For Your Enrichment

Integrating Technology and Human Interaction with a Library Reference Chatbot

Robin Fowler and Kelly Handy

For Your Enrichment is an occasional column that offers a platform for pieces that are of interest to the work of librarians, but that might not otherwise fit in one of the traditional *RUSQ* column areas. I was intrigued here by the concept of maintaining the human connection while utilizing AI.

Based on the Florida Association of College and Research Libraries (FACRL) Poster Presentation, "'Are You a Bot?' Implementing a Chatbot While Keeping the Human Connection in Virtual Reference," October 25, 2024, and the Professional Development Alliance of Library Consortia online seminar, "'Are You a Bot?' Merging Technology and Human Interaction Through the Implementation of a Library Reference Chatbot," (invited) March 27, 2025.—*Editor*

Introduction

Virtual chat reference has been a point of service within the George A. Smathers Libraries at the University of Florida (UF) since 2000,[1] adopted then as an emerging technology. Even after joining the state of Florida Ask a Librarian (AAL) cooperative hosted by the Tampa Bay Library Consortium (TBLC) in 2006 and expanding the coverage provided to the UF community and other state libraries in Florida, live chat remained an underutilized service for several years. With the onset of the global pandemic in early 2020, when lockdown sent everyone home and online for work, school, and research, the virtual reference service became the primary service point in the library. The staffing of the virtual reference desk had to be tripled to keep up with the demands of the university community's needs. Following lockdown and return to campus, virtual reference at the Smathers Libraries maintained a steady flow of business, requiring staffing of the virtual desk to remain at pandemic levels. And because "virtual reference, especially chat, is highly appealing to patrons for its convenience and immediacy,"[2] the virtual reference desk had evolved into a key service point in the Smathers Libraries. As the libraries returned to some sense of normalcy post-pandemic, addressing ways to further improve the virtual reference model in the Smathers Libraries was logical.

Through the implementation of a rule-based chatbot, the Smathers Libraries would be able to continue offering a high standard of reference service[3]; the chatbot would be able to interact directly with patrons, "providing instant access to information about books, journals, and other resources available at the library,"[4] as well as provide answers to directional and circulation queries.

Robin Fowler, Physical Sciences and Engineering Librarian, University of Florida, email: rfowler73@ufl.edu, ORCID ID: 0009-0001-2042-0577; and **Kelly Handy**, Assistant User Services Manager and Public Services Associate, University of Florida, email: kl.handy@ufl.edu, ORCID ID: 0009-0007-2429-1278. Correspondence concerning this column should be addressed to Barry Trott, RUSQ Editor, email: barry.trott1@gmail.com.

Patrons would be able to receive help from reference services in the mode that makes them the most comfortable: in person, via live chat, or by clicking through the options provided by the chatbot.

Background and Objectives

In August 2023, Springshare, the platform that hosts UF's chat service and many other library services, announced the arrival of its rule-based chatbot product. Though the usage volume of the UF AAL service had begun to taper off slightly in 2023 (still higher than prepandemic usage), it was brisk enough to consider incorporating the Springshare chatbot product into the existing service. The potential benefits of a chatbot are clear; the bot could intercept more commonly asked reference queries and answer them, allowing patrons with more complex reference questions to spend time with online librarians getting help. The AAL service could be available to the UF community twenty-four hours a day, seven days a week, without having to physically staff the virtual desk, a difficult undertaking during later hours when fewer library workers are available. The benefits to the broad range of patrons served are also numerous; students often view live chat reference favorably as a more personalized and often conversational experience. To achieve this goal, according to Mawhinney, "service providers can personalize interactions by using their first names to identify themselves during interactions and/or provide their contact information at the end in case users have follow-up questions."[5] Further, certain students who feel library anxiety, or a hesitancy to approach a physical service desk out of fear of feeling inadequate or seeming unfamiliar with the library, benefit from the anonymity that live chat and, to more of an extent, a chatbot offers.[6] With the implementation of a rule-based chatbot, both goals could be achieved. There is, however, one drawback to a rule-based chatbot: the creators of the chat flows have to accurately predict the direction of the "conversations" the chatbot will have with patrons, which would be specific to the type of query it is being asked, with the knowledge that the chatbot will be the manager of the interaction with the patron.[7]

For the purpose of this discussion, it is important to examine the differences between a rule-based chatbot and an AI (artificial intelligence) chatbot. A rule-based chatbot operates as a flow chart, following sets of preprogrammed question strings and extracting responses from a preloaded question bank.[8] An AI chatbot uses machine learning and can understand a user's question as it is manually typed into the product. It can seem more conversational and more nuanced, providing layered responses to more complex queries. The Springshare rule-based chatbot is not interactive in the way that an AI chatbot is. Patrons cannot type their questions into the product to receive feedback. They must follow the predefined options set in the chat flows.[9] The chatbot will never stray from the predetermined language of the creators of the chat flows and will never "learn" to say anything other than what it was "taught" to say. As a fully customized product, building the bank of questions and series of answer flows for the UF chatbot was no small undertaking. Because the UF Libraries were already paying for a long-term subscription for Springshare services that included a rule-based chatbot, building an AI chatbot was never a real consideration. One of the most appealing features about the Springshare rule-based chatbot was the readily available platform from which to launch this additional service. Additionally, artificial intelligence was a bit of an undiscovered frontier in late 2023 and had not yet asserted itself into all corners of the library world. Machine learning was somewhat intimidating, and the notion that a generative AI product would potentially distribute incorrect information seemed too much of a risk to take at that time.

In January 2024, a team of experienced virtual reference operators (located at various on- and off-campus library locations) came together to create a rule-based chatbot using the Springshare

chatbot product. In the initial planning discussions of incorporating a digital assistant into the virtual reference service, one goal was principal above all others: the chatbot would not completely erase the human component from the virtual chat reference experience. The human interaction and the point-of-need availability of trained library operators remained paramount to the service. The role of the chatbot would be to make simple library queries answerable immediately for patrons and at all times of day or night. Using Springshare's chatbot tool, the implementation team was able to successfully integrate this extension of service into the existing LibChat platform through survey, analysis, chat flow creation, implementation, and continuing assessment of the tool through statistics.

Methodology

First Steps and Implementation Team Formation

On August 1, 2023, TBLC began beta testing on the chatbot feature, employing the help of several libraries that participate in the consortium. The UF AAL site coordinator chose not to take part in beta testing due to a few factors. In the beginning stages of researching the feasibility of integrating the chatbot, it appeared the best way to create the chat flows would be to use the existing Frequently Asked Question (FAQ) modules. Because UF's FAQ section had historically been open to all operators for contribution, there were many incorrect, duplicate, and outdated entries. The time commitment to reconstruct the FAQ section on top of creating chat flows would have been prohibitive, but the chatbot was not off the table for the future. The decision to finally implement the chatbot feature was made in December 2023 after attending training sessions where the AAL site coordinator learned of alternative (though no less simple) methods of implementation. The site coordinator invited four experienced virtual reference operators to collaborate on the project to create the George A. Smathers Libraries' first virtual reference chatbot.

Prior to the first Chatbot Implementation Team meeting, team members viewed applicable Springshare training videos to get a sense of the work ahead. The AAL site coordinator (and leader of the implementation team) also reached out to a few of the institutions that took part in TBLC's beta testing and received useful feedback and advice on where to begin in the process, confirming suspicions that using FAQs would work best only if they were already updated and robust. The team met on January 10, 2024, to begin the work toward creating the chat flows for the new chatbot. The first steps included formalizing the goals for the chatbot, holding preliminary discussions on creating layered chatbot flows to guide patrons to the answers they were searching for, and learning how to create the question bank from which the chatbot would extract those answers.

Survey

In discussions about potential topics to include for the chatbot, the team decided that the best source from which to glean that information would be the virtual chat operators who staff the UF Libraries virtual desk daily. Using a Google Forms survey, the team asked for the UF AAL chat operators' aid in identifying the more commonly asked and easily answered questions they were fielding during their chat shifts. The team distributed the survey to 91 Smathers Libraries employees on the AAL chat operator distribution list on January 18, 2024, and were asked to submit completed surveys within one week.

The survey consisted of the following questions:

- How many hours per week do you work on Ask a Librarian?
- What times are your shift(s)?

- What types of questions do you get asked most often? (This survey question was enabled with multiple prechosen options and included space for open-ended responses.)
- If you are comfortable with us reaching out to you to ask further questions about the chats you answer on Ask a Librarian shifts, please leave your name below.

Of the thirty-five responses received, thirty-two chat operators listed one of their most frequently asked questions during their virtual desk shifts as "Can you help me find this journal article/book/ebook?" This response, along with the other top recurring queries identified in the survey (Table 1), helped the team begin to design chat flow categories that would be most beneficial to the UF community using the chat service.

Initial Categories and Chat Flows

The team created an initial set of categories based on feedback from the surveys:

- Complaints and Concerns
- Reservation Questions
- Building Information Questions
- How Do I?
- Other

Table 1. Top Responses from Chat Operator Survey

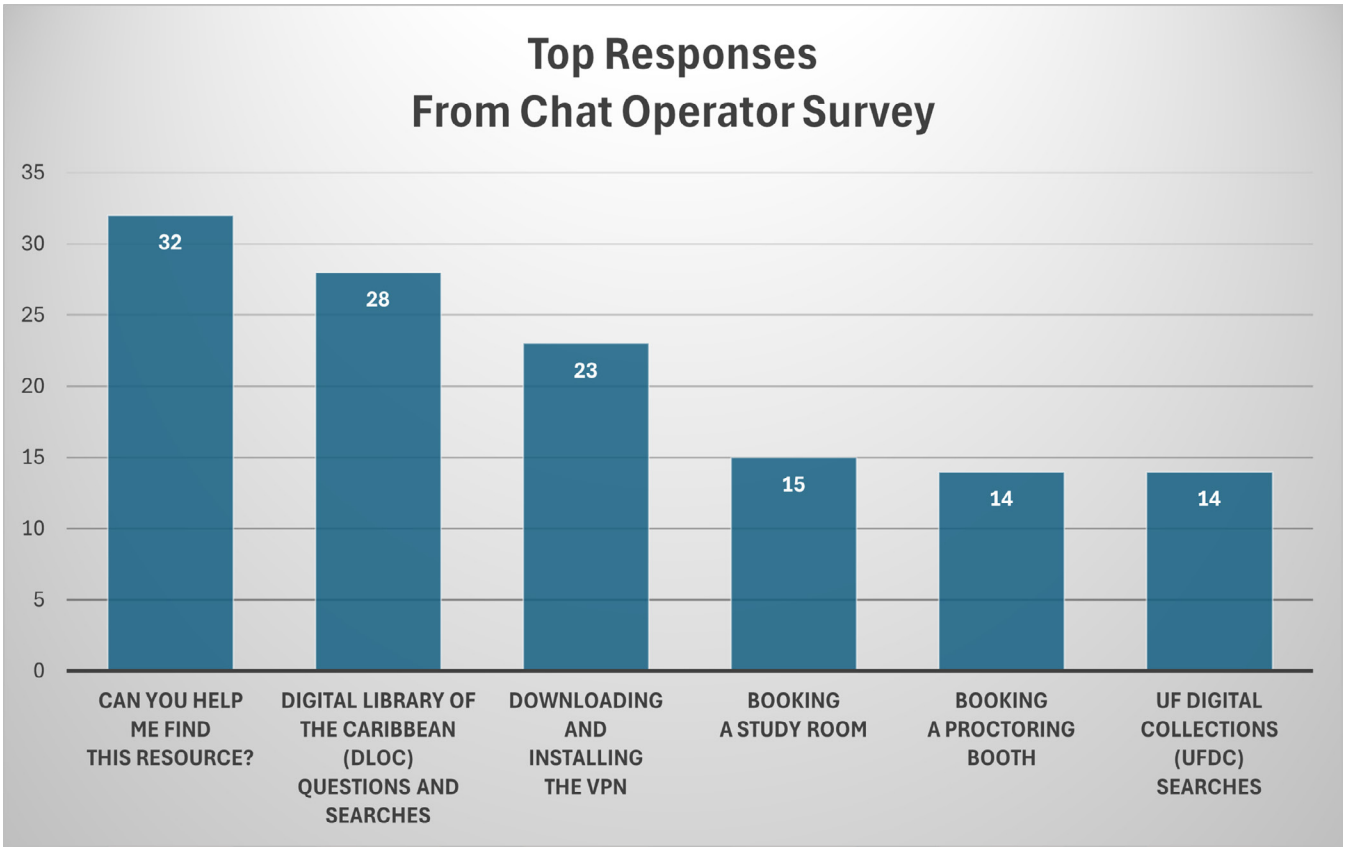

Table 1 shows the top six most commonly asked questions reported by the thirty-five Ask a Librarian chat operators who responded to the survey designed by the chat implementation team.

Through brainstorming sessions, the team determined what specific answers should be located under each category. From there, each team member selected a category and began to map out the question bank and chat flows that would live within that category. It was at this point that the team agreed that embedding the option of speaking to a live chat operator (if available) in every step of every chat flow was necessary. The team used a Microsoft Teams group to create lists of categories with commonly asked questions related to those categories (each team member had one or two categories to work on), followed by appropriate responses to those questions. Each team member could view all work being done on these early chat flows and offer feedback or corrections. During regular meetings through this phase of the project, the team reviewed each flow together and edited for correct information when necessary. Team members also began exploring the chatbot features in LibChat to gain a deeper understanding of how the chatbot platform worked and how chat flows would be linked to the question bank.

In these first stages of development and discussion, it quickly became clear that some of the chat flows were too convoluted, with too many pathways that could potentially confuse or frustrate patrons. As the team started to understand the capacity and limitations of the chatbot product and its creation process, they took a simpler approach to the chat flows to maintain the scope of the project: quick answers for the more commonly asked questions. Shortening the flows would alleviate the potential for a patron to be led down the rabbit hole, become frustrated, and leave the chatbot without receiving the help they needed or opting to transfer to a live chat operator.

Final Categories and Constructing the Chat Flows

To better visualize the chat flows, a team member constructed a flowchart of the specific chat flows to show how each topic would build off previous responses and how many levels deep to take the initial chat flow (Figure 1). Throughout the development of the final chat categories, some of the original categories were broken down even further in the interest of streamlining the length of the chat flows. The final categories were:

- Access Library Materials
- Complaints and Concerns
- Study and Testing Rooms
- Printing and 3D Printing
- Another Problem

The team included the additional category of "Another Problem" to help direct patrons to the public service desks at the library branches for urgently needed assistance. The goal was to launch a chatbot that would be usable but not confusing, and that would still offer quick answers to patrons' queries and an option to speak with on-shift live chat operators (when available) or submit a ticket that a library affiliate would address as soon as possible. Ongoing maintenance of the chat flows over time by the AAL site coordinator was also on the minds of the implementation team members; the coordinator should be able to easily make corrections, updates, and additions through monitoring of the statistics and transcripts in the interest of continuous improvement and accuracy.

Final construction of the primary chat flow and question bank occurred in early August 2024. During this session, the team focused on standardizing the language of each chat flow worked on by individual team members. This was for consistency throughout the flows and to ensure there was no "library jargon" or library acronyms being used that a patron might be confused by or not understand. The team also tried to limit the number of flows per question, striving for no more than three under each branched-off section of the main categories, when possible. To satisfy the goal

of keeping a connection between patrons and live chat operators, the option to connect to one of them (if available) was included in the answer selections for each level of the chat flow. It was also added later as an option with the final categories when patrons activate the widget for the chatbot. Further, the team included a way to return to the beginning once the patron reached the end of the existing flow if they indicated they had not received the answer they were looking for. If the patron received help and had no further questions, they would be thanked for using the service and the session would end, prompting them to rate the chat experience.

In terms of aesthetics, the implementation team thought the chatbot should take on a friendly persona. The team proposed a "friendly gator librarian" avatar (the UF students are The Gators, after all) for the chatbot, and affectionately named her "Alice," after a campus landmark (Figure 2). Using an AI image–generating program and some customization by the Libraries Communications team, the team brought Alice to life. Once a patron clicks on the chat widget to initiate an interaction with Alice, they receive a friendly greeting from her along with the chat flow options to begin their virtual reference journey.

Chatbot Testing and Launch

Beta testing was a crucial step in the final stages of development of the chatbot. A team member created a chatbot testing zone through an internal TBLC webpage to allow a secure space to extensively test the chat flows. Once the implementation team concluded their testing and made minor adjustments in the flows, the AAL chat operators who responded to the survey were invited to review the chat flows and try to "break" Alice. The implementation team encouraged the chat operators to submit feedback and suggestions within one week.

Among the comments received after the chat operator testing period began were minor updates and clarifications to the language used, adding nonlibrary spaces on campus for some resources, and a way to easily return to the current chatbot flow to find additional answers on the same general topic. Unfortunately, due to limitations in the Springshare chatbot platform, the team was not able to complete the suggested flow changes within the same general topic. Alice went live on August 21, 2024, as the first George A. Smathers Libraries chatbot.

Discussion

Developing the chat flows and getting Alice ready for launch was not without obstacles.

Limitations of the Springshare interface compelled the team to scale back the initial intended scope of the project. The team found the platform used to design the chatbot flows to be difficult to navigate and unnecessarily complicated when creating flows with multiple levels. The design of the Springshare chatbot product forced the team to create a separate and unique chat flow for each topic. In the interest of time, the team scaled back and created a series of chat flows that were more simplified with a limited number of flow levels. On the plus side, this calibration of the chat flows made it an easier product for patrons to navigate, keeping the chances of getting lost in too many options at a minimum. Further, due to being part of the TBLC cooperative, the team had limited access to the full Springshare product (and the UF portions of the LibChat dashboard) and had limited creativity in terms of naming conventions and chat flow language. The team wanted to create unique names for each of the flows. Because other participating institutions could view the chat flow information in the admin dashboard, the naming convention of each flow required a more traditional and simple institution identifier. The limitation of the statistics available to the team

has also proven to be a source of frustration. More robust and thorough data when running usage reports would help the AAL site coordinator make continuous substantive improvements to Alice.

Another example of limitations met by the team was the ambiguous interpretation of each type of chatbot status within the chatbot statistics. While "Sent to Live Chat" and "Ticket Created" statuses have proven themselves to be straightforward when reviewing reports, the difference between "Incomplete," "Successful," and "Resource Clicked" are often unclear. "Resource Clicked" (Figure 3) and "Incomplete" (Figure 4) chat flow statuses regularly look similar when reviewing the transcripts, making it difficult to determine whether patrons are truly getting the help they need, where they are being directed to after their last chatbot interaction, or if the resource information has helped them at all. Because these chats do not officially end, the team is also not able to get any patron feedback on the service. One benefit of chat flows deemed "Successful" is that they show the chat ended. A truly "Successful" chat flow shows the patron located the resources or information they needed, and all their questions had been answered to their satisfaction (Figure 5). However, some of the chat flows categorized as "Successful" by Springshare do not have satisfactory conclusions. They either end immediately without any resources being clicked or with the patron stating their questions were not fully answered (Figure 6). These seemingly ineffectual interactions hinder the review of statistical data.

The time commitment of the project, from first design to final launch, was more than the team originally estimated. The team first sought to launch the chatbot within a few months of the project's start. In total, the project took nine months to complete. Due to the disparate schedules and locations of the team members, virtual meetings were the standard way the group collaborated, which made both scheduling to work together and creating cohesive and consistent language among the chat flow drafts difficult. When the team gathered in the same room, they completed the remaining tasks much faster and more efficiently.

Early analysis of transcripts and other data points shows that Alice is doing her job (Table 2). The information Alice delivers to patrons on chat is accurate and current, in keeping with the implementation team's work to populate the flows. Patrons that need more in-depth assistance are able to receive it from the live library workers staffing the virtual reference desk. Very few patrons have expressed dissatisfaction or made recommendations to the service; to date there has been a relatively low rate of feedback (24 total comments out of 6,719 interactions), with only two patrons (or roughly .03%) requesting the ability to type their specific questions rather than following the predetermined flow of questions. Further, 127 ratings (or 1.89%) of Alice have been provided by patrons, with an average rating of 3.5 out of 5. Two other patrons typed desired research queries into the comment box, which offers valuable insight into types of question flows that can be added to the chatbot. Enough queries have been intercepted and sufficiently addressed that the current AAL site coordinator, who handles building the shift schedule for the service, has been able to decrease the depth of hourly shifts from three operators to two. Since the launch of Alice, the current AAL Site Coordinator has continued to update and improve the chat flows based on statistics and feedback received from patrons.

Conclusion

The intent behind implementing a rule-based chatbot was never to replace the already robust virtual reference service provided by experienced live chat operators at the Smathers Libraries. The chatbot's role was to expand the service by creating another contact point for patrons seeking library assistance. Information seekers sometimes need help when there is no live chat operator

available, and the rule-based chatbot, Alice, can fill those gaps by providing answers to more basic queries. Information about library hours or printing in the libraries can be answered via website links or by directing the patron to submit a help ticket that can be addressed as soon as a librarian is available. During live chat operational hours, Alice fields those less complex questions and invites all patrons to speak to a live operator to assist with more intricate reference queries.

Surveying the current AAL operators was a helpful first step in developing the chat flows for the chatbot. Their experience provided valuable "boots on the ground" data about what types of information patrons are looking for. Based on that feedback, the team developed an initial list of categories from which to create efficient chat flows. Through trial and error, along with beta testing, the implementation team finalized a series of chat flows and a robust question bank using the Springshare LibChat product. Alice directs patrons to the appropriate library websites for answers to commonly asked questions while always providing the patron with the option to speak to a live chat operator, submit a ticket, or restart the interaction. The combination of simple but thorough chat flows and Alice's pleasant and whimsical appearance puts a friendly face and tone on a service that could be perceived as cold or impersonal by some patrons (and curtailing the often-asked question, "Are you a bot?" when a patron is speaking with a live chat operator). Since its launch, the chatbot has been a successful addition to the AAL virtual reference service at UF, intercepting and answering enough patron queries to allow for a reduction in the volume of physical staffing. The process of determining the types of questions Alice would address, the length and depth of the chat flows encountered by patrons, and the creation of those chat flows in the LibChat chatbot interface was sometimes frustrating. Nevertheless, the implementation team accomplished their goal of merging technology with human interaction in the virtual space of the library, giving patrons an added point of service whenever they need help from a librarian.

Acknowledgments

The authors wish to extend their sincerest gratitude to Adam Chalmers, Sarah Coates, and Garfield Simmons, without whose knowledge, diligence, and patience Alice would never have come to life. And a special thank you goes out to CJ Ellward for continuing and improving upon the initial work of the Alice Implementation Team.

References

1. Jana Smith Ronan, "Staffing a Real-Time Reference Service: The University of Florida Experience," *Internet Reference Services Quarterly* 8, no. 1-2 (June 16, 2003): 33–47, https://doi.org/10.1300/J136v08n01_05.
2. Tara Mawhinney, "User Preferences Related to Virtual Reference Services in an Academic Library," *The Journal of Academic Librarianship* 46, no. 1 (January 2020): 102094, https://doi.org/10.1016/j.acalib.2019.102094.
3. Jana Smith Ronan, "Staffing a Real-Time Reference Service."
4. Krishna Mojjada and Harihararao S. Siva, "Artificial Intelligence Based Academic Libraries: A Review Artificial Intelligence Based Academic Libraries: A Review," *Futuristic Trends in Artificial Intelligence 3*, no. 12 (2024): 153–72.
5. Tara Mawhinney, "User Preferences Related to Virtual Reference Services in an Academic Library," *The Journal of Academic Librarianship 46*, no. 1 (2020): 102094.
6. Constance A. Mellon, "Library Anxiety: A Grounded Theory and Its Development," *College & Research Libraries* 76, no. 3 (March 1, 2015): 276–82, https://doi.org/10.5860/crl.76.3.276.

7. Michal Černý, "Educational Psychology Aspects of Learning with Chatbots without Artificial Intelligence: Suggestions for Designers," *European Journal of Investigation in Health, Psychology and Education* 13, no. 2 (2023): 284–305, https://doi.org/10.3390/ejihpe13020022.
8. Jagdish Singh, Minnu Helen Joesph, and Khurshid Begum Abdul Jabbar, "Rule-Based Chabot for Student Enquiries," *Journal of Physics: Conference Series* 1228, no. 1 (May 1, 2019): 012060, https://doi.org/10.1088/1742-6596/1228/1/012060.
9. Thomas L. Reinsfelder and Katie O'Hara-Krebs, "Implementing a Rules-Based Chatbot for Reference Service at a Large University Library," *Journal of Web Librarianship* 17, no. 4 (October 2, 2023): 95–109, https://doi.org/10.1080/19322909.2023.2268832.