I just want my website to talk to me: implementation and trial of a chatbot at the University of Wolverhampton

John Dowd
Hybrid Services Manager
Learning and Information Services
University of Wolverhampton

There is an irreversible trend in academic libraries towards the migration of more and more of its content from print to electronic. Alongside this are the challenges of meeting the changing needs and expectations of users in a web-based world where library services are increasingly virtual. In this context a library's website inexorably moves towards a central focal point. The website is both the 'glue' holding together and presenting a library's resources and services, and also a 'front line' with regard to interacting with its users.

This article describes a number of parallel and overlapping activities which led Learning and Information Services (LIS) at the University of Wolverhampton to investigate chatbot technology as a tool to support its users and as a way of improving access to web-based information. It reviews the implementation and the initial findings following its launch in January 2010.

What is a chatbot?

The term 'chatbot' is created from the words 'chat' and 'robot'. Wikipedia defines a chatbot as a 'computer program designed to simulate an intelligent conversation with one or more human users'[1]. It can be placed on a website to help visitors find the information they need by answering specific questions which are pre-programmed into it and/ or by directing them to relevant pages on the site. The user types in questions and the response is displayed as text and/or audio output. Normally

the front-end of the chatbot, which the user interacts with, is an avatar, a 'graphical representation of a person or character in a computer generated environment'[2].

A PROJECT WITH MANY STARTS

In December 2006 the University of Wolverhampton was one of the first UK academic libraries to launch a live chat service, branded ASSIST, an acronym of 'Available Simple Supportive Information Service Today', enabling users to use instant messaging to engage with staff. The platform for the service is OCLC's QuestionPoint[3]. One of the useful by-products of the service is that chats are recorded and stored in the form of transcripts. For the first time, we were able to analyse objectively what questions people were asking. A closer examination of transcripts showed that a significant percentage of enquiries were for basic and routine information about our services; in answering the enquiries, librarians would often support the response by page-pushing relevant pages from our website.

In the summer of 2008 a usability study of the department's website was also conducted. The findings noted that:
- the library catalogue and electronic resources pages were associated with LIS; other than that the department's website had no real identity, being visually similar to the other school and department websites
- students do not navigate top down or browse; they access sections directly and have tunnel vision, focusing on 'getting the task done'; they do not spontaneously explore the site and are not aware of all the information and help available to them
- these problems had a tendency to push users to Google, not just for electronic resources, but also general LIS information.

The immediate effect of the study was to investigate refreshing the website, including a review of our web metrics. The metrics supported the usability report and also highlighted the importance of our website in that it was (a) the most heavily used website in the university, (b) the OPAC was the most heavily used element, followed by the electronic resources page, (c) hits on other pages in comparison dropped significantly and (d) the LIS homepage was actually one of the least used pages.

The combination of quantitative and qualitative findings about our website, alongside the chats

from the ASSIST service, suggested that no matter how much we rationalised and reorganised our website, a collection of passive, linked html pages would not be enough to help users find the information they need. At the 2007 TICER summer school 'Digital libraries à la carte', a presentation was given by Anne Christensen, Web Services Librarian, State and University Library Hamburg[4], on their chatbot Stella[5]. This chatbot tries to automate the reference interview, which was in itself a very broad scope for the chatbot, and as a result required the development of a complex and large knowledgebase. However, even though it appeared to be a difficult tool to create and there were potential cost vs. benefit issues, it did show a potential solution to the issues we had identified.

WHAT SHOULD THE CHATBOT DO?

After some internal discussion and lobbying, a project proposal was signed off in December 2008 to trial a chatbot. The project was initially linked to a parallel project refreshing our website (although subsequently decoupled due to some internal delays in moving the site to a new content management system). The main drive in moving forward was that, with the ASSIST chat transcripts, we had a clear evidence base as to not only what was being asked but *how* it was being asked (in particular the phrasing and terminology). The project was entitled HAL, an acronym standing for 'Helping ASSIST Librarians'; of course it had nothing at all to do with Arthur C Clarke's fictional *Space Odyssey* saga and the sentient computer with the same name.

- As a supplement to the ASSIST Service, the chatbot's scope was defined as being:
- supportive with regard to the Web Refresh project, and aiming to:
  - develop an evidence based FAQ
  - a 'push' to web site content not being fully utilised
- an alternate front-end/access method to content on LIS website, by offering an interactive service to enhance the passive web pages.

This scope was encapsulated in a statement of purpose: 'LIS's chatbot is an interface which will provide real-time interaction and navigation to resources and services provided by LIS. It supplements the ASSIST online chat service, by dealing with routine information enquiries.'

This initial scoping and definition of purpose proved critical to the project's success, especially during the development stages when the project

team, as they became more familiar with the technology, were tempted to utilise more functionality and deviate from our initial intent.

How should the chatbot do it?

Our main criteria for the end solution was that it would be lightweight to implement and maintain, easily controllable, require minimal staff training and be low cost.

Based on these criteria we identified several options on how a potential chatbot could be implemented on our website. The team decided that the best option was for the chatbot to be configured to push relevant pages of the website in response to queries, instead of encapsulating specific responses within the chatbot. This solution would avoid building a complex knowledgebase of responses to specific questions and instead it would become more of a supportive search tool for the website; it would also better meet the aim of promoting content on our website and would also avoid duplication between the website and the chatbot knowledgebase. The user, through interaction with the chatbot, would see web pages of relevance to them and so promote serendipitous learning. Figure 1 shows the chatbot responding to an enquiry on how to reference a website with a context specific response and the appropriate page being pushed.

maintained locally and uploaded into it. Keeping the knowledgebase and chatbot separate meant that the knowledgebase would be easily reusable should the project prove unsuccessful; it also meant that it was easier to provide quality control. As the chatbot would be pushing web pages, it was decided that it would not hold specific answers but respond with a short message (also appearing in text below the bot) and push the relevant page. Once it was launched by the user, it would remain viewable whilst pages were pushed using a split screen, so allowing for follow-up questions. As the bot would be based on keyword matching within the knowledgebase, if the results being retrieved for a specific user were continually low, it would push the main contacts page and suggest to the user that it may not be helping them.

Which chatbot to use?

Early on in pre-project discussions a chatbot vendor, Daden[6], was invited in to find out more about the technology. Daden was approached because the company had been name-checked in the TICER presentation and it had been runner-up in a Talis 'mashing up the library' competition in 2006 with a chatbot called Lillian the Librarian. Daden's hosted solution was then selected via a competitive process, and offered best value for money.

Figure 1

It was agreed that the chatbot's initial knowledgebase would be based on an analysis of ASSIST chat transcripts. The knowledgebase would not be created within the chatbot, but would be

During the project Daden was receptive to suggestions with regard to development and introduced additional tools such as a 'confidence meter', which gave the user a visual indicator of how confident the chatbot was in its responses to the questions being asked of it.

Over a two-month period two librarians reviewed the ASSIST chat transcripts. The review included two passes of the data, the first being a cut and paste of the summary of questions and associated responses along with any URLs used. This information was put into an Excel spreadsheet (see figure 2). Whilst this was being done they created a folksonomy to categorise query types together.

Figure 2

A second pass (see figure 3), reduced the summary of the question to its kernel keyword components.

Figure 3

The department's systems team then applied the vendor's syntax to the keywords. A standard response was added to each of these cases along

with a URL to be pushed. The file was then was then uploaded into the chatbot engine (see figure 4).

Figure 4

Once loaded, a reiterative process of testing was then was carried out. This was done by our dedicated telephone and online chat enquiry service team, Learning Centre Direct (LCD). Subsequent revisions then increased the accuracy of results. To support this process, the systems team created several automated procedures to convert files to the Daden import format, perform basic validation of the contents, execute a batch run to check the test questions and email the results and validate the status of the web page.

OTHER CHATBOT CONSIDERATIONS

For the front-end for the chatbot, the avatar, we followed Daden's recommendation to use Sitepal[7]. Initially the project team thought of running an internal department competition to decide the features of the chatbot; this was not pursued as it was difficult to get a consensus on what the 'face' of the department should be. The focus switched to concentrating on the development of the knowledgebase. Prior to launch, several mock-ups of faces were created and the LIS Executive team were asked to make a selection (see figure 5 for the 'face').

Figure 5

A name was not specifically given to the chatbot as there had been some initial discussion on

routinely switching its gender/race/age. This was discounted as it would have been impossible for the Chabot to answer specific personality questions, for example, 'How old are you?', if the age was constantly changing. Even though a 'face' was developed a name was not initially assigned; however, in a reply to query 'Who are you?', the response is 'I am Learning and Information Services Assistant'; the acronym of the response being LISA, a name was informally adopted by chance.

A 'small talk' file was built in parallel to its core knowledgebase, the only difference being that it was held within the chatbot engine itself. It did demonstrate the project teams' and other colleagues' creative flair in providing some interesting responses to 'small talk' users were engaging it with. The 'small talk' file is an essential supplement to the chatbot in that it provides a level of 'social' interaction, designed to make the chatbot more appealing to use, and providing a 'fun side' to it.

RESULTS OF BETA TESTING

LISA was tested against a small group of students and against a sample of ASSIST enquiries, and the results were promising enough to move it onto the website for a phase of live beta testing. The chatbot went live on the departmental website in January 2010 and was available for the start of the second semester 2. Usage was closely monitored by LCD. For the first two months the chatbot was available only on the LIS homepage in order to test robustness; it was then subsequently rolled out to the two most heavily used pages: OPAC homepage and electronic resources. It was not rolled out further. On reflection this should have been done, as the queries asked were often based on the page itself.

At launch, its knowledgebase had 733 cases (keyword/response/URL links). During the period this grew to 974 cases, representing a growth of 25%. It is estimated that another 200 of the original cases were amended.

During semester 2 it received 1351 queries; for the same period ASSIST received 826. However, it should be noted that the chatbot records each query whilst the ASSIST service records only a user session, which may include more than one question.

Pattern-matching traffic on the chatbot against ASSIST showed that the ebb and flow of usage is similar on a daily basis in terms of peak hours, with a sharp uptake in use from 09:00, staying constant during the day and tailing off after 21:00. Monthly comparisons also saw both services mirror each other in terms of traffic.

The quantitative review (see figure 6) has seen an overall success rate of 59%, 'success' being defined as the chatbot pushing the appropriate page. With revisions to the knowledgebase the success rate grew over the semester from mid to high 50s in January through to March, moving into the early to mid 60s in April/May.

some problems, though there is a text translation file. The main problem is users using the 'wrong' word to describe something; wrong in the sense of not pre-defined by the project team. For example, instead of asking how they find out when their books are due for *renewal* they ask what the *expiry date* of their books is. Another problem is use of homophones, for example asking for '*passed*' papers rather than '*past*' exam papers. It is rare that the chatbot is asked a question on a topic not programmed for it to answer; rather, it is how the user asks the question that can cause problems.

Reviewing the logs has seen most users engage to some degree with the chatbot. Most say 'hello', 'thank you' and 'goodbye', but there is little added conversation beyond this. Though 'small talk' is minimal, when a 'small talk' conversation does take place, users are often impressed. There were some instances of abusive and banned language (swear words) when the chatbot first went live, but now there is very little. On occasions when this happens now, it is normally caused by the chatbot not knowing the answer to the question and asking the user to contact a member of staff. Many compliments have been noted in the logs, for example, 'you are good', 'this is a great service', etc.

Figure 6

The final month of the evaluation period, June 2010, saw the chatbot achieving a 72% hit rate.

Most questions are asked using short sentences or by using one keyword; the latter will often result in the chatbot not knowing the answer, as not enough information is provided. Only a few users have typed in lengthy questions.

Spelling has been slightly problematic, but not as much as anticipated. Users have at times used text speak when asking questions, which has caused

Some users have thought the chatbot was a live person talking through the bot or have asked if it is a real person and asked them to do things like renew their books until a certain date, etc.

Others have missed the web page that is pushed. For example, if asked a referencing query the bot will say 'This Harvard referencing webpage should have the information you require' and then the user asks 'Which web page?', indicating that they have not noticed that the web page has changed.

Some users have lost the web page given to them by the chatbot by the simple fact of them saying 'thank you' to the bot. 'Thank you' triggers a response that pushes the main web page, causing the page they needed to disappear. The user then has to find the page again as there is no 'back' button. As no URLs are given when a page is pushed, the option to bookmark would be useful.

Throughout semester 2 there was a link to a passive survey for users to provide qualitative feedback; passive in the sense that there was a link to the survey should the user be motivated enough to respond, but this was not a requirement during the process of using the chatbot. The survey asked for their opinion on:

- helpfulness of the chatbot
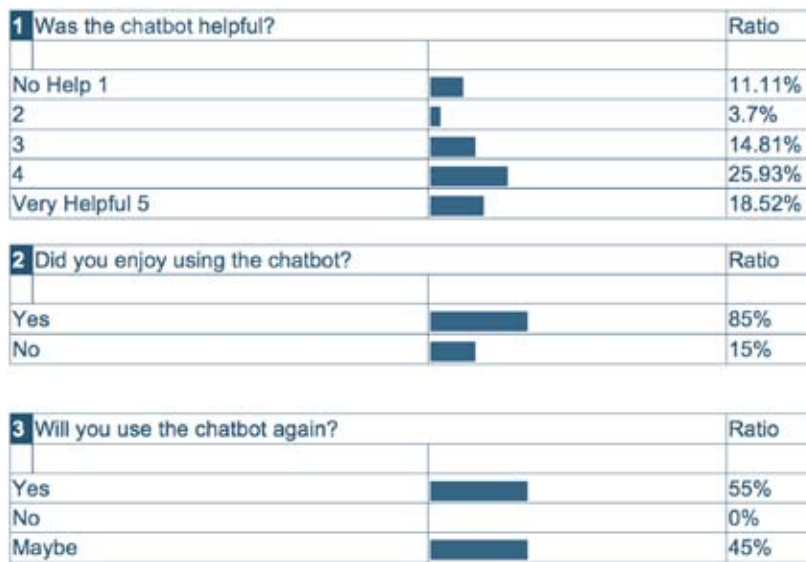- enjoyment in using the chatbot
- whether they would use it again

1 Was the chatbot helpful?		Ratio
No Help 1		11.11%
2		3.7%
3		14.81%
4		25.93%
Very Helpful 5		18.52%

2 Did you enjoy using the chatbot?		Ratio
Yes		85%
No		15%

3 Will you use the chatbot again?		Ratio
Yes		55%
No		0%
Maybe		45%

Figure 7

Figure 7 shows that (a) an even spread found it helpful to a varying degree of satisfaction, with a weighting towards 'very helpful'; (b) most enjoyed using it and (c) more than half would use it again; no one said that they would not use it again .

They were also asked for their general comments; the majority were positive and included:

- 'very good idea!'
- 'the response was quick' and 'it was easy to use'
- 'the website sadly is very awkward to use in terms of finding important information such as opening time; however the chat box gave the answer, fantastic!!!'

Negative comments were mainly with regard to incorrect matching of keywords to a response or URL problems:
'in response to "I need a journal article on engineering" the chat bot found a site error page'
'I asked 'who is the librarian for deaf studies?" and it gave me a visitor information page. Other questions answered well'.

FINAL THOUGHTS

The chatbot is achieving a good success rate in terms of responding to queries; and the user surveys, both pre-launch testing and beta, show that it is popular with those who use it. Is it suitable for all types of user? The short answer is no, but it was never meant to be. Some people like using chatbots, others do not. What we have seen is that if a user uses our chat service, they are more likely to engage with the chatbot, and what it does provide is an additional interactive tool to ask questions. Even though the project was called HAL for 'Helping ASSIST Librarians', it has not impacted on ASSIST usage. Rather, ASSIST has helped the chatbot by providing an evidence-based knowledgebase. The fact that users have spent time engaged in 'small talk' with the bot aids in the site's overall stickiness; this 'fun' element means they stay longer on the site and by simple interaction with the chatbot see more of what is available.

To conclude: a comment from one of our survey respondents: '…the more available options we have to find the answers to our questions, the better'. This summarises nicely what our chatbot is all about.

NOTES

1. Wikipedia, http://en.wikipedia.org/wiki/Chatterbot [accessed 19 November 2010].
2. Oxford English Dictionary Online, http://dictionary.oed.com [accessed 19 November 2010].
3. QuestionPoint 24/7 Reference Services, http://www.questionpoint.org [accessed 4 November 2010].
4. Tilburg Innovation Centre for Electronic Resources, http://www.tilburguniversity.nl/

ticer/07carte/publicat/07achristensen.pdf
[accessed 4 November 2010].

5   Bibliothekssystem Universität Hamburg,
     http://www.sub.uni-hamburg.de/bib-
     liotheken/projekte/chatbot-stella.html
     [accessed 4 November 2010].

6   Daden Limited, http://www.daden.co.uk/
     [accessed 4 November 2010].

7   Sitepal, http://www.sitepal.com/ [accessed 4
     November 2010].

5   Bibliothekssystem Universität Hamburg,
     http://www.sub.uni-hamburg.de/bib-
     liotheken/projekte/chatbot-stella.html
     [accessed 4 November 2010].