPLE (cont.)	• A new class of professionals will be required in the Library. New roles will be required that blend the curatorial aspects of librarianship with data science skills such as visualisation, feature selection, and model building. That is, these skills are not additive; they are in combinations that are qualitatively different. • For the NLA's audiences, MAC technologies enable more of the collection to be accessible in new ways. This means more complex search capabilities on the digital collection, new possible ways of interacting with the collection and new ways of thinking about how the collection might be used for research and analysis enabled by computation. • New roles will be required that blend the curatorial aspects of librarianship with data science skills such as visualisation, feature selection and model building.
DATA	• Data needs to be in (usually structured) machine readable form, which is a transition away from (usually unstructured) human readable data. • Discoverability, interactive interfaces, and other systems that process the collection as data need to be trained on data that is relevant to the materials at hand and the intent for use of the collection. • There is a wide degree of variation in the applicability of models to enable discoverability, interactive interfaces, and other data analytic solutions, given different types of multimedia data as well as the range of script styles and speech practices across languages and time periods. • Datasets need to be made interoperable with one another and with other data sets. • Clear communication is needed for recognising and thinking about how to both manage and communicate the fact that NLA datasets are partial, potentially fragmented, and incomplete. • Considerations are needed around which data will be harvested, for which purpose, by whom and to go in which context.

PROCESSES	• Human review of the output is needed for MAC used for discoverability, interaction with the collection, and also of the new objects created via machines acting on the collection. • In the post-processing of outputs, review of new objects created from the collection, including data visualisations may be required. • Some control and responsibility will be needed over how the collection is used in new contexts through sets of guidelines, and how this is communicated to audiences of these objects may be required. • Training is required to ensure that the various data models used in Machine Transcription are fit-for-purpose and regularly monitored and updated. • Workflow process documentation will need to be reviewed and updated to include MAC technologies. • Guidelines for interactions between Library staff and audience online will need to be reviewed and updated, particularly to include practices around using the collection as a data set. • User rights must be managed alongside the increased accessibility of collection items. • Considerations are needed around compressing data, while sustaining the quality of digital collection items. • Considerations are needed around ways the digital collection can be used as a cohesive data collection, which is actionable as a data set irrespective of the original format. • Considerations are needed around making the collection interoperable with other data sets. This includes data sets that are created within the NLA from the collection or from data (like metadata) associated with the collection.
INFRASTRUCTURE	• Existing infrastructure and workflows will need to be adjusted to accommodate machine-derived metadata. • As digital collections scale, storage capacity will need to be continuously managed, including storage for the metadata that describes it. • Considerations are needed around making the digital collection accessible via cloud computing to allow NLA audiences to use the collection as data with oversight by the NLA. • Energy intensive workloads might drive new operating expense costs.

3.5.3 Salient General Dynamics for MAC Capabilities

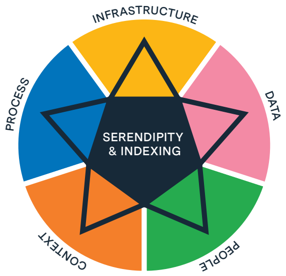

MAC transform each digital item in the collection from a *representation* to an *actuator*. An item is no longer only a digital object or digital facsimile of the physical item, but an actuator that can cause a machine to operate in ways mediated by models and APIs. The action that machines can take on the collection create new pathways through the collection. Commonly, these paths are links drawn between fragments of the collection: parts of an object such as the images in a newspaper are connected to other parts of the collection; audio sounds can be linked by correlating similar melodies; and sets of text that can be connected in multiple ways. This process happens at multiple data scales; at the level of the collection as a data set, and at the level of collection items that can become datasets in and of themselves, and at the level of parts or aspects of collection items that link to other parts or aspects of other items and form new data sets.

Computing activity on the collection also has impact on the ways that subjects of the collection are represented. The ability for more complex search and use of the collection as data to create new objects means that the Library may have less control over the ways the collection is used and the meaning that is found within it. For instance, looking with a contemporary lens across history generates narratives: narratives that fit with the concerns of the moment. History is always a narrative of some type: mechanising this process may limit the potential to read across the grain of history, for discovery that transforms how people think rather than discovery that aligns with how people think *now*.

Computerisation of access to the collection creates data structures that shape and constrain the way meaning can be created and interpreted. For instance, statistical methods of mapping semantic nets of words through natural language processing organises the collection, in terms of search outputs, to a particular hierarchy of associations. While finding associative meaning between collection items and parts of the collection suggests the potential for serendipitous discovery, these connections are constrained by an indexing mechanism. This serves to yoke potential discovery to only that which is “like” —by suggesting like-words, like-images, like-audio and other like-objects – instead of providing a thematic funnel of increasing specificity.

The intent here is that search can be more interpretative; that a research question might bring a wider array of search results based on associative meanings. The serendipity that is inherent in navigating a physical space for information, however, remains elusive to this type of retrieval. While it allows for searching across the collection, through a lateral dimension, rather than just tunnelling down toward specificity, this is still constrained to a net of like-connections defined by an natural language processing system, as one example.

MAC change how audiences can engage with the collection. One potential change is a bifurcation into *Agents & Audiences*. Some audience members will act as creators who use the collection to form new representations of the collection as data, and other audience members will then engage with these new objects. The opportunity of MAC means that long-established ways of understanding how audiences interact with the collection become mediated to some extent by representations of the collection.

Another important dynamic is the extent to which humans are agents on the collection. Representations of the collection change as they are set in new contexts, which includes different ordering systems. Computing access to digital collections changes the ways that these contexts are created/curated. Important to understand in this context is that underlying computing capabilities such as these are organising principles that typically use statistical methods to generate outputs i.e. responses to queries etc. The question becomes, “who is curating the collection for both access and for meaning?”.

MAC also ask more from the Library in terms of digital and data skills. While some traditional IT skills – such as application programming interfaces, database administration, server administration and so on – are still foundational to operations, they will not provide the transformational capabilities required to harness NLA's cybernetic, Machine-Actionable future (Brown 2020, Jones 2020). Examples of roles that may be required include:

- **Data engineer:** Fluent in data acquisition, data cleaning, reformatting, MAC data exchange formats, data interoperability, and data extract, transform and load (ETL), this role ensures the right data is in the right format at the right time for the right purpose. For example, the data engineer might ensure that two MAC datasets are able to interoperate and share data between them or transform large datasets into a different format to be used in, say, virtual reality. Existing roles that might retrain into this skillset would include database administrator, or systems administrator.
- **Curatorial engineer:** Skilled in statistical analysis and machine learning algorithms, as well as curation and collections, this *hybrid* role uses data as an input to create models that guide how collections data is categorised and classified as part of larger cyber-physical systems that generate or create audience experiences. For example, the curatorial engineer might identify patterns in collection data usage and make predictions about collection item access, volume of borrowings, or anticipated foot traffic in a physical site, and shape curatorial activities, in real time, in response. It is possible that staff from either a technological background or a curatorial background could transition into this type of role.
- **Cyber-physical conductor:** Just like the conductor of an orchestra, this role has responsibility for ensuring that the portfolio of systems within the Library of the future harmonise to a given rhythm, providing a delightful experience for the audience. For example, this oversight and strategic role would plan for, deliver, validate, and monitor the many heterogenous computing systems to ensure overall NLA mission alignment. This role generally requires strategic leadership capabilities, a customer experience focus, and computing systems experience.

Additionally, these types of skills currently command a steep premium in the labour market; NLA may not be able to hire directly and may need to identify a digital and data service delivery partner or panel of partners to mitigate against this risk.

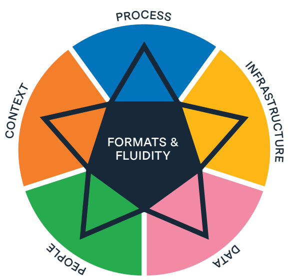

Collections as data suggests a transition to new ways of thinking about the relationship between an analogue and digital collection. The ideal end state for deploying the types of capabilities that are made possible by MAC is a fully digital collection, one in which the entire collection is a data set of interoperable formats where collection items, irrespective of form or era can be seen as a set of like data items. For born-digital content, this means creating uniformity across the collections. It also means considering ways to preserve the longevity of born-digital content, transitioning it to formats that work with data models. For the physical collection, this means creating a digital twin: a facsimile, which in its creation changes the nature of the collection, flattening objects (including books) to digital representation, replicating images in different digital formats, and in doing this, untethering the collection from the systems that index them in physical collections and layering over different organising mechanisms, which are managed by machines.

In reality, the ideal — a collection of data in standardised format (i.e. a single data set) — is atypical at this point in time. Instead, digital collections are frequently multiple sets of data, segregated into sets of likeness. At this point in time, MAC are being applied to just parts of a digital collection.

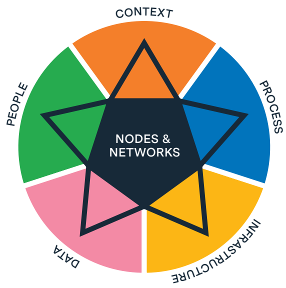

Making MAC available for computation puts libraries in a new, more open context. While access to physical collections will always be constrained by the time and labour costs of moving items, digital libraries have enabled increasingly easier access to broader swathes of content. MAC, however, do not only suggest open access to large amounts of collection data, but enable researchers to make new objects from the collection. These objects are a new representation of the collection or a subset of it; one which is not a facsimile or version of the collection, but a reading of those items. These objects relate in interesting ways to what we might distinguish as the boundary of a library collection. In one way, they extend a library's collection into a larger corpus of information, enabling the collection to link with other data sets beyond the library's hold. In another way, they make the boundaries of curatorial accessibility more permeable by allowing people beyond the library to have open access to their collection. As the number of MAC scales across the global library sector, new relationships between collections are also possible. This raises questions on how open the NLA wishes to be in the context of increasing paths toward networks of libraries and collections as data. We contend that openness, access, and controls for such will be *contingent* upon contexts, use cases, audiences, actors and intents; again, underscoring rather than discounting the value of curation.

3.5.4 Current and Projected Applications of Machine-Actionable Collections Capabilities

Using two current examples and projecting forward to mid- and far-future applications, including those in Scenario 3, *Great Granddaughters Take Back the Library,* we consider the potential new dynamics that will emerge if these capabilities is integrated into the current known dynamics of the NLA collections system and the attendant possibilities, risk, pitfalls, and issues each presents.

Example 1: Infrastructures for enabling collaboration and collective nurturing of new ideas: Library of Congress Computing Cultural Heritage in the Cloud Project

The American Library of Congress Lab (LC Lab), supported by the Andrew W. Mellon Foundation, has been developing a "Computing Cultural Heritage in the Cloud Project", designed to pilot a new workflow and system model for creative use of the Library of Congress collections (Ferriter 2019). The project brings together cloud infrastructures and Machine-Actionable collections to provide opportunities for researchers and artists to work with Library data in novel ways. Historically, researchers have tended to download data from a library website, or access data via an API, to work with on their local machine. In the model proposed by this pilot project, researchers work with data directly in a shared cloud computing environment. This example speaks to the infrastructural and processual opportunities and costs that may be associated with MAC.

The project and its new infrastructures and workflows seek to overcome a few current challenges for the Library and the people engaging with the collections. Firstly, by bringing people onto a carefully controlled shared cloud environment, these researchers can overcome the potentially limiting affordances of their own personal storage and compute capacities for their projects. They can instead work directly on the infrastructure provided by the Library, noting that this requires the user to have reasonable digital skills and capabilities. Secondly, their research work is expected to flow back into and be retained by the Library community. This new research itself becomes a data source accessible through the Library systems to accumulate shared community knowledge. It helps to overcome the challenge for libraries of being able to access results of research in a timely fashion or at all, leading to that data and knowledge created from it being lost to the broader library community.

The proposed governance and funding model is also worth noting. In it the Library and research institutions would share the responsibility for, and cost of, the cloud computing environment. This could equally extend to the digitisation of the collections used in this shared context. This LC Lab novel pilot example may inspire the NLA to identify new business and community collaboration models. These models could include sharing the agency between the Library and people in its broader ecosystem, as well as providing novel curation opportunities and outcomes of national value. Should new cooperative business models be created between libraries, research funders, researchers and potentially those benefitting from the research, then where relevant there may be a financial opportunity to increase the volume and quality of the collection over time. We are already seeing the beginning of some cooperative models of system development and financing such as in the new service model for Trove Partner organisations, Trove Collaboration Services. Over time these could be supported by further diversification.

In the future scenario *Great Granddaughters Take Back the Library*, we point to a Virtual Reading Room that could be a technological evolution from the collaborative cloud environments of the LC Lab pilot. Such infrastructures could afford opportunities for new contextualisation of data and questions that have not previously been asked, in our example the bringing together of disparate Machine-Actionable collection resources that underpinned the knowledge required to recreate the March4Justice. These new collaborations enabling infrastructures, and the business models underpinning them, will be key to value creation from MAC for the Library and the nation.

Example 2: Cultural complexity of understanding canonical spaces: Insights from the Stanford Literary Lab

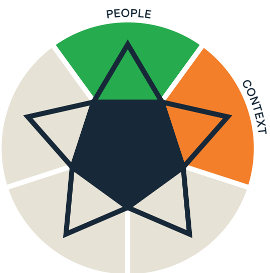

While the MAC cluster may seem to impact most significantly on "expert" library users, like academic researchers, in this example we explore one possible impact of MAC on the broader cultural context and the Australian public at large. This example is taken from the Stanford Literary Lab, a digital humanities lab that includes the literary historian and theorist Franco Moretti among its members (Moretti 2013). Moretti was one of the first to articulate the concept of computational literary analysis with his coinage of the phrase "distant reading" (as opposed to 'close reading') in 2000. One of the proposed benefits of "distant reading", and computational approaches to literature, and historical archives more generally (enabled by MAC), is that these methods enable researchers to move outside of the established literary and historical canons. From this perspective, the creation of MAC may be seen to be a crucial step in diversifying the stories Australia tells about itself with the help of the NLA collection.

Here, we have chosen to focus on an experiment conducted by J.D. Porter of the Stanford Literary Lab in 2018, which suggests how MAC can allow us to (re)define canonical spaces in multiple ways, providing potential opportunities for discoverability systems and the meta-data of the collection items that support them.

In globally connected library ecologies, information on collection items is spread far beyond libraries in a variety of near-real time updated systems (e.g., Goodreads for book recommendations) and those that work on longer time delay (e.g. academic and literary review in the peer-reviewed/professionally edited literature). Porter (2019) gives an example of mapping large collections based on easy-to-access indicators of popularity like Goodreads data against their prestige in literary terms—here a score based on statistics from the MLA International Bibliography. The logarithmic 2D output provides a canvas for interpretation of the canonical space and why different texts might be organised in similar areas of the graph.

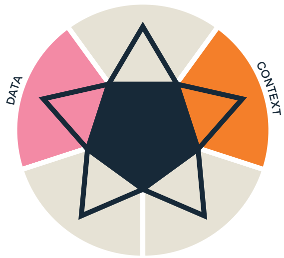

This example highlights the common challenges of incommensurability and the partial (e.g. differently time-bounded) nature of indicators. It stems, in part, from differently curated datasets and the extent of the archive they cover (in these cases often English and/or Western-centric literatures). They bring with them, and remained underpinned by, their user communities and their cultural preferences. However, acknowledging these challenges of current MAC configurations gives libraries and the people who use them the opportunity to not remain content with the current system, but to create new classification and interpretative systems to apply to the data in the MAC. This is certainly the approach of the Syracuse School of Information Studies' "Critical Catalogue" (Clarke and Schoonmaker 2020) that is actively developing metadata, values, and organisational structures that promote and expose diversity (e.g. content from people of colour, women, people of non-cisgenders) in library catalogue searches.

Understanding the values and limits behind the indicators and metadata tags we use to define and provide access to our own national canonical spaces through the NLA will thus be important as we (re)construct and manage MAC over time. These issues are pointed to in Scenario 3, *Great Granddaughters Take Back the Library*, in terms of the value of location metadata, as well as the potential balkanisation and power struggles over preferred indicators, search algorithms and recommendation engines that underpin the ability to access and benefit from locally produced and stored archival material.

3.5.5 Summary of Opportunities, Risks, Pitfalls, and Issues with Machine-Actionable Collection Capabilities

The opportunities, risks, pitfalls, and issues (ORPI) identified through this Cybernetic Systems analysis of Machine-Actionable Collections include:

OPPORTUNITIES

- Create new artefacts useful for research purposes.
- Provide new data sources to popular research sites.
- Improve search and sense-making capabilities.
- Drive national, regional, and international standards for Linked Open Data interoperability standards to support MAC.

RISKS

- Focus on researchers to the exclusion of the public.
- Fail to manage the security risks inherent in large open datasets.
- Reproduce rather than resist patterns of colonialism inherent in collections through digitisation and linkage process.

PITFALLS

- Association with researchers or other organisations using the collection for nefarious purposes.
- Computational analysis may generate readings of the collection which amplify distasteful perspectives from historic and contemporary lenses.

ISSUES

- Data trails created by users of MAC may breach the privacy expectations of NLA's audiences, and individuals or groups named or included in collection items.
- Increased digital and data capabilities and skillsets are required to leverage and capture opportunities.

Scenario 4:
Waking up Ancestors and Walking with the Dead[1]

The National Library of Australia was brought before the Indigenous Voice to Parliament today to explain how a partially fabricated hologram of Aunty Ellen Holsworthy has been held by the library for the last 10 years without the knowledge or consent of the Dharug people.

The hologram, which was brought to light by the Australian historian Martin Bost, was originally created as an experiment in 2022 by the Indigenous technology developer Indigital, making it one of the earliest known examples of the medium. Indigital worked with Aunty Holsworthy and other Dharug elders to develop a hologram that could answer questions about Dharug experiences of the Blacktown Natives Institute.

A spokesperson for the Dharug people, Jasmine Griffiths, said she remembered her mother pioneering hologram technology in the early 2020s, and addressing the Dharug cultural protocols for the spatial web together with Dharug elders at the time before holograms were created.

However, the original hologram was subsequently repurposed by a non-Indigenous developer, who augmented the original interactive model with additional training data sourced from the National Library of Australia, and it is this version that found its way into their digital archives.

"This holographic representation and the information spoken by the hologram of 'Aunty Holsworthy' is colonial," Ms Griffiths said. "The data the hologram was trained on was from resources stored at the NLA relating to the Blacktown Natives Institute which were written by white people. There are much better examples of Dharug holograms being created by our own people. The hologram does not represent Dharug perspectives and in fact breaks cultural law by representing our Peoples without our consent. This hologram is nothing more than colonial voyeurism and should be destroyed."

Bost, who intends to showcase the hologram in an exhibition on the Blacktown Natives Institute, acknowledged Ms Griffith's concerns and those of her community.

1 Michaela Jade, a Cabrogal Woman from the Dharug-speaking Nations of Sydney, the founder of Indigital, and a graduate of the Masters of Applied Cybernetics (2021) created this scenario in collaboration with the authors of this report.

However, he argues that "from a technological perspective, Aunty Holsworthy's hologram is one of the earliest examples of the medium we have. She is a national treasure."

The Indigenous Voice to Parliament is expected to advise on whether or not the hologram should be destroyed by the end of the week.

In the year 2040, interactive three-dimensional holographic videos (known simply as 'holograms') have emerged as a popular new storytelling medium. Holograms, which can be viewed through smart glasses or in dedicated holographic parlours, combine volumetric video capture with artificial intelligence to create what could be described as interactive live action films. Volumetric video capture entails the filming of a performance in three dimensions using hundreds of cameras. The first volumetric video capture studio, Imagine Room, was launched in Australia in 2022.

Cabrogal woman and technology entrepreneur, Mikaela Jade, has been a pioneer in extended reality and artificial intelligence since at least 2014, when she founded Australia's first Indigenous Edu-tech company, Indigital. In 2020, she began working with Dharug elders to develop Dharug cultural protocols for the 'spatial web', a phrase that describes the convergence of our digital and physical worlds in extended reality technologies. In particular, the guidelines addressed Dharug elders' concerns about cultural representations of their ancestors breaking cultural law when – as they rightly predicted – holograms were created by AI, rather than guided by Indigenous people themselves.

Jade worked with the Dharug elders and Aunty Ellen Holsworthy to create one of the first interactive holograms in 2023, which featured Aunty Holsworthy truth-telling about Dharug experiences of the Blacktown Natives Institute. Unlike contemporary holograms, this early example of the medium had very limited interactive features. The only AI element used was natural language processing for parsing user's questions about Dharug experiences of the Institute, which were responded to with authentic answers given by Aunty Ellen and other Dharug elders speaking 'through' Aunty Ellen's holographic image. However, the cost of hosting this hologram was prohibitively expensive and so it quickly disappeared from public view. In 2025, Jade turned her attention to co-developing an Indigenous-owned Dharug language model, that would allow users to interact with future holograms in language.

It is unclear how another technology developer gained access to the original data for the Aunty Holsworthy hologram in 2027, or what their intent was in building a more interactive version that contravened Dharug cultural protocols. The start-up company quickly folded after it failed to raise a Series B round of capital. All that the National Library of Australia knows is that the developers used library data to augment the interactive model and that a copy of the file was submitted via the National E-Deposit Service several years later, when it began accepting audio-visual and holographic content.

The holograph went undiscovered for several years, until 2040, when it was discovered by historian Martin Bost, who had received a $50,000 grant to develop an exhibition on the Blacktown Natives Institute. Bost assumed that showcasing the hologram through a public exhibition at the NLA would bring great joy to Dharug People, and all people across Australia who need to hear the 'real' stories of Australia's past interactions with First Peoples.

3.6 From Future to Present with Transparent AI

This speculative future highlights an emerging need to develop new standards for the use of AI-enabled systems in libraries. As Machine-Transcription and Machine-Actionable collections create new ways for audiences to interact and engage with the collection, people have a right to know and understand how these technologies inform their experiences. This means telling the origin stories of the data which make up a part of AI-systems. It also means making explicit the ways in which AI-systems use this data to tell a set of different stories, which are informed by the internal logic of a system, and its way of interpreting and responding to the context it sits in.

> As Machine Transcription and Machine-Actionable collections (MAC) create new ways for audiences to interact and engage with the collection, people have a right to know and understand how these technologies inform their experiences. Transparent AI is not a technology or a technology capability per se. Rather it is an agenda that aims to empower people to make informed decisions about when and how they engage with an AI-system (Fjeld et al. 2020). This agenda is set against a backdrop of emerging standards and principles coming from both government and industry.

3.6.1 What is Transparent AI?

Transparent AI is not a technology or a technology capability per se. Rather it is an agenda that aims to empower people to make informed decisions about when and how they engage with an AI-system (Fjeld *et al.* 2020). This agenda is set against a backdrop of emerging standards and principles coming from both government and industry. For example, Fjeld *et al.* (2020) have developed a framework on building and implementing AI responsibly that is derived from a survey of literature published by developers and regulators (see Section 2.3). One of the eight principles the authors set out suggests that AI-systems should be designed so that users and audiences understand the intent with which these systems act on in the world (ibid.). In Australia, the Commonwealth Government has determined that people should know when they are being "significantly impacted" by an AI system, and should also be able to find out when an AI system is engaging with them (*Data Availability and Transparency Bill 2020* (Cth)).

The requirement to consider, design, implement, secure, and manage Transparent AI and Transparent AI systems is a complicated one. Transparent AI implicates everything from data (including data collection, data quality, date lineage, and data hygiene) to machine learning techniques and their explainability, to sensors, algorithms and security. As such, Transparent AI requires processes, procedures, and organisational engagement, in addition to existing and new technological features. Even as one of the most frequently cited means of achieving Transparent AI is the notion of explicability or explainability, the questions of to whom something should be explicable and how explicability would be determined remain remarkably opaque.

At this point in time, standards for transparency in AI-enabled systems are still emerging and are not yet fixed as regulation, law, or product certifications. However, regulations and court proceedings such as those linked to the European Union's *Regulation 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)* may rapidly set new global norms and expectations for transparent AI. In Australia, there are already a series of existing AI governance frameworks. These frameworks support improved data practices and promote public understanding of technologies, which may inform how the NLA might seek to shape their own practices around transparent AI. Relevant examples include:

- The Commonwealth Government's own AI Ethics Principles within Australia's Artificial Intelligence Ethics Framework (Department of Industry, Science, Energy and Resources 2020).
- The Fives Safes framework (Australian Institute of Health and Welfare 2020).
- The Data Sharing Principles in the *Data Availability and Transparency Bill 2020* (Cth). These principles are based on the Five Safes framework and have been used by the Australian Institute of Health and Welfare (AIHW) and Australian Bureau of Statistics.

Specific to library contexts, there is also increasing engagement with attempts to establish best-practices around AI-governance. For example, the Machine Learning in Libraries community has developed numerous reports calling for working groups and symposia to surface historical and contemporary approaches to managing bias in library systems, which include a focus on protecting library audiences (Padilla *et al.* 2019). There are also calls to draft and publish a Statement of Values for the use of Machine Learning in Libraries (Cordell 2020). These initiatives reflect and acknowledge a need to approach AI governance pro-actively rather than re-actively in libraries (ibid.)

Recent events suggest that beyond designing these principles, culturally embodying these principles might remain a frequent challenge. Notably, the technology industry has recently set up several boards, taskforces, and toolkits in an attempt to develop and deploy better standards around AI (e.g. Rossi 2020, Google PAIR 2021, Microsoft 2021). The lack of consistent outcomes of these entities suggests that *having* principles and *applying* these principles require different sets of conversations and processes that need to respond to the context within which they sit. The obvious tensions between transparent AI-enabled systems and profit margins continues to play out in quite public ways with legal and regulatory consequences.

Even as Transparent AI remains a diffuse agenda without a clear implementation path (human, technological, or ecological), there are several critical questions the NLA should consider in thinking about developing their own transparent AI practices, such as:

- How can transparent AI influence existing workflows (see: *Quality & Quantity* dynamic)?
- How will transparent AI fit within the broader services provided by the library (see: *Agents & Audiences* dynamic)?
- What are some of the influencing factors for transparent and explainable AI (see: *Nodes & Networks* dynamic)?
- And, how can transparent and explainable AI support the development of a more robust discovery interface (see: *Serendipity & Indexing* dynamic)?

3.6.2 Integrating Transparent AI into the NLA's cybernetic system

From this high-level overview of Transparent AI, its history and how it is used in libraries, we turn to a cybernetic analysis of these capabilities in the system of the Library. We outline the aspects of each system component – context, people, data, infrastructure, and process. We then consider the salient general dynamics among them that we consider important to understand in relation to Transparent AI at the NLA. Following this, we consider the emerging dynamics among system components through examples of the application of Transparent AI capabilities. We also look at the potential changes to dynamics along a continuum from the internal-facing (implementation issues, knowable risks, automation and workforce) through to more audience-facing (ethical concerns, unintended consequences, impact on users). Finally, we consider potential changes to the dynamics and feedback loops at the NLA that occur across multiple time rhythms, which inform collections work at the NLA.

CONTEXT	• Transparent AI is a set of technologies and practices designed and implemented in such a way that oversight of their operations is possible at scale. • Transparent AI will be heavily influenced by external governing bodies and regulations. • Advancements in broader industry sectors, including Google, Microsoft, etc. will influence regulations and standards. • There is currently a lack of existing best practice standards around Transparent AI. • Governing and regulatory bodies progress at a slower pace than AI technology development and implementation. • In the library sector, questions of transparency focus more on the processes and practices of curation, discoverability, accessibility, data quality of collections.
PEOPLE	• Users increasingly desire a better understanding of how services and tools function at a high level. • NLA staff actively work towards supporting and maintaining the provision of a robust search engine. • NLA audiences are evolving, and the way people interact with the collections and each other is changing.
DATA	• Processes around data acquisition, preservation, curation, and access can be influenced by broader governing frameworks. • Mechanisms for surfacing transparency around data will influence existing work processes within collections. • Best practice processes can lead to a reduction in potential risks associated with accessibility of digitised information.
PROCESSES	• Processes around data will be influenced by broader governing frameworks around Transparent AI. • Transparent AI encourages improved data practices and open and transparent communication to promote public understanding and trust in technologies. • Mechanisms for surfacing transparency around data sources, data access and permissions, etc. will influence existing work processes.
INFRASTRUCTURE	• Information technology services represent the core link between technology capabilities and services for the Library. • Transparent AI will look different for back-end services versus services applicable to end users. • Underpinning infrastructure is needed to support back-end and user services.

3.6.3 Salient General Dynamics for Transparent AI Capabilities

Processes around data acquisition, preservation, curation, and access are often influenced by broader governing frameworks around Transparent AI. Mechanisms for surfacing transparency around data sources, data access, and permissions will influence existing work processes, most notably within collections.

There exists an opportunity for the NLA to implement mechanisms around data and information in support of best practice processes which can lead to a reduction in the potential risk the NLA is exposed to with accessibility of digitised information.

The requirements for meeting standards for Transparent AI, or any other standards, have the potential to slow down or speed up processes. The outcome will depend on how mechanisms around transparency are implemented, maintained, and upheld. Transparent AI can function in libraries in a multitude of ways. Libraries are considered a trusted source of well-curated information and knowledge. A lack of transparency around processes and practices of curation, discoverability, accessibility, and data quality will impact how a general audience perceives the quality of library services.

Issues regarding privacy and ownership of data emerge with more accessible digitised materials and search capabilities. The contract between the Library and its audiences becomes more challenging to navigate with the implementation of more complex AI-enabled capabilities.

The evolving contexts and expanding networks that are created through the development and implementation of AI-enabled technologies will require the NLA to reassess its contract with its audiences relative to the evolving ways in which people interact with the collection and each other. Increased accessibility and digitisation, enabled by AI technology capabilities, have transformed the Australian public from from audiences to collaborators who will have strong opinions about collection practices.

Understanding and knowledge of systems is increasingly important for users of library systems. The NLA will have to consider how Transparent AI fit within the broader services provided by the library, from both a user and workforce perspective. This will include thinking about what kind of training should be provided around the implementation and use of technology capabilities, from the perspective of NLA workforce and that of the collaborative Australian public.

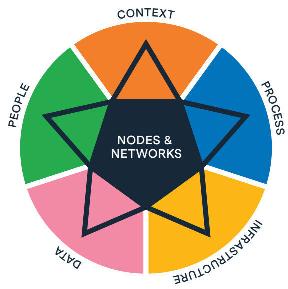

The NLA operates in an open system with multiple dependencies with other systems and organisations. Governing and regulatory bodies impose boundaries through requirements and standards around data sharing processes, accessibility, services, and more.

The complexity and challenges associated with AI-enabled technology capabilities will breed new standards and regulations for libraries. This may mean that the NLA will need to consider the tension between choosing to follow legacy processes versus choosing to adopt standards by other organisations.

Governance and regulation standards around mechanisms around transparent and explainable AI will be heavily influenced by external governing bodies and advancements in broader industry sectors.

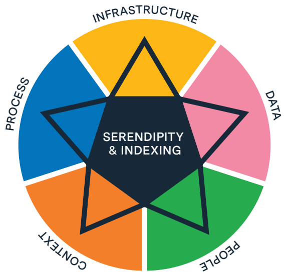

The NLA's goal is not only to collect and catalogue information, but to create opportunities for people to connect information together into new strings of meaning.

There is a need to consider the balance between maintaining room for audiences to create new connections among collection items and generate new knowledge and, describing, organising and creating indexing tools to enable easy discovery of any given item.

Transparent AI and algorithmic transparency can provide an opportunity for the NLA workforce to develop a more robust discovery interface. Transparent and explainable AI mechanisms can support users in better understanding how the service they are using functions at a high-level and consequently, why certain results are being generated. These mechanisms can go a step further than just making information accessible as they can act as a vessel for enhancing search and user experience options through increased public knowledge and understanding of library services.

Fostering discovery and new ideas and insights will require the design of new kinds of search tools. Algorithmic transparency can increase awareness for NLA workforce around potential biases that may be present in the design and implementation of a technology capability, an outcome that is difficult to avoid with human-designed tools. These biases, in addition to the ways that emerging technologies continue to shape search, will fundamentally influence the ways that the collection is curated.

3.6.4 Current and Projected Applications of Transparent AI Capabilities

Multiple reports on AI and libraries have identified that research library staff, drawing upon their expertise in information science, are well positioned to shape policies that ensure accountable, transparent, and ethical AI-enabled systems (Henry 2019, Padilla *et al.* 2019, Cordell 2020). The design of oversight of future ML processes will benefit from library staff's expertise with considering data quality and curation issues. Library professionals are well positioned to offer guidance on the privacy and ethical use questions of personal information policies, and help students develop research skills (International Federation of Library Associations 2020).

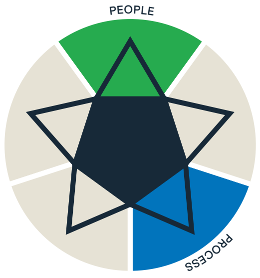

Library staff with contextual knowledge of the collection are also well suited to assist with the task of making automated systems in the library more transparent. For example, in a project where ML is used to generate subject headings, AI development teams can work with library staff to identify existing biases in subject headings developed by cataloguers (International Federation of Library Associations 2020).

Ensuring that the library's people, processes, data, and infrastructure are prepared to support transparency in the library requires understanding how historical biases influence existing systems. This includes having methods and processes in place for understanding the library's ML tools, be they sourced from external vendors or internally developed. It also includes knowing how a vendor has trained their AI, understanding both the provenance of the data used to train the model, as well as the indicators applied to verify and validate the AI system. At this stage, standardisation around these activities is lacking. As such, these methods and processes should be considered, developed, and implemented alongside technology development and implementation within the library.

As the NLA seeks out upgrades or new models of AI-enabled systems, it can understand how the processes developed around such systems can be open and transparent. In doing so, the NLA should understand how it views transparency and oversight in its existing processes and understand how these are impacted by the AI-enabled technologies it is considering. Approaching transparency through the framework provided also creates an opportunity for the NLA to engage other stakeholders in the private and public sector that share an interest in transparent processes. How are other branches of the Commonwealth Government considering transparency and oversight in their ML processes? What experiences and expertise can the library offer to this endeavor?

Transparent AI versus Algorithmic Transparency

Transparency is a multifaceted concept. When it is applied to AI, this concept can broadly be divided into two categories; transparency as it relates to AI and transparency as it relates to algorithms. Algorithmic transparency refers to oversight of the purpose, structure, and underlying logic and actions of the set of steps that comprise an algorithm. Comparatively, transparency in AI encompasses more of a socio-legal notion linked to knowing and understanding. It is a more nuanced concept in which the conceptual framework of transparency exceeds specific technological details and expands to include context and process.

Algorithmic transparency is essential for library employees working directly with technology capabilities. To implement and use these capabilities effectively, library employees will require algorithmic transparency, as well and the tools and processes to manage and maintain it. In contrast, a more general audience interacting with library services, from a user perspective, will not need to understand the intricate and nuanced construct of the underlying algorithm. Rather, they will require a high-level understanding of the capability and how they can best engage with the external interface to achieve a desired result. Alternately, as has been seen in other sectors, it might simply be enough to provide a consistent and trusted certification process that gives users comfort in the technological tools. However, this kind of standardization seems a long way off. Does that need to change and if so, what role should the NLA play in driving such a change?

Algorithmic transparency can be overwhelming to absorb for those not well versed in software development. Additionally, a general audience will not need to understand how a capability sits within the broader system of the library, interacting with the various internal processes and procedures. What they will want to know is how a particular capability influences core services provided by the Library. For example, a general audience may want to know that OCR uses algorithms to detect and extract letterforms and other written markers from images which are then converted into machine-readable text. This information may help them in interpreting information and understanding the nuances of that information source. However, the intricate details of the construct of algorithms that make up OCR and how that process sits within the broader ecosystem of processes and procedures within the NLA will most likely be too granular in detail for consumption by a general audience.

Developed mechanisms that align with existing AI governance frameworks can assist in improving Transparent AI for library users. These mechanisms, such as the example given in Figure 9, should intersect the back-end processes and procedures that deal with the specificities of the technology application, through to the front-end interfaces that a general audience will interact with.

Figure 9: An example of what a Transparent AI mechanism might look like for the NLA.

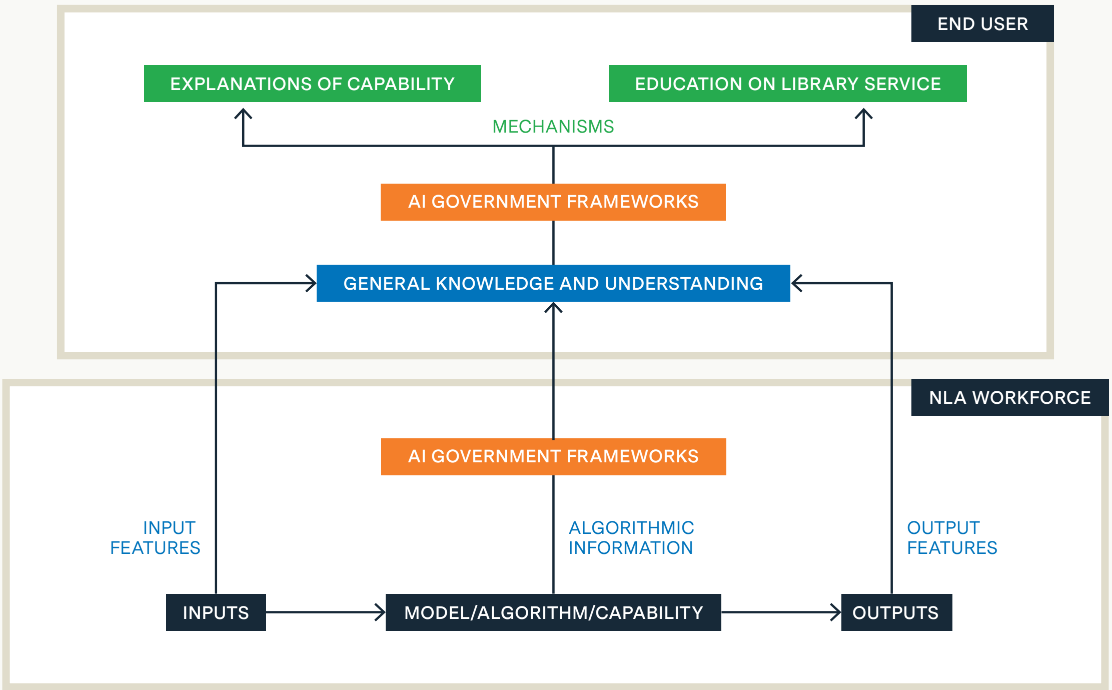

When technology meets humans

In 2016, Microsoft released an AI chat bot in Twitter named Tay. The intention of the chatbot was to conduct research on conversational understanding. It was an experimental program that combined explorations of machine learning, natural language processing and social networks. The integration of these technologies was anticipated to result in a chatbot capable of conducting conversations on any topic through a developed understanding of language based on anonymised public data. The outcome of this experiment was a notable amount of abusive and offensive messages coming from Tay, leading to the removal of the chatbot in less than 24 hours following its release. It was discovered that Twitter users were deliberately sending Tay messages with offensive topics and abusive language. These messages influenced the learning element of the algorithm, resulting in skewed, offensive behavior.

Tay is a poignant example of the unintended outcomes that can arise alongside public participation with technology capabilities, specifically those obtaining data from the public through a mechanism lacking good data quality measures (Schwartz 2019). Like the 3D holographic representation of Aunty Ellen Holsworthy described in Scenario 4, the considerations and thinking behind the data sourcing, development and release of the capability did not fully encompass the potential harm it could – and did – cause. While Tay resulted in direct forms of harm that are more clearly identifiable, i.e., abusive and confronting language, the hologram elicited more complex forms of harm that resulted in the breaking of cultural law and deep offense to Dharug people.

The hologram is also an example of the consequences that can arise from a lack of transparency around processes and practices of preservation and accessibility. Had the process around preservation been more transparent, the Dhurag people would have had the knowledge of the preservation of that information a decade prior. This could have elicited a more positive relationship between the NLA and the Dhurag people, encouraging a collaborative release of the hologram through formal permission. The data quality could also have been improved; thereby upholding the integrity of information made accessible by the Library.

Both capabilities were developed with good intentions; however, they failed to consider the technology within the broader ecosystem of its intended application. Technology capabilities exist within broader systems that encompass people and the environment; influencing factors that shape how technology is utilised and received, which can often be disparate from the intentions behind the development of a capability. Evaluation of technology capabilities should extend beyond the development stage and well into implementation to effectively assess its impact on people and the environment. A "set and forget" approach fails to consider how technology applications can evolve with human interaction.

Here transparency would have had the potential to reduce unintended consequences through an increased awareness of automated decision making. While standards and best practice are continuing to be developed and refined, it can be challenging to understand how to incorporate these elements into capabilities such that they are visible to a general audience.

Bias: a perennial challenge

The integration of AI-enabled capabilities in libraries influences the broader library sector. Accessibility is scaled dramatically, thereby expanding and diversifying the audience of library services. Changing societal needs have the potential to influence how library services are shaped over time. These influences combine with others to weave a complex web of interests, intents, assumptions and biases through the system.

Safiya Noble provides a thorough exploration of some of the underlying biases in search engines in her book Algorithms of Oppression (Noble 2018). In the case of Google, a dominant search engine, commercial interests shape the system significantly, moreso than other factors, such as trustworthiness or credibility of information.

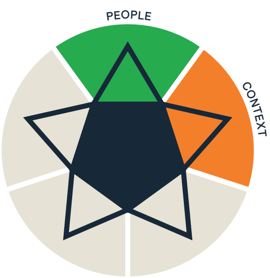

The algorithms developed for search engines, such as Google, are developed by humans. As humans we have different histories, values, assumptions and perspectives. These are ultimately embedded in the systems we develop, with the intricacies of these perspectives evident in the, at times, unintended consequences of technology capabilities. What these biases look like changes with time. Values of the past are deemed unethical in the present and norms of the present may be deemed inappropriate or inaccurate in the future. What is consistent over time is its presence, as bias is perennial.

As such, standards around technology, accessibility, data, etc. will change alongside the change in societal landscapes over time. It is almost impossible to predict every potential bias that may be present or issue that may arise with technology applications. This is made more challenging when the implementation of technology capabilities is considered in isolation of the various factors that influence a system (such as people and ecology). For the NLA, technology capabilities should be considered over the time scales and dynamics that characterise the Library as a cybernetic system.

The NLA benefits from its position as a public institution not motivated by profit, thereby being able to shape capabilities that create public value. However, while the business model of the Library may differ to those of other search engines, such as Google, the Library is not immune to the complexities of the unavoidable human element in technology development. While Scenario 4 is a speculative piece, the notions driving that narrative are reflected in the complexities experienced in the present. Martin Bost assumed showcasing the hologram would bring great joy to Dharug People, and all people across Australia who need to hear the "real" stories of Australia's past interactions with First Peoples. He assumed that augmenting information that was preserved 10 years ago, would be received positively in a present setting. This perspective did not account for the concerns around data quality and transparency of processes and practices of curation and accessibility.

Biases have a direct impact on how AI is both implemented and used. As the collection grows, there is limited capacity to catch and fix errors and often Library staff are not aware of a problem until they are informed. The library sector will benefit from a greater awareness of technology capabilities, most specifically around automated decision making. Understanding the underlying logic behind automated decision making can aid in reducing the impact of biases within algorithms that have the potential to negatively impact Library staff and audiences.

3.6.5 Summary of Opportunities, Risks, Pitfalls, and Issues with Transparent AI Capabilities

The opportunities, risks, pitfalls, and issues (ORPI) identified through this Cybernetic Systems analysis of Transparent AI include:

OPPORTUNITIES

- Develop a robust search interface motivated by a desire to create public value.
- Increased awareness of potential biases in a technology capability.
- Increased knowledge and understanding for a general audience.
- Reduce the risk associated with accessibility of digitised information.
- Reduce unintended consequences through an increased awareness of automated decision making.

RISKS

- Transparency of technology capabilities can make them easier to manipulate and to be utilised in unintended ways that may pose harm towards Library staff and audiences.
- Increased accessibility of digitised information changes the way people interact with that information.

PITFALLS

- Transparency of technology capabilities can make the systems vulnerable to threats and attacks.
- Biases are unavoidable, and always embedded in systems. They can never be eliminated, only managed over time. This management starts, where possible, with clearly articulating, documenting and disclosing existing bias in the system and its component pieces.

ISSUES

- Standards and regulations around transparent and explainable AI are still relatively immature.
- Legacy or newly implemented process may need to be re-evaluated following all future releases of governing frameworks around transparent and explainable AI for libraries.

3.7 Emerging threads and directions of the NLA

In the next half century, the AI capability clusters explored in this section will potentially radically transform the NLA, enabling a more open, more complex, and more inclusive library than has been possible to date. The scope and edges of the cybernetic library will be more diffuse, and its publics, as agents and audiences, will engage with the NLA in settings far removed from the shores of Lake Burley Griffith, and far from the computers, tablets and smartphones that mediate their digital library experience today. The collection will engage Australians in streets, in courts, in bush and more.

From the threads of today's AI capabilities, we have woven visions of the NLA's roles in the social fabric of the future. Today's threads are tomorrow's fabric, though only achievable when considered as parts of complex systems. The cybernetic analysis presented in this section explicates some of the dynamics – the dynamics that combine the components of the system including data, infrastructure, people, data, and context into the specific weft and warp that characterise the library as a cybernetic system. In analysing these components and their emergent dynamics with the introductions of four clusters of AI capabilities, we have identified opportunities, risks, pitfalls, and issues the NLA must address.

Evident in the analysis is that enhanced discoverability does not always correlate to improved sense-making, indeed the rigidity of current applications that fall within the OCR capability cluster can result in inflexible networks in the layers of data ascribed to collection items. Machine Transcription capabilities increase this complexity by enabling new processing of diverse collection items. The broadened diversity and scale of collected materials requires the integration of new pre-processing and post-processing workflows within the NLA. New approaches to discoverability and accessibility for digital collections are engendered through applications that sit within the Machine-Actionable Collections cluster. The collation of these advancements, and the known potential for their continued evolution, requires libraries to incorporate standards and regulations around transparent and explainable AI.

In evaluating the capability clusters against the salient general dynamics and the most likely dynamics and across multiple time scales, we have explicated some of the critical questions and conversations that these new capabilities necessitate. The conducted analysis should inform debate and considered decision making among Library staff around proposed technology capabilities.

"Without libraries what have we? We have no past and no future."

Ray Bradbury
(1998) (cited in Fjeldsted 2016)

"…the Library collects documentary resources relating to Australia and the Australian people so that the Australian community can discover, learn and create new knowledge… We create public value by increasing economic, social and intellectual wellbeing."

National Library of Australia
(2020b, pp. 1 & 7)

04

The Future: A Reference Guide

In Section 4, we offer a cybernetically-grounded technology reference guide for evaluating prospective and emerging AI applications the NLA may consider incorporating into its collections system. This Cybernetic Star Guide (the Guide) comprises both a conceptual perspective on the Library as a cybernetic system as outlined in Section 1, as well as a tool for supporting confident decision making that leverages the general dynamics and the AI capability-specific dynamics among the NLA's system components explored in Section 3.

The Guide is designed to enable a comprehensive examination of the opportunities, risks, pitfalls, and issues with a proposed AI application from the perspective of the internal work processes of the Library (implementation issues, knowable risks, automation, and workforce) through to more audience-centric perspectives (ethical concerns, unintended consequences, impact on users).

Taken as a whole, the Guide complements existing approaches and processes for evaluating proposed technology applications already used by the NLA's Digital Branch and Collections Branch, by supporting decision makers' capacities to clarify the potential new dynamics that will emerge if, and when, the NLA integrates tools and processes enabled by promising AI technology applications into the organisation's existing systems to collect, preserve and make accessible library materials.

4.1 From Evaluating AI Capabilities to Evaluating AI Applications

In our analysis in Section 3, we examined clusters of AI *capabilities* – the services a class of AI technologies provide – as well as knowledge that can be applied to specific technology products and processes. This allowed us to outline a pathway for the NLA to evaluate ORPI with AI over the time scales and rhythms that characterise the Library as a cybernetic system. In this section, we provide a process and set of tools the NLA can use to evaluate specific AI *applications* instantiated as products and processes in the market or deployable now, whilst still attendant to the time scales and rhythms that characterise the library as a cybernetic system.

How do we propose to clarify potential new dynamics introduced by AI applications? The Guide follows the same logic that informed our analysis of AI capability clusters in the Section 3 (Figure 10), and has a range of possible applications (Figure 11).

The process starts with defining the broad capability and the technical specifications of the proposed applications to the Library (Figure 10). This is followed by careful consideration of those aspects of each system component that initially appear to be most relevant; as the analysis unfolds, these aspects will be updated as needed.

Work then proceeds to elucidate the dynamics among these components. If decision makers get stuck, they may consider using future casting (as was deployed in Section 3) to think about longer term outcomes and dynamics. Synthesis and analysis of these complementary inputs is then used to generate ORPI for the proposed application. This process will be further detailed in this section, and a full set of supporting materials, including worksheets and templates, is included in Appendix C: Cybernetic Star Guide.

The process is useful in early roadmap planning activities – before decisions for resourcing are made. It is also an important tracking tool to update periodically when a new application is introduced into the NLA and stabilised. Finally, we imagine it is also useful as an application reaches end of life or is significantly upgraded (see #2 in Figure 11), as a means of documenting lessons and best-known methods for future applications. The following step-by-step walkthrough of this process is based on these scenarios, though we acknowledge that the tool has broader applicability for decision making when considering the NLA as a cybernetic system.

Figure 10: Process of analysis applied in preparing this report, indicating the relationship between identifying research areas, capturing component information, and generating future scenarios to explore general dynamics and OPRI.

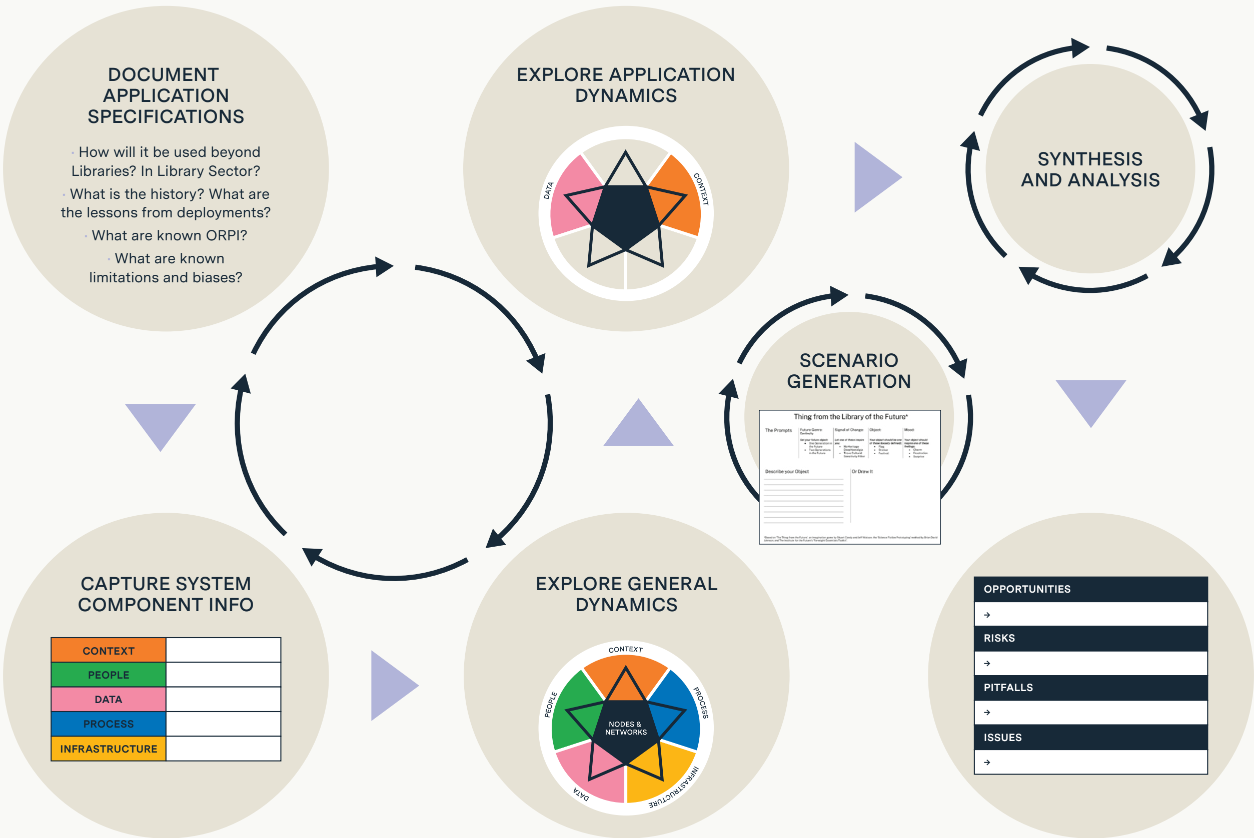

1. Upgrade Application/ Existing Process	2. New Application/ Existing Process	3. New Application/ New Process
Upgrading infrastructure and data resources with an AI-enabled application to enhance existing processes and/or services.	Deploying an AI-enabled application that is new to the NLA in order to enhance existing processes and/or services.	Deploying an AI-enabled application that is new to the NLA in order to create new processes and/or services.

Figure 11: Three Scenarios in which the Cybernetic Star Guide will be useful at NLA.

4.2 A Walk Through of Each Step in the Cybernetic Star Guide

4.2.1 Document Application Specifications

The process starts with identifying the broad AI capabilities, down to the technical specifications of the proposed applications. This type of documentation is already part of the technology evaluation process that the NLA's Digital team undertakes when taking on a new project and/or library collection capability update. For the cybernetic analysis, particular attention is paid to contextualising the application with desktop research and potentially through conversations with other organisations who have experimented, and/or successfully applied and scaled the same or similar applications in their organisations. Information useful to gather at this stage that will inform exploration of dynamics in later steps include:

- *What is the historic trajectory of this application and of the capabilities it includes?*
- *Where and why did they originate, and how have they been used elsewhere? How have they been used in other libraries?*
- *What are the known limitations and how are they accounted for in deployments?*
- *What are the known ORPI from deployments beyond the library sector? The ORPI in the library sector?*
- *If an example of a deployment of the application is available, what can we learn about the system components (context, data, people, process, infrastructure) that enabled or hindered the success of the application?*

4.2.2 Capture System Component Information

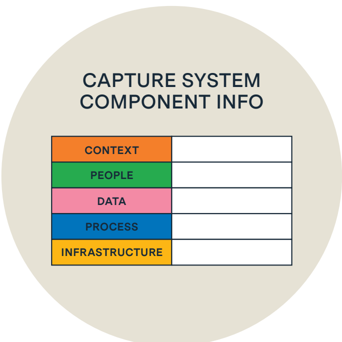

Having documented the application specifications, the next step is to capture the system components at the NLA that are most relevant in planning the integration of the application into the existing Library system. This initial capture of the aspects of the system components should be informed both by the information gathered in the preceding Application Specification step, and by information gathered internally from the NLA. Strategy documents, process guidelines, and other documentation are one source of insight; inputs from staff from multiple branches within the Library are equally important. Input and feedback on the system components should be sought from multiple sources within a branch, both at the level of strategy (leadership) and the level of operations (staff directly responsible and involved in current processes that will be affected by the new application).

Input from staff can be generated in a number of ways, including through working meetings or workshops, through one-on-one meetings, through a digital platform, or within existing workflows. Approaches that allow for cross-branch staff discussion (between Collections, Digital, and Collaboration branches, for example) about the components and the dynamics among them are useful to surface perspectives and dynamics that may not be obvious in siloed conversation, or among staff all working within the same branch. A worksheet that can be used in workshops or conversations to gather inputs on system components is included in Appendix C.3.1.

4.2.3 Explore General Dynamics

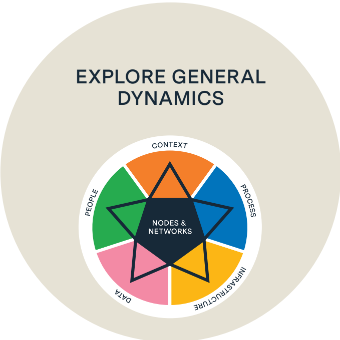

Work should now proceed to elucidate the dynamics among these components, including the signals, flows of work, and data that will be added, altered, or erased by the proposed application. In this elucidation, NLA decision makers will need to consider both the general dynamics that characterise the collections system, and the dynamics specific to the application. They should consider how dynamics will play out over time – not just how the solution will impact or improve a process now, but in the mid and longer terms. This is an analytical skill at which NLA staff are particularly adept.

Informed by the characteristics of the system components identified as relevant in the preceding step, decision makers now gather inputs on the potential dynamics that may be generated by the proposed application. We recommend starting with the general dynamics of the collections system identified in Section 1, including: *Quality & Quantity*, *Formats & Fluidity*, *Agents & Audiences*, *Nodes & Networks*, and *Serendipity & Indexing*. For each of these dynamics, we have identified a generative prompt, in the form of a question, that is a simple starting point for moving from the general dynamic to clarifying how it manifests for a specific application.These can be found in Appendix C.1.1.

For each general dynamic, we also provide sub-questions to the overarching prompt to help decision makers think through potential emerging dynamics, and suggested methods for answering these questions. These are given in Appendix C.1.2.

4.2.4 Explore Application-Specific Dynamics

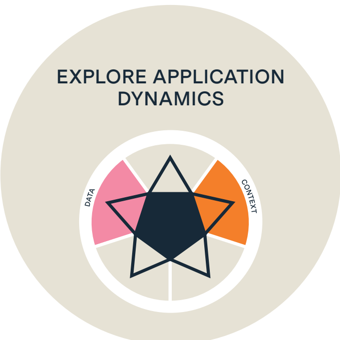

Following an exploration of potential changes to the general dynamics of the collections system, a next set of overarching prompts and sub-questions focused on dyads of system components are used to generate application-specific dynamics. Similar to the general dynamics, we provide a series of generative prompts for the ten possible dyads of the system components. These are summarised in Appendix C.2.1. For each dyad, sub-questions and proposed activities are provided in Appendix C.2.2.

As in identifying system component information, input from a range of staff will be useful here as well. Ideally, these activities will happen in quick succession, at the same workshop, working meeting, or one-on-one meeting. Again, approaches that allow for cross-branch staff discussions are preferable as they will generate perspectives on dynamics from multiple points in the collections system.

As dynamics are identified, the System Component Information should be updated. As application-specific dynamics are generated, the general dynamics should be updated as well.

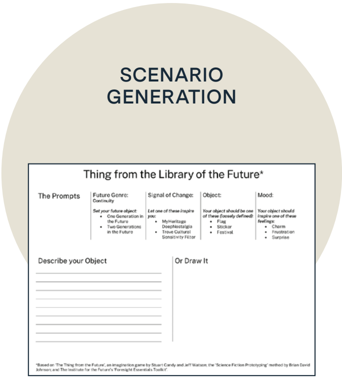

4.2.5 Generate Future Scenarios to Inform Strategies around Potential Dynamics

If, while generating relevant system component information and the subsequent general and application-specific dynamics, decision makers or staff from whom they are soliciting inputs feel 'stuck' or unsure of how to generate mid and longer time horizon dynamics, they may consider using speculative future scenarios – as was deployed in Section 3 – to think about longer term outcomes and dynamics.

A full description of the process, including instructions on how to generate scenarios, is included in Appendix B.

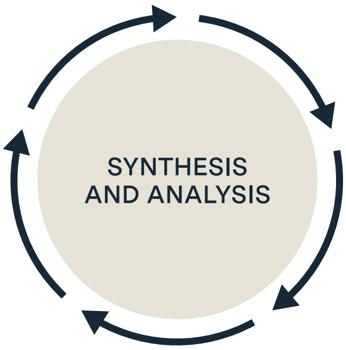

4.2.6 Synthesis and Analysis

Having gathered information from multiple sources within and beyond the Library, the task remains to analyse the information that has been created. Whilst it is tempting to think that the information gathering phase will generate answers, it is necessary to apply a set of analytic lenses to the material that has been gathered. This will take time.

Analysis of materials should take into account the perspectives from which they were generated. Component descriptions and dynamics may be described differently depending on where the creator sits within the collections system. In other words, the perspective of staff from the Digital branch may differ from those in the Collaboration branch; both perspectives are important, even if contradictory or at odds. Triangulating data and finding supporting evidence across data types (desktop research, case studies, internal workshops, etc.) will also be important. The exact nature of this analysis will depend on the materials generated, and we leave the details of this analysis to the expertise of the highly trained staff of the NLA. We note that this cybernetic analysis is intended to complement the library's existing practices around strategic planning, technology and collections roadmaps, and technology readiness analyses.

Periodic reviewing and updating of this process is recommended in tandem with these existing planning workstreams. As we have noted, the NLA is a dynamic open system – we expect that conditions will change, and updates of these analyses should be ongoing.

4.3 Cybernetic Star Outputs

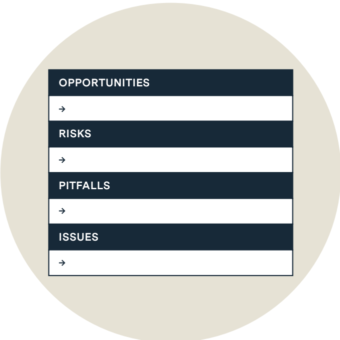

By working through the process outlined in the Guide, the NLA Collections branch can produce a summary of the opportunities, risks, pitfalls, and issues associated with implementing a particular AI-enabled technology. During the course of analysis, decision makers should be attentive to data that helps identify:

Opportunities – *Possibilities afforded by the technology capability.* These will surface strongly in the first step of the Guide (see Section 4.2.1). Measured critique of these opportunities through the following stages will generate more realistic framings of possibilities for the NLA.

Risks – *Potential negative consequences of the intended application of a technology capability.* Risks will begin to surface strongly as general and application-specific dynamics are explored (Sections 4.2.3 and 4.2.4 above) and the characteristics of the NLA's cybernetic system become better understood – in other words, when the abstract opportunities afforded by the application meet the realities of the system components.

Pitfalls – *Potential negative consequences of the unintended application of a technology capability*. Like risks, these will come into focus as systems dynamics are mapped. Particular attention should be paid to mid- and longer-term dynamics.

Issues – *Known challenges with implementing the technology capability.* These will surface strongly in the first step of the Guide (see Section 4.2.1) and are also captured in other technology evaluation and suitability processes that the NLA's Digital branch already uses. The Cybernetic Star Guide can complement these processes by bringing into focus those issues that emerge through interactions between *infrastructure* and *data* system components on the one hand, and *people*, *process* and *context* components on the other.

The output can serve as a communication tool to summarise the findings from the Guide and facilitate discussions between the NLA Collections branch and other parts of the NLA – or beyond – regarding implementation of the proposed technology. A template for summarising ORPI is included in Appendix C.3. With this process, the NLA can make informed decisions about AI applications available now, while keeping a clear eye on the future implications of adopting a broader set of AI capabilities. Taken as a whole, the Cybernetic Star Guide enables the NLA to create "powerful descriptions" of the future dynamics of the library as a cybernetic system – the most important result of any cybernetic analysis.

“The trust or confidence that can be established between writer and reader is real, though entirely mental; on both sides it consists in the willingness to animate, to project one’s own thinking and feeling into a harmony with a not-yet-existent reader or a not-present and perhaps long-dead writer …. When the power of the relationship is not abused, when the trust is mutual, as when a parent tells a bedtime story or a teacher shares the treasures of the intellect or a poet speaks both to and for the listener, real community is achieved; the occasion is sacred.”

Ursula LeGuin
(1985)

05

Conclusions

Library collections are vast and growing. As physical materials are digitised and libraries acquire born-digital materials, library resources will continue to be constrained relative to the materials that the Library wishes to collect and curate. Libraries around the world are also responding to changes in how people create and access information. At the same time, the development of AI-enabled technologies is catalysing different applications based on the new world of information. These technologies are proposed to increase the rate of processing, enrich collections to aid discoverability and serendipity, and develop new modes for audience interaction with the collection. Whilst meta-data will sometimes flatten out the unevenness of collecting practices, time itself will continue to be uneven: the timeliness of when materials are collected and made accessible is not straightforward, depending on the material type and the speed by which it can be described in the catalogue.

The new AI technologies and capabilities have the capacity to shape our futures. Today, and in the past, library professionals live with the decisions made by their predecessors about what and how to collect. This has frequently resulted in heterogeneity, unevenness and unexpected blind and blank spots from the vantage point of their present. Library professionals in the NLA grapple with making good decisions about what to collect and how to make it knowable to audiences who will want to use it and make sense of it in the future. What might seem appropriate in one era can be wildly inappropriate later. In this way, the NLA is always in dialogue with the future.

It isn't just about the future. Library staff, and especially those whose work touch upon collections, must think in past, present and future tenses at once. While the Library has a responsibility to provide access to current generations, it also has a responsibility to preserve things for future generations. Items in the collection have an audience now, and an equally valued audience in 20, 50, and 100 years' time. Current and future audiences need to be accounted for in decisions made about the collection today. Thus, the items the NLA collects in the present are curated through the lenses of possible "future histories" and possible "futures-to-be".

Items that have already been collected are frequently revised by new audiences and their meanings are reinterpreted through an evolving contemporary lens.

History is, as we know, cultural, contextual, and contingent; so is our sense-making of it. AI-enabled technologies, within the Library's cybernetic system, will ensure that meaning and sense-making remain emergent properties. They will change and adapt and reconfigure over time as new materials are added to collections and as people interact with the collections to make new meaning.

In this final section, we provide a glimpse into the emerging socio-technical, business, and regulatory trends that may be relevant to the National Library in coming years. There is always a temptation to read these future glimpses literally, but we would caution to regard them more as pointing to outlines of possible futures, not necessarily about its precise shape. Think of them as signals, akin to the ones used as inputs for the speculative stories throughout this report (see Appendix B). The future, in its indeterminacy, invites our action: "action takes over where episteme fails" (Candy 2021, p. 6). We hope to prompt questions around what the future *could* look like, before asking what it *should* look like.

Library staff, and especially those whose work touch upon collections, must think in past, present and future tenses at once. While the Library has a responsibility to provide access to current generations, it also has a responsibility to preserve things for future generations. Items in the collection have an audience now, and an equally valued audience in 20, 50, and 100 years' time.

5.1 Glimpses into the Future/s

> "Doing digital preservation requires thinking like a futurist. We don't know the tools and systems that people will have and use in the future to access digital content. So, if we want to ensure long term access to digital information we need to, at least on some level, be thinking about and aware of trends in the development of digital technologies. This is a key consideration for risk mitigation. Our preservation risks and threats are based on the technology stack we currently have and the stack we will have in the future so we need to look to the future in a way that we didn't need to with previous media and formats."
>
> **Trevor Owens**
> (2018, p. 9)

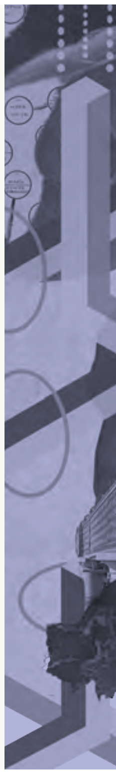

In 2003, American author and futurist William Gibson was interviewed by the *The Economist* and asked his views about the coming decades. One imagines that the journalist hoped for a vivid description of novel and unique technologies, and complex new behaviour and social systems. Instead, Gibson offered the following statement: "the future is already here, it is just unevenly distributed" (Gibson 2003, p. 152). On the face of it, this feels like a deflection – a clever way to avoid the trap of making predictions about the future that will later be seen as very foolish indeed. However, if you dwell with that statement just a little while longer, you start to see it as both a provocation and an invitation. It is a pointed reminder to pay more attention to the present, and to see it slightly differently.

In what follows, we describe five moments in which the future might have been on display in the present. Whilst none of these moments are immediately about libraries, they have threads that pull through to the library's role in the social fabric of the future.

Moment 1: The Portrait of Edmond de Belamy

On the 25th of October 2018 and less than a ten-minute walk from the New York Museum of Modern art, the Portrait of Edmond de Belamy was sold for US$432,500. It is a strange image; a blurred figure in a dark suit with a white collar stares out of the canvas; there are shadows and an odd light. The face appears misshapen, and the jaw swollen. But until you search for the signature block and read a line of mathematical symbols: "*min G max D x [log (D(x))] + z [log(1 – D (G(z)))]*" (see Vincent 2018), it is not immediately clear that this is not human-made. Instead, the Portrait of Edmond de Belamy is the first piece of AI-generated art to go to auction; it seems unlikely to be the last.

It was made by an art collective in Paris using a Generative Adversarial Network (GAN). A GAN is a class of ML techniques (Goodfellow et al. 2014) that pits two neural networks against one another, with the aim to create new data that can pass as the old data. It functions akin to a two-player game: on one side is the "generator" which creates a new data object based off the training data that it has been given, and on the other side is a "discriminator" which tries to determine if the data is real or was generated by its opponent. For the generator, the aim of the game is to trick the discriminator into thinking that their generated data is real (ibid.). In the case of the Portrait of Edmond de Belamy, the GAN was borrowed from an open source library (Vincent 2018), and the training data set consisted of 15,000 portrait paintings created between the fourteenth and twentieth century (Christie's 2018). The Portrait of Edmond de Belamy is the image that the discriminator imagined was real.

Edmond de Belamy is a born-digital object, created from code challenging code. Unlike so-called "deep fakes" which use ML and AI to deliberately trick or fool us for malicious intents (Karnouskos 2020, Mirsky and Lee 2021), or commercial applications like ''Deep Nostalgia" (Hern 2021), which uses very similar computational techniques to delight and surprise humans, the Portrait of Edmond de Belamy is a digital object that code believed was real. Do motives matter – trickery, wonder, simulacrum? And what does it mean to contemplate a world not just of collections of born-digital objects but of their complex adversarial reproduction? And our reactions and engagements and relationships to, with and through them? About training data, and its possible constraints, consents, and limitations and desires? And how might we think about those digital kin and their creation and curation?

The Portrait of Edmond de Belamy generates a number of critical questions for libraries:

- Do motives for the creation of AI-generated collection items matter in how they are handled by libraries?
- What will it mean to contemplate a world not just of collections of born-digital objects, but of their complex adversarial reproduction through neural networks and other future AI approaches?
- How will we engage, react, and relate to, with, and through such collection items?
- How will we think about training data, it's possible constrains, consents, limitations and desires?
- How will we think about these digital kin and their creation and curation?

Moment 2: "Buying myself back"

In May 2021, Emily Ratajkowski sold a photo of herself standing beside a photo of herself, the latter photo taken, she claims, without her explicit consent and reproduced in ways she does not endorse (Ratajkowski 2020, Dwyer 2021, Elan 2021). The complicated appropriation/re-appropriation image was purchased at a Christie's auction for US$145,000. Ratajkowski said she was challenging the ways in which images were commodified, circulated, and consumed. This image found its way into the broader conversation, not so much for her critique of the male gaze and its sometimes-violent enactment, but because she sold the image as a non-fungible token, or NFT.

NFTs are units of data stored on digital ledgers – unique and non-interchangeable. They are effectively one-of-a-kind certificates of ownership for various digital assets, which can then be bought and sold like property. NFTs operate under similar cryptographic protocols to technologies like Bitcoin, in that a record of this ownership is stored on a secure blockchain sustained by thousands of computers (Greenfield 2021, Jones 2021).

While the first NFT was prototyped in 2014, the market for NFTs have experienced a boom in 2021, particularly in the realms of digital art, sport, and celebrity experiences. Notably, American artist Mike Winkelmann, also known as Beeple, sold an NFT of his digital artwork entitled "The First 5000 Days" for US$69.3 million in March 2021 at auction (Kastrenakes 2021). A month later, Kate Moss sold NFTs for three short videos designed to capture moments of reality and authenticity in the supermodel's life (Yotka 2021).

NFTs are not without their critics: some point out that they are highly energy intensive and therefore pose a significant environmental cost, while others argue the NFT bubble will soon pop, rendering most NFTs worthless in a few years (Cao 2021, Greenfield 2021). This latter group point out the ways in which this market is highly speculative and volatile.

The market economics of NFTs are clearly unstable; and the notion of a well-calibrated ecosystem seems unlikely in the short term. The future that glimpses through here could be less about the matter and more about the medium. After all, NFTs would surely raise interesting challenges for libraries: should - and if so how - might a library contemplate collecting, storing and curating these kinds of born-digital objects? What other kinds of networks and technologies would that implicate? How do NFTs shape and reshape underlying ideas about copyright and IP; and how do they flow, or not flow, across jurisdictions and borders? The ways in which NFTs might be assembled and reassembled into other configurations is also worthy of contemplation, including how such assemblages could occur and what new kinds of economic, social, and regulatory forms will be created.

NFTs surely raise interesting challenges for libraries:

- Should - and if so how - might a library contemplate collecting, storing, and curating these kinds of born-digital objects?
- What other kinds of networks and technologies would such collection practices implicate?
- How do NFTs shape and reshape underlying ideas about copyright and IP; and how do they flow, or not flow, across jurisdictions and borders?
- How might NFTs be assembled and reassembled into other configurations and what new kinds of economic, social, and regulatory forms will be created?

Beyond the specifics of this current generation of NFTs, and their very particular technical instantiations, it is interesting to consider how this orientation to unique, stable, digitally assigned data relates to equally constructed digital simulacrums and what sets of cultural anxieties are being enacted, and how those might flow through to the broader imaginings of library collections and sense-making.

Moment 3: "Our goal is to make money, and not creating problems for society"

It is the pictures of plastic bags full of petrol that make it clear there is a crisis unfolding in the United States. It starts out with a news story reporting on yet another cyber-attack on yet another American company; it is familiar news in 2021. In the preceding two years, there have been repeated, high profile, front page news stories of attacks on power companies, and hospitals, and even whole cities. Baltimore, for instance, has seen attacks on its city administration and its schools (Chokshi 2019, Paybarah 2020).

This time it was not a city school district, or a power plant, but a company that owned a pipeline that supplied fuel to the whole of the eastern seaboard. Within days, the United States Government's Federal Motor Carrier Safety Administration had issued a regional emergency declaration, designed to keep fuel supply lines open during the largest ransomware attack of its kind to date. Ultimately the outage would last for five days and there would be fuel shortages for far longer. The company, Colonial Pipeline, paid a ransom to the hacker organisation Dark Side, who said their goal in the attack was purely financial (Menn and Satter 2021). A month later, the United States Department of Justice announced that it had recovered nearly two thirds of the Bitcoin ransom it had facilitated paying one month earlier. As executives from the Colonial Pipeline have continued to unfold this story, it is has also become clear that the data stolen and the systems compromised were customer-facing, not infrastructural.

Many of these cyber security incidents of the last two years have been ransomware attacks where critical data or files are seized, encrypted, and held remotely pending payment of a ransom, at which time a decryption key might be proffered and the process of decrypting the data and reinstalling the systems can begin. Of course, it is rarely that simple – disrupted systems can take days, weeks or months to be re-established; data can be corrupted through the encryption and decryption process; backups can be out-of-date or absent, or even when present can take significant amounts of time to reload; ransom payments can be impossible to make, or immoral or illegal to consider. The demand for payment in various forms of cryptocurrencies can make policing difficult, and enforcement almost impossible.

The sheer volume of these kinds of ransomware attacks has been significant, with growth in the size and quantity and scale of the operations, and the clear articulation of a business model. As the RobinHood note read to the city of Baltimore, "We won't talk more, all we know is MONEY!"' (Chokshi 2019). Some analysts in this space conjecture that whereas cyber-attacks once required a degree of computational literacy and skills, a whole new ecosystem has grown up around this space; it is possible to imagine ransomware as a service.

This model works because of the increasing reliance on data, digital systems, and networked infrastructure, and because of the paucity of good security protocols that can be feasibly administered in complex, heterogenous organisations, and systems. It is less about how to prevent such attacks, and more about resilient systems and the plans to re-establish the organisation afterwards. Libraries have had to worry about fires, and floods, and the damage to their physical buildings, and they have long considered how to secure valuable physical objects – behind glass doors, padlocks, chains and even guards (Bouwman 2018). Now they need to consider a whole different set of vulnerabilities and perceived values and economics.

As collections of digital and digital-born objects increases, and the reliance on digital tools and digital infrastructure is normalised, there are new challenges. At least one artist, Phillip David Stearns, has found a way to transform these attacks into unexpected new objects and to re-materialise the digital-born artefacts into physical ones (Stearns 2019). In 2019, he launched a textile collection that recalled jacquard weaving as a way to visualise malware samples *(ibid.)*. The digital is born as physical once again.

Ransomware attacks work because of the increasing reliance on data, digital systems, and networked infrastructure, and because of the paucity of good security protocols that can be feasibly administered in complex, heterogenous organisations, and systems. Recent examples raise interesting challenges for libraries:

- How will libraries build resilience in the face of such attacks?
- How will libraries plan to re-establish their organisations afterwards?
- Libraries have had to worry about fires, and floods, and the damage to their physical buildings; how will they now consider a whole different set of vulnerabilities and perceived values and economics?

Moment 4: The Pigs are Hiding in the Tall Grass

Despite all the remarkable changes in technology in the last five years, the sound of drones is still very precise and distinctive; a persistent metallic buzzing, a little bit like a mechanical blowfly somewhere just out of reach, and surprisingly loud. This sound has been known to vex all manner of animals, and there are many videos on the internet of animals attacking drones, from wedge-tail eagles to long-horn sheep. Of course, this is complicated because, for the most part, drones are there to help. They represent the next step in long-standing local and global activities to track and monitor animals using various kinds of technology. Sometimes this monitoring has been used with wildlife: to help establish population numbers, habitat, movements, and migrations, and also to provide forms of protection from predations and changes in the landscape (for instance Benson 2010). In other places, tracking devices have been hacked and the location data exploited to hunt endangered animals rather than protect them. In yet other places, smart tags and location-sensitive devices have been used to automate things like herding and milking, and make it possible, for instance, for cows to milk themselves.

Here in Australia, feral pigs are raising different questions. On the beaches of Queensland's northern Cape York region, feral pigs are responsible for finding and eating vast numbers of turtle eggs, placing the survival of these species in jeopardy (Felton-Taylor and Claughton 2015). In the past, companies have used drone technology and thermal imaging to identify feral pigs through thick layers of forest (Felton-Taylor and Claughton 2015). For the past six years, scientists from the CSIRO and Indigenous rangers have using drones as part of a collaboration to manage these feral populations and monitor turtle nests in order to reduce the harmful impacts of predation (Felton-Taylor and Claughton 2015). This can be time-consuming and laborious. Through a recent partnership between CSIRO, Indigenous Rangers, and Microsoft, AI-enabled image detection algorithms have been applied to aerial survey videos to automate and accelerate these processes (The Commonwealth Scientific and Industrial Research Organisation 2021). Initially this seemed successful. However, it has been noted that some pigs have learned to associate the sound of the drones with both danger and the presence of turtle eggs, thereby circumventing the original intent of the technology.

In all these instances, non-humans are a part of broader complex systems; they are the objects and subjects in a technologically determined world. In many cases, new and emerging technologies and capabilities have rendered such non-humans visible in unanticipated ways. We love these moments – real animals, real lives, unvarnished, unadorned. We are seduced by the idea that we can *really* see the world. This did not mean, however, that we were ready for the lessons that were then in plain sight – loss of habitat, pollution of waterways, near-extinction events.

In the worlds we are building and the ones we imagine, it is tempting to give constant primacy to human actors; the control in the cybernetic system resting in the hands of people. However, the world is full of other kinds of actors and actions. How we account for them as part of collection practices – "can Naruto leave his selfie to the NLA" – would just be the first question we might ask, not the last. How do we manage the appearances of non-human actors in the margins and the peripheries of so many collection habits and activities? And could we imagine them as part of our sense-making apparatus? Or, indeed, should we?

Queensland's feral pigs, self-milking cows, and crested macaque photographers remind us that the world is full of actors and actions beyond humans, and raise interesting questions for libraries:

- How can we account for non-human creators as part of collection practices?
- How will libraries manage the appearances of non-human actors in the margins and the peripheries of so many collection habits and activities?
- Can we imagine non-human actors as part of library sense-making apparatus? Or, indeed, should we?

Moment 5: The Internet Smells Like Old Books

Libraries and their collections used to smell; some of them still do. In fact, there is a particular smell in used bookstores—and sometimes in the library stacks – a kind of dusty decay. In some places that smell is given off by the slow disintegration of an organic material used in bookbinding in the middle of the nineteenth century (Godfrey 2018). One source of that organic material was the gutta percha trees of Sumatra, the Malay Peninsula, Singapore and Borneo. They are tall trees, up to forty meters tall, and one meter in diameter with dark, glossy green leaves with glowing yellow undersides set against dark red bark and clusters of white flowers. Their sap has been harvested and used for many things. As it turns out, it doesn't make particular good bookbinding adhesive; but it does make for an excellent electrical insulator (Wong 2016).

In the mid-nineteenth century, Victorian engineers struggled to identify a material that could adequately insulate the first submarine electric telegraph cables. They wanted to submerge live wires in the ocean and keep them safe; many of their attempts failed. Then they found gutta percha, and that inert rubber was responsible for accelerating the development of a global telecommunications network, which in turn laid the foundation for today's fibre-optic network and the global web (Tully 2011, Wong 2016). Gutta percha was eventually replaced by polyethylene in the 1950s, but there are some cables that are still wrapped in it, underneath layers of other more modern things, and deep in the ocean. Those cables that once carried telegrams, now carry the internet.

In a strange kind of way, that means the internet smells like old books, although it doesn't smell at all. The future that is glimpsed here is one of the memories of a smell, or perhaps more complicatedly its utter absence.

Do born-digital assets smell? And should they? How would we render that kind of full sensory experience? AI capabilities function best when objects can be rendered digital. It is possible to argue that the digitisation of collections and artefacts presupposes a particular way of engaging and interacting that is always and already detached from the physicality of touch or smell. And yet, we know those senses matter, and those kinds of encounters likewise. How will the senses be engaged in the libraries of the future? Will there be alternative ways of sensing and engaging with electronic collections, in which the physical and the digital are always necessarily interdependent? What will be their smell and their lingering decay?

Digital collections presuppose engaging and interacting with collections detached from the physicality of touch or smell, and raise interesting questions for libraries:

- How will the senses be engaged in the libraries of the future?
- Will there be alternative ways of sensing and engaging with electronic collections, in which the physical and the digital are always necessarily interdependent?
- What will be their smell? And their lingering decay?

5.2 A Final Word, and a Final Story from the Future

Straddling the dual role of custodian of documentary resources relating to Australia and the Australian people, and of midwife to the discovery, learning, and creation of new knowledge by the Australian Community, the National Library of Australia is a constantly regenerating and emergent cybernetic system. It continuously ingests new resources and grapples with making confident, informed decisions about collections in the context of multiple time frames and feedback loops.

Throughout this report, we have taken the library to be a complex open system with multiple goals and built on the work of Heinz von Foerster in describing some of the dynamics of the Library as a cybernetic system. What makes it complex is that its goals can be multiple and can differ depending on position/perspective and on the level of granularity in accounting for people, processes, infrastructure, data, and context. What makes it open is that it sits at a nexus of relationships with one thousand other institutions, big and small, and within government, and with materials that are constantly streaming in, in many formats and types.

The NLA has a complicated future; it remains a complex open system, a cybernetic system. Many of the biggest questions with which it must grapple are ones that technologies alone will not solve: what does “relating to Australia and the Australian people” mean? Who gets to decide? How can the NLA collect in ways that enable discoverability, learning and creation of new knowledge? What is seen to be of “public value”?

Just like all work that the NLA does, this report reflects a particular moment in time, always situated within the knowledge that past decisions and knowledge inform what we do now, and that what NLA employees and their audiences will do/want/need in the future is both critical and unknown.

Our goal has been to inform you of what is on your horizon in terms of tech capabilities that have the potential to transform collections and Library practices as fundamentally as indexing did in the nineteenth century, and as the internet and search did in the last forty years, and to empower you to make the best-informed decisions you can now for your future audiences. As many NLA staff said in interviews, what you want your successors to say in forty years is “they made good decisions in 2021.”

To help make those good decisions, we want to leave you with one last speculative story about a possible future; a story which we would like to imagine is as generative as it is hopeful . . .

Scenario 5:
Serendipity Engine Accused of Sparking Unplanned Joy

A young woman was rescued on the streets of Albury-Wodonga today after suffering what appeared to be a fit of uncontrollable laughter.

Nyala Kassa claims to have been given information by the National Library of Australia's Serendipity Engine that was neither convenient nor immediately relevant to her current location or status.

Ms Kassa, who was enroute to her annual retinal scan verification appointment, found herself listening to the English poet and comedian Pam Ayres addressing the Australian National Press Club in 1987. The address, which dwells primarily on the subject of Ms Ayres' marriage, bears no apparent relation to optical security or any of the geographical points of interest identified by her anticipation agent on Ms Kassa's planned route.

All requests for an interview with Ms Kassa have been denied, as physicians warn she is showing signs of unpredictable curiosity. There are reports that she smiled at someone she didn't know, sparking fears of a repeat of the stranger-hugging phenomenon of 2036.

Meanwhile, a spokesperson for the National Library of Australia has confirmed that the institute is aware of Ms Kassa's experience and that the Collections team is "exploring whether this was a chance occurrence or some other happenstance."

In 2040, the line between search and recommendation has disappeared. What we know as 'search engines' in 2020 are better described as 'anticipation engines' that pre-empt rather than respond to inquiries. At first, anticipation engines played the role of ubiquitous personal assistant. Running out of milk or forgetting a birthday were things of the past for many Australians in 2030. Soon though, machine anticipation took an executive role, deciding on, as much as serving, our wants and needs. By 2040, searching for information is considered an arcane past time of the very bored or very old, since everything we need to know is always already at our fingertips.

One of the technologies driving these anticipation engines are Graph Neural Networks (GNNs). GNNs, which were used by scientists to discover new particles in physics, and drugs in medicine, have been

put to work by corporations to discover new correlations between user attributes and actions, and a huge range of contextual factors, from time of day to left- or right-handedness. Whereas traditional neural networks take a sequence (like a sentence) or a grid (like an image) as their input, GNNs take graphs – a cloud of related entities. Rather than predicting the next word in a sentence, or detecting a cat in an image, a GNN predicts how a range of items relate: if its 2 am during Mardi Gras and a teenager in downtown Sydney is hungry, they don't want a salad.

Another factor driving anticipation engines was the move away from text-based search that began in the late 2010s and early 2020s. Voice-, location-, and image-based search had collectively overtaken the search bar by 2025. Voice-based search started as structured queries ("Hey, Siri") before evolving into more conversational inquiries, often prompted by Siri themselves. Mobile location-based search responds to a person's immediate environment to give relevant information. The launch of Google Lens in 2017 kicked off the trend towards image-based search, which allows users to capture an image of a pair of boots, plate of food, or nifty widget, and have Google tell them how to make it their own, as well as other secondary information-only queries. As smart glasses overtook smart phones by 2030, image-based search became the primary search medium.

Developments in 'affective computing' saw these anticipation agents grow increasingly sensitive to their users' apparent moods, picking up on changes in vocal intonation, mobility, and visual variety, which the agents responded to by suggesting light exercise and shielding their wards from unpleasant news. By 2030, a new person-based search paradigm was emerging, which pre-emptively delivered information on, and interaction tips for, other humans in a user's vicinity. Stranger-hugging was briefly a thing in 2036, when anticipation agents learned that physical contact was the fastest way to restore their human to equilibrium but was quickly quashed by regulation.

Based on GNNs, anticipation agents see through individuals. Not in the sense of intimacy, but transparency. Anticipation agents look through their users like a window to a statistical landscape, a graph, of the common traits and behaviours they represent. While this comes with the benefit of comfort and convenience, it also comes at the cost of individuation, the process by which a human comes to feel like themselves. It turns out that the stranger-hugging phenomenon was a sign of things to come, as anticipation agents found themselves having to work harder and harder to make their users happy. Fortunately, they were very, very clever, and by 2040 we still had no idea what was making us so miserable.

The National Library of Australia, like all information-based services in the 2020s, embraced GNNs in their catalogue search engine. Graph-based search allowed library readers, listeners and lookers to find relevant items that they *weren't* looking for. The same technology that informed Australians that, if they liked the Pad See Ew, they'd also enjoy the Tom Kha Gai, also helped them discover photographs of the farm where their grandfather grew up, and recordings of a lyre bird that haunted the area. The NLA's 'Serendipity Engine' became a source of great pride.

However, by around 2035, curators at the library noticed that, although the Serendipity Engine was encouraging individual users to explore more items in the collection, a decreasing number of items overall were being accessed. Unexplored corners of the catalogue were remaining unexplored. It appeared that the Serendipity Engine had reached a kind of saturation point, such that users were being increasingly directed towards the same statistical peaks in the catalogue topography. One particularly astute library staff member observed that the name 'Serendipity Engine' was actually a bit of a furphy, since the engine wasn't driving happy accidents but calculated surprises. The library set about creating a new serendipity engine that would be worthy of the name …

"The role of the archivist and manuscript curator, thus transformed by electronic recording and transmission equipment of millisecond speed and virtually unlimited storage capacity, becomes as a consequence an increasingly dynamic one"

Fisher and Evans
(1967, p. 335)

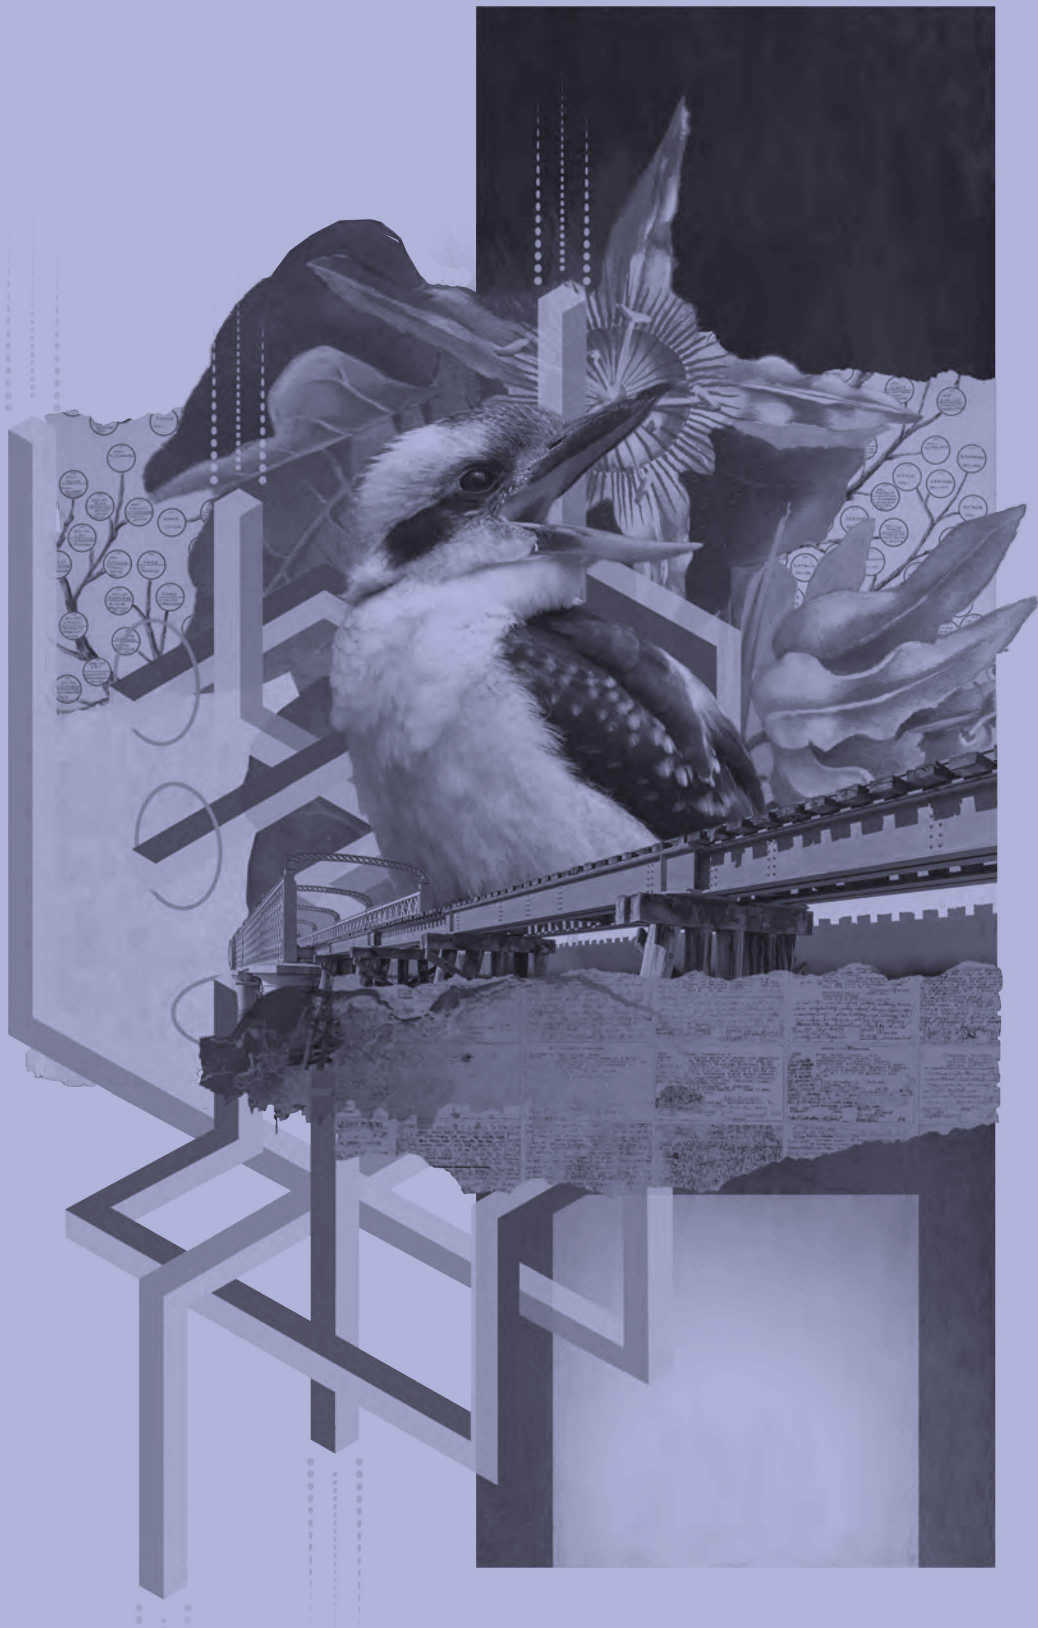

Appendices

Appendix A.
Materials Informing this Report

The insights and recommendations contained in this report are the result of a matrixed analysis of the following materials. By matrixed, we refer to an analysis in which materials of each type are read against and in the context of other materials, rather than in isolation or as wholly independent content types.

A.1 Existing NLA materials

The research team reviewed materials provided by the NLA as a starting point for understanding the structure, mission, goals, and dynamics of the organisation. These materials were critical in developing our understanding of the NLA as a complex system, and informed the shape and depth of the Literature Review (Appendix A.2) and our preparation for Expert Interviews (Appendix A.3).

NLA materials included:

- National Library Act 1960 (Cth)
- NLA IT Division Strategic Plan 2018–2022
- NLA Mobile Strategy
- NLA Collection Digitisation Policy
- NLA Digital Preservation Policy, 4th Edition
- NLA Digital Library Infrastructure Replacement website
- NLA Corporate Plan 2020–2021
- NLA Temporary Employment Register 2021
- NLA Annual Report 2019–2020
- NLA Statement of Intent 2020–2021
- NLA Organisation Chart 2019
- NLA Statement of Expectation, Letter to the Minister for Communications, Cyber Safety and the Arts 2020-2021

A.2 Desktop Research

The research team conducted a holistic literature review of resources related to AI and libraries. This included resources that described library practices that generally relate to AI (e.g. practices associated with the digitisation of materials), resources that discussed applications of AI in the library, and resources that considered the implications of AI for the Library. The review also included deep consideration of resources on cybernetics, and referred to principles, frameworks and codes that relate to the NLA as a Commonwealth Government institution. The consideration of these resources was critical in developing our analysis of the broad landscape of AI capabilities in the library sector, and to the presentation of the opportunities, and the analysis of the risks, pitfalls, and issues with each of the four capabilities clusters discussed in Section 3.2, as well as the emerging socio-technical, business, and regulatory trends that might be relevant to the Library in the coming years (Section 5.1).

From the literature review, we produced an annotated bibliography focused on resources discussing AI and libraries, particularly resources focused on the implications of AI for libraries and principles, frameworks, and codes that relate to AI and the NLA. This bibliography has been provided to the NLA under separate cover. Also under separate cover is a Zotero library of the broader set of literature that informed this report, including materials related more generally to artificial intelligence, libraries, and cybernetics.

A.3 Expert Interviews

The research team conducted semi-structured, qualitative expert interviews with NLA team members responsible for strategy and operations in multiple parts of the organisation including Collections, Engagement, Collaboration, and Digital. The expert interviews were critical to informing the dynamics of the NLA as a cybernetic system in Section 1.3, as well as the questions in the Cybernetic Star Guide in Section 4.

In this report, we have drawn on the insights from these interviews, and have chosen to synthesise the outputs into generalised statements from NLA experts, rather than using direct quotes or identifying individual speakers. We have adopted this approach as the primary outputs of the report are not an ethnographic analysis of the workings of the Library or the viewpoints of stakeholders in different parts of the organisation. The original transcripts have been shared with the NLA should it wish to explore them further or extract specific quotes.

Interviews were conducted on-premises at the NLA in Canberra by teams of School of Cybernetics researchers. The interviews were audio recorded and transcribed, and the transcripts made available to interviewees to check for accuracy and desired edits before being confirmed as research data.

We wish to thank the following NLA employees for their patience and generosity in sharing their expertise with us:

- Marie-Louise Ayers, Director-General
- Elizabeth Baillie, Director, Collections Management
- Kevin Bradley, Assistant Director-General, Collections Branch
- Heather Clark, Director, Reader Services
- Alison Dellit, Assistant Director-General, Collaboration Branch
- Kent Fitch, Programmer, Contractor
- Julia Hickey, Assistant Director, Platform Transition
- Terence Ingram, Director, Technology Operations
- Paul Koerbin, Assistant Director, Strategy & Product
- Cathy Pilgrim, Assistant Director-General, Engagement Branch
- Qi Fan, Director of Development, Digital
- Megan Williams, Assistant Director, Archival Processing
- David Wong, Chief Information Officer

A.4 Speculative Futures and Science Fiction Prototyping

Building on methodologies developed at Intel Corporation, we use science fiction prototyping as a methodology for exploring the systemic impacts of prospective technologies (see Johnson 2011a). As Johnson puts it, science fiction prototyping "is a short story, movie or comic based specifically on a science fact for the purpose of exploring the implications, effects and ramifications of that science or technology" (2011a, p. 25). For this report, we developed a series of five speculative future scenarios, based on signals and drivers of change at the National Library of Australia (as identified in NLA-specific documents like annual reports) and in the broader library and technology landscape. These inputs fed into a School of Cybernetics internal workshop, in which a range of possible futures were generated, with five selected for further development. These scenarios continued to develop throughout the stakeholder interviews and as the substance of the report was drafted. Many alternative futures may occur, depending on our actions in the present. The merit of the future scenarios shown in this report should not be judged on their veracity, but on their effectiveness in raising meaningful questions regarding the choices we make today.

Each scenario is accompanied by a photo collage that illustrates the main story. The design aesthetic of these collages, which feature images and documents created throughout the nineteenth through twenty-first centuries, aims to reflect the multiple temporal rhythms that underscore the work of Collections (as discussed in Section 1.3 of this report). Each image features at least one item sourced from the NLA's catalogue. For example, each collage includes a native flower illustration by Ellis Rowan, and the collage illustrating the Deakin scenario includes the digitised version of one of the official records of the debates of the Australasian Federal Convention, which are held by the NLA and directly referenced in that story.

A.5 Additional Materials

Finally, the School of Cybernetics acknowledges that the insights created here are informed by prior experience of the research team in studying libraries as sites of innovation and as innovations in themselves. While these materials are not directly referenced in the report, we recognise that our experience thinking about libraries precedes our engagement with the NLA. Specifically, we wish to draw attention to prior work that has informed our present thinking.

Work done by Genevieve Bell and Alex Zafiroglu between 1999-2010 while at Intel corporation, including the following published materials:

- Dourish, P. and Bell, G., 2007. *The Infrastructure of Experience and the Experience of Infrastructure: Meaning and Structure in Everyday Encounters with Space*. Environment and Planning B: Planning and Design, 34 (3), 414–430.
- Bell, G. (2010). 'U are Happy Life: Making sense of new technologies', *Innovation Ideas Forum 2010*, National Library of Australia.
- Bell, G. (2007). 'Going to the Library? a glimpse into the future.' Presidential Panel, Reference and User Service's Association, American Library Association, Washington, DC.

Work that Genevieve Bell undertook as a Thinker in Residence for South Australia in 2008-2010, including the final published report: *Getting Connected, Staying Connected: Exploring South Australia's Digital Futures* (2009).

From 2015-2021, Genevieve Bell has also been part of several expert working groups tackling technology (including AI) and its roles in Australian society. This work also informs the analysis contained here. It includes:

- Australian Government Department of Industry, Science, Energy and Resources, 2020. Australia's Artificial Intelligence Ethics Framework.
- Lattimore, F., O'Callaghan, S., Paleologos, Z., Reid, A., Santow, E., Sargeant, H., and Thomsen, A., 2020. *Using artificial intelligence to make decisions: Addressing the problem of algorithmic bias*. Australian Human Rights Commission, Technical Paper.
- Farthing, S., Howell, J., Lecchi, K., Paleologos, Z., Saintilan, P., and Santow, E., 2021. *Human Rights and Technology Final Report*. Sydney: Australian Human Rights Commission.
- Walsh, T., Levy, N., Bell, G, Elliot, A., Maclaurin, J., Mareels, I., & Wood, F. (2019). *The effective and ethical development of artificial intelligence: An opportunity to improve our wellbeing*. Australian Council of Learned Academies.
- Williamson, R. C., Raghnaill, M. N., Douglas K. & Sanchez, D. (2015). *Technology and Australia's future: New technologies and their role in Australia's security, cultural, democratic, social and economic systems*. Melbourne: Australian Council of Learned Academies.

Appendix B.
Speculative Glimpses

This appendix contains the full text of the five speculative future scenarios included in the report, along with a selection of the signals of change that informed those scenarios. We have also provided a brief description of how the National Library of Australia (NLA) might develop its own speculative scenarios along with a worksheet used during the School of Cybernetics (SoCy) Report Back session to the NLA.

Building on methodologies developed at Intel Corporation, we use science fiction prototyping as a methodology for exploring the systemic impacts of prospective technologies (see Johnson 2011a). As Johnson puts it, science fiction prototyping "is a short story, movie, or comic based specifically on a science fact for the purpose of exploring the implications, effects and ramifications of that science or technology" (Johnson 2011a, p. 5).

For this report, we developed a series of five speculative future scenarios, based on signals and drivers of change at the NLA (as identified in NLA-specific documents like annual reports) and in the broader library and technology landscape. These inputs fed into a SoCy internal workshop, in which a range of possible futures were generated, with five selected for further development. These scenarios continued to develop throughout the stakeholder interviews and as the substance of the report was drafted. Many alternative futures may occur, depending on our actions in the present. The merit of the future scenarios shown in this report should not be judged on their veracity, but on their effectiveness in raising meaningful questions regarding the choices we make today.

Each scenario is accompanied by a photo collage that illustrates the main story. The design aesthetic of these collages, which feature images and documents created throughout the nineteenth through twenty-first centuries, aims to reflect the multiple temporal rhythms that underscore the work of Collections (as discussed in Section 1.3 of this report). Each image features at least one item sourced from the NLA's catalogue. For example, each collage includes a native flower illustration by Ellis Rowan, and the collage illustrating the Deakin scenario includes the digitised version of one of the official records of the debates of the Australasian Federal Convention, which are held by the NLA and directly referenced in that story.

Each speculative vision has also been provided to the NLA as an A3 Poster. The following images were used under Creative Commons license, with full license details and accreditation to be found at the supplied URL.

- High Court of Australia: https://commons.wikimedia.org/wiki/File:High_Court_Australia02JAC.jpg
- Girl Marching: https://www.flickr.com/photos/matthrkac/51039335446/in/photostream/
- Parliament House: https://aus01.safelinks.protection.outlook.com/GetUrlReputation
- National Library of Australia: https://commons.wikimedia.org/wiki/File:National_Library_of_Australia_(2192098809).jpg
- Indigenous Flag: https://commons.wikimedia.org/wiki/File:Indigenous_and_Rainbow_Flags_(15741501997).jpg
- Cursive Diary: https://commons.wikimedia.org/wiki/File:Indenture_between_Daniel_D._Ridenhour_and_Adam_Haffley_-_DPLA_-_8f4855c9bfd61180a59c9365fd90b035_(page_1).jpg
- Svalbard Seed Vault: https://upload.wikimedia.org/wikipedia/commons/9/97/Entrance_to_the_Seed_Vault_%28cropped%29.jpg
- Man with Binoculars: https://www.shutterstock.com/image-photo/passengers-on-ship-duringcruise-696651739
- Old Card Catalogue: https://commons.wikimedia.org/wiki/File:Old_card_catalog_(3089541727).jpg
- Kookaburra: https://www.istockphoto.com/photo/australian-kookaburra-perched-andlaughing-gm1176553747-328077158
- Albury-Wodonga Rail Bridge: https://www.istockphoto.com/photo/australian-kookaburraperched-and-laughing-gm1176553747-328077158

B.1 Scenario 1: Deakin Speaks, World Listens (OCR Capability Cluster)

Deakin Speaks, World Listens

A lawyer representing the experimental neural implant company Affect Link in the High Court of Australia is trending worldwide today after calling a digital avatar of Alfred Deakin, the second Prime Minister of Australia, as an expert witness in a constitutional law case against the Commonwealth.

The lawyer claims that the avatar, trained on notes from the constitutional convention debates and Deakin's personal records held by the National Library of Australia, can represent the original intention of the founders of the Australian Federation.

The lawyer for the Commonwealth argues that there is no way of assuring either the accuracy of the avatar or the underlying data. "We have no way of knowing if that machine dotted every 'i' and crossed every 't' while reviewing those documents," the lawyer told the court. "A stray apostrophe could render unto Caesar that which is due elsewhere."

This news follows unverified claims leaked from Parliament earlier this week that the current Prime Minister, Karen Wong, has consulted with the National Library's Robert Menzies avatar on policy issues.

Affect Link is opposing the Commonwealth's attempt to legislate against communication via neural implant. The case turns on whether the Affect Link device, which allows one user to share a physical sensation with another, falls under the Commonwealth's constitutional power to regulate "postal, telegraphic, telephonic, and other like services".

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

Deakin Speaks, World Listens
How did we get here?

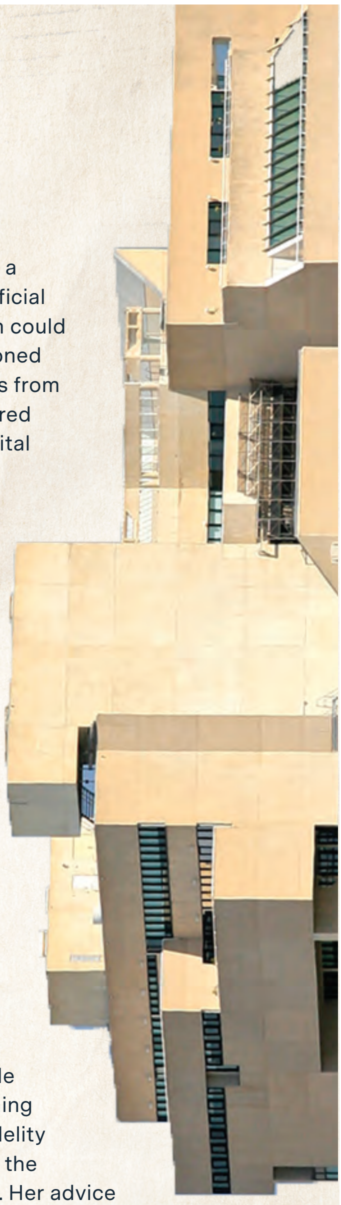

In the year 2035, artificial agents trained on the personal data of individuals are a compelling proxy for interacting with the person themself. Of course, these artificial agents were only ever as good as their data sets, and tiny errors in transcription could have significant ripple effects. And yet, we still feed them everything: old-fashioned text- and image-based data like emails and social media posts, voice recordings from smart speakers, gesture tracking from wearables and even dreamscapes captured by experimental neural implants were among the data sources used to train digital avatars that stand in for a real person.

This trend began with small, personal projects. In the late 2010s, a handful of technologists were reported to have trained text-based 'chatbots' on the personal data of deceased friends and family members. In 2020, Microsoft was granted a patent for a similar method, though they expanded the range of targets to include any "past or present entity … such as a friend, a relative, an acquaintance, a celebrity, a historical figure," who would be rendered in 2D or 3D animation. Squeamishness over the ethics of such an undertaking, though, meant that it was a start-up, Gemini, who offered the first commercial avatar building service in 2025. Gemini's offering was extremely expensive and restricted to relatively young people for whom a vast corpus of born-digital data was available. But it launched a market that, over the next 10 years, drove significant improvements in avatar technology. This, combined with improvements in machine learning and synthetic data techniques, means that developers now claim to be able to build realistic avatars from pre-digital data like printed newspapers and handwritten notes.

By the late 2020s, the most common form of avatar was the so-called 'grief bot' - an interactive memorial for the living to commune with the dead. There are now a range of platforms offering this service. Grief bots are relatively affordable compared to other forms of avatar because they require less data and less training to be made fit for their private purpose. At the higher end of the market, high fidelity avatars are built to represent their living subjects publicly. For example, in 2028 the CEO of Amazon began offering executive coaching services through her avatar. Her advice to young entrepreneurs? Invest early in your own avatar so you can be in two places at once. In 2031, a tech billionaire combined the two concepts of grief bot and digital representative, appointing his avatar as an executor on his will. The legality of the appointment is yet to be tested.

The NLA was among the first libraries to experiment with offering digital avatars as a novel way of interacting with their collections, partnering with a successful Australian high fidelity avatar provider, Perpetua. This caught the attention of an aging Australian novelist, who offered his personal data to the library along with an exclusive license to his identity by way of bequest, with the stipulation that it only be made publicly available after his own passing, and those of his six children. Exclusive identity licenses have become a promising source of alternative revenue for the NLA, who have sought and established similar agreements with other public figures. The current prime minister has made it known that they plan to donate their identity to the library for this purpose. Already, the NLA and Perpetua have created digital avatars of 5 deceased Australian prime ministers.

Deakin Speaks, World Listens

Signals that contributed to this scenario

Bringing the past to life: Services like Deep AI's image colorisation (Antic2021) allows users to convert black and white images to supposedly realistic colour images. MyHeritage's DeepNostalgia service animates still photographs (Animate your family photos n.d.). The process of Deep Image Inpainting uses AI to fill in the blanks on damaged or partial photographs (Li 2020). Services like these may help bring library content to life for users. However, they clearly render an interpretive role to machines that cannot be checked by humans without significant expertise. While non-expert humans-in-the-loop might be used to check the veracity of OCR, one needs significant historical contextual knowledge to check the colourisation of a photograph.

Early griefbots: There are already at least two examples of technologists creating 'griefbots' to talk with their deceased loved ones. Russian technology developer and entrepreneur Eugenia Kuyda built a chatbot based on the personal data of her deceased friend Roman Mazurenko (Newton 2016). And a team of researchers at MIT Media Lab, led by Canadian academic and technology entrepreneur Hossein Rahnama have been developing a similar commercial offering named 'augmented eternity' (Rahnama 2021).

Cutting edge language models: A Deakin-specific language model has been trained on the data of his papers and speeches captured in the records of the constitutional debates. However, this domain-specific language model will be built on top of a larger, pre-existing model, known as a 'large language models' (or LLM). Example LLMs, which have been trained on text scraped from the Internet, include like Open AI's GPT-3, Google's Switch Transformer, and Facebook's M2M-100.

B.2 Scenario 2: Magnifica Emerges from the Web (Transcription Capability Cluster)

Magnifica Emerges from the Web

Amateur botanists across Australia are celebrating after one of their own, Jarli Grange, discovered a small grove of supposedly extinct Australian orchids growing in the wild, aided by a series of clues left in historical maps, oral histories and other resources held online by the National Library of Australia.

The caladenia magnifica, or magnificent spider orchid, is endemic to Central Victoria but has not been seen there since 1979. "I couldn't believe what I was seeing," Grange told In Media Res. "I've been bush-bashing my way across Clydesdale for years, but it turns out I should have been looking online the whole time!"

Grange's trail of discovery began when he came across a reference to a sloped area of stony soil, the orchid's ideal habitat, on an early pastoral map of Victoria in Avoca -an area not previous associated with the flower.

The map, and many others like it, had been carefully read and annotated by a new kind of librarian: an AI that finds information 'buried' in archival documents and brings it to the surface, transcribing handwritten, audio and visual information into searchable catalogue records.

Having found the map, Grange picked his way through a series of additional clues also made accessible by AI. One notable find was a 1923 letter from Ms. Dymphna Marston to her fiancée, which contained a fragile pressed flower–none other than the caladenia magnifica –and a detailed description of the pleasant afternoon she spent walking near Avoca, where she found the bud nestled under a tree on a hill.

"It notices things and connections we don't," said Grange. "It found that same damn tree again in an interview done with a local after the 1985 bushfires. Poor bloke. Lost most of his farm, but all he wanted to talk about was how Avoca's one tree hill kept standing."

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

Australian National University

NATIONAL LIBRARY OF AUSTRALIA

Magnifica Emerges from the Web
How did we get here?

In the year 2032, the National Library of Australia has established an international reputation as a leader in applying Machine Transcription technologies to their collections. After several tentative years of experimenting with speech-to-text translation of oral histories in the mid-2020s, the library branched out to object recognition in photographs, and text extraction from maps, among other AI services. In each case, the extracted information was attached to items first as metadata, and gradually made more public as the library gained confidence in the technology and relevant governance processes.

The case of pastoral and other 19th century maps is a particular source of pride to the NLA. Inspired by the success of its AUSTLANG code-a-thon, which saw more than 8,000 records in the Australian Bibliographic Database coded with a unique Indigenous Australian language identifier, the library took a similar approach to de-colonising the records for its colonial era maps. After using Machine Transcription processes to pull text out of maps into a searchable database, the NLA recruited a pool of expert volunteers to apply new subject headings to the map records, which had been developed by the Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS) to better serve the needs of Indigenous Australians.

It was one of these human-applied subject headings – not the AI alone – that led Grange to his map. In fact, without these AIATSIS subject headings, Grange may not have found the information that led to his discovery, as the NLA struggled to adjust its catalogue search algorithms to cope with the influx of new metadata and information generated by its Machine Transcription efforts.

While its initial experiments were achieved with small digital transformation grants, the NLA required a more sustainable business model to fund its expanded Machine Transcription activities. After contemplating an exclusive partnership with a major machine learning platform provider, the NLA eventually opted to develop a shared Machine Transcription roadmap with the state libraries, who were excited by the results of the NLA's experiments. Seed funding for that project was provided by a philanthropic campaign targeting successful tech entrepreneurs who also formed an advisory council for the roadmap, with additional ongoing funding provided by the states. Together, the libraries developed a suite of Machine Transcription tools, built on top of a range of commercial and open-source platforms and data models, uniquely suited to the broader Australian collection.

Unfortunately, Grange's remarks on the future of Australian biodiversity proved all too prescient. Extinction rates continued to increase with global warming throughout the twenty-first century. By 2056, Synthetic Biology labs have started mining the library's collections for reference material, to help re-create lost plant- and eventually animal life. Drawing inspiration from the Svalbard Global Seed Vault, the NLA and Te Puna Mātauranga o Aotearoa, the National Library of New Zealand, co-found an Oceanic Life Vault (OLV) in 2078 to store images and sound recordings relating to endemic flora and fauna to aid in their eventual reconstruction. One unexpected but significant source of data for the OLV was animal and bird calls caught in the background of other audio and video items held by the library, identified through audio event detection.

Magnifica Emerges from the Web

Signals that have contributed to this scenario

Enabling new kinds of research: The Commonwealth Scientific and Research Organisation (CSIRO) Climate Researcher, Dr Robert Godfree, used Trove's digitised newspapers to reconstruct the impact of the historic, continent-wide 'megadrought' of 1891-1903, known as the Federation Drought, on Australia's flora, fauna and landscape (National Library of Australia 2020a). This kind of research is only imaginable with the Trove search engine and digitised newspapers. What new forms of research will be enabled by ML-powered Machine Transcription?

Speech-to-text for search support: It has been reported that Google Podcasts automatically creates transcripts of podcasts on its platform to aid with user search (Bhagyashree 2019). The full transcript is not made public but logged as metadata on the web portal. A similar approach has been suggested to aid in the discovery of oral histories held by the NLA.

De-colonising maps: NLA Curatorial Fellow Bess Moylan has explored how a better understanding of historical maps can support research into Indigenous Australian cultural landscapes (Moylan 2019). She describes how place names and other features in Cadastral maps give clues to Indigenous ownership and land use. In this scenario, we have combined Moylan's suggestion of applying AIATSIS subject headings to maps with the NLA's previous experiments with crowdsourcing the application of AUSTLANG Indigenous Australian language codes to the Australian Bibliographic Database to suggest a hybrid Machine Transcription + human expertise model for realizing Moylan's vision at scale.

A changing planet: the NLA, along with the rest of Canberra, suffered the effects of a severe hailstorm and bushfires in 2019 and 2020. While the bushfires were undoubtedly frightening and destructive, the hailstorm was notable as a comparatively unexpected climate event that did significant damage to the library's heritage roof. How do speculative future scenarios like those shared in this report help us prepare for both probable and possible futures?

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

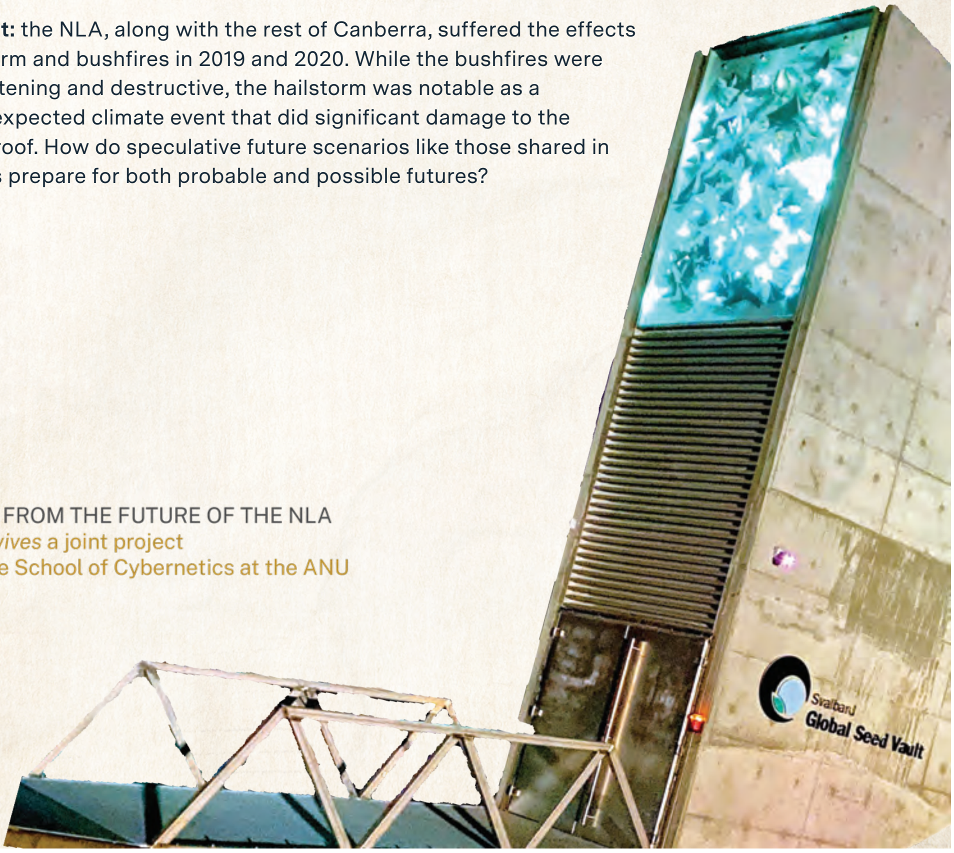

B.3 Scenario 3: Great Granddaughters Take Back the Library (Machine-Actionable Collections Cluster)

Great Granddaughters Take Back the Library

A group of young Australians braved the elements in Ngunnawal country this morning to follow in the footsteps of their great-grandmothers, re-enacting the landmark Women's March 4 Justice that occurred there 30 years ago today.

Each participant followed the path taken by an individual protestor, which had been recreated for them from extremely rare location data made accessible to the public by the National Library of Australia.

Marchers lifted replica signs as they followed the directions projected through their smart rigs and listened to a play-by-play account of the march pieced together by organisers from a cache of data pulled from a 'micro-blogging website'. The website data, which comprises short user-generated text content, was retrieved from a vintage physical storage device found at the bottom of an Australian National University researcher's closet after they died in 2046.

As the marchers gathered on Capital Hill, a young woman appeared on their rigs. "Change took a long time," march organiser Amelia Chen told the crowd, "But it came."

When asked what inspired her to organise the re-enactment, Ms Chen told In Media Res that she "wanted people to remember the good things our great-grandparents did for us. I know we have reason to blame them for a lot, but they weren't all bad."

"I don't have anything physical that belonged to my great-grandmother, just some low-definition Internet-era stuff. And I'm lucky, most people don't even have that. Walking in her footsteps today, it made me feel physically connected to her for the first time."

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

Australian National University

NATIONAL LIBRARY OF AUSTRALIA

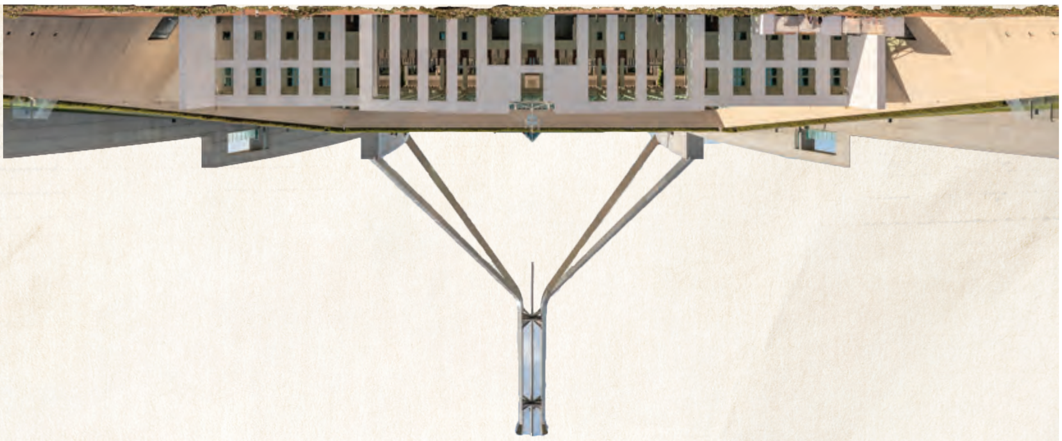

Great Granddaughters Take Back the Library

How did we get here?

In the year 2050, smart eye- and ear-wear have replaced smart phones as the ubiquitous personal computing device. 'Smart rigs', as they're called, generally comprise one or two contact lenses for receiving and displaying visuals, bone-conduction pads for audio input and output, a discreet biometric sensor, and a central computing and networking hub, usually implanted in a molar. Users receive a constant stream of carefully calibrated visual and audio information, projected over and into the real world. Meanwhile, a constant stream of user data is taken in return. This blending of the cyber and physical worlds is referred to as the 'spatial web' or Web 3.0.

Before smart rigs, smart glasses overtook smart phones from around 2030. This transition, combined with the shift from Web 2.0 to 3.0, and balkanisation of the internet by competing global and corporate powers, meant that much of the content produced online in the 2010s and 20s became very difficult to access. Internet-era powerhouses, repositioning themselves in the new AI-everything environment were loathe to share their private stores of machine learning training data. Publicly accessible caches like that found in ANU researcher's closet are rare treasures.

Like the smart phone before it, the smart rig generated new ways of interacting with information, and conducting research, including embodied explorations like the March 4 Justice re-enactment. Researchers can literally feel their way through a collection, exploring twisting pathways through real and virtual landscapes of images, sounds and text. In the early days of the 2030s, the major limit to public participation in such experiences was access to the computational power required to render such detailed information worlds. In 2035 the NLA launched its first public Virtual Reading Room, a cloud computing environment in which users can readily work with the library's Machine-Actionable collections in a virtual space.

The location data used to guide each individual participant in the re-enactment was collected by the COVIDSafe app, a primitive crowd sensing application promoted by the Australian Government during the COVID-19 Pandemic of 2020-2022. The COVIDSafe data was donated to the National Library of Australia in 2032 and placed under indefinite restricted access. In addition to containing personal information, the COVIDSafe dataset was huge, making it both cognitively and computationally inaccessible to the general public. The NLA has only recently begun supporting public use of the dataset by providing access via a data visualisation tool, which offers access to anonymised data in the form of representative, rather than actual, trajectories.

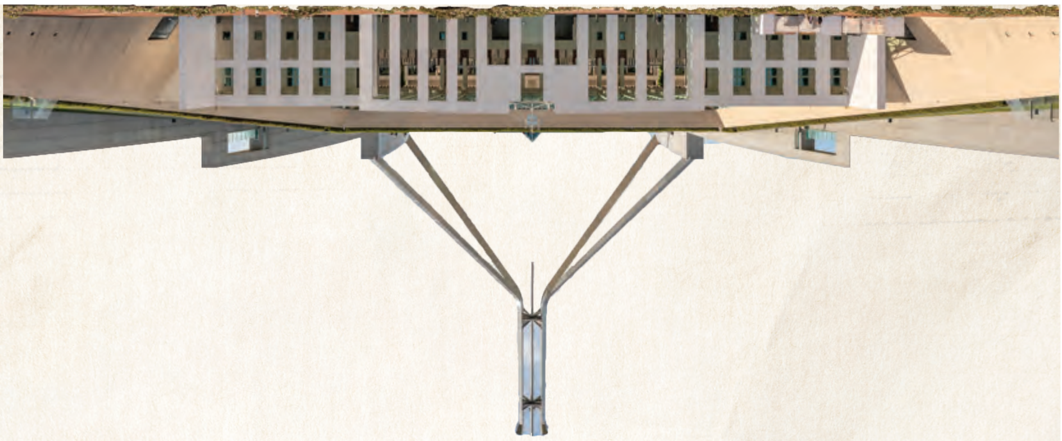

Great Granddaughters Take Back the Library

Signals that have contributed to this scenario

The challenge of funding: A new philanthropy team was founded within Office of the Director-General at the NLA in 2019-20, with the goal of raising $30m by 2030 (National Library of Australia 2020c). At the same time, the current giving trend in Australia is towards a smaller proportion of Australians donating larger amounts of money (Giving in Australia: the fast facts 2020). This trend may favour projects that attract funding from wealthy individuals, and/or enable to the NLA to pursue ambitious goals that government funding alone cannot sustain. The challenge, as always, will be funding the ongoing cost of any new technology investments.

Low code/no code machine learning: Companies like Apple (Create ML - Machine Learning 2021) and Google (Cloud AutoML Custom Machine Learning Models 2021) platforms that allow users to build ML models using little to no code. At present, these are designed to be used by developers without ML expertise but might also be used by consumers. This suggests a future in which an AI-literate general public works with its own AI tools to assist in their family genealogy project, for example, or in processing the big data collections held by the library, like the location data example explored in this scenario.

Computing Cultural Heritage in the Cloud: this is the name of a pilot program run by the Library of Congress exploring infrastructures for supporting access to and analysis of large-scale data in the cloud. Rather than having researchers and artists download and work with small data sets on local computers, this pilot allowed interested users to access and work directly with data in the cloud, where the library had some oversight over what was being produced (Ferriter 2019).

New forms of information: location data represents just one example of the new kind of information emerging from our increasingly cyber-physical world. Information of this kind is not just used by data scientists, but is increasingly made accessible to, and valued by, a broader community of users. For example, The Strava Global Heatmap (Strava 2021) buffs explore where people are running and cycling in their area, while Google's COVID-19 mobility maps summarise pedestrian traffic trends in target zones like supermarkets and pharmacies, public transport and workplaces for public health workers (COVID-19 Community Mobility Reports 2021).

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

B.4 Scenario 4: Waking up Ancestors and Walking with the Dead (Transparent AI Cluster)

Waking Up Ancestors and Walking with the Dead

The National Library of Australia was brought before the Indigenous Voice to Parliament today to explain how a partially fabricated hologram of Aunty Ellen Holsworthy has been held by the library for the last 10 years without the knowledge or consent of the Dharug people.

The hologram was originally created as an experiment in 2022 by the Indigenous technology developer Indigital, making it one of the earliest known examples of the medium. Indigital worked with Aunty Holsworthy and other Dharug elders to develop a hologram that could answer questions about Dharug experiences of the Blacktown Natives Institute.

A spokesperson for the Dharug people, Jasmine Griffiths, said she remembered her mother pioneering hologram technology in the early 2020s, and addressing the Dharug cultural protocols for the spatial web together with Dharug elders at the time before holograms were created.

However, the original hologram was subsequently repurposed by a non-Indigenous developer, who augmented the original interactive model with additional training data sourced from the National Library. It is this version that found its way into their digital archives.

"This holographic representation and the information spoken by the hologram of 'Aunty Holsworthy' is colonial voyeurism and should be destroyed," Ms Griffiths said. "The data the hologram was trained on was from resources written by white people. The hologram does not represent Dharug perspectives and in fact breaks cultural law by representing our Peoples without our consent."

The Indigenous Voice to Parliament is expected to advise on whether or not the hologram should be destroyed by the end of the week.

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

Australian National University

NATIONAL LIBRARY OF AUSTRALIA

Waking Up Ancestors and Walking with the Dead

How did we get here?

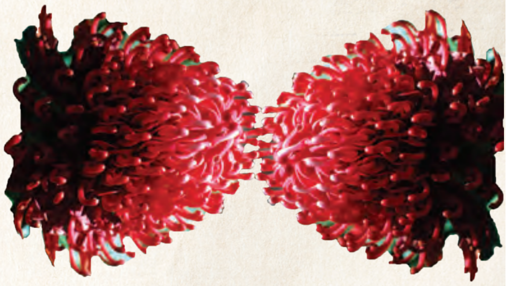

In the year 2040, interactive three-dimensional holographic videos (known simply as 'holograms') have emerged as a popular new storytelling medium. Holograms, which can be viewed through smart glasses or in dedicated holographic parlours, combine volumetric video capture with artificial intelligence to create what could be described as interactive live action films. Volumetric video capture entails the filming of a performance in three dimensions using hundreds of cameras. The first volumetric video capture studio, Imagine Room, was launched in Australia in 2022.

Cabrogal woman and technology entrepreneur, Mikaela Jade, has been a pioneer in extended reality and artificial intelligence since at least 2014, when she founded Australia's first Indigenous Edu-tech company, Indigital. In 2020, she began working with Dharug elders to develop Dharug cultural protocols for the 'spatial web', a phrase that describes the convergence of our digital and physical worlds in extended reality technologies. In particular, the guidelines addressed Dharug elders' concerns about cultural representations of their ancestors breaking cultural law when – as they rightly predicted – holograms were created by AI, rather than guided by Indigenous people themselves.

Jade worked with the Dharug elders and Aunty Ellen Holsworthy to create one of the first interactive holograms in 2023, which featured Aunty Holsworthy truth-telling about Dharug experiences of the Blacktown Natives Institute. Unlike contemporary holograms, this early example of the medium had very limited interactive features. The only AI element used was natural language processing for parsing user's questions about Dharug experiences of the Institute, which were responded to with authentic answers given by Aunty Ellen and other Dharug elders speaking 'through' Aunty Ellen's holographic image. However, the cost of hosting this hologram was prohibitively expensive and so it quickly disappeared from public view. In 2025, Jade turned her attention to co-developing an Indigenous-owned Dharug language model, that would allow users to interact with future holograms in language.

It is unclear how another technology developer gained access to the original data for the Aunty Holsworthy hologram in 2027, or what their intent was in building a more interactive version that contravened Dharug cultural protocols. The start-up company quickly folded after it failed to raise a Series B round of capital. All that the National Library of Australia knows is that the developers used library data to augment the interactive model and that a copy of the file was submitted via the National E-Deposit Service several years later, when it began accepting audio-visual and holographic content.

The holograph went undiscovered for several years, until 2040, when it was discovered by historian Martin Bost, who had received a $50,000 grant to develop an exhibition on the Blacktown Natives Institute. Bost assumed that showcasing the hologram through a public exhibition at the NLA would bring great joy to Dharug People, and all people across Australia who need to hear the 'real' stories of Australia's past interactions with First Peoples.

Waking Up Ancestors and Walking with the Dead

Signals that have contributed to this scenario

Indigenous Language Models: Te Hiku Media, a small non-profit radio station in New Zealand, have captured international attention by creating a speech-to-text engine for Te Reo Māori from a very small pool of crowd-sourced training data (Coffey 2021). They have subsequently been approached by both commercial and open-source developers, vying for access to their code. However, Te Hiku Media has rejected almost all requests, based on the principle that the only people who should financially benefit from Te Reo are Māori people themselves.

Data Sovereignty: Te Hiku Media's approach is a specific example of the more general principle of Indigenous Data Sovereignty, defined by the Australian Maiam nayri Wingara Indigneous Data Sovereignty Collective as "the right of Indigenous people to exercise ownership over Indigenous Data" where 'Indigenous Data' refers to "information or knowledge, in any format or medium, which is about and may affect Indigenous peoples both collectively and individually" (Key Principles 2021).

Cultural Sensitivity Filter: The NLA has recently added a cultural sensitivity filter to Trove, based on the recognition that much material of significance to First Australians held by the library was recorded in disrespectful ways by non-Indigenous Australians. The adoption of the Cultural Sensitivity Filter places the onus of work managing culturally sensitive material on the library, rather than requiring Indigenous Australians to self-censor. How could this principle be extrapolated out to other projects?

Indigital: Mikaela Jade and her company Indigital have already begun work on the Dharug cultural protocols for the spatial web referenced in this scenario (Indigital 2021). Jade also recently joined the World Economic Forum's Global Future Council on Augmented Reality and Virtual Reality (Global Future Council on Augmented Reality and Virtual Reality 2021), where she is recognised as a world leader in working with Indigenous and remote communities to leverage the potential benefits of extended reality technologies. Jade is just one prominent example of the many Indigenous and other Australians, like Lynette Wallworth (Collisions 2016), exploring the role this technology might play in preserving and sharing Indigenous culture and knowledge.

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

B.5 Scenario 5: Serendipity Engine Accused of Sparking Unplanned Joy

Serendipity Engine Accused of Sparking Unplanned Joy

A young woman was rescued on the streets of Albury-Wodonga today after suffering what appeared to be a fit of uncontrollable laughter.

Nyala Kassa claims to have been given information by the National Library of Australia's Serendipity Engine that was neither convenient nor immediately relevant to her current location or status.

Ms Kassa, who was enroute to her annual retinal scan verification appointment, found herself listening to the English poet and comedian Pam Ayres addressing the Australian National Press Club in 1987. The address, which dwells primarily on the subject of Ms Ayres' marriage, bears no apparent relation to optical security or any of the geographical points of interest identified by her anticipation agent on Ms Kassa's planned route.

All requests for an interview with Ms Kassa have been denied, as physicians warn she is showing signs of unpredictable curiosity. There are reports that she smiled at someone she didn't know, sparking fears of a repeat of the stranger-hugging phenomenon of 2036.

Meanwhile, a spokesperson for the National Library of Australia has confirmed that the institute is aware of Ms Kassa's experience and that the Collections team is "exploring whether this was a chance occurrence or some other happenstance."

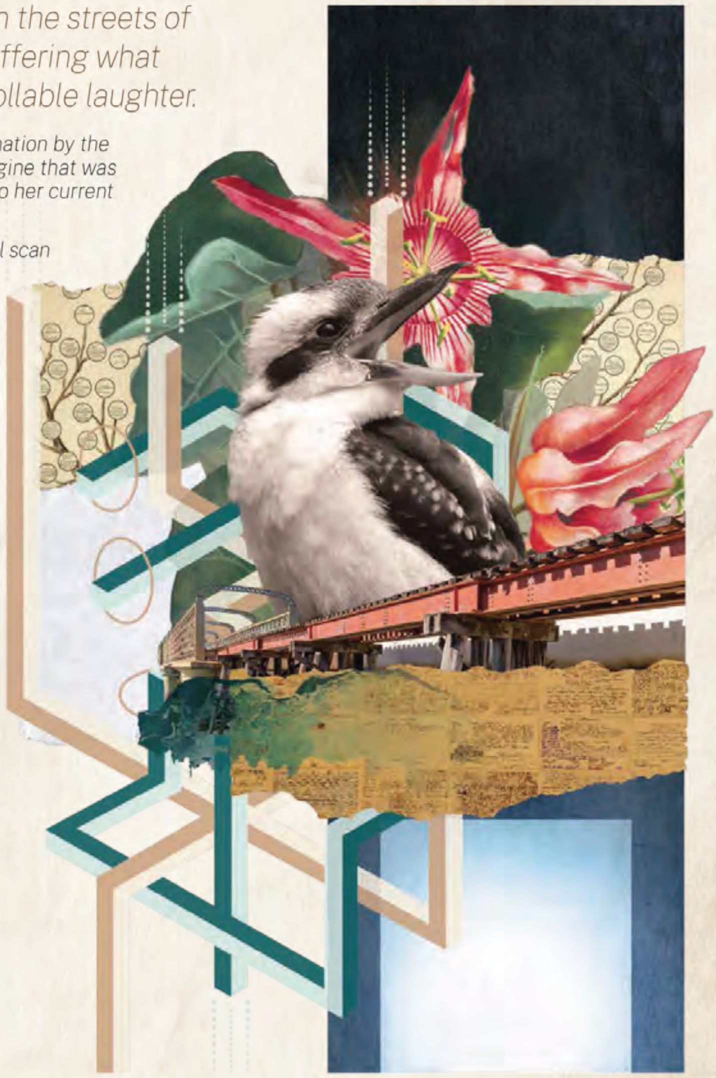

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

Australian National University

NATIONAL LIBRARY OF AUSTRALIA

Serendipity Engine Accused of Sparking Unplanned Joy

How did we get here?

In 2040, the line between search and recommendation has disappeared. What we know as ‘search engines’ in 2020 are better described as ‘anticipation engines’ that pre-empt rather than respond to inquiries. At first, anticipation engines played the role of ubiquitous personal assistant. Running out of milk or forgetting a birthday were things of the past for many Australians in 2030. Soon though, machine anticipation took an executive role, deciding on, as much as serving, our wants and needs. By 2040, searching for information is considered an arcane past time of the very bored or very old, since everything we need to know is always already at our fingertips.

One of the technologies driving these anticipation engines are Graph Neural Networks (GNNs). GNNs, which were used by scientists to discover new particles in physics, and drugs in medicine, have been put to work by corporations to discover new correlations between user attributes and actions, and a huge range of contextual factors, from time of day to left- or right-handedness. Whereas traditional neural networks take a sequence (like a sentence) or a grid (like an image) as their input, GNNs take graphs – a cloud of related entities. Rather than predicting the next word in a sentence, or detecting a cat in an image, a GNN predicts how a range of items relate: if its 2am during Mardi Gras and a teenager in downtown Sydney is hungry, they don’t want a salad.

Another factor driving anticipation engines was the move away from text-based search that began in the late 2010s and early 2020s. Voice-, location-, and image-based search had collectively overtaken the search bar by 2025. Voice-based search started as structured queries (“Hey, Siri”) before evolving into more conversational inquiries, often prompted by Siri themselves. Mobile location-based search responds to a person’s immediate environment to give relevant information. The launch of Google Lens in 2017 kicked off the trend towards image-based search, which allows users to capture an image of a pair of boots, plate of food, or nifty widget, and have Google tell them how to make it their own, as well as other secondary information-only queries. As smart glasses overtook smart phones by 2030, image-based search became the primary search medium.

Developments in ‘affective computing’ saw these anticipation agents grow increasingly sensitive to their users’ apparent moods, picking up on changes in vocal intonation, mobility, and visual variety, which the agents responded to by suggesting light exercise and shielding their wards from unpleasant news. By 2030, a new person-based search paradigm was emerging, which pre-emptively delivered information on, and interaction tips for, other humans in a user’s vicinity. Stranger-hugging was briefly a thing in 2036, when anticipation agents learned that physical contact was the fastest way to restore their human to equilibrium but was quickly quashed by regulation.

Based on GNNs, anticipation agents see through individuals. Not in the sense of intimacy, but transparency. Anticipation agents look through their users like a window to a statistical landscape, a graph, of the common traits and behaviours they represent. While this comes with the benefit of comfort and convenience, it also comes at the cost of individuation, the process by which a human comes to feel like themselves. It turns out that the stranger-hugging phenomenon was a sign of things to come, as anticipation agents found themselves having to work harder and harder to make their users happy. Fortunately, they were very, very clever, and by 2040 we still had no idea what was making us so miserable.

The National Library of Australia, like all information-based services in the 2020s, embraced GNNs in their catalogue search engine. Graph-based search allowed library readers, listeners and lookers to find relevant items that they weren't looking for. The same technology that informed Australians that, if they liked the Pad See Ew, they'd also enjoy the Tom Kha Gai, also helped them discover photographs of the farm where their grandfather grew up, and recordings of a lyre bird that haunted the area. The NLA's 'Serendipity Engine' became a source of great pride.

However, by around 2035, curators at the library noticed that, although the serendipity engine was encouraging individual users to explore more items in the collection, a decreasing number of items overall were being accessed. Unexplored corners of the catalogue were remaining unexplored. It appeared that the Serendipity Engine had reached a kind of saturation point, such that users were being increasingly directed towards the same statistical peaks in the catalogue topography. One particularly astute library staff member observed that the name 'serendipity engine' was actually a bit of a furphy, since the engine wasn't driving happy accidents but calculated surprises. The library set about creating a new serendipity engine that would be worthy of the name…

Serendipity Engine Accused of Sparking Unplanned Joy

Signals that have contributed to this scenario

Graph Neural Networks: GNNs use multi-label systems (graphs) as inputs. Unlike other forms of neural networks, such as a convolutional neural network, where inputs can be usefully classified according to single labels (e.g. 'Miniature Schnauzer'), GNNs are useful in situations of irreducible complexity. GNNs have been used by UberEats (Jain et al. 2019) and Alibaba (Zhu et al. 2019) to drive their product recommendation engines, by chemical engineers to discover new ways of storing green energy (Facebook AI and Carnegie Mellon University 2021), and biological engineers to discover new antibiotic molecules (Stokes et al. 2020).

Recommendation Engines: Recommendation engines used by Google, Spotify, and Netflix, for example, are generally based on either item-based ratings, and/or user- or item-based similarity. The measure of success for such systems is predictive accuracy: have we accurately identified the most popular item or most significant shared user trait? Of course, this is a source of great convenience to users. However, designing for serendipity in information discovery presents an opportunity for the NLA to offer a unique service and experience to its constituents.

Opportunity for Library as Leader: Ryan Cordell has argued that libraries have the opportunity to be leaders in the practice of 'Responsible AI' (Cordell 2020). The NLA, as a public institution and Chair of the National State Libraries Australia, is already recognised as a leader. Cordell's call to action suggests that, rather than focus on early adoption, the NLA has an opportunity to be a leader in the field by drawing on its expertise in responsible information management.

The Critical Catalogue: this is the name of a Research through Design prototype developed at the Syracuse School of Information Studies to explore metadata elements, values and organizational structures that would enable diversity in library catalogue search results (Clarke and Schoonmaker 2020). The Critical Catalogue is a typical library catalogue in most respects, except that it intentionally advocates for diversity and exposes library users and readers to resources from populations traditionally marginalised by floating content produces by people of colour, women, people of non-cisgenders and so on to the top of the search results.

ANOTHER STORY FROM THE FUTURE OF THE NLA
Custodians & Midwives a joint project
by the NLA and the School of Cybernetics at the ANU

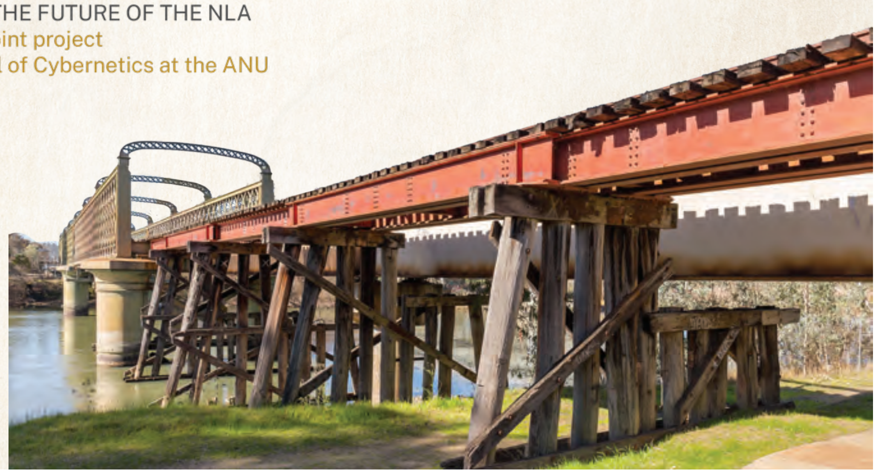

B.6 How to Create Your Own Future Scenarios

We briefly outlined the methodology used to generate these future scenarios in the introduction to this appendix. Here, we provide the next level of detail, describing the activities the NLA could undertake to produce its own future scenarios. This process combines elements of 'The Thing from the Future', an imagination game by Stuart Candy and Jeff Watson (Candy and Watson 2021); the 'Science Fiction Prototyping' method by Brian David Johnson (Johnson 2011b); and The Institute for the Future's 'Foresight Essentials Toolkit' (IFTF Foresight Essentials 2021).

Step One:
Identify Drivers and Signals of Change

The goal of this activity is to generate a pool of drivers and signals of change that combine to inform a forecast in the next step.

> **Drivers of change** are large scale forces impacting the Library and the world at large. For example, we identified declining trust in public institutions as a provocative driver of change that may impact the NLA. Challenge yourself to identify drivers from different contexts, including social, technological, economic, environmental, and political. When recording a driver, also note its anticipated impact on the NLA. For example, declining trust in public institutions may result in the loss of public support for government funding of the NLA.

> **Signals of change** are concrete examples of futures in the making. For example, we identified the Semantic Scholar academic search engine's TL;DR (Too Long; Didn't Read) function as a signal of where automatic content summary is headed (Hao 2020). When recording a signal, also note what direction that signal points in: what are you going from and to? For example, the TL;DR function suggests that we are going from expert-written abstracts and literature reviews to machine-generated content summaries. The question arises as to where these can be treated as serving equivalent needs, and where there is a need for both.

Create a database of drivers and signals of change that you can easily refer to. We worked in an Excel spreadsheet, but post-it notes are another good option.

Step Two:
Draft a Forecast

The goal of this activity is to cluster the drivers and signals of change that you identified in the first step to form the basis of a forecast – a brief description of a future object, pattern, or situation.

Start by clustering signals of change. If you are working as a group, take turns at sharing a signal from the pool and explain how you think it relates (or not) to the other signals that have already been shared.

Next, add drivers of change that relate to your signal clusters.

Once you have a healthy selection of clusters, divide the work of drafting forecasts based on those clusters. Forecasts should be provocative, plausible and internally consistent. Give them a snappy title that indicates the direction of change the forecast suggests.

Here is an example forecast drafted by us, that was not further developed:

Indigenous Languages Revived

In the year 2030, researchers from across Australia have built 15 Indigenous Australian language models capable of supporting speech-to-text and text-to-speech applications, and even some text-generative applications. Data held by the NLA has been used in the production of some of these models. This was made possible by the NLA prioritising access to this data. Applications built on top of these models have helped support an Indigenous language revival. There has been a four-fold increase of speakers of Indigenous Australian languages and some Australian councils have adopted local Indigenous languages as an official language alongside English. Statements issued by the First Nations Voice to parliament are translated into a language chosen by their sponsor. However, some controversy has emerged over the access to the language models. Some have been made open source, others have been commercialised. Languages represented in open-source models appear to be flourishing more widely, but data sovereignty advocates argue that this unfairly benefits non-Indigenous application developers who profit from applications built on top of the open-source models.

Step Three:
Draft a Scenario

The goal of this activity is to transform your forecast into an inspiring, provocative story. In our case, we used our forecasts as the basis of the 'background' section of each scenario, then created a shorter news article to sit on top of that forecast.

Some things to consider:

- What kind of future is your scenario set in, and when? One popular way of conceiving of prospective futures is according to four genres: Growth (assuming extrapolation of current trends), Constraint (where a core value organises society and governs behaviour), Collapse (following a major breakdown), and Transform (where society has fundamentally transformed).
- In addition to the drivers and signals of change that informed the seed forecast, what are the other key story elements? Who is your protagonist? What do they value? What is the status quo and what change happens?
- Give yourself a snappy 'headline from the future' to spark your interest and that of your eventual readers.

A Short, Sharp Version: Thing from the Library of the Future

The following truncated version of the process and attending worksheet was developed for the SoCy Report Back session to the NLA on 21 June 2021. The purpose of this exercise is to give participants a taste of the process used to generate the scenarios that appear throughout this report, though it may also be adapted to generate new scenarios.

The goal of this activity is to quickly generate a description or drawing of an object or event from the Library of the future. Participants are provided with four prompts to aid in their speculation:

- A future genre. Choose from the four discussed in Step Three above – Growth, Constraint, Collapse, or Transform. You should also suggest a time frame, which could range from 5 to 100 years.
- A signal or signals of change. Choose one or two from the selection you have generated in Step One above.
- An object.
- A mood that the object evokes.

Participants combine these prompts to imagine an object or event from the future.

Take 5 minutes to explain the activity, then give your participants a further 5 minutes to respond using the worksheet shared below.

Thing from the Library of the Future*

The Prompts

Future Genre:
Continuity

Set your future object:
- One Generation in the Future
- Two Generations in the Future

Signal of Change:

Let one of these inspire you:
- MyHeritage DeepNostalgia
- Trove Cultural Sensitivity Filter

Object:

Your object should be one of these (loosely defined):
- Flag
- Sticker
- Festival

Mood:

Your object should inspire one of these feelings:
- Charm
- Frustration
- Surprise

Describe your Object

Or Draw It

*Based on 'The Thing from the Future', an imagination game by Stuart Candy and Jeff Watson; the 'Science Fiction Prototyping' method by Brian David Johnson; and The Institute for the Future's 'Foresight Essentials Toolkit'.

Appendix C.
Cybernetic Star Guide Materials

This appendix complements the overview and instructions for using the Cybernetic Star Guide.

C.1 Generative Prompts for the Cybernetic Star Guide

C.1.1 Summary of Overall Prompts for General Dynamics

What is the right balance to strike between the quality of collection item processing and quantity of processing that can be done in a given time interval?

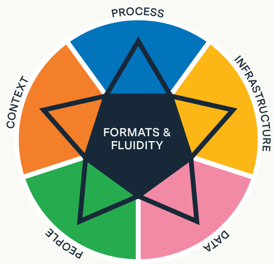

How do we strike a balance between assured longevity and risky ephemerality of collection item formats when making decisions about which items to collect, and with what urgency and frequency?

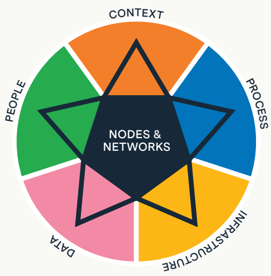

What is the right balance to strike between following our specific values, processes, infrastructure, and data standards, versus standards adopted by other organisations with which we collaborate or to which we are obligated?

How do we strike a balance between interacting with Australian publics as an audience and negotiating with them as collaborators who will have strong opinions about our collections practices and will want to influence and act upon them, enabled by new tech applications?

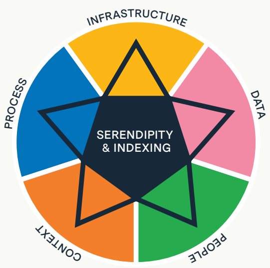

What is the right balance to strike between, on the one hand, maintaining room for audiences to create new connections among collection items and generate new knowledge and, on the other hand describing, organising, and creating indexing tools to enable intuitive discovery of any given item? Between some things that are fixed and perhaps ‘‘true’’ and others that are accumulating or changing?

C.1.2 Detailed Questions & Approaches for General Dynamic Exploration

In the table below are questions useful to explore general dynamics of the NLA's collections system as they may apply to a specific AI application. Questions in bold are the primary questions for consideration, with sub-questions providing context or stimulus to answer the primary question. Possible approaches to answering these questions are provided.

GENERAL DYNAMIC	Questions for consideration	Possible Approaches
	**All Components** **What is the right balance to strike between the quality of collection item processing and quantity of processing that can be done in a given time interval?** • What are current expectations for quality of collections processing, and for quantity of processing? Do these vary by collection, artefact or item type (i.e. born-digital versus converted; photo versus manuscript etc.)? • Who is responsible for managing expectations and delivery of quality of collections practices that underpin items available to audiences? • Who is responsible for managing expectations and delivery of quantity of collections available to audiences? • How does this application affect that balance? • Where are the decision-making forums that balance quality and quantity? • Who are the decision makers? • How is this balance implemented across workstreams?	*Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.* *This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside an acceptable range? What are our health indicators for the balance?*

GENERAL DYNAMIC	Questions for consideration	Possible Approaches
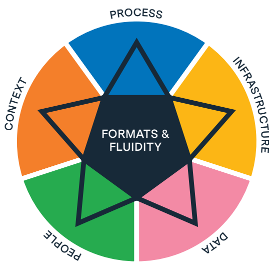	**All Components** **How do we strike a balance between assured longevity and risky ephemerality of collection item formats, when making decisions about which items to collect, and with what urgency and frequency?** • How does this application affect that balance? • How is this balance implemented across workstreams? • Who are the decision makers? • How do we build into our application plan the inevitability that infrastructure changes may make today's preferred formats more difficult to work with or even obsolete? • How does this application affect time frames for updating other dependent processes, data, infrastructure, and staff training? • Should some items be allowed to decay or resist further collection or meaning making (i.e. an extension of collection rules in various manuscript collections that mean things cannot be seen or opened for fixed periods)?	*Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.* *This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside and acceptable range? What are our health indicators for the balance?*

<table>
<tr><th>GENERAL DYNAMIC</th><th>Questions for consideration</th><th>Possible Approaches</th></tr>
<tr><td>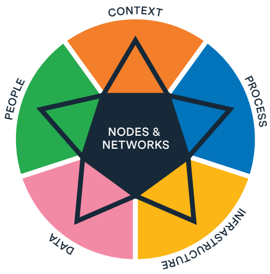
</td><td>

All Components

What is the right balance to strike between following our specific values, processes, infrastructure, and data standards, versus standards adopted by other organisations with which we collaborate or to which we are obligated?

- How does this application affect that balance?
- Where are the decision-making forums that balance NLA-specific versus library and government sector standards?
- Who are the decision makers?
- How is this balance implemented across workstreams?
- What happens when objects might sit physically in one place, and be accessed digitally in another? Or sit digitally in one cultural frame and be accessed in another (i.e. how do we manage cultural protocols over different regulatory, legal, geographic domains?)
- Should meta-data also do cultural/indexing work?
- Should certain objects/collections be able to resist traveling on/through the network?

</td><td>

Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.

This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside and acceptable range? What are our health indicators for the balance?

</td></tr>
</table>

<table>
<tr><th>GENERAL DYNAMIC</th><th>Questions for consideration</th><th>Possible Approaches</th></tr>
<tr><td>
</td><td>

All Components

How do we strike a balance between interacting with Australian publics as an audience and negotiating with them as collaborators who will have strong opinions about our collections practices and will want to influence and act upon them, enabled by new tech applications?

- What are the implicit contracts that we have with audiences now and how will they be affected by this application, including growing or losing this audience, or gaining new audiences?
- How will the types and amount of interactions that audiences have with what have historically been our internal processes be affected by this application?
- What audience feedback loops into collections will the application strengthen, weaken, transform, or create?
- Where are the decision-making forums that address audience contracts?
- Who are the decision makers for audience contracts?
- How will this application be implemented across process, infrastructure, and data workstreams?
- What happens when the same audience wants something different over time? Or the same audience that was hiding can now afford to be seen, or wants to be seen?
- What if certain collections or products of collections make audiences visible who have otherwise been invisible, either by choice or by force?

</td><td>

Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.

This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside and acceptable range? What are our health indicators for the balance?

</td></tr>
</table>

<table>
<tr><th>GENERAL DYNAMIC</th><th>Questions for consideration</th><th>Possible Approaches</th></tr>
<tr>
<td>
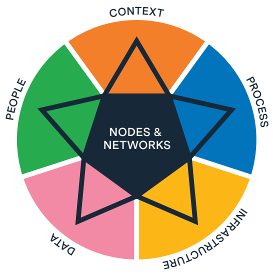

</td>
<td>
All Components
What is the right balance to strike between following our specific values, processes, infrastructure, and data standards, versus standards adopted by other organisations with which we collaborate or to which we are obligated?

How does this application affect that balance?
Where are the decision-making forums that balance NLA-specific versus library and government sector standards?
Who are the decision makers?
How is this balance implemented across workstreams?
What happens when objects might sit physically in one place, and be accessed digitally in another? Or sit digitally in one cultural frame and be accessed in another (i.e. how do we manage cultural protocols over different regulatory, legal, geographic domains?)
Should meta-data also do cultural/indexing work?
Should certain objects/collections be able to resist traveling on/through the network?

</td>
<td>
Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.
This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside and acceptable range? What are our health indicators for the balance?
</td>
</tr>
</table>

<table>
<tr><th>GENERAL DYNAMIC</th><th>Questions for consideration</th><th>Possible Approaches</th></tr>
<tr>
<td>
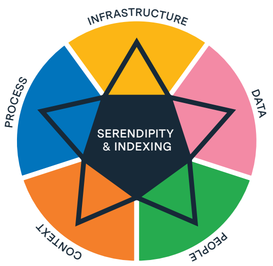

</td>
<td>
All Components
What is the right balance to strike between, on the one hand, maintaining room for audiences to create new connections among collection items and generate new knowledge and, on the other hand describing, organising, and creating indexing tools to enable easy discovery of any given item? Between some things that are fixed and perhaps “true” and others that are accumulating or changing?

How does this application affect that balance?
Where are the decision-making forums that balance serendipity and indexing?
Who are the decision makers?
How is this balance implemented across workstreams?
How do we allow for emergent meaning for individual audiences/agents that does NOT need to scale – serendipity might be intensely personal? How do we do that kind of mass customisation?

</td>
<td>
Consider developing and implementing a consistent tool (framework, methodology, technique) to monitor the balance in this dynamic.
This tool should allow us to track dynamic balance; how will we know the state of the balance? How will we know when the balance is outside and acceptable range? What are our health indicators for the balance?
</td>
</tr>
</table>

C.2 Detailed Questions & Approaches for Application-Specific Dynamic Exploration

C.2.1 System Dyad Generative Prompts

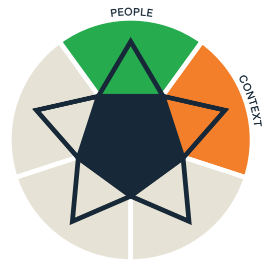

People ↔ Context

[Audience/Role]

What audience(s) does the proposed technology serve and how does that shape our roles and responsibilities when implementing the proposed application?

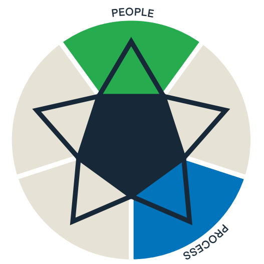

People ↔ Process

[Library staff]

How can we equip the next generation of Library staff with the skills and knowledge to work with and through the proposed application?

Process ↔ Infrastructure

[Scale/Lifecycle]

How will we plan for the proposed application at scale and throughout its full lifecycle?

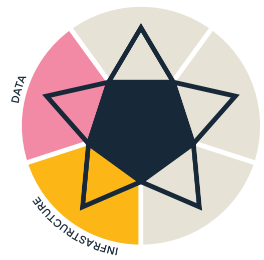

Data ↔ Infrastructure

[Accuracy]

How does the scale afforded by the proposed application affect accuracy for different types of materials?

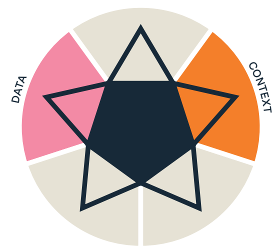

Context ↔ Data

[New connections in collection]

How could the proposed application connect items that were not previously connected in the catalogue?

People ↔ Infrastructure

[Potential harms]

What potential harms might the proposed application create or amplify?

Context ↔ Process

[Organisational considerations]

How does implementing the proposed application align with our current strategies?

Context ↔ Infrastructure

[Collaboration/Influence]

Does it make sense for us to implement the proposed application or are others better placed to do so?

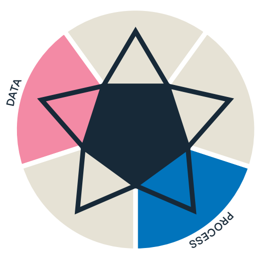

Process ↔ Data

[Errors]

What errors could occur as a result of implementing the proposed application? What will we do about them?

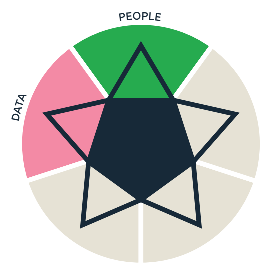

People ↔ Data

[Transparency]

How will we ensure transparency when implementing the proposed application?

C.2.2 Extended Questions and Suggested Methods for Examining System Dyads

In the table below are questions useful to explore general dynamics of the NLA's collections system as they may apply to a specific AI application, organised by the interactions of two points on the Cybernetic Star. These questions may relate to other points of the Cybernetic Star but have been connected to two key points. Questions in bold are the primary questions for consideration, with sub-questions providing context or stimulus to answer the primary question. There is also guidance on how to answer the questions for NLA consideration.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
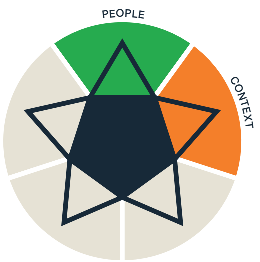	**Context ↔ People** *[Audience/role]* **What audience(s) does the proposed application serve and how does that shape our roles and responsibilities when implementing the proposed application?** • How could the proposed application improve accessibility to the collection? For which groups? • What audience(s) does the proposed application exclude? • What are the benefits of scale to our audience? • What are the benefits of scale to our staff? • How does the proposed application change the Library's responsibility for the collection to others? Who are the 'others'?	*Internal workshop:* Facilitate meeting with staff from the Collections branch and other branches affected by the proposed application. Agenda to be populated using relevant questions.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
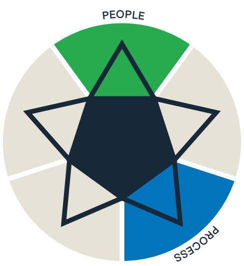	**People ↔ Process** *[Library staff]* **How can we equip the next generation of Library staff with the skills and knowledge to work with and through the proposed application?** • How can practices associated with the proposed application be combined with Library expertise (e.g. how could the library add to its existing curation practices with online 'curation' practices such as like, subscribe, thumbs up etc.)? • How can our contextual knowledge (as Library staff and other members of the library) be augmented by the proposed application? • What will the next generation of Library staff need to learn to use the proposed application?	*Capability and workforce planning:* Identify capabilities required today and in 5-10 years to develop, maintain, and monitor the proposed application. Compare to staff capabilities today, and plan workforce development activities accordingly.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
	**Process ↔ Infrastructure** *[Scale/lifecycle]* **How will we plan for the proposed application at scale and throughout its full lifecycle?** • What is the plan for scaling of the proposed application? • What is the plan for de-commissioning the application? • How do we know when we have achieved the desired scale (metrics)? How do those metrics change over time? • What resources (human, financial, and material) does the proposed application require for development, maintenance, and monitoring? How should this be modelled to understand full lifecycle costs? • What staff capabilities does the proposed application require for development, maintenance, and monitoring? Do our staff have these capabilities?	*Resources forecast:* Model resources required for development, maintenance, and monitoring over 1-, 5- and 20-year timelines.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
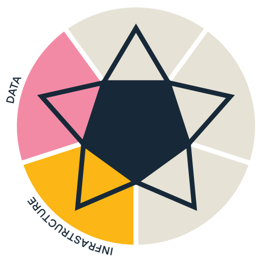	**Data ↔ Infrastructure** *[Accuracy]* **How does the scale afforded by the proposed application affect accuracy for different types of materials?** • What metrics will we use to determine if accuracy is decreasing? • Will these metrics look different for different types of materials?	*Experiment and evaluation:* Test proposed application on a subset of materials and assess (quantitative and qualitative measures) for the relevant question, prior to scale deployment.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
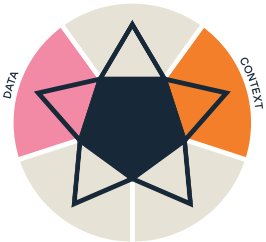	**Context ↔ Data** *[New connections in collection]* **How could the proposed application connect items that were not previously connected in the catalogue?** • What data must be added to materials in the collection to implement the proposed application (e.g. chronological, relevance, language)? • How do we determine if the Library has ascribed excessive information to a material using the proposed application? How do we course-correct if we get it wrong? • What new questions does the proposed application enable our staff to ask when using the collection? • What new questions does the proposed application enable our audience to ask when using the collection? How and when would we measure that? • How does the proposed application enable our staff to identify relationships and 'gaps' between materials within the collection? • How could the proposed application improve the breadth or depth of the collection?	*Internal workshop:* Facilitate meeting with staff from the Collections branch and other branches affected by the proposed application. Agenda to be populated using relevant questions. *Experiment and evaluation:* Test proposed application on a subset of materials and assess (quantitative and qualitative measures) for the relevant question, prior to scale deployment.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches
	**People ↔ Infrastructure** *[Harms]* **What potential harms might the proposed application create or amplify?** • How will we identify the harms created or amplified by the proposed application? • Which audiences are likely impacted? • Who will be responsible for monitoring the potential harms? • What metrics could we use to measure and monitor the impact of the potential harms created or amplified by the proposed application?	*Internal workshop:* Facilitate meeting with staff from the Collections branch and other branches affected by the proposed application. Agenda to be populated using relevant questions. Draw on Microsoft's 'Harms Modelling' practice guide (Microsoft 2020) where appropriate.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches

Context ↔ Process

[Organisational considerations]

How does implementing the proposed application align with our current strategies?

- How does implementing the proposed application align with:
  - Our Corporate Plan?
  - Our collections strategy?
  - Our IT strategy?
  - Other strategies?
- How might implementing the proposed application enable new revenue streams for us?
- How might implementing the proposed application strengthen our position as leader in the GLAM sector?

Internal workshop:

Facilitate meeting with staff from the Collections branch and other branches affected by the proposed application. Agenda to be populated using relevant questions.

SPECIFIC DYNAMIC	Questions for consideration	Possible Approaches

Context ↔ Infrastructure

[Collaboration/influence]

Does it make sense for us to implement the proposed application or are others better placed to do so?

- What do our options look like in terms of implementing the proposed application? (e.g. will we build it in-house? Will we procure an off-the-shelf solution? Will we procure a provider to build it for us? Will we champion others to build it so we can partner with them?)
- What can we learn from others who have implemented the proposed application?
- What are the benefits to us of implementing the proposed application?
- What are the costs to us of implementing the proposed application?
- If we don't implement the proposed application, who will? How might we encourage others to act?
- If we don't implement the proposed application, what will we miss out on?

Environmental scan:

Conduct desktop analysis of the market to answer relevant questions.

Internal workshop:

Facilitate meeting with staff from the Collections branch and other branches affected by the proposed application. Agenda to be populated using relevant questions.

SPECIFIC DYNAMIC | **Questions for consideration** | **Possible Approaches**

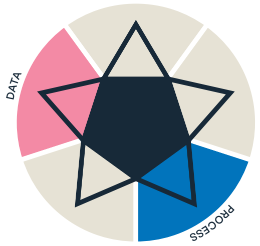

Process ↔ Data

[Errors]

What errors could occur as a result of implementing the proposed application? What will we do about them?

- What errors could be made by humans with the proposed application in place? At what parts of the process will humans be in-the-loop?
- What errors could result in our workflows because of implementing the proposed application?
- What is the process when our staff encounter an error caused by the proposed application?
- What is the process when our audiences encounter an error caused by the proposed application?

Process map:

Draw a map of the NLA processes that would affect and be affected by the application post-implementation.

SPECIFIC DYNAMIC | **Questions for consideration** | **Possible Approaches**

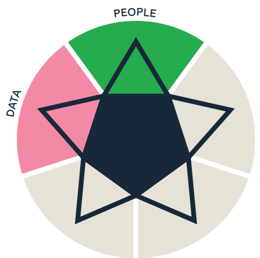

People ↔ Data

[Transparency]

How will we ensure transparency when implementing the proposed application?

- Who designed the proposed application? What influences and biases should we be aware of?
- What needs to be made explicit about our collections practices because of implementing the proposed application?
- How can we make collections practices associated with the proposed application auditable?
- How do we make our staff aware of the role of algorithmic curation that has emerged as a result of implementing the proposed application?
- How do we make our audiences aware of the role of algorithmic curation that has emerged as a result of implementing the proposed application?

Process map:

Draw a map of the NLA processes that would affect and be affected by the application post-implementation.

C.3 Cybernetic Star Guide worksheets

C.3.1 System component information

Add a description of the proposed technology and its use case in the library (2–3 sentences). Then consider context, people, data, process and infrastructure as they relate to the scenario.

PROPOSED APPLICATION:

COMPONENT	DESCRIPTION	YOUR INPUTS
CONTEXT	The external and intersecting systems, which may include global information ecosystems that cross commercial markets, regulation, funding and more, which impact the NLA's operation.	
PEOPLE	NLA audiences, workforce, and third parties referenced in library materials.	
DATA	The information collected and stored by the NLA, including the records of the collection as well as information on NLA audiences.	
PROCESSES	The chains of activity that happen in the NLA to deliver services. For example, the process that occurs when a collection item is accessioned and eventually accessed by NLA audiences.	
INFRASTRUCTURE	The information technology services in place to support NLA activities and services. The capability embedded in the NLA's infrastructure underpins the links between technologies and services.	

C.3.2 General Dynamics Worksheet

Question	Discussion	ORPI Observations
 What is the right balance to strike between the quality of collection item processing and quantity of processing that can be done in a given time interval?		

Question	Discussion	ORPI Observations
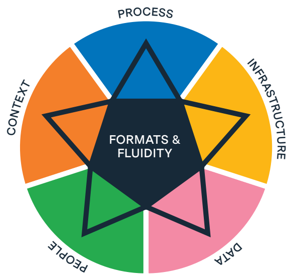 How do we strike a balance between assured longevity and risky ephemerality of collection item formats, when making decisions about which items to collect, and with what urgency and frequency?		

Question	Discussion	ORPI Observations
 How do we strike a balance between interacting with Australian publics as an audience and negotiating with them as collaborators who will have strong opinions about our collections practices and will want to influence and act upon them?		

Question	Discussion	ORPI Observations
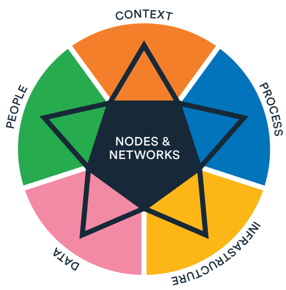 What is the right balance to strike between following our specific values, processes, infrastructure, and data standards, versus standards adopted by other organisations with which we collaborate or to which we are obligated?		

Question	Discussion	ORPI Observations
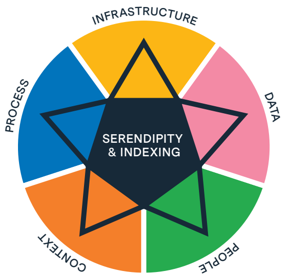 What is the right balance to strike between, on the one hand, maintaining room for audiences to create new connections among collection items and generate new knowledge and, on the other hand describing, organising, and creating indexing tools to enable easy discovery of any given item? Between some things that are fixed and perhaps “true” and others that are accumulating or changing?		

C.3.3 Application-Specific Dynamics Worksheet

Question	Discussion	ORPI Observations
 *[Audience/role]* What audience(s) does the proposed application serve and how does that shape our roles and responsibilities when implementing the proposed application?		

Question	Discussion	ORPI Observations
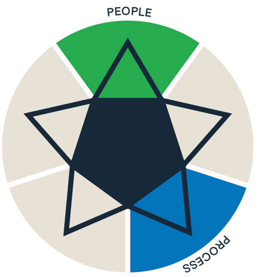 *[Library staff]* How can we equip the next generation of Library staff with the skills and knowledge to work with and through the proposed application?		

Question	Discussion	ORPI Observations
 *[Scale/lifecycle]* How will we plan for the proposed application at scale and throughout its full lifecycle?		

Question	Discussion	ORPI Observations
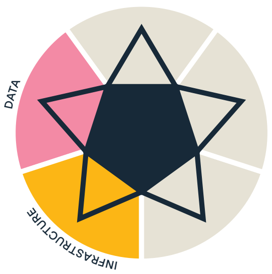 *[Accuracy]* How does the scale afforded by the proposed application affect accuracy for different types of materials?		

Question	Discussion	ORPI Observations
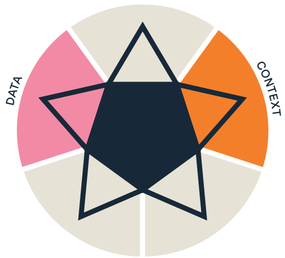 *[New connections in collection]* How could the proposed application connect items that were not previously connected in the catalogue?		

Question	Discussion	ORPI Observations
 *[Harms]* What potential harms might the proposed application create or amplify?		

Question	Discussion	ORPI Observations
 *[Organisational considerations]* How does implementing the proposed application align with our current strategies?		

Question	Discussion	ORPI Observations
 *[Collaboration/influence]* Does it make sense for us to implement the proposed application or are others better placed to do it?		

Question	Discussion	ORPI Observations
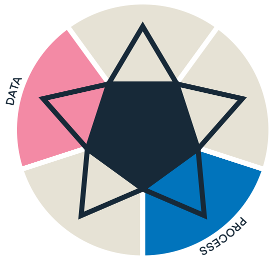 *[Errors]* What errors could occur as a result of implementing the proposed application? What will we do about them?		

Question	Discussion	ORPI Observations
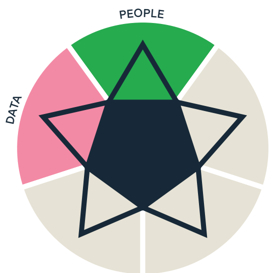 *[Transparency]* How will we ensure transparency when implementing the proposed application?		

C.3.4 ORPI Summary Worksheet

Description of proposed technology	**Activities undertaken**
	People involved

ORPI SUMMARY

Opportunities advantageous possibilities afforded by technology	
Risks potential negative consequences of the intended application of technology	
Pitfalls potential negative consequences of the unintended application of technology	
Issues known challenges with implementing the technology	
Overall	

Appendix D.
References and Resources

D.1 Overview of Resource Lists

The research team conducted a holistic literature review of resources related to AI and libraries. This included resources that described library practices that generally relate to AI (e.g., practices associated with the digitisation of materials), resources that discussed applications of AI in the library and resources that considered the implications of AI for the library. The review also included deep consideration of resources on cybernetics, and referred to principles, frameworks and codes that relate to the NLA as a Commonwealth Government institution. The consideration of these resources was critical in developing our analysis of the broad landscape of AI-capabilities in the library sector, and to the presentation of the opportunities, and the analysis of the risks, pitfalls, and issues with each of the 4 capabilities clusters discussed in Section 3.2, as well as the emerging socio-technical, business, and regulatory trends that may be relevant to the National Library in coming years (Section 5.1) .

From the literature review, we produced an annotated bibliography focused on resources discussing AI + libraries, particularly resources focused on the implications of AI for libraries and principles, frameworks and codes that relate to AI + the NLA. This bibliography has been provided to the NLA under separate cover. Also under separate cover is a Zotero library of the broader set of literature that informed this report, including materials related more generally to artificial intelligence, libraries and cybernetics.

D.2 References Cited in this Report

Abdilla, A., 2018. Beyond Imperial Tools: Future-proofing technology through Indigenous Governance and Traditional Knowledge Systems. In: J. Harle, A. Newman, and A. Abdilla, eds. *Decolonising the Digital: Technology as Cultural Practice*. Sydney: Tactical Space Lab, 67–81.

Alpert-Abrams, H., 2016. Machine Reading the Primeros Libros. *Digital Humanities Quarterly*, 10 (4).

Alveo, 2021. About [online]. *Alveo*. Available from: https://alveo.edu.au/about/ [Accessed 29 Jun 2021].

American Civil Liberties Union, 2020. Emails Show Microsoft Tried to Sell Face Recognition System to Federal Law Enforcement [Press release]. Available from: https://www.aclu.org/press-releases/emails-show-microsoft-tried-sell-face-recognition-system-federal-law-enforcement [Accesed 22 June 2021]

Ananny, M., 2016. Toward an Ethics of Algorithms: Convening, Observation, Probability, and Timeliness. *Science, Technology, & Human Values*, 41 (1), 93–117.

Ananny, M. and Crawford, K., 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20 (3), 973–989.

Andreopoulos, A. and Tsotsos, J.K., 2013. 50 Years of object recognition: Directions forward. *Computer Vision and Image Understanding*, 117 (8), 827–891.

Animate your family photos [online], 2021. *MyHeritage*. Available from: https://www.myheritage.com/deep-nostalgia [Accessed 22 Jun 2021].

Antic, J., 2021. Image Colorization API [online]. *DeepAI*. Available from: https://deepai.org/machine-learning-model/colorizer [Accessed 22 Jun 2021].

Antons, D., Grünwald, E., Cichy, P., and Salge, T.O., 2020. The application of text mining methods in innovation research: current state, evolution patterns, and development priorities. *R&D Management*, 50 (3), 329–351.

Arlitsch, K. and Newell, B., 2017. Thriving in the Age of Accelerations: A Brief Look at the Societal Effects of Artificial Intelligence and the Opportunities for Libraries. *Journal of Library Administration*, 57 (7), 789–798.

Australian Institute of Health and Welfare, 2020. The Five Safes framework. Available from: https://www.aihw.gov.au/about-our-data/data-governance/the-five-safes-framework.

Ayres, M.-L., 2013. Singing for their supper: Trove, Australian newspapers and the crowd. Presented at the IFLA World Library and Information Congress, 17-23 August 2013 Singapore. IFLA.

Baker, J., 2015. The British Library Machine Learning Experiment. *Digital Scholarship Blog*. Available from: https://blogs.bl.uk/digital-scholarship/2015/04/the-british-library-machine-learning-experiment.html. [Accessed on 22 Jun 2021]

Barnes, J., 1982. *The Presocratic Philosophers. Revised Edition*. London: Routledge & Kegan Paul.

Baró, A., Chen, J., Fornés, A., and Megyesi, B., 2019. Towards a Generic Unsupervised Method for Transcription of Encoded Manuscripts. *In*: *Proceedings of the 3rd International Conference on Digital Access to Textual Cultural Heritage*, 8 May 2019 Brussels Belgium. Association for Computing Machinery, 73–78.

Basaran, D., Essid, S., and Peeters, G., 2018. Main Melody Extraction with Source-Filter NMF and CRNN. *In*: E. Gómez, X. Hu, E. Humphrey, and E. Benetos, eds. *Proceedings of the 19th International Society for Music Information Retrieval Conference*, 23-27 September 2018 Paris France. 82–89.

Beer, S., 1959. *Cybernetics and Management*. 2nd ed. London: English Universities Press.

Beer, S., 1972. *Brain of the Firm: The Managerial Cybernetics of Organization*. London: Allen Lane The Penguin Press.

Beer, S., 1985. *Diagnosing the system for organizations*. New York: Wiley.

Bell, G., 2007. Going to the Library? A Glimpse into the future [Presidential Panel]. Reference and User Service's Association, American Library Association. Washington, DC.

Bell, G., 2010. U are Happy Life: Making sense of new technologies. *Innovation Ideas Forum 2010*. National Library of Australia.

Bell, G., 2021a. Talking to AI: An anthropological encounter with artificial intelligence. *In*: L. Pedersen and L. Cliggett, eds. *The SAGE Handbook of Cultural Anthropology*. Thousand Oaks, CA: SAGE, 428–443.

Bell, G., 2021b. Touching the Future: Stories of systems, serendipity and grace. *Griffith Review*, 71 (Remaking the Balance).

Bell, G., 2021c. We need mass surveillance to fight covid-19—but it doesn't have to be creepy. *MIT Technology Review*, (The coronavirus issue: navigating a world reshaped by covid-19).

Bell, G. and McLennan, A., 2020. The New Cybernetics: Lessons from the last great re/design. *In*: M. Schrader and V. Martens, eds. *The Great ReDesign: Frameworks for the Future*. Next Factory Ottensen, 192–209.

Benson, E., 2010. *Wired Wilderness: Technologies of Tracking and the Making of Modern Wildlife*. Baltimore: Johns Hopkins University Press.

Berners-Lee, T. and Fischetti, M., 1999. *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*. 1st ed. San Francisco: HarperSanFrancisco.

Berthon, H. and Hickie, J., 2018. *Digitising journals on Trove: a national approach to sharing* *content, engaging communities and collaboration in the digital world.* National Library of Australia. Available from: https://www.nla.gov.au/our-publications/staff-papers/digitising-journals-on-trove-a-national-approach-to-sharing-content [Accessed 15 April 2021]

Besteman, C.L. and Gusterson, H., eds., 2019. *Life by Algorithms: How Roboprocesses Are Remaking Our World*. The University of Chicago Press.

Bhagyashree, R., 2019. Google Podcasts is transcribing full podcast episodes for improving search results. *Packt Hub*. Available from: https://hub.packtpub.com/google-podcasts-is-transcribing-full-podcast-episodes-for-improving-search-results/ [Accessed 22 June 2021]

Biskup, P., 1983. *Library Models and Library Myths: The Early Years of the National Library of Australia*. Sydney: History Project Inc.

Biskup, P., 1994. State Libraries in Australia. *Alexandria: The Journal of National and International Library and Information Issues*, 6 (2), 115–131.

Biskup, P. and Goodman, D.M., 1982. *Australian Libraries*. 3rd ed. London: Bingley.

Biskup, P. and Henty, M., eds., 1991. *Library for the Nation*. Canberra: Australian Academic & Research Libraries, National Library of Australia.

Blair, A.M., 2011. *Too Much to Know: Managing Scholarly Information before the Modern Age*. New Haven, CT: Yale University Press.

Blom, P., 2005. *Enlightening the World: Encyclopédie, the Book that Changed the Course of History*. 1st ed. New York: Palgrave Macmillan.

Bode, K., 2016. Thousands of Titles Without Authors: Digitized Newspapers, Serial Fiction, and the Challenges of Anonymity. *Book History*, 19, 284–316.

Borgman, C.L., 1999. What are digital libraries? Competing visions. *Information Processing & Management*, 35 (3), 227–243.

Bouwman, A., 2018. Mundus Inversus in a Chained Library. *Leiden Special Collections Blog*. Available from: https://leidenspecialcollectionsblog.nl/articles/mundus-inversus-in-a-chained-library [Accessed 22 June 2021]

Bowker, G.C., 1993. How to be Universal: Some Cybernetic Strategies, 1943-1970. *Social Studies of Science*, 23 (1), 107–127.

Bowman, M., 2017. Ladies and libraries in nineteenth-century Melbourne. *Victorian Historical Journal*, 88 (2), 281–291.

Brill, J.H., 1998. Systems Engineering – A Retrospective View. *Systems Engineering, 1 (4)*, 258–266.

Brown, J., 2020. An examination of the Skills Framework for the Information Age (SFIA) version 7. *International Journal of Information Management*, 51 (102058).

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D., 2020. Language Models are Few-Shot Learners. *In*: H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, eds. *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, 6-12 December 2020 virtual. Curran Associates, 1877–1901.

Brownlee, J., 2019. What Are Word Embeddings for Text? *Machine Learning Mastery*. Available from: https://machinelearningmastery.com/what-are-word-embeddings/ [Accessed 22 Jun 2021].

Bryan, H., 1991. The National Library Of Australia: An Historical Perspective. *Australian Academic & Research Libraries*, 22 (4), 163–179.

Bucholtz, M., 2000. The politics of transcription. *Journal of Pragmatics*, 32 (10), 1439–1465.

Buolamwini, J., 2016a. *How I'm fighting bias in algorithms* [transcript]. TEDxBeaconStreet. Available from: https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms/transcript?language=en.

Buolamwini, J., 2016b. InCoding – In The Beginning Was The Coded Gaze. *Medium: MIT Media Lab*. https://medium.com/mit-media-lab/incoding-in-the-beginning-4e2a5c51a45d [Accessed 22 Jun 2021].

Buolamwini, J., 2018. When the Robot Doesn't See Dark Skin. *The New York Times*, 21 Jun.

Buolamwini, J. and Gebru, T., 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *In*: S. Friedler and C. Wilson, eds. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. 23-24 February 2018 New York: PMLR, 77–91.

Burrows, T.N. and Davidson, A., 2015. Enriching the HuNI Virtual Laboratory with Content from the Trove Digitized Newspapers Corpus. Presented at the Digital Humanities 2015: Global Digital Humanities, 29 June – 3 July 2015 Sydney, Australia.

Bush, V., 1945. As We May Think. *The Atlantic*.

Cai, Y. and Sun, J.-T., 2009. Text Mining. *Encyclopedia of Database Systems*.

Calhoun, K., 2014. Emergence and definitions of digital libraries. *In: Exploring Digital Libraries: Foundations, Practice, Prospects*. London: Facet, 1–26.

Candela, G., Sáez, M.D., Escobar Esteban, M.P., and Marco-Such, M., 2020. Reusing digital collections from GLAM institutions. *Journal of Information Science, 1-17*.

Candy, S., 2021. Adding Dimensions to Development Futures with UNDP. *The Sceptical Futuryst*. Available from: https://futuryst.blogspot.com/2021/03/experiential-futures-undp.html [Accessed 21 Jun 2021].

Candy, S. and Watson, J., 2021. The Thing From The Future [online]. *Situation Lab*. Available from: http://situationlab.org/project/the-thing-from-the-future/ [Accessed 22 Jun 2021].

Cao, S., 2021. 90% of NFTs Will Be Worthless in 3-5 Years, Coinbase Cofounder Warns. *Observer,* 18 Jun.

Cathro, W. and Collier, S., 2010. Developing Trove; the policy and technical challenges. *eLucidate*, 7 (3).

Ceruzzi, P.E., 1998. *A History of Modern Computing*. Cambridge, MA: MIT Press.

Chan, L.M., 2007. *Cataloguing and Classification: An Introduction*. Lanhamn, MD: Scarecrow Press.

Chanod, J.-P., 1999. Natural Language Processing and Digital Libraries. *In*: *Information Extraction: Towards Scalable, Adaptable Systems*. Berlin, Heidelberg: Springer-Verlag, 17–31.

Chokshi, N., 2019. Hackers Are Holding Baltimore Hostage: How They Struck and What's Next. *The New York Times*, 22 May.

Chowdhury, G., 2003. Natural language processing. *Annual Review of Information Science and Technology*, 37 (1), 51–89.

Christie's, 2018. Is artificial intelligence set to become art's next medium? *Christie's*.

Clapp, V., 1951. Archivists and Bibliographical Control: A Librarian's Viewpoint. *The American Archivist*, 14 (4), 305–311.

Clapp, V., 1964. *The Future of Research Libraries*. First ed. University of Illinois Press.

Clarke, R.I. and Schoonmaker, S., 2020. The Critical Catalog: Library Information Systems, Tricksterism, and Social Justice. *In*: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 25-30 April 2020 Honolulu. Association of Computing Machinery, 1–13.

Cleveland, G., 1998. *Digital libraries: Definitions, issues and challenges*. Ottawa: International Federation of Library Associations and Institutions.

Cloud AutoML Custom Machine Learning Models [online], 2021. *Google Cloud*. Available from: https://cloud.google.com/automl [Accessed 22 Jun 2021].

Coffey, D., 2021. Māori are trying to save their language from Big Tech. *Wired UK*.

Collisions: Lynette Wallworth, 2016. *ACMI*.

Conley, S.N. and York, E., 2020. Public engagement in contested political contexts: reflections on the role of recursive reflexivity in responsible innovation. *Journal of Responsible Innovation*, 7 (sup1), 1–12.

Cordell, R., 2020. *Machine Learning + Libraries: A Report on the State of the Field*. Boston, MA: LC Labs Library of Congress.

COVID-19 Community Mobility Reports [online], 2021. *Google*. Available from: https://www.google.com/covid19/mobility?hl=en [Accessed 22 Jun 2021].

Crawford, K., 2019. Halt the use of facial-recognition technology until it is regulated. *Nature*, 572, 565–565.

Crawford, K., Dobbe, R., Dryer, T., Fried, G., Green, B., Kaziunas, E., Kak, A., Mathur, V., McElroy, E., Nill Sanchez, A., Raji, D., Rankin, J.L., Richardson, R., Schultz, J., Myers West, S., and Whittaker, M., 2019. *AI Now 2019 Report*. New York: AI Now Institute.

Crawford, K. and Whittaker, M., 2016. *The AI Now Report: The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term*. New York: AI Now Institute.

Create ML – Machine Learning [online], 2021. *Apple Developer*. Available from: https://developer.apple.com/machine-learning/ [Accessed 22 Jun 2021].

Cultural AI Lab [online], 2021. *Innovation Center for Artificial Intelligence*. Available from: https://icai.ai/cultural-ai-lab/ [Accessed 5 Jun 2021].

Datig, I. and Whiting, P., 2018. Telling your library story: tableau public for data visualization. *Library Hi Tech News*, 35 (4), 6–8.

Department of Industry, Science, Energy and Resources, 2020. Australia's Artificial Intelligence Ethics Framework.

Desk Set, 1957. Film. 20th Century Fox.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K., 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *In*: J. Burstein, C. Doran, and T. Solorio, eds. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2-7 June 2019 Minneapolis USA. Association for Computational Linguistics, 4171–4186.

Dewey, M., 1876. *A Classification and Subject Index for Cataloguing and Arranging the Books and Pamphlets of a Library*. Amherst, MA.

Diebold, J., 1966. Computers, Program Management and Foreign Affairs. *Foreign Affairs*, 45 (1), 125–134.

Dignum, V., 2019. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer International Publishing.

Disability Discrimination Act 1992 (Cth), Act 135.

Dixon, S., Gabbard, K., Crayencour, H.-C., Peeters, G., Pfeiderer, M., Solis, G., and Weyde, T., 2019. Dig That Lick: Analysing large-scale data for melodic patterns in jazz performances. About. [online]. *UK Research and Innovation*. Available from: http://dig-that-lick.eecs.qmul.ac.uk/Dig%20That%20Lick_About.html [Accessed 4 Jun 2021].

Dourish, P. and Bell, G., 2007. The Infrastructure of Experience and the Experience of Infrastructure: Meaning and Structure in Everyday Encounters with Space. *Environment and Planning B: Planning and Design*, 34 (3), 414–430.

Draxler, C., van Hessen, A., and van de Heuvel, H., 2019. *Transcription Chain*. Presented at the DH2019 Workshop: The Case of Interview Data, Utrecht, The Netherlands.

Dubberly, H., Esmonde, P., Geoghegan, M., and Pangaro, P., 2014. Notes on the Role of Leadership & Language in Regenerating Organizations. *In*: M. Shamiyeh, ed. *Driving Desired Futures*. Berlin: Birkhäuser, 400–418.

Duke, C., 2014. *The History of Modern Japanese Education: Constructing the National School System, 1872-1890*. Rutgers University Press.

Dwyer, K., 2021. Emily Ratajkowski Is Selling an NFT at Christie's. *The New York Times*, 23 Apr.

Eaton, M., 2017. Seeing Library Data: A Prototype Data Visualization Application for Librarians. Journal of Web Librarianship, 11 (1), 69–78.

Egan, M.E. and Shera, J.H., 1949. Prolegomena to bibliographic control. *Journal of cataloging and classification*, 5 (2), 17–19.

Elan, P., 2021. Kate Moss auctions Sleep With Kate video as non-fungible token. *The Guardian*, 20 May.

Elizalde, B., Shah, A., Dalmia, S., Lee, M.H., Badlani, R., Kumar, A., Raj, B., and Lane, I., 2017. An Approach for Self-Training Audio Event Detectors Using Web Data. *In: 2017 25th European Signal Processing Conference (EUSIPCO)*. August 28 - September 2 2017 Kos, Greece. Piscataway, NJ: IEEE, 1863–1867.

EuroCC National Competence Centre Sweden, 2021. National Library of Sweden Now Has Access to VEGA. *ENCCS News*. Available from: https://enccs.se/news/2021/06/kb-vega/ [Accessed 22 Jun 2021].

European Parliament, Council of the European Union, 2016. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal, L119*.

EuropeanaTech AI for GLAMs task force, 2020. *Interim Analysis of EuropeanaTech AI in Relation to GLAMs Survey*. Europeana Network Association.

Facebook AI and Carnegie Mellon University, 2021. Open Catalyst Project: Home [online]. *Open Catalyst Project*. Available from: https://opencatalystproject.org/ [Accessed 22 Jun 2021].

Fagen, M.D., Joel, A.E., Schindler, G.E., and O'Neill, E.F., 1975. *A History of Engineering and Science in the Bell System: National Service in War and Peace* (1925 - 1975). New York: Bell Telephone Laboratories.

Farthing, S., Howell, J., Lecchi, K., Paleologos, Z., Saintilan, P., and Santow, E., 2019. *Human Rights and Technology Discussion Paper.* Sydney: Australian Human Rights Commission.

Farthing, S., Howell, J., Lecchi, K., Paleologos, Z., Saintilan, P., and Santow, E., 2021. *Human Rights and Technology Final Report*. Sydney: Australian Human Rights Commission.

Fatourechi, M., 2015. The Evolving Landscape of Recommendation Systems. *Tech Crunch*. Available from: https://techcrunch.com/2015/09/28/the-evolving-landscape-of-recommendation-systems/ [Accessed 22 Jun 2021].

Felton-Taylor, A. and Claughton, D., 2015. Drones to help manage invasive pest animal species in Australia. *Australian Broadcasting Corporation*, 22 Jul.

Ferriter, M., 2019. Introducing the Computing Cultural Heritage in the Cloud Project. *The Signal.* Available from: https://blogs.loc.gov/thesignal/2019/11/introducing-the-computing-cultural-heritage-in-the-cloud-project/ [Accessed 22 June 2021].

Fielding, D., 1991. Changing roles: an outsider's view of the National Library. *Special issue: Library for the Nation, Australian Academic & Research Libraries*, 22 (4), 151–162.

Finn, E., 2017. *What Algorithms Want: Imagination in the Age of Computing*. Cambridge, MA: MIT Press.

Fisher, B. and Evans, F.B., 1967. Automation, Information, and the Administration of Archives and Manuscript Collections: Bibliographic Review. *The American Archivist*, 30 (2), 333–348.

Fitzgerald, B., 2020. *NSW Public Libraries- the COVID-19 response: Report of a survey into how libraries adapted to a pandemic*. State Library of NSW.

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., and Srikumar, M., 2020. *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI.* Berkman Klein Center for Internet & Society at Harvard University.

Fjeldsted, S., 2016. Library Display Honors Ray Bradbury Upon His Induction into the California Library Hall of Fame. *City of South Pasadena*. Available from: https://www.southpasadenaca.gov/Home/Components/News/News/749/714?arch=1&npage=20 [Accessed 22 June 2021].

Floridi, L. and Chiriatti, M., 2020. GPT-3: Its Nature, Scope, Limits, and Consequences. *Minds and Machines*, 30 (4), 681–694.

von Foerster, H., 1984. Technology: What will it mean for Libraries. *In*: *Observing Systems*. USA: Intersystems Publications, 211–230.

Foley, B., Arnold, J., Coto-Solano, R., Durantin, G., Ellison, T.M., van Esch, D., Heath, S., Kratochvil, F., Maxwell-Smith, Z., Nash, D., Olsson, O., Richards, M., Stoakes, H., Thieberger, N., and Wiles, J., 2018. Building Speech Recognition Systems for Language Documentation: The CoEDL Endangered Language Pipeline and Inference System (Elpis). In: *Proceedings of The 6th International Workshop on Spoken Language Technologies for Under-Resourced Languages*, 29-31 August 2018 Gurugram India. 205–209.

Gallus, G. and Regoliosi, G., 1974. A decisional model of recognition applied to the chromosome boundaries. *Journal of Histochemistry & Cytochemistry*, 22 (7), 546–553.

Gibson, W., 2003. The future is already here – it's just not evenly distributed. *The Economist*, 4 (2), 152.

Giving in Australia: the fast facts [online], 2020. *Philanthropy Australia*. Available from: https://www.philanthropy.org.au/tools-resources/fast-facts-and-stats/ [Accessed 22 Jun 2021].

Global Future Council on Augmented Reality and Virtual Reality [online], 2021. *World Economic Forum*. Available from: https://www.weforum.org/communities/gfc-on-virtual-reality-and-augmented-reality/ [Accessed 22 Jun 2021].

Global Indigenous Data Alliance, 2018. CARE Principles for Indigenous Data Governance. Available from: https://www.gida-global.org/care [Accessed 22 Jun 2021].

Godfrey, H., 2018. *Submarine Telegraphy and the Hunt for Gutta Percha: Challenge and Opportunity in a Global Trade*. Submarine Telegraphy and the Hunt for Gutta Percha. Brill.

Goodchild, M., 2021. Relational Systems Thinking: That's How Change Is Going to Come, from Our Earth Mother. *Journal of Awareness-Based Systems Change*, 1 (1), 75–103.

Goode, H.H. and Machol, R.E., 1957. *System Engineering: An Introduction to the Design of Large-scale Systems*. McGraw-Hill.

Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., 2014. Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]*.

Google PAIR, 2021. People + AI Guidebook [online]. Available from: https://pair.withgoogle.com/guidebook [Accessed 4 Jun 2021].

Gray, M.L. and Suri, S., 2019. *Ghost Work: How to Stop Silicon Valley From Building a New Global Underclass*. Boston, MA: Houghton Mifflin Harcourt.

Greene, J., 2020. Microsoft bans police from using its facial-recognition technology. *The Washington Post*, 12 Jun.

Greenfield, A., 2021. Non-fungible tokens aren't a harmless digital fad – they're a disaster for our planet. *The Guardian*, 29 May.

Griffen, A.M., 1987. Images of Libraries in Science Fiction. *Library Journal*, 112 (14), 137–142.

Gu, F., Sarkar, S., Sun, Y., Wu, H., and Wu, K., 2020. NLP: Word Embedding Techniques for Text Analysis. *NLP: Word Embedding Techniques for Text Analysis*.

Hajkowicz, S., Karimi, S., Wark, T., Chen, C., Evans, M., Rens, N., Dawson, D., Charlton, A., Brennan, T., Moffatt, C., Srikumar, S., and Tong, K., 2019. *Artificial intelligence: Solving problems, growing the economy and improving our quality of life*. Australia: CSIRO Data61.

Hall, A.D., 1962. *A Methodology for Systems Engineering*. New York: D. Van Nostrand Company.

Hang, X., 2017. Microsoft researchers achieve new conversational speech recognition milestone. *Microsoft Research*.

Hao, K., 2020. An AI helps you summarize the latest in AI. *MIT Technology Review*.

Hayles, N.K., 1999. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. University of Chicago Press.

Heims, S.J., 1993. *Constructing a social science for postwar America: the cybernetics group, 1946-1953*. Cambridge, MA: MIT Press.

Henry, G., 2019. Research Librarians as Guides and Navigators for AI Policies at Universities (RLI 299, 2019). *Research Library Issues*, (299), 47–65.

Hern, A., 2021. Deep Nostalgia: ‘creepy’ new service uses AI to animate old family photos. *The Guardian*, 2 Mar.

Hervieux, S. and Wheatley, A., 2021. Perceptions of artificial intelligence: A survey of academic librarians in Canada and the United States. *The Journal of Academic Librarianship*, 47 (1).

Hider, P., 2019. From Anderson to ORDAC: A History of Bibliographic Policy Discussion in Australia. *Journal of the Australian Library and Information Association*, 68 (1), 3–17.

Hilton, A.M., 1963. *Logic, computing machines, and automation*. Washington: Spartan Books.

Hilton, A.M., 1973. Cybernetics and Cybernation. *The Science Teacher*, 40 (2), 34–40.

Hoare, P. and Black, A., eds., 2006. *Cambridge History of Libraries in Britain and Ireland*. Volume 3, 1850-2000. Cambridge University Press.

Holley, R., 2009. How Good Can It Get?: Analysing and Improving OCR Accuracy in Large Scale Historic Newspaper Digitisation Programs. *D-Lib Magazine*, 15 (3/4).

Holley, R., 2010. Trove: Innovation in Access to Information in Australia. *Ariadne*, July (64).

Hunt, T.D., Ryan, G., and Ryan-Pears, C., 2017. An investigation and comparison of speech recognition software for determining if bird song recordings contain legible human voices. *Journal of Applied Computing and Information Technology*, 21 (1).

IFTF Foresight Essentials [online], 2021. *Institute for the Future*. Available from: https://www.iftf.org/foresightessentials/ [Accessed 22 Jun 2021].

Indigital [online], 2021. *Indigital*. Available from: https://indigital.net.au/ [Accessed 22 Jun 2021].

International Federation of Library Associations, 2013. IFLA Trend Report 2013 – Literature Review.

International Federation of Library Associations, 2020. IFLA Statement on Libraries and Artificial Intelligence.

Isaacson, K., 1982. Machine-Readable Information in the Library. *RQ*, 22 (2), 164–170.

Jade, M., 2021. *A discussion paper to propose Cultural Protocols for Dharug Communities in the Spatial Web*. Canberra: School of Cybernetics, Capstone project discussion paper.

Jain, A., Liu, I., Sarda, A., and Molino, P., 2019. Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations. *Uber Engineering Blog*. Available from: https://eng.uber.com/uber-eats-graph-learning/ [Accessed 22 June 2021].

Jakeway, E., 2020. Machine Learning + Libraries Summit: Event Summary now live! *The Signal*. Available from: https://blogs.loc.gov/thesignal/2020/02/machine-learning-libraries-summit-event-summary-now-live/ [Accessed 22 June 2021].

Johnson, B.D., 2011a. *Science Fiction Prototyping: Designing the Future with Science Fiction*. San Rafael, CA: Morgan & Claypool.

Johnson, B.D., 2011b. Science Fiction Prototyping: Designing the Future with Science Fiction. *Synthesis Lectures on Computer Science*, 3 (1), 1–190.

Jones, D.J., 2005. Public library development in New South Wales. *The Australian Library Journal*, 54 (2), 130–137.

Jones, N., 2021. How scientists are embracing NFTs. *Nature*, 594, 481–482.

Jones, R., 2020. Digital opportunities: what digital skills do library staff need in a changing world? And do they have them? *In*: M. Weaver and L. Appleton, eds. *Bold minds: library leadership in a time of disruption*. London: Facet Publishing, 211–226.

Kangas, P., 2019. AI use cases [online]. *AI for librarians*. Available from: https://www.aiforlibrarians.com/ai-cases/ [Accessed 29 Jun 2021].

Karnouskos, S., 2020. Artificial Intelligence in Digital Media: The Era of Deepfakes. *IEEE Transactions on Technology and Society*, 1 (3), 138–147.

Kastrenakes, J., 2021. Beeple sold an NFT for $69 million. *The Verge*.

KB Labs [online], 2021. *Det Kgl. Bibliotek*. Available from: https://labs.kb.dk/ [Accessed 29 Jun 2021].

Keller, E.F., 2004. Marrying the premodern to the postmodern: computers and organisms after WWII. *In*: *Prefiguring Cyberculture: An Intellectual History*. Cambridge, MA: MIT Press, 52–65.

Kelly, K., 1992. In the Library of Form. *In*: *Out of Control: The New Biology of Machines, Social Systems, and the Economic World*. Basic Books, 221–239.

Kelly, T., 1952. The Origin of Mechanics' Institutes. *British Journal of Educational Studies*, 1 (1), 17–27.

Kemp, K. and Greenleaf, G., 2020. The COVIDSafe bill doesn't go far enough to protect our privacy. Here's what needs to change. *The Conversation*. Available from: http://theconversation.com/the-covidsafe-bill-doesnt-go-far-enough-to-protect-our-privacy-heres-what-needs-to-change-137880 [Accessed 19 Jun 2021].

Kesserwan, K., 2018. How can Indigenous knowledge shape our view of AI? *Policy Options*.

Kettunen, K., Kervinen, J., and Koistinen, M., 2018. Creating and using ground truth OCR sample data for Finnish historical newspapers and journals. *In*: E. Mäkelä, M. Tolonen, and J. Tuominen, eds. *Proceedings of the Digital Humanities in the Nordic Countries 3rd Conference*. March 7-9 2018 Helsinki, Finland, 162–169.

Kettunen, K., Koistinen, M., and Kervinen, J., 2020. Ground Truth OCR Sample Data of Finnish Historical Newspapers and Journals in Data Improvement Validation of a re-OCRing Process. *LIBER Quarterly*, 30 (1), 1–20.

Key Principles [online], 2021. *Maiam nayri Wingara*. Available from: https://www.maiamnayriwingara.org/key-principles [Accessed 22 Jun 2021].

Keyes, O., 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proceedings of the ACM on Human-Computer Interaction*, 2 (CSCW), 1–22.

Kilgour, F.G., 1970. History of Library Computerization. *Information Technology and Libraries*, 3 (3), 218–229.

Kleiner, A. and Kurzweil, R., 1977. A description of the Kurzweil reading machine and a status report on its testing and dissemination. *Bulletin of prosthetics research*, 1027, 72–81.

Kline, R.R., 2015. *The Cybernetics Moment: Or Why We Call Our Age the Information Age*. JHU Press.

Koehn, P. and Knowles, R., 2017. Six Challenges for Neural Machine Translation. *In*: *Proceedings of the First Workshop on Neural Machine Translation*. 30 July – August 4 Vancouver. Association for Computational Linguistics, 28–39.

Konch, 2020. Konch – Powerful transcription platform made for researchers [online]. Available from: https://www.konch.ai/product [Accessed 22 Jun 2021].

Krajewski, M., 2011. *Paper machines*. Cambridge, MA: MIT Press.

Kukutai, T. and Taylor, J., eds., 2016. *Indigenous Data Sovereignty: Toward an agenda*. Canberra: ANU Press.

Kumar, A. and Raj, B., 2016. Audio Event Detection using Weakly Labeled Data. *In: Proceedings of the 24th ACM international conference on Multimedia*, 15-19 October 2016 Amsterdam. New York: Association for Computing Machinery, 1038–1047.

Lee, B.C.G., 2020. Compounded Mediation: A Data Archaeology of the Newspaper Navigator Dataset. *Humanities Commons*.

Lee, H., 2012. Epistemic foundation of bibliographic classification in early China: A Ru classicist perspective. *Journal of Documentation*, 68 (3), 378–401.

Leedham, C., 1994. Historical Perspectives of Handwriting Recognition Systems. *In: IEEE Colloquium on Handwriting and Pen-Based Input*, 11 March 1994 London. Institution of Engineering and Technology.

LeGuin, U., 1985. *Always Coming Home*. 1st ed. New York: Harper & Row.

Lewis, J.E., 2019. An orderly assemblage of biases: Troubling the monocultural stack. *In*: I. Schweitzer and G. Henry, eds. *Afterlives of Indigenous Archives*. Hanover, NH: Dartmouth College Press, 218–233.

Li, C.-T., 2020. 10 Papers You Must Read for Deep Image Inpainting. *Towards Data Science.* Available from: https://towardsdatascience.com/10-papers-you-must-read-for-deep-image-inpainting-2e41c589ced0 [Accessed 12 June 2021]

Liddy, E.D., 1998. Enhanced Text Retrieval Using Natural Language Processing. *Bulletin of the American Society for Information Science and Technology*, 24 (4), 14–16.

Liker, J.K., Ettlie, J.E., and Campbell, J.C., eds., 1995. *Engineered in Japan: Japanese Technology-management Practices*. Oxford University Press.

Ling, T. and Mclean, A., 2004. Taking it to the People: Why the National Archives of Australia Embraced Digitisation on Demand. *Australian Academic & Research Libraries*, 35 (1), 2–15.

Lorang, E., Soh, L.-K., Liu, Y., and Pack, C., 2020. Digital Libraries, Intelligent Data Analytics, and Augmented Description: A Demonstration Project. *Faculty Publications, UNL Libraries*, 396.

Lowden, B., 2006. *Mechanics' institutes, schools of arts, athenaeums, etc. : an Australian checklist*. 1st ed. Donvale, Victoria: Lowden Publishing.

Lynch, C., 2002. Digital Collections, Digital Libraries and the Digitisation of Cultural Heritage Information. *First Monday*, 7 (5).

Lyons, M., 2011. *Books: A Living History*. Farnborough, Hampshire: Thames & Hudson.

Maguire, C., 2016. Ifould, Remington and Metcalfe: three mostly wise men behind the New South Wales Library Act 1939. *The Australian Library Journal*, 65 (1), 41–49.

Malapi-Nelson, A., 2017. *The Nature of the Machine and the Collapse of Cybernetics*. Cham, Switzerland: Springer International Publishing.

McCallum, S., 2003. 40 Years of Technology in Libraries: A Brief History of the IFLA Section on Information Technology, 1963/64 – 2003. International Federation of Library Associations and Institutions.

McCarthy, J., Minsky, M., Rochester, N., and Shannon, C., 1955. *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. Hanover, NH: Dartmouth College, Research proposal.

McCue, T.J., 2019. Amazon Alexa Wants To Guard Your House And Listen For Broken Glass, Smoke Alarm. *Forbes*.

McInerney, T. and Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1 (2), 91–108.

Melville, P. and Sindhwani, V., 2010. Recommender Systems. *Encyclopedia of Machine Learning*.

Menn, J. and Satter, R., 2021. Pipeline hackers say their aim is cash, not chaos. *Reuters*, 10 May.

Microsoft, 2020. Foundations of assessing harm: Microsoft Azure. Available from: https://docs.microsoft.com/en-us/azure/architecture/guide/responsible-innovation/harms-modelling/ [Accessed 22 June 2021].

Microsoft, 2021. Our approach to responsible AI at Microsoft [online]. *Microsoft*. Available from: https://www.microsoft.com/en-us/ai/our-approach [Accessed 4 Jun 2021].

Middleton, M., 2006. Scientific and Technological Information Services in Australia: I. History and Development. *Australian Academic & Research Libraries*, 37 (2), 111–135.

Minto, J., 1932. *A history of the public library movement in Great Britain and Ireland*. London: George Allen and Unwin.

Mirsky, Y. and Lee, W., 2021. The Creation and Detection of Deepfakes: A Survey. *ACM Computing Surveys*, 54 (1), 1–41.

Mohamed, S., Png, M.-T., and Isaac, W., 2020. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology*, 33 (4), 659–684.

Mohammed, H., Märgner, V., and Ciotti, G., 2021. Learning-free pattern detection for manuscript research:: An efficient approach toward making manuscript images searchable. *International Journal on Document Analysis and Recognition (IJDAR)*.

Moretti, F., 2013. *Distant Reading*. New York: Verso.

Morris, A., 2017. After the Facts. *Los Angeles Review of Books*.

Moylan, B., 2019. Decolonising Historical Maps. *National Library of Australia Behind the Scenes*.

Muehlberger, G., Seaward, L., Terras, M., Ares Oliveira, S., Bosch, V., Bryan, M., Colutto, S., Déjean, H., Diem, M., Fiel, S., Gatos, B., Greinoecker, A., Grüning, T., Hackl, G., Haukkovaara, V., Heyer, G., Hirvonen, L., Hodel, T., Jokinen, M., Kahle, P., Kallio, M., Kaplan, F., Kleber, F., Labahn, R., Lang, E.M., Laube, S., Leifert, G., Louloudis, G., McNicholl, R., Meunier, J.-L., Michael, J., Mühlbauer, E., Philipp, N., Pratikakis, I., Puigcerver Pérez, J., Putz, H., Retsinas, G., Romero, V., Sablatnig, R., Sánchez, J.A., Schofield, P., Sfikas, G., Sieber, C., Stamatopoulos, N., Strauß, T., Terbul, T., Toselli, A.H., Ulreich, B., Villegas, M., Vidal, E., Walcher, J., Weidemann, M., Wurster, H., and Zagoris, K., 2019. Transforming scholarship in the archives through handwritten text recognition: Transkribus as a case study. *Journal of Documentation*, 75 (5), 954–976.

Mühling, M., Meister, M., Korfhage, N., Wehling, J., Hörth, A., Ewerth, R., and Freisleben, B., 2019. Content-based video retrieval in historical collections of the German Broadcasting Archive. *International Journal on Digital Libraries*, 20 (2), 167–183.

Murison, W., 1971. *The Public Library: its origins, purpose and significance*. 2nd ed. London: Harrap.

National Library Act 1960 (Cth).

National Library of Australia, 2013. Digital Preservation Policy 4th Edition (2013).

National Library of Australia, 2017. DLIR Program details [online]. *National Library of Australia*. Available from: https://www.nla.gov.au/dlir/project-details [Accessed 28 Jun 2021].

National Library of Australia, 2020a. Collecting Strategy 2020-21 – 2023-24.

National Library of Australia, 2020b. Corporate Plan.

National Library of Australia, 2020c. *Annual Report 2019-2020*. Parkes Place, Canberra: National Library of Australia.

National Library of Australia, 2021. Digitisation of Library collections [online]. *National Library of Australia*. Available from: https://www.nla.gov.au/digitisation [Accessed 28 Jun 2021].

Neilson, O., 2010. Sharing Our Nation's Treasures: Digitisation at the National Library of Australia. *Access*, 24 (2), 10–15.

Neudecker, C., Baierer, K., Federbusch, M., Boenig, M., Würzner, K.-M., Hartmann, V., and Herrmann, E., 2019. OCR-D: An end-to-end open source OCR framework for historical printed documents. *In*: *DATeCH2019: Proceedings of the 3rd International Conference on Digital Access to Textual Cultural Heritage*, 8-10 May 2019 Brussels Belgium. New York: Association for Computing Machinery, 53–58.

Newton, C., 2016. Speak, Memory: When her best friend died, she used artificial intelligence to keep talking to him. *The Verge*.

Noble, S.U., 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press.

Ong, W.J., 1982. *Orality and Literacy: The Technologizing of the Word*. New York: Methuen.

Owens, T., 2018. *The Theory and Craft of Digital Preservation*. Baltimore: Johns Hopkins University Press.

Padilla, T., 2019. Responsible Operations: Data Science, Machine Learning, and AI in Libraries. *OCLC Research*.

Padilla, T., Allen, L., Frost, H., Potvin, S., Roke, E.R., and Varner, S., 2019. *Always Already Computational: Collections as Data: Final Report*. OSF.

Paglen, T., 2016. Invisible Images (Your Pictures Are Looking at You). *The New Inquiry*.

Pandey, P., 2019. The Remarkable World of Recommender Systems. *Towards Data Science*. Available from: https://towardsdatascience.com/the-remarkable-world-of-recommender-systems-bff4b9cbe6a7 [Accessed 22 June 2021].

Pangaro, P., 2006. Cybernetics: A Definition. *Paul Pangaro PhD*. Available from: https://pangaro.com/definition-cybernetics.html [Accessed 22 June 2021].

Partnership on AI, 2020. *Understanding Facial Recognition Systems*. San Franscisco: Partnership on AI.

Paybarah, A., 2020. Ransomware Attack Closes Baltimore County Public Schools. *The New York Times*, 30 Nov.

Pias, C., 2016. *The Age of Cybernetics*. Zurich, Switzerland: Diaphanes AG.

Porter, J., 2019. Popularity/Prestige: A New Canon. *Stanford Literary Lab*. Available from: https://litlab.stanford.edu/pop-pres-ny/ [Accessed 22 June 2021].

Prizeman, O., 2012. *Philanthropy and Light: Carnegie Libraries and the Advent of Transatlantic Standards for Public Space*. Routledge.

Huang, B.H. and Rabiner, L., 2004. Automatic Speech Recognition – A Brief History of the Technology Development.

Radebaugh, A., 1959. Closer Than We Think: Electronic Home Library. *Chicago Sunday Tribune*, 1 Feb, p. 10.

Rahnama, H., 2021. Augmented Eternity and Swappable Identities. *MIT Media Lab*.

Ratajkowski, E., 2020. Buying Myself Back. *The Cut*.

Raven, J., 2006. Libraries for sociability: The advance of the subscription library. In: G. Mandelbrote and K.A. Manley, eds. *The Cambridge History of Libraries in Britain and Ireland.* Volume 2 1640–1850. Cambridge University Press, 239–263.

Remington, G.C. and Metcalfe, J., 1945. The Free Library Movement—1935-1945. *The Australian Quarterly*, 17 (2), 87–97.

Rid, T., 2016. *Rise of the Machines: The Lost History of Cybernetics*. Brunswick, VIC: Scribe Publications.

Rigaud, C., Doucet, A., Coustaty, M., and Moreux, J.-P., 2019. ICDAR 2019 Competition on Post-OCR Text Correction. *In*: *2019 International Conference on Document Analysis and Recognition (ICDAR)*, 20-25 September 2019 Sydney Australia. Los Alamitos: IEEE Computer Society, 1588–1593.

Rochester, M.K., 1990. Australian Library History Research. *Libraries & Culture*, 25 (1), 115–129.

Rossi, F., 2020. How IBM Is Working Toward a Fairer AI. *Harvard Business Review*. 5 Nov.

Rowe, J.S. and Ryan, M., 1974. Library Service from Numerical Data Bases: The 1970 Census as a Paradigm. *College and Research Libraries*, 35 (1), 7–15.

SAS Denmark, n.d. Demand forecasts keep readers engaged. *SAS Institute*. Available from: https://www.sas.com/no_no/customers/singapore-national-library-board.html [Accessed 22 June 2021].

Schantz, H.F., 1982. *The history of OCR, optical character recognition*. Manchester Center: Recognition Technologies Users Association.

Schjønhaug AS: Home [online], 2021. *Schjønhaug AS*. Available from: https://en.schjonhaug.com [Accessed 22 Jun 2021].

Schlager, J., 1956. Systems engineering: key to modern development. *IRE Transactions on Engineering Management*, EM-3 (3), 64–66.

Schwartz, C., 2000. Digital Libraries: An Overview. *Journal of Academic Librarianship*, 26 (6), 385–93.

Schwartz, O., 2019. In 2016, Microsoft's Racist Chatbot Revealed the Dangers of Online Conversation. *IEEE Spectrum*.

Scott, A.F., 1986. Women and Libraries. *Journal of Library History*, 21 (2), 400–405.

Selyukh, A., (Pres) 2018, 'Microsoft Urges Congress to Regulate Facial Recognition Technology', *All Things Considered*, M.L. Kelly (Prod), Radio program transcript, National Public Radio, 6 December 2018. Available from: https://www.npr.org/2018/12/06/674310978/microsoft-urges-congress-to-regulate-facial-recognition-technology [Accessed 22 June 2021].

Shahid, S.M., 2005. Use of RFID technology in libraries: A new approach to circulation, tracking, inventorying, and security of library materials. *Library Philosophy and Practice*, 8 (1), 1–9.

Siwecka, D., 2018. Knowledge Organization systems used in European national libraries towards interoperability of the semantic Web. *In*: F. Ribeiro and M.E. Cerveira, eds. *Challenges and Opportunities for Knowledge Organization in the Digital Age: Proceedings of the Fifteenth International ISKO Conference*, 9-11 July 2018 Porto Portugal. Baden Baden: Ergon, 633–643.

Slotkin, J., 2017. 'Monkey Selfie' Lawsuit Ends With Settlement Between PETA, Photographer. *NPR The Two-Way*.

State Library New South Wales, 2020. Amplify: audio transcription [online]. *State Library of NSW*. Available from: https://www.sl.nsw.gov.au/public-library-services/amplify-audio-transcription [Accessed 22 Jun 2021].

StatsNZ, 2020. *Algorithm Charter for Aotearoa New Zealand*. Wellington: New Zealand Government.

Stearns, P., 2019. The Honeypot Collection now on Kickstarter [online]. *Glitch Textiles*. Available from: https://glitchtextiles.com/blog/2019/8/9/the-honeypot-collection-now-on-kickstarter [Accessed 25 Jun 2021].

Stokes, J.M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N.M., MacNair, C.R., French, S., Carfrae, L.A., Bloom-Ackermann, Z., Tran, V.M., Chiappino-Pepe, A., Badran, A.H., Andrews, I.W., Chory, E.J., Church, G.M., Brown, E.D., Jaakkola, T.S., Barzilay, R., and Collins, J.J., 2020. A Deep Learning Approach to Antibiotic Discovery. *Cell*, 180 (4), 688-702.

Strava, 2021. Strava Global Heatmap [online]. *Strava*. Available from: https://www.strava.com/heatmap [Accessed 22 Jun 2021].

van Strien, D., Beelen, K., Ardanuy, M., Hosseini, K., McGillivray, B., and Colavizza, G., 2020. Assessing the Impact of OCR Quality on Downstream NLP Tasks: *In*: A. Rocha, L. Steels, and J. van den Herik, eds. *Proceedings of the 12th International Conference on Agents and Artificial Intelligence*. Presented at the Special Session on Artificial Intelligence and Digital Heritage: Challenges and Opportunities, 22-24 February 2020 Valletta Malta. SCITEPRESS Science and Technology Publications, 484–496.

Suominen, O., 2019. Annif: DIY automated subject indexing using multiple algorithms. *LIBER Quarterly*, 29 (1), 1–25.

Swann, T., 2018. Towards an anarchist cybernetics: Stafford Beer, self-organisation and radical social movements. *Ephemera*, 18 (3), 427–456.

Talib, R., Kashif, M., Ayesha, S., and Fatima, F., 2016. Text Mining: Techniques, Applications and Issues. *International Journal of Advanced Computer Science and Applications*, 7 (11).

The Commonwealth Scientific and Industrial Research Organisation, 2021. Teaming up for turtles. *CSIRO Media Releases and Statements*. Available from: https://www.csiro.au/en/news/news-releases/2021/teaming-up-for-turtles [Accessed 22 June 2021].

The Laughing Room (November 16-18, 2018), 2018. *metaLAB Harvard*. Available from: https://metalabharvard.github.io/projects/laughingroom/ [Accessed 22 Jun 2021].

The Parliament of the Commonwealth of Australia, 2020. *Data Availability and Transparency Bill 2020*.

Thorvaldsen, G., Pujadas-Mora, J.M., Andersen, T., Eikvil, L., Llados, J., Fornes, A., and Cabre, A., 2015. A Tale of Two Transcriptions: Machine-Assisted Transcription of Historical Sources. *Historical Life Course Studies*, 2, 1–19.

Too, Y., 2010. *The Idea of the Library in the Ancient World*. Oxford University Press.

Transkribus: Where AI meets historical documents [online]. 2021. REAS-COOP SCE. Available from: https://readcoop.eu/transkribus/ [Accessed 28 May 2021]

Tully, J., 2011. *The Devils Milk: A Social History of Rubber*. New York: NYU Press.

Turner, J., 2019. *Robot Rules: Regulating Artificial Intelligence*. Cham, Switzerland: Springer International Publishing.

Uttarwar, M.L., Kumar, A., and Chong, P.H.J., 2017. BeaLib: A Beacon Enabled Smart Library System. *Wireless Sensor Network*, 9 (8), 302–310.

Vincent, J., 2018. How three french students used borrowed code to put the first AI portrait in Christie's. *The Verge*.

Waldrop, M.M., 2001. *The Dream Machine: J.C.R. Licklider and the Revolution That Made Computing Personal*. New York: Viking.

Walsh, T., Levy, N., Bell, G., Elliott, A., Maclaurin, J., Mareels, I., and Wood, F. 2019. *The effective and ethical development of artificial intelligence: an opportunity to improve our wellbeing*. Melbourne: Australian Council of Learned Academia (ACOLA). Available from: https://acola.org/hs4-artificial-intelligence-australia/

Watson, P.D., 1994. Founding Mothers: The Contribution of Women's Organizations to Public Library Development in the United States. *The Library Quarterly: Information, Community, Policy*, 64 (3), 233–269.

Web Accessibility Initiative, 2018. Web Content Accessibility Guidelines (WCAG2): Overview [online]. Available from: https://www.w3.org/WAI/standards-guidelines/wcag/ [Accessed 22 Jun 2021].

Webb, K., 1991. Twenty Years of Systems in the National Library of Australia. *Australian Academic & Research Libraries*, 22 (4), 142–150.

Weiner Grotta, S., 2012. Is the Keyboard Going the Way of the Typewriter? Transcript of interview. Available from: https://spectrum.ieee.org/podcast/geek-life/tools-toys/is-the-keyboard-going-the-way-of-the-typewriter.

Wetjen, F. and Brygfjeld, S.A., 2018. *Machine learning for production of Dewey Decimal.* Presented at the 84th International Federation of Library Associations and Institutions World Library and Information Congress, 24-30 August 2018 Kuala Lumpur, Malaysia.

Wheatley, A. and Hervieux, S., 2020. Artificial intelligence in academic libraries: An environmental scan. *Information Services & Use*, 39 (4), 347–356.

Whittaker, M., Alper, M., Bennett, C., Hendren, S., Kaziunas, L., Mills, M., Ringel Morris, M., Rogers, E., Salas, M., and Myers West, S., 2019. *Disability, Bias and AI*. New York: AI Now Institute.

Wiener, N., 1948. *Cybernetics or Control and Communication in the Animal and the Machine*. 1st ed. Paris: Hermann & Cie.

Wiener, N., 1950. *The Human Use of Human Beings: Cybernetics and Society*. Boston: Houghton Mifflin.

Wilkinson, T. and Brun, A., 2015. Visualizing Document Image Collections Using Image-Based Word Clouds. *In*: G. Bebis, R. Boyle, B. Parvin, D. Koracin, I. Pavlidis, R. Feris, T. McGraw, M. Elendt, R. Kopper, E. Ragan, Z. Ye, and G. Weber, eds. *Advances in Visual Computing*. Cham: Springer International Publishing, 297–306.

Williamson, R.C., Raghnaill, M.N., Douglas, K., and Sanchez, D., 2015. *Technology and Australia's future: New technologies and their role in Australia's security, cultural, democratic, social and economic systems*. Melbourne: Australian Council of Learned Academies.

Wischmeyer, T. and Rademacher, T., eds., 2020. *Regulating Artificial Intelligence*. Cham, Switzerland: Springer International Publishing.

Wittmann, R., Neatrour, A., Cummings, R., and Myntti, J., 2019. From Digital Library to Open Datasets: Embracing a 'Collections as Data' Framework. *Information Technology and Libraries*, 38 (4), 49–61.

Wong, J.I., 2016. The story of gutta percha, the humble tropical tree sap, which laid the foundation for telecommunications. *Quartz 5 Oct.*

Wood, B.A. and Evans, D., 2018. Librarians' Perceptions of Artificial Intelligence and Its Potential Impact on the Profession. *Computers in Libraries*, 38 (1), 26–30.

Writbeat, n.d. Norwegian speech to text at the National Library of Norway. *Schjønhaug AS*. Available from https://en.schjonhaug.com/projects/the-national-library-of-norway [Accessed 22 June 2021].

Yotka, S., 2021. Kate Moss Gets Into the NFT Art Market With Three New Video Works. *Vogue 9 April*.

Zhu, R., Zhao, K., Yang, H., Lin, W., Zhou, C., Ai, B., Li, Y., and Zhou, J., 2019. AliGraph: A Comprehensive Graph Neural Network Platform. *arXiv:1902.08730*.

For further information on this project, please contact The School of Cybernetics in the College of Engineering and Computer Science at the Australian National University.

The Australian National University
Level 3, Building 35, 35 Science Rd
Acton ACT 2601
02 6125 8121
https://3ainstitute.org/
director.3ainstitute@anu.edu.au
@ANUcybernetics

Who we are

In 2021, ANU established the School of Cybernetics, based within the College of Engineering and Computer Science. Our mission is to reappraise and refit cybernetics for the 21st century – to connect people with diverse perspectives, equip and encourage them to address tomorrow's challenges and ensure a more hopeful future for Australia and the world.

Our beginnings

In 2017, the 3A Institute was created as the first of the innovation institutes at ANU, intended as structures adjacent to established research and teaching programs to explore new education models and pathways towards the application of research.

The School of Cybernetics is now the home of the 3A Institute. It is building upon the 3A Institute's foundational work and carrying forward its mission to create a new branch of engineering and its mantra of keeping humanity in technology.

Contacts

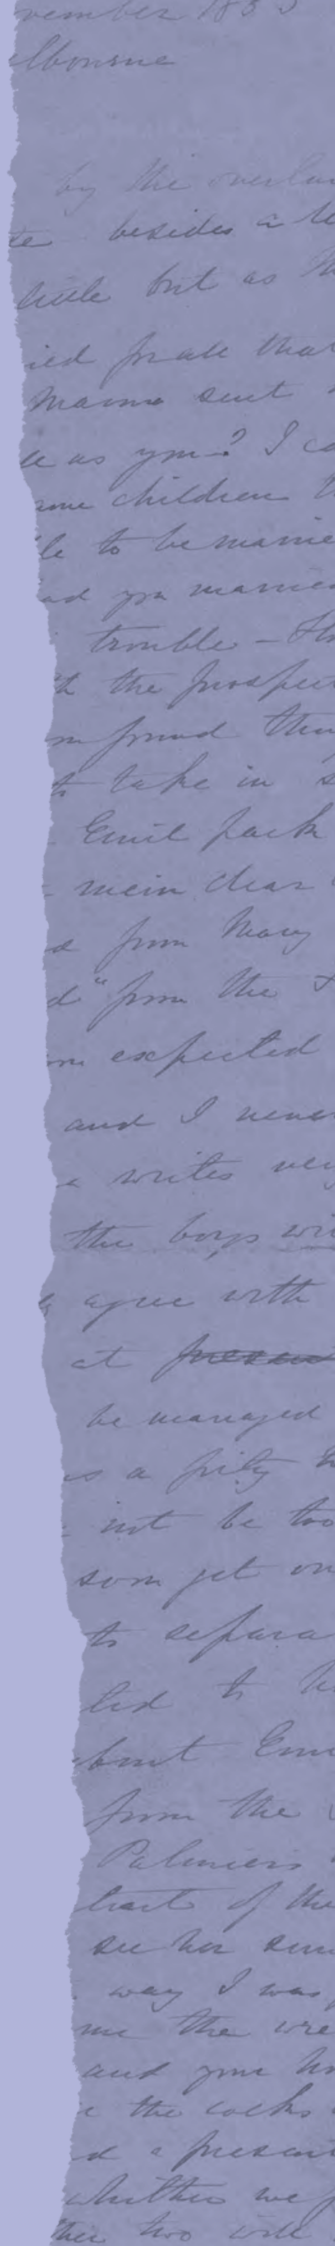

School of Cybernetics
The Australian National University
Canberra ACT 2601 Australia
www.anu.edu.au