


~Mark Lucente

CONVERSATIONAL INTERFACES *for* E-COMMERCE APPLICATIONS

These e-commerce shopping agents are knowledgeable, patient, and affable during their conversational interaction with shoppers — while helping generate sales for site owners.



CONSUMERS SPEND BILLIONS ONLINE, THOUGH MOST ARE UNHAPPY WITH their shopping experiences [1]. My colleagues and I at Soliloquy, Inc., are creating conversational natural language interfaces for e-commerce applications that appear on Web sites as sales “Experts” conversing with shoppers, answering their questions, and helping find products they can buy. Each Soliloquy Expert is a software agent custom-developed for a particular e-commerce application. The metaphor reflected in an Expert-based e-commerce experience is natural conversation, as if the shopper were talking to a human salesperson. The shopper and the Expert co-produce knowledge and understanding through the conversation—a shopping experience more effective, fun, and human than its conventional counterpart of point-click-download.

By constraining an Expert’s knowledge domain to a particular product line, say, laptop PCs, restaurants, or mutual funds, they can be made robust and helpful. Domain constraint facilitates several features of the personal shopping experience:

- Rich contextual understanding;
- Robust speech recognition;
- High-quality text-to-speech synthesis;
- Effective Web site navigation; and
- A business model based on spoken interfaces.

The first two features benefit from the disambiguation, or derivation of precise contextual meaning, made possible by the limited size of a particular domain—its concepts and concomitant words—that is, of course, significantly smaller than all the concepts and all the words of any particular human language. For the company behind the Web site, profitability is emphasized; for spoken interfaces to flourish and ultimately meet popular expectations, they must support a “killer app” that is both compelling and profitable—in this case, conversational e-commerce.

Each Expert combines several technologies suitable for mass-market e-commerce applications: speech recognition, speech synthesis, natural-language understanding, and the Web. Early use over the past two years has yielded many insights from this confluence of technologies while prompting several important questions about the evolving online shopping experience. Especially notable are the multimodal nature of the interaction and the natural human perception built into each Expert’s “personality.”

Experts appear to the shopper as a small conversation window on an e-commerce Web page (see Figure 1). A human-like Expert (which can include a cartoon face) appears eager to help the shopper understand what is for sale and find the most appropriate items to buy. An Expert begins the conversation by greeting and prompting the shopper for questions. The shopper speaks to the Expert or types into the window. The Expert responds in several ways:

- Spoken (if audio output is enabled);
- Text; or
- Hyperlinked multimedia, including images, such as those of products available for purchase.

An Expert’s main goal in a conversation is to help a shopper find the right item. Its response aims to satisfy several subgoals:

- Answer the shopper’s questions;
- Proffer items the shopper may wish to buy;
- Ask for clarification; and
- Prompt the shopper for information, enabling itself to refine the shopper’s search.

Each Expert uses a proprietary natural language understanding system developed by Soliloquy built atop a structured knowledge base. Natural-language subsystems convert the incoming text (from the keyboard or from a speech recognizer) into conceptual representations, add these concepts to the conversational context, determine subgoals, and finally motivate actions, such as constructing an answer,

searching a product database, navigating a Web page, or prompting the shopper for more information. A built-in conversation manager orchestrates the input and output while updating the context representation. The structured knowledge base in each Expert contains the following components:

- Functional representations of general concepts (relevant to e-commerce shopping) and their relationships to the words and phrases of a particular language (English for now);
- Functional representations of domain knowledge (such as the often-specialized language of PC technology);
- A semantic representation of the essential structure and function of domain products and what they are used for (such as a laptop computer with a certain amount of RAM and at least one hard disk or a laptop that doesn’t weigh much, necessitated, perhaps, by the user’s traveling schedule);
- A database listing the products for sale and their features (such as price, speed, and size) on a given e-commerce Web site.

The knowledge-engineering approach facilitates a rich representation of context, allowing the shopper-computer conversation to be natural, nonlinear, and efficient [1]. Our object-oriented approach to conceptual representation makes it possible to quickly scale up the knowledge base of an Expert and create additional Experts in a multitude of domains.

Inherently Multimodal

The interaction between a human user and a shopping Expert differs from human-human conversation in important ways that we have now begun to understand and exploit. Because each Expert is designed to appear on a Web page, the shopper-Expert interaction is inherently multimodal. The Expert senses the shopper’s mouse and keyboard activity. The shopper sees images and hypertext generated on-the-fly. Moreover, each Expert can offer graphical interface elements (such as buttons or hypertext words) to augment the conversation, like a human-human conversation involving physical objects the participants might gesture toward or use iconically.

Multimodal interaction allows for rich conversation and contextual development often exceeding that of human-human speech-only conversation. Although each Expert’s natural-language understanding is not, of course, as advanced as that of a typical human, the presentation of multimedia responses (such as pictures of the items for sale) facilitates complex conversation and delights users. For example, if a

Amazon/VB Assets
Exhibit 1010
Page 3