

Beyond Time and Space: using AI to Solve Client Service Challenges Now and Into the Future

Rory L.L. Sie & Debra Nolan

University of Wollongong Australia

Northfields Avenue

Wollongong NSW 2522

Abstract

One of the biggest challenges for libraries today is fulfilling client service requests 24/7. We have skilled staff in the Library during the day and evening, and LiveChat for those not able to visit in person, but what happens after hours? Who is there to help those who work shifts, have children, live on the other side of the planet or just prefer studying at 2am?

Artificial intelligence (AI) promises an affordable solution, but will it provide the quality, flexibility and authenticity needed to ensure client satisfaction? Importantly, can the accuracy and veracity of responses be assured? At the University of Wollongong, the Library has teamed up with Academic AI specialists, Engineering and Information Science Faculty and a team of exceptionally gifted students to solve this problem.

Using a combination of Agile Scrum and design thinking methodology to ensure client satisfaction and service needs were met, the team commenced building a

chatbot designed to converse, provide assistance, and refer clients when needed – in other words: to create, rather than recruit, a new Library staff member.

This paper will explore the journey, noting the challenges, breakthroughs and methodologies used to create and evaluate the chatbot.

Introduction

The University of Wollongong (UOW) is a young and vibrant university with global reach. The UOW has several regional, metropolitan and international campus locations as well as active partnerships with other educational institutions across the globe. The size of the University enables the ability to move relatively rapidly when testing new ideas of implementing change, however, it also brings its own challenges, especially in relation to budget and staff capacity to develop and sustain a new service provision.

The UOW Library has a reputation of strong collaboration and innovation within the University. A future-ready strategy, recently written by the Director Library Services and driven by the Library's executive team, embraces change and demands agility and flexibility, building on a culture that welcomes innovation and is supportive of risk.

Current strategy includes an "online first" mindset that recognises the need to move into the online world beyond webpages, and to explore new technologies such as artificial intelligence (AI) to facilitate student agency regardless of location or time.

This article lays out how a team of exceptional Computer Science students, librarians, and AI experts led the successful delivery of a first version of the Library

chatbot called *Moodji*, that automatically answers student questions on a broad range of topics.

We start off with the background into our current service model and the importance of the introduction of a chatbot. We give a basic explanation of what AI is, and particularly, why it has gained attention in the past few years, then we share how we went about creating Moodji, given a diverse team from distinct work cultures with diverse work methodologies.

Moving to embrace new technology to support 24/7 access

The challenges faced when developing 24/7 interactive library assistance are complex due to factors such as multiple time zones, language barriers and varied location-specific information. The Library explored various options such as introducing shift work and joining global university library alliances, but the solutions were either too expensive or not sustainable due to the size of our institution.

In recent years, AI has shed new light on alternate possibilities. International projects such as the Jill Watson virtual assistant (Goel & Polepeddi 2016) have successfully expanded Georgia Tech's ability to provide quality support to students anywhere, anytime. While interest in artificial intelligence solutions at UOW Library was strong, development remained stagnant due to lack of expert knowledge within the Library. It was not until AI expertise was recruited to the Learning, Teaching & Curriculum unit at UOW that the Library was able to realise the potential of developing a Library chatbot or conversational agent. This opportunity was very attractive to the Library, especially given our quest to provide 24/7 enquiry services and having already identified that 32% of queries to existing systems were recorded outside of our

LiveChat's operational hours. When a subsequent literature scan failed to uncover any similar projects within Australian academic libraries, we were reasonably sure we would be cutting new ground nationally.

The Library's LiveChat service works within a three-tiered service model (Figure 1). Tier 0 is unmediated, with routine reference questions, Tier 1 is pathway questions that require action from Library staff, and Tier 2 refers complex questions to Reference or Specialist Librarians. The LiveChat service is provided by the SpringShare system (Figure 2), and operates during the hours of 9am to 5pm AEST. The service is backed up by an extensive knowledge base with basic keyword search capability that students can use to search for answers 24/7. It was the combination of this established knowledge base and LiveChat service that attracted the attention of the new AI expert and provided the catalyst for a series of discussions which eventually led to the collaborative project to create *Moodji*, the Library chatbot.

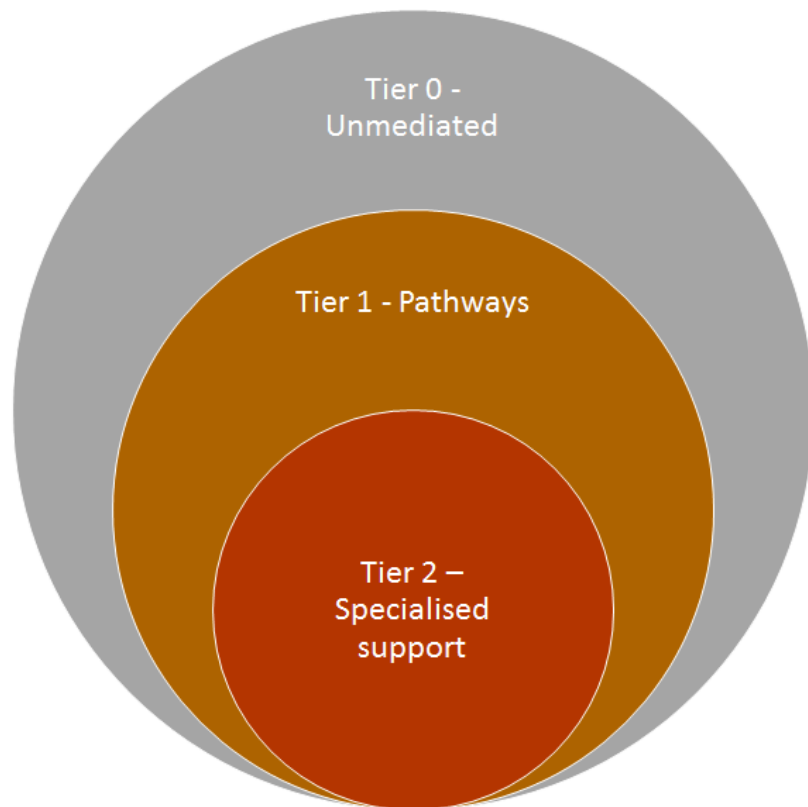


Figure 1: Three-tiered service model

The Library intended to augment the existing LiveChat service rather than replace it, by extending hours of chat operation and possibly taking LiveChat overflow at busy times. Just as the current LiveChat service refers complex questions, the chatbot would be designed to refer complex queries to Reference or Specialist staff and extend the reach of the triage service in a cost-effective way.




Figure 2: Current LiveChat service

Artificial intelligence

Artificial intelligence (AI) is a field of research that enables computers to mimic or even outperform human intelligence. It spans from software that captures knowledge in a structured way (Semantic Web), to intelligent agents that act autonomously and communicate among themselves. AI has recently drawn the attention of the public through vehicles that drive autonomously.

The recent advancements in the space of autonomous vehicles (Fagnant & Kockelman, 2015) are the result of two main developments: the amount of data that can be collected and stored, and recent successes in computer vision. Computer vision, as the name suggests, revolves around computers detecting objects from images or video. They do so by processing an image (the input) to detect distinct

objects, followed by a process to determine what that object is (the classification) (Figure 3). Nowadays, AI algorithms can tap into huge volumes of manually labelled images that help the computer learn how to correctly classify objects and thus mimic human vision.



Figure 3: A simplified example of how a car detects objects from image or video. The original photo was taken by Herry Lawford. (<https://www.flickr.com/photos/herry/>).

Some of the underlying AI techniques that are applied to computer vision can also be used to understand language, and this is where it becomes interesting for library professionals. Librarians often receive questions as to where to find certain books, or where to borrow a laptop. As the customer base grows, so do the number of questions, and often these (Tier 0) questions have routine answers. Similar to the computer vision example, the computer can learn the associations between questions (input) and answers (classification) from past questions and answers.

Chatbots can be built in various ways (from simple to complex): a chatbot can compute how similar a given question and a question in the database are (semantic similarity; Li, McLean, Bandar, O'Shea, & Crockett, 2006) and responds with the answer associated with the question in the database; it can learn what the underlying intent of the question is (intent recognition; Holtgraves, 2008), which can be mapped to a broader range of responses based on the context; and it can learn how a good response is generated by example, which allows it to generate its own completely new answer (generative model; Vinyals, & Le, 2015). Obviously, a generative model is preferred, but as complexity grows, so does the need for data. *Moodji* would follow this path of increasing complexity, and as the development of *Moodji* matured, so would the knowledge base it could tap into, allowing for more complex and more accurate responses, similar to the product-users-data-product virtuous circle proposed by Andrew Ng, the most well-known expert in AI (Ng, 2017).

How *Moodji*, the Library chatbot, came to life

Several factors were explored by the project team prior to developing the chatbot. These included whether we would try to emulate a human experience or to let our clients know from the outset that they were communicating with a machine, whether the chatbot was male or female, and if it should have a name. Feedback from students suggested the chatbot should be identified as such and should not try to emulate a person, but may take the physical look of a robot, and the name should be gender neutral.

Choosing a name for the chatbot quickly became important when the UOW Media unit requested information for a potential news article. The project group and several

creative minds within the Library considered a number of gender neutral names, but it was the idea of using an Indigenous word from a local language group that seemed to fit best. “Moodji” means “friend” in the Dharawal language. The word sounded contemporary, while being derived from an ancient language steeped in local tradition and custom.

The most interesting aspect of naming the chatbot was an almost magical transformation of what was a “thing” into the prospect of an entity that may one day be able to reason for itself. The name also helped Library staff better understand that AI could be an ally rather than a threat – almost as if we were training a new staff member. This line of reasoning particularly helped when the chatbot was found to be lacking data to learn from. The project team asked their colleagues at the Library to contribute to *Moodji*’s knowledge: “What would a new staff member need to know about the services we offer as the Library?”

User interface and functionality

Moodji tapped into a knowledge base of 300+ questions provided by Library staff members. When a chat session started, *Moodji* offered a greeting by asking how you were (Figure 4). After the user responded, it showed it was ready to answer the question. All of *Moodji*’s responses could be played through a text-to-speech button and *Moodji*’s speech-to-text module ensured that it understood speech: for example, when microphone access was enabled, the spoken keyword ‘Duck’ ensured *Moodji* listened to your verbal question.

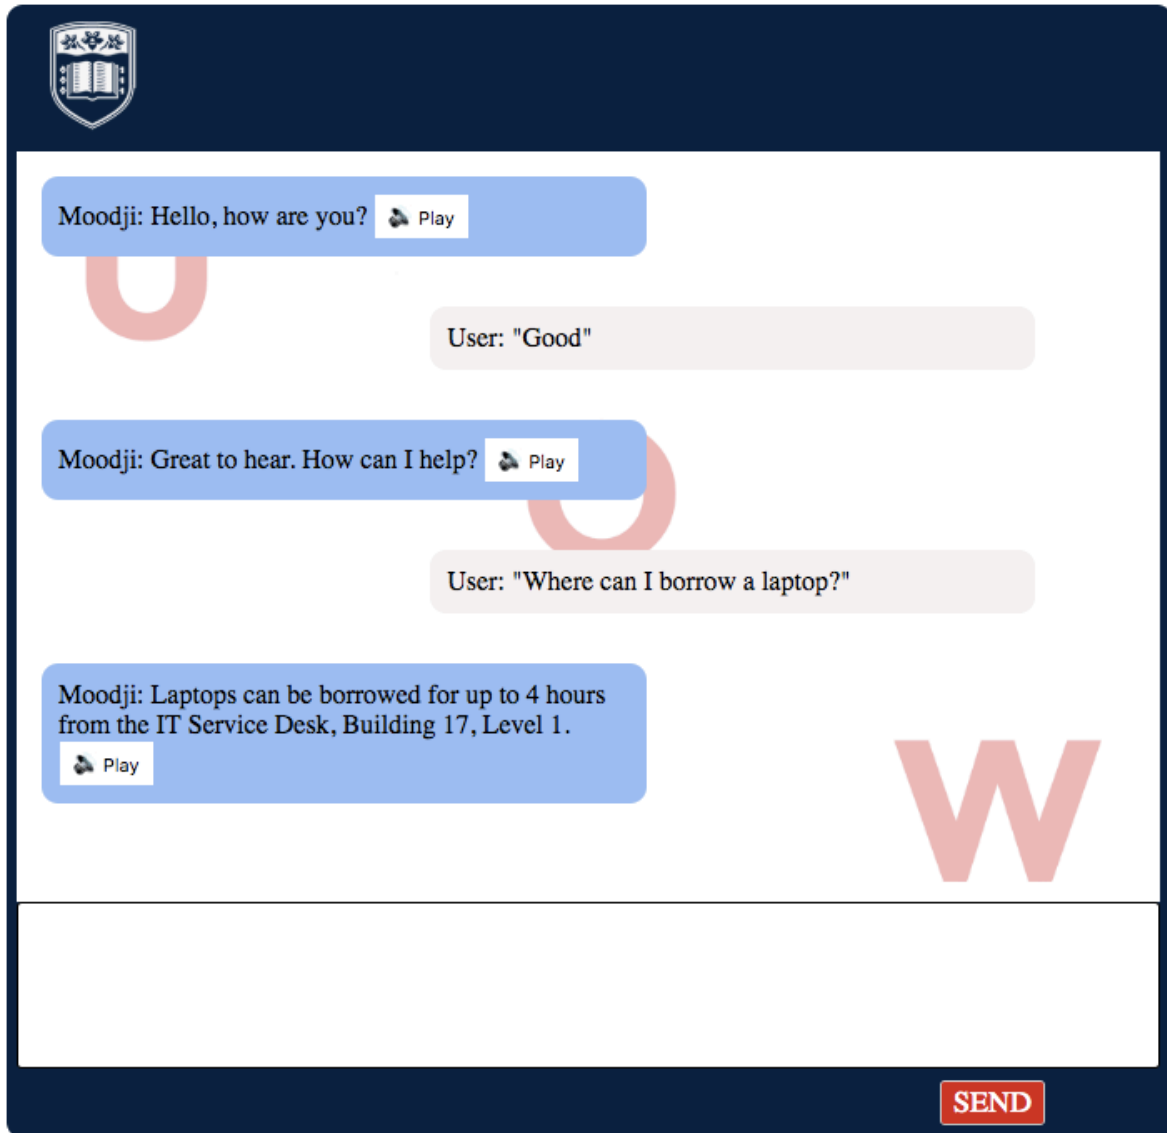


Figure 4: Screenshot of *Moodji*'s current user interface.

That version of *Moodji* was purely text-matching from the knowledge base and still far from production ready. From an AI perspective, 300 questions was not sufficient for *Moodji* to interact with the user and create accurate responses, let alone generate entirely new responses by itself (generative model). The extent of data needed was something that would need to be tested by developing predictive models.

Nevertheless, that early stage system was demonstrated during an internal Library staff conference when the project was three months old and generated lots of

questions and enthusiasm. Importantly, it provided the project team with buy-in from the Library staff and the opportunity to clarify what was needed for *Moodji* to learn. Subsequently, a system was built to register new questions and answers, especially in the context of what might be relevant to “a new staff member”.

This allowed *Moodji* to quickly develop into a minimum viable product (MVP) that users could interact with, which in turn generated more data and thus a better product, along the lines of Ng’s virtuous circle of AI (Ng, 2017).

Lessons learnt: work methodologies

The challenge we faced during this project was how to merge the diverse methodologies that the Library and software development employ. Although the Library and software development teams showed some similarities in the structured way they organised and presented information (the *processes*), they seemed to differ greatly in collaborating with their customers (the *people*), and how new *products* were developed. Nevertheless, the UOW Library has a proven openness to innovation with their focus on learning analytics (Jantti & Heath, 2016), a MakerSpace (<https://uow.libguides.com/uowmakerspace>) and adoption of modern work methodologies such as design thinking.

The Library way of working: design thinking

Design thinking methodology was adopted by the UOW Library in 2016 and is one of the main tools used when developing and evaluating services. Design thinking was chosen because of its human-centred design that combines creative and analytical mindsets and methods to empathise with people. Design thinking enables the Library

to design products and systems that meet client needs. UOW Library uses two models: the Stanford model (<https://dschool.stanford.edu/>) and the IDEO model (<https://www.ideo.com/>).

The human-centred approach enables the client's voice to be heard throughout the process, ensuring innovations or solutions are firmly based on the client's needs. The ideation process suspends the urge to simply solve, or Band-Aid, issues and enables creativity to flourish. New ideas or prototypes are tested prior to and after implementation to ensure continuous improvement is achieved (Figure 5).




Figure 5: Design thinking process, graphic by d.school, Hasso Plattner Institute of Design at Stanford.

Design thinking techniques can be used by individuals or by teams. The crucial element is ensuring a high level of client interaction throughout the process.

The software developer way of working: Agile Scrum

Agile Scrum (Scrum Alliance, 2018) is a well-established work methodology to develop software in teams. It revolves around stakeholder commitment and agility, or flexibility. Products are developed in iterations of two to four weeks, called "sprints". Each sprint starts with a task breakdown and estimation of effort (sprint planning). A

scrum meeting is held every day to determine what has been done and what is left to be done, and to assess potential hurdles to successful software development. At the end of the sprint, a sprint demo is held to showcase the product to the customer, who can decide on new functionality or determine new priorities for the project. However, during a sprint all tasks are “frozen”, which means they cannot be changed. This provides the development team with more reliable insight into their progress and relieves pressure for dealing with ad hoc tasks. This caters for transparency and flexibility that is valued by customers.

Agile Scrum results in developing “shippable products” that are ready to be used and reflected on by the customer within two to four weeks of iterations (sprints). The first critical weeks in a project are when the first shippable product is to be built, which is why functionality is very limited at first. As the project progresses, new functionality is added to an already working product, which ensures there is always a “production-ready” product at all stages.

The individuals needed to utilise Scrum depend on the task at hand. Usually, there are three main roles: the product owner (who represents the customer), the scrum master, and the team member. However, these are not distinct roles since multiple roles can be fulfilled by one person.

Combining work methodologies

The two methodologies did not seem too different in terms of processes: the client is continually involved in the process, which builds commitment from either side. The only difference, perhaps, was that in Agile Scrum, the development requirements are

frozen during a sprint: no new tasks can be added, nor can they be changed. Given the highly experimental type of project at hand, we employed a milder approach: the software development team aimed for concrete tasks, but the software development team (mainly students) was inexperienced in the field of natural language processing which was the underlying technology. A flexible mindset was required to cope with any skill deficiencies in the team, changing demands or ideas from the Library, or new insights that derived from working with different technologies.

Moodji's incremental development did not differ among the two work methodologies. The Library was always interested and available to demo new products and to offer feedback to the software development team. The Library was also very supportive of new ideas generated by the students, such as speech-to-text and text-to-speech functionalities, and an accessible mobile interface, which sparked the idea of creating a cardboard cut-out of a robot with a tablet in it, that students could talk to.

Developing the Product

We worked with a group of students who were available 10-12 hours each week, which entailed a slight change in the frequency of Scrum meetings. These were held weekly, with sprints taking up to six weeks. The core team consisted of three Librarians to represent the customer, two AI specialists (academics) to provide guidance and consultancy to the students, a DevOps (student) to develop software while maintaining the server, a front-end developer (student), back-end developer (student), user interface designer (student), and a technical writer (student). Regular discussions were held with the IT department within the University, as well as other potential customers such as the Student Services Division. Rather than making

Moodji a bespoke, Library-only project, discussions with potential customers ensured an underlying infrastructure that could cater for new types of questions and also conduct service-based requests such as “book a study room”.

Challenges, recommendations and future work

The current version of *Moodji* was funded largely by in-kind efforts from the Librarians and AI specialists. The team’s students worked on *Moodji* as part of their third-year capstone project across two academic sessions. But what would happen when the students finish their capstone project? This was a key concern occupying the team. The Library and the Learning, Teaching & Curriculum unit currently lack the skills to work with the specific back-end technologies (Python/Django) and infrastructure (Amazon Web Services server, known as AWS). This is also a new area for the University’s IT department to explore, making it more difficult to take the software into production within the University’s current ecosystem.

The cross-institution collaboration that made this project so unique also proved to be its Achilles heel. The AWS server was expected to cost \$2000 to \$3000 on a yearly basis, and it remains unclear as to whether it should be funded by the Library, the Learning, Teaching & Curriculum unit, the IT department or the University. Apart from this, capable DevOps would be fundamental to keeping *Moodji* up and running. Our recommendation would be that if one were to start this project at another university, one would need to think about the maintenance costs after the project ends. In our case, a bare minimum to keep *Moodji* up and running for another year would be approximately \$25,000 (AWS and DevOps). However, it is assumed that *Moodji* is far from production-ready by the end of 2018. An additional investment

would be necessary to increase its functionality and accuracy, which would need the following team composition, based on our experience (the FTEs are an approximation):

- *DevOps engineer (0.6 FTE)*: a DevOps engineer to set up the server infrastructure is an often-forgotten prerequisite, but crucial to the initial phase of the project. Each hour the DevOps spends in getting the server architecture correct (e.g. server instances, storage, security, load balancing, containerisation and orchestration) in the start-up phase, will pay off when deploying the product and maintaining the server architecture.
- *Back-end developer (1 FTE)*: back-end developers are crucial in their role to store and retrieve data from the database. They also typically release data from the database in a secure manner through an Application Protocol Interface (API) with authentication and authorisation.
- *Front-end developer/UX designer (0.4 FTE)*: in the initial phase, the front-end developer usually takes on multiple roles, such as designing and developing the user interface and connecting the user interface to the back-end (via the API).
- *Natural language processing (NLP)/AI engineer (1 FTE)*: the chatbot needs to understand what its users are asking. An NLP engineer is a combination of someone who is knowledgeable on how to train a natural language processing model to recognise the user's intents, and at the same time knows how to incorporate this into the overall architecture of the chatbot (to "engineer" it).

- Chatbot “trainer” (0.4 FTE): the chatbot needs a domain expert that can register questions and answers, and the underlying intents and entities that are specific to the university.

Some of the roles, such as the front-end developer and UX designer, could be split when the project expands its scope. A separate role would be necessary for mobile development.

Last but not least, data was a vital requirement in the development of a chatbot. The first versions of *Moodji* were only required to answer routine questions, but what if we were to take *Moodji* to a more active service-based tier? For instance, what if *Moodji* could book a study room, or borrow a book on behalf of the student or staff member? This would involve identity verification of the user against the University’s databases, which would involve asking users for their student ID, triggering questions regarding ethics, privacy and security. Identity verification is only achievable if access to the University database is granted.

Another phase of skill development for *Moodji* would be to recognise its user, and adapt its responses to the user’s specific circumstances and profile. Google already has a long history of adapting its search results to its users based on recent search activity or, for example, YouTube videos viewed. Similarly, *Moodji* could adapt its service to the individual user, making the user experience more personal, more nuanced. For instance, it could detect that a student is currently enrolled in a cybersecurity subject and if that student asks “what is a hash?”, it could distinguish the term “hash” as being a way to encrypt data from the way a topic on Twitter is defined by the # sign.

Evaluation: customer/client journey mapping (CJM)

Once the minimal viable product version of *Moodji* is complete, the Library intends to evaluate student interaction with *Moodji* by employing the client journey mapping (CJM) technique. This technique has recently been chosen by the Library because of its ability to identify and interpret authentic client experiences, adopt the user perspective, to adjust the environment to meet user expectations, and provide a better overall experience (Marquez, Downey & Clement 2015). The technique has been used successfully in a range of library and commercial situations over the past several years.

In their 2013 article on the implementation of CJM at Birmingham University Library, Andrews and Eade described CJM methodology as:

“plotting a process or service by using business process mapping symbols within swimlanes to produce a visual representation of a transaction from the point at which the customer accesses a service to the goal the customer is aiming to reach” (Andrews & Eade 2013).

In the case of the Library chatbot, emojis will replace business process mapping symbols, with swim lanes becoming touch points highlighting the student’s journey or experience. This method will provide a tool to gain in-depth understanding of the client’s perspective before, during, and after their interaction with a service or product. For example, in the case of the Library chatbot, students will be interviewed using open-ended questions regarding why they chose to interact with the chatbot; what happened during their interaction and what they gained from the experience. Unlike many survey or focus group methods, the interview will focus on enabling the student to tell their story rather than prompting them with direct questioning, which

will provide the Library with: an authentic and unbiased understanding of student motivation and expectation (before); their level of satisfaction with the interaction (during); and the value, if any, taken from the interaction (after).

Conclusion

This article reported on our experiences building a conversational agent, *Moodji*, for the University of Wollongong Library, with the potential of rolling out this service across the University's services, as a single triage system. We demonstrated the first version of *Moodji* and shared some of the important considerations for universities when commencing the development of a chatbot as:

- being future-ready and ready to take advantage of technological opportunities. This involves collecting and releasing current data, and catering for a technological infrastructure (servers, storage, computation capacity) to train a chatbot
- a willingness and support to take risks
- being open to collaboration and embracing new ideas
- a willingness to drive change regardless of the challenges, pushing through adversity, and not holding back and waiting for the ideal time or funds etc.
- an awareness of the maintenance costs involved after a chatbot, or any software in general, is deployed.

This highlighted the importance of combining an open, innovative mindset with technical expertise while developing connections that enable creative solutions.

Moreover, our continuing challenge will be to create a common ground between the Library and its technological partners, thereby developing a collective vision and commitment to grow *Moodji's* abilities into future.

References

- Andrews, J, & Eade, E 2013, Listening to Students: Customer Journey Mapping at Birmingham City University Library and Learning Resources, *New Review Of Academic Librarianship*, 19, 2, p. 161-177.
- Fagnant, D. J., & Kockelman, K. (2015). Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice*, 77, 167–181. <https://doi.org/10.1016/j.tra.2015.04.003>
- Goel A. K. and Polepeddi L, Jill Watson: A virtual teaching assistant for online education. School of Interactive Computing Technical Reports, Georgia Institute of Technology, 2016.
- Holtgraves, T. (2008). Automatic intention recognition in conversation processing. *Journal of Memory and Language*, 58(3), 627–645. <https://doi.org/10.1016/j.jml.2007.06.001>
- Johnson, C. (2015). New Toolkit Brings Patron Centered Design to Libraries. Retrieved August 2, 2018, from <https://www.shareable.net/blog/new-toolkit-brings-patron-centered-design-to-libraries>
- Li, Y., McLean, D., Bandar, Z. A., O’Shea, J. D., & Crockett, K. (2006). Sentence similarity based on semantic nets and corpus statistics. *IEEE Transactions on Knowledge and Data Engineering*, 18(8), 1138–1150. <https://doi.org/10.1109/TKDE.2006.130>
- Marquez Joe, J., Downey, A. & Clement, R. (2015) Walking a Mile in the User's Shoes: Customer Journey Mapping as a Method to Understanding the User Experience. *Internet Reference Services Quarterly*, 20:3-4, 135-150, doi: 10.1080/10875301.2015.1107000
- Ng, A. (2017). Artificial Intelligence is the New Electricity. Retrieved July 28, 2018, from <https://medium.com/syncedreview/artificial-intelligence-is-the-new-electricity-andrew-ng-cc132ea6264>

Serban, I. V., Sordani, A., Bengio, Y., Courville, A. C., & Pineau, J. (2016). Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. *Proceedings of the Thirtieth {AAAI} Conference on Artificial Intelligence*, 3776–3784.

Scrum Alliance. (2018). Learn About Scrum. Retrieved August 3, 2018, from <https://www.scrumalliance.org/learn-about-scrum>

Vinyals, O., & Le, Q. (2015). A Neural Conversational Model. In *Proceedings of the 31st International Conference on Machine Learning*(Vol. 37). https://doi.org/10.1007/978-3-319-19291-8_22