

ALife-WebGuide: An Intelligent User Interface for Web Site Navigation

Paolo Gaudiano

Artificial Life, Inc.
Four Copley Place, Suite 102
Boston, MA 02116 USA
+1 617 266 5542

Paolo.Gaudiano@artificial-life.com

Klaus Kater

Artificial Life, Inc.
Four Copley Place, Suite 102
Boston, MA 02116 USA
+1 617 266 5542

Klaus.Kater@artificial-life.com

ABSTRACT

This article describes Artificial Life, Inc.'s *WebGuide*, an intelligent software bot that helps users navigate a Web site using natural language. The article describes the technology behind ALife-WebGuide, discusses some of the issues involved in commercialization of this type of user interface, and summarizes possible enhancements for future versions of this product.

Keywords

Web site interface, Web navigation, intelligent navigation, software bots, natural language.

1. INTRODUCTION

The dramatic growth of the World Wide Web has brought about profound changes in the way people find information, interact with each other, and purchase goods. The proliferation of personal and commercial Web sites has created a substantial problem for users, who must sift through enormous amounts of information to find what they want. Likewise, companies are struggling to make their Web sites highly visible, so as to attract more visitors.

A significant amount of research has gone into collaborative filtering (e.g., [5],[6]), personal Web assistants [2],[3], and other technologies that can help users find information on the Web. Most of these approaches focus on simplifying navigation for individual users, for instance by learning a user's navigation preferences, and then matching the user's preferences against the content or other features of Web pages, or against the preferences of other users.

An alternative approach, preferred by the owners of commercial Web sites, is to focus the effort on the content and appearance of the site itself, increasing its value or visibility by making it more informative and entertaining, easier to access, and generally more palatable to users. Typical methods for increasing the accessibility of Web sites include site maps and local search engines. However, these techniques still require a significant amount of effort on the part of the user, who must navigate a complex map with terse page titles, or learn what keywords might be appropriate for the search engine, and then sift manually through multiple search engine results.

The research that focuses on learning user preferences has one important drawback: companies might be reluctant to support tools that make it easier for a user to surf *any* Web site, including competitors' Web sites. Ultimately, the commercial viability of a technology will have a profound influence on its widespread acceptance.

This article describes ALife-WebGuide [1], an intelligent user interface that helps users navigate a Web site using natural language. ALife-WebGuide, a commercial product by Artificial Life, Inc. (NASDAQ:ALIF), is part of a suite of *smart bots* designed to facilitate many typical tasks of Internet users. Our product takes into account the needs of the Web site owner, not only by making navigation simpler for users, but also by making it possible to focus the bot's knowledge on the contents of a specific Web site, thereby increasing the value of this product for Web site owners.

As a commercial application of intelligent user interfaces, ALife-WebGuide is of interest for several reasons. First, it shows how intelligent user interfaces can have commercial value by offering benefits both to the owners of Web sites and to Web site users. Second, it is an example of how an intelligent user interface can improve the way people visit Web sites, making their experience more similar to what they would expect from visiting a real store. Third, we have made some informal observations about Web site visitors who use ALife-WebGuide. These observations confirm the usefulness of friendly and intelligent user interfaces, they show what strategies seem particularly successful, and they offer a heuristic for determining how much artificial intelligence (AI) is "enough."

The next section provides an overview of ALife-WebGuide, followed by a more detailed description of two versions of this intelligent user interface. The article closes with a brief discussion of our experience with ALife-WebGuide, and a summary of some of the improvements that we are planning for future versions of this product.

2. AN OVERVIEW OF ALIFE- WEBGUIDE

ALife-WebGuide was designed as a tool to allow user-friendly, intuitive navigation of a Web site. For the user visiting a Web site, ALife-WebGuide typically appears as a *bot*, a static or animated character that accepts input from a text entry box. The bot's reply typically appears in a text balloon. Figure 1 shows an example in which ALife-WebGuide is used as a sales representative for a fictitious shoe shop.

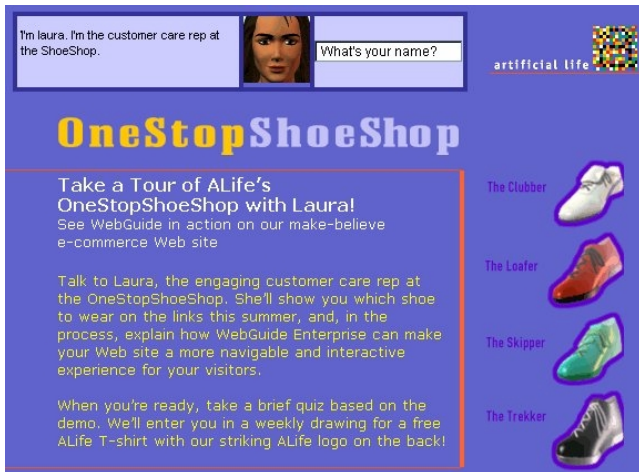


Figure 1: The *OneStopShoeShop* ALife-WebGuide demo (www.artificial-life.com/shoeshop).

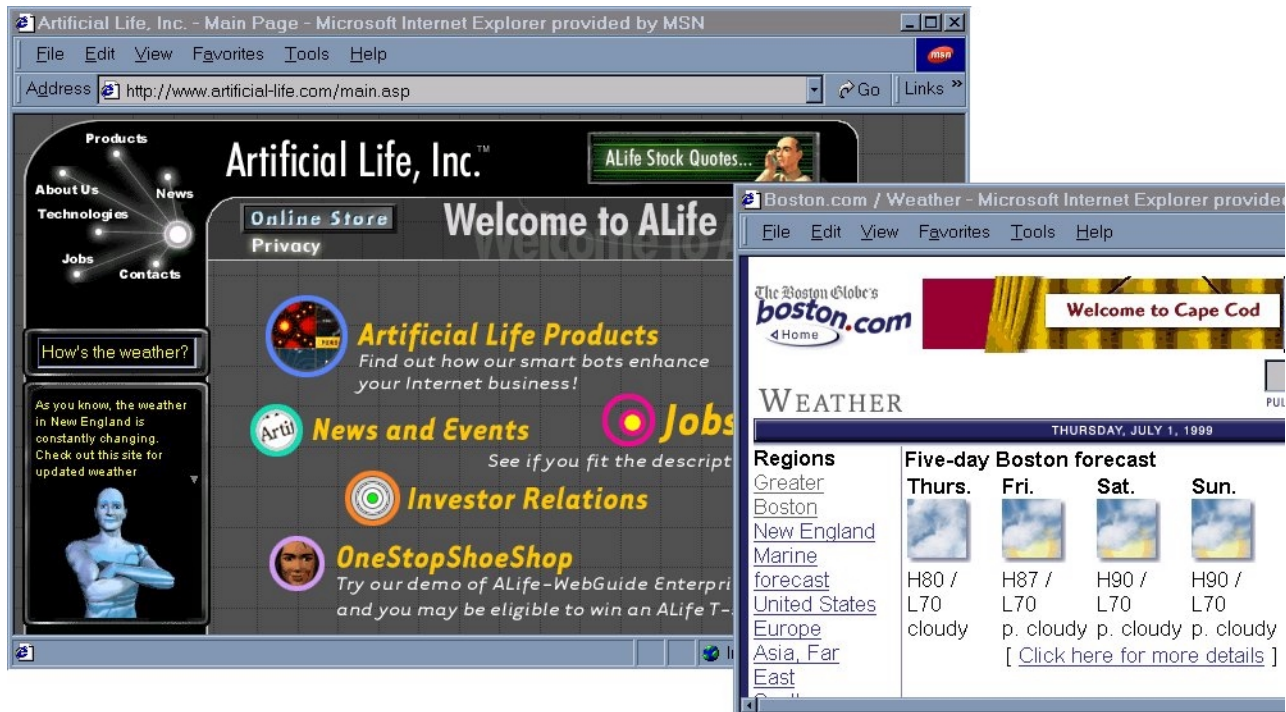
Customers visiting the ALife *OneStopShoeShop* can ask the bot questions using natural language. The bot can answer general questions (e.g., “what is your name?”) and can provide specific information about the available shoe models, sizes, and so on. In response to each user query, the bot responds with text, it can show some animated facial expressions, and it can load different content in the main frame of the browser.

Figure 2 shows another sample application of ALife-WebGuide, in this case from our company’s Web site (notice that our site is likely to change by the time of the conference). The character in the left frame of the browser has an extensive Knowledge Base that covers general topics as well as specific issues pertaining to the company. As with the first example, the bot can reply with text, animations, and Web pages. In addition, the bot can choose to display information in a new browser window. For instance, when asked about the weather, the bot gives a relevant verbal response and opens a new browser window, loading a URL that shows the current weather in Boston.

2.1 Simplifying Web site navigation

Users visiting a WebGuide-enabled Web site can navigate the site much as they might visit a regular store: they can simply browse by clicking on hyperlinks, or they can ask the bot for guidance. The latter navigation method offers several advantages:

- Users need not have any knowledge about the structure of the site or even the nature of the site. The bot is capable of answering general and specific questions about the Web site and other topics.
- Unlike search engines, the bot returns at most only a single page in response to a single query. Note that the bot can give multiple, different answers to the same or similar queries, but provides only one answer at a time.
- The information returned by the bot is likely to be highly relevant as it is based on a Knowledge Base designed by the Web site owner (see below).
- It is possible with a single question to reach a page that is several levels deep within the site’s hierarchy.



2.2 Increasing customer tracking and satisfaction

The extent of the bot's knowledge, and thus the bot's usefulness for the site user, depends on the amount of detail that the Web site owner builds into ALife-WebGuide's Knowledge Base. From the point of view of a commercial product, it is important to show how the benefits obtained by the Web site owner can exceed the cost of creating and maintaining the bot's Knowledge Base. Benefits fall into at least three categories:

- Users are more likely to find the information they seek, and will thus be more satisfied with their visit, more likely to purchase items, etc.
- The Web site owner can bias the bot's conversation to favor particular topics. For instance, in response to chit-chat, the bot can shift the topic to a particular product and present relevant pages.
- By interacting with the bot, users leave behind much richer information than the information available by merely tracking mouse clicks.

We will return to some of these issues after a more detailed description of ALife-WebGuide.

3. ALIFE-WEBGUIDE TECHNICAL DETAILS

Currently two versions of ALife-WebGuide exist. The 1.0 version, which is available in "Professional" and "Enterprise" releases, is based primarily on pattern-matching techniques that match user input sentences against information contained in the bot's Knowledge Base. The two releases use virtually the same Knowledge Base format and language understanding techniques, so they are treated here as a single version. Both releases are written in Java, and both include a Knowledge Editor based on Microsoft Access forms.

At the time of this writing, version 2.0 is used for Artificial Life's Web site (as shown in Figure 2), but is not yet commercially available. The newer version includes a much more powerful, modular *SmartEngine* that allows for context recognition, memory about the conversation, emotional processing, and several other features. Both versions are described in the rest of this section.

3.1 ALife-WebGuide 1.0

The "intelligence" of bots based on ALife-WebGuide 1.0 is captured in a Knowledge Base constructed by the owner of the Web site. Information about the Web site (or whatever knowledge the bot should have) is first divided into a hierarchy consisting of *classes*, *topics*, and *subjects*. Classes organize the content of the Knowledge Base at the highest level. For example, a company's Knowledge Base might have classes called *products*, *technology*, *company information*, etc. Topics further subdivide each class into finer categories. For instance, the class *company* might include topics called *address* and *staff*. Topics in turn contain several subjects: for instance, *WebURL* and *e-mail_address* might be subjects in the topic *address*.

Subjects are the first "functional" level of a Knowledge Base: each user query, after a preprocessing step, is matched against the names of all subjects. All matching subjects are ranked based on a user-specified priority level.

Each subject contains a collection of *user input patterns* (or *userins*) which, when matched, elicit one of possibly several *responses* from the bot. Each *userin* is a pattern combining text and the * (asterisk) wildcard. For instance, the pattern

** what* company* do**

matches sentences such as "tell me what your company does" or "what does your company do?".

After the initial matching against subject names, ALife-WebGuide starts with the highest priority subject and tries to match the user query against all *userins* specified for that subject. If no match is found, the subject with the next highest priority is tested for matching, and so on. If all matching subjects are exhausted, a *catch-all* subject is selected, which returns a "generic" response.

Each subject, including the catch-all, specifies at least one response. The response can include text (which will be displayed in the bot's text balloon), the URL of any Web page to be displayed in the user's browser, a specification of where the Web page should be displayed, and the name of an action corresponding to an animation for the bot. When multiple responses are specified for a given *userin*, ALife-WebGuide cycles through each response in order. Defining multiple responses is especially important in giving the bot a more "human" quality, because repeating the same sentence does not yield the same response.

ALife-WebGuide includes additional features, such as the ability to specify *filters* (for instance to normalize misspelled words or different forms of irregular verbs), *synonyms*, and a *goto* response directive that can force selection of an alternative subject in the Knowledge Base. The Enterprise version of the program also includes administrative tools to analyze individual conversations or statistics about the topics chosen by Web site visitors.

Clearly the amount of "intelligence" of version 1.0 is limited. However, even with this relatively simple pattern matching and priorities it is possible to construct complex Knowledge Bases that can handle many kinds of user queries. What is more important is that this product offers users a novel, interactive way to navigate Web sites. For version 1.0 it was judged important to provide a user-friendly, intuitive interface and other features to help gain user acceptance for a type of Web navigation that is relatively unknown to the majority of Web users.

In our experience this approach has been successful: we have monitored user-bot conversations, and found that users can become involved in lengthy and highly informative conversations that would not be possible with traditional navigation schemes. Furthermore, user-bot conversations contain much more information than is available through log files resulting from traditional click-based navigation. The Web site owner who uses ALife-WebGuide can have a much better sense of what the customers are really interested in, and can use this information to improve the Web site or to influence product marketing decisions.

3.2 ALife-WebGuide 2.0

There are many ways in which ALife-WebGuide 1.0 could be improved, but arguably the most important is to add memory. Version 1.0 only considers one sentence at a time, without taking into account context from past sentences in a given conversation (the exception is that a *userin* with multiple responses will cycle through each response if it is selected repeatedly). This means that ALife-WebGuide 1.0 might present the same information more than once to a given user during a single conversation. Also, ALife-WebGuide 1.0 cannot have a true conversation in which the user and the bot take turns in a continuous conversation. The design of ALife-WebGuide 2.0 makes it possible to overcome these shortcomings, and adds several other desirable features.

The main technical difference lies in the new *SmartEngine 2.0* technology. The SmartEngine is the “brain” of ALife-WebGuide and of several other bots that we are developing. The SmartEngine 2.0 is able to follow conversations, keeping track of context across multiple sentences. For instance, suppose a user asks “Who is the boss at ALife?” and the bot replies “Eberhard Schöneburg is the CEO of Artificial Life, Inc.” while showing the relevant Web page. If the user then asks “What is his phone number?” the SmartEngine will realize that the pronoun “his” refers to the subject of the previous sentence, in this case the company CEO.

Another feature of the SmartEngine 2.0 is the ability to manage discourse by specifying “goal topics”. When the user and the bot have exchanged a few sentences, the bot can switch topics and discuss specific items that the Web site owner wishes to present to customers. The SmartEngine can select multiple goals that are chosen dynamically based on the entire conversation. It is also possible to select how “pushy” the bot will be in switching topics. Discourse management is also used to restrict the focus of the conversation so that, for instance, if the bot asks the user “would you like to hear about our latest product”, it will know how to handle answers such as “yes” or “no” that would be meaningless in an unconstrained context.

Other changes include the ability to influence the bot’s responses based on the emotional content of the user’s sentences. This capability is embedded in the SmartEngine but has not yet been applied in a commercial setting.

We have already used the SmartEngine 2.0 for ALife-Messenger [1], an automated e-mail response program that interprets e-mail messages. We have also used the SmartEngine 2.0 for the WebGuide on our company’s site, and for a number of custom projects. The commercial version of ALife-WebGuide 2.0 should be available around the time of this conference.

4. DISCUSSION

We have found that ALife-WebGuide is a useful product because it offers clear benefits both to the end-user and to the owner of a Web site who is trying to attract and retain more customers. Our focus on the Web site owner is crucial from a commercial standpoint, as this is what makes the product viable. This is one main difference between ALife-WebGuide and other types of personal navigation assistants proposed in the scientific literature (e.g., [2]-[3]), which focus almost entirely on the end-user. We also differ from other commercial offerings, such as those of Extempo, Virtual Personalities and Haptek, by virtue of our emphasis on a common core technology that works on the Web, but also through e-mail, standard call centers and other media.

It is interesting to note that none of the commercial companies, ourselves included, utilize a particularly deep or sophisticated form of Natural Language Processing (NLP). There is a vast literature on NLP, and we ourselves are pursuing this area of research. However, we have not found a rigorous approach to NLP that works satisfactorily in the unconstrained, realistic situations that our bots encounter on a daily basis. Furthermore, our low-level reliance on pattern matching makes the underlying bot Knowledge Bases largely language independent. In fact, we have already developed bots in English, German and Italian.

Given that computers are unlikely in the foreseeable future to really appear “human” to human users, we believe that a shallow

approach that minimizes linguistic accuracy but maximizes flexibility gives much better returns than an approach that really tries to replicate human-level conversations. Nonetheless we have found, as have others, that many users have a strong preference for interacting with an ostensibly artificial character in plain language than for using keyword-based search engines or other traditional navigation methods.

Another important observation is that some customers seem to really enjoy the verbal interactivity, even when they clearly realize that they are interacting with an artificial creature. We have logged thousands of conversations with visitors to our Web site. The average conversation involves six turns (question/answer), lasts 280 sec., and covers three different topics. The longest conversation we have logged took over 240 turns and lasted approximately two hours. Users like to discuss many topics, often drifting to personal and even offensive topics. However, we can see that the Web pages most frequently seen by our visitors correspond to those identified as goals in the bot’s conversation.

In summary, we find that a little bit of AI has a great positive effect, but that one quickly encounters diminishing marginal returns when trying to adopt complex AI or NLP techniques.

There are many ways in which we intend to improve our products. One is to increase the flexibility of the language parsing, for example by using a built-in thesaurus and spell-checker. Another area of planned improvement is in the development of tools to facilitate the creation and maintenance of Knowledge Bases, tasks that currently require a significant investment of human resources. Finally, we are going to use the content of the conversations to improve our ability to recognize Web site users and their preferences. We believe that considering user preferences is an important element of all future Web transactions, as it can offer users a much more personalized experience, and it can give the Web site owner a clearer picture of its customers.

5. REFERENCES

- [1] *Alife Solutions*. White paper available on the Web at www.artificial-life.com/whitepapers/ALifeSolutionsWhitePaper.zip, 1999.
- [2] Chen, L. and Sycara, K. WebMate: A Personal Agent for Browsing and Searching. *Proceedings of the International Conference on Autonomous Agents*, Minneapolis, MN, May 1998. ACM Press, 132-139.
- [3] Lieberman, H. Letizia: An Agent That Assists Web Browsing. *Proceedings of the International Joint Conference on Artificial Intelligence*. Montreal, CA, August 1995.
- [4] Shavlik, J., Calcari, S., Eliassi-Rad, T. and Solock, J. An Instructable, Adaptive Interface for Discovering and Monitoring Information on the World Wide Web. *Proceedings of the 1999 International Conference on Intelligent User Interfaces*, ACM Press, 157-160.
- [5] Van Dyke, N.W., Lieberman, H. and Maes, P. Butterfly: A Conversation-Finding Agent for Internet Relay Chat. *Proceedings of the 1999 International Conference on Intelligent User Interfaces*, ACM Press, 39-41.
- [6] Wasfi, A.M.A. Collecting User Access Patterns for Building User Profiles and Collaborative Filtering. *Proceedings of the 1999 International Conference on Intelligent User Interfaces*, ACM Press, 57-64.