Cobot in LambdaMOO: A Social Statistics Agent

Charles Lee Isbell, Jr. **Michael Kearns**
Dave Kormann **Satinder Singh** **Peter Stone**
AT&T Shannon Labs
180 Park Avenue
Florham Park, NJ 07932-0971

Abstract

We describe our development of Cobot, a software agent who lives in LambdaMOO, a popular virtual world frequented by hundreds of users. We present a detailed discussion of the functionality that has made him one of the objects most frequently interacted with in LambdaMOO, human or artificial.

Introduction

The internet is a medium where large groups of people build social communities. This presents both a challenge and an opportunity for artificial intelligence and software agent researchers. In such communities, agents may do more than filter mail, and retrieve price quotes for consumer items; they may be legitimate, if still limited, participants in close-knit social environments.

This paper presents Cobot, a software agent that lives in an active online community frequented by several hundred users (LambdaMOO, which we describe in requisite detail in the next section). His goal is to interact with other members of the community and to become a vital, useful and accepted part of his social fabric. Toward this end, Cobot tracks actions taken by users, building statistics on who performs what actions, and on whom they use them. For example, Cobot tracks which users converse with each other most frequently. Using his chatting interface, Cobot can answer queries about these and other usage statistics, and describe the statistical similarities and differences between users. This information also provides Cobot with a user model that may be used for learning, imitation and conversation. Cobot's chat abilities include a simple and novel method inspired by information retrieval that allows him to choose appropriate utterances from very large documents.

During Cobot's months in LambdaMOO, he has become a member of the community. As we will see in subsequent sections, users interact with Cobot more than with any other user (human or artificial), take advantage of social statistics he provides, converse with him, and discuss him.

Following the pioneering studies of (Foner, 1993; Foner, 1997), we present transcripts establishing the sociological impact of Cobot. We compare the techniques that we have used to implement Cobot's set of social skills to those of previous MUD agents, such as Julia (Foner, 1997), and discuss how these techniques affect user expectations. In addition to the more anecdotal evidence provided by transcripts, we provide quantitative statistical support for our belief that Cobot has not only become part of the social environment in which he resides, but has significantly altered it as well.

The paper begins with a brief history of LambdaMOO. We then detail the two major components of Cobot's functionality—his ability to provide social statistics, and his conversational abilities—and quantify and discuss their impact. After an examination of privacy issues and our approach to them, we discuss future plans for Cobot.

LambdaMOO

LambdaMOO, founded in 1990 by Pavel Curtis at Xerox PARC, is one of the oldest continuously-operating MUDs, a class of online worlds with roots in text-based multiplayer role-playing games. MUDs (multi-user dungeons) differ from most chat and gaming systems in their use of a persistent representation of a virtual world, often created by the participants, who are represented as characters of their own choosing. The mechanisms of social interaction in MUDs are designed to reinforce the illusion that the user is present in the virtual space. LambdaMOO is a MOO: a MUD that uses an object-oriented programming language to manipulate objects in the virtual world.

LambdaMOO appears as a series of interconnected rooms (modeled as a mansion), populated by users and objects who may move from room to room. Each room provides a chat channel shared by just those users in the room (users can also communicate privately), and typically has an elaborate text description that imbues it with its own "look and feel." In addition to speech, users express themselves via a large collection of *emotes,* allowing a rich set of simulated actions, and the expression of emotional states:

(1) Buster is overwhelmed by all these paper deadlines.
(2) Buster begins to slowly tear his hair out, one strand at a time.
(3) HFh comforts Buster.
(4) HFh [to Buster]: Remember, the mighty oak was once a nut like you.
(5) Buster [to HFh]: Right, but his personal growth was assured. Thanks anyway, though.
(6) Buster feels better now.

Lines (1) and (2) are initiated by emote commands by user Buster, expressing his emotional state, while (3) and (4) are examples of emote and speech acts, respectively, by HFh.

Lines (5) and (6) are speech and emote by Buster. (In our transcripts the name of the user initiating an action always begins the description of that action or utterance.) Though there are many standard emotes, such as the use of "comfort" in line (3) above, the variety is essentially unlimited, as players have the ability to create their own emotes.

The rooms and objects in LambdaMOO are created by users themselves, who devise descriptions, and control access by other users. Users can also create objects with methods (or *verbs*) that can be invoked by other players.[1] LambdaMOO is thus a long-standing, ongoing experiment in collective programming and creation, with often stunning results that can only be fully appreciated firsthand. Inventions include technical objects, such as the *lag meter*, which provides recent statistics on server load; objects serving a mix of practical and metaphorical purposes, such as elevators that move users between floors; objects with social uses, such as the *birthday meter*, where users register their birthdays publicly; and objects that just entertain or annoy, such as the *Cockatoo*, a virtual bird who occasionally repeats an utterance recently overheard (often to amusing effect). There is also a long history of objects that can be viewed as experiments in AI, as we will discuss below.

LambdaMOO's long existence, and the user-created nature of the environment, combine to give it with one of the strongest senses of virtual community in the on-line world. Many users have interacted extensively with each other over a period of years, and many are widely acknowledged for their contribution of interesting objects. LambdaMOO has been the subject of articles and books in many different literatures, including the popular press (Dibbell, 1999), linguistics and sociology (Cherny, 1999), computer science, and law.[2] The complex nature of the LambdaMOO community goes a long way towards explaining why it is difficult to simply characterize what users "do" on LambdaMOO. As in real life, users engage in a wide variety of activities, including social activity, programming, and exploring.

LambdaMOO is an attractive environment for experiments in AI. The population is generally curious and technically savvy, and users are interested in automated objects meant to display some form of intelligence (called "puppets"). There is a rich history of automated agents and constructions: Markov chainer, an object that builds a Markov model from the conversations in a room; Dudley, a well-known agent with simple chatting abilities; an automated bartender who provides virtual drinks and small-talk; and many others. There are also object classes allowing users to specialize and create their own AI agents. The agent Julia, a descendant of Colin (created by Fuzzy Mauldin (Mauldin, 1994)), who once resided in a different MUD, is perhaps the closest ancestor of Cobot. We will discuss both Julia and her analysis by Foner (Foner, 1997), which has strongly influenced our thinking, throughout the paper where appropriate.

[1] Everything in LambdaMOO is an object, and every event is the invocation of a verb on some object, including speech (usually invocations of the **tell** verb). The LambdaMOO server maintains the database of objects, and executes verbs. As of this writing, the database contains 118,154 objects (not all are valid), including 5158 active user accounts.

[2] LambdaMOO has its own rather intricate legal system.

Cobot

Most of Cobot's computation and storage occurs off-server. He is built using *the Cobot platform*, an agent architecture that uses a directed graph-based metaphor to define a standard for describing a wide-range of virtual environments, including MUDs. A complete discussion of the platform is beyond the scope of this paper.

Cobot appears to be just another user. Once connected, he usually wanders into the LambdaMOO Living Room, where he spends most of his time.[3] The Living Room is a central public place, frequented by many regulars. It is also located next to the Linen Closet, where guests tend to appear, so it is also frequented by users new to LambdaMOO. There are several permanent objects in the Living Room, including a couch with various features, a cuckoo clock, and the aforementioned Cockatoo. The Living Room usually has between five and twenty users, and is constantly busy. Over a three month period there, Cobot has counted over 550,000 separate events (roughly one event every eleven seconds).

As a regular of the Living Room over several months, Cobot has sought to engage other users. His social development can be divided into three distinct stages: inanimate object, social statistics engine, and conversationalist.

In the beginning, Cobot was socially inept: he sat in the Living Room and did nothing but answer one or two basic questions about why he was there. When spoken to in an unanticipated way, he did not respond. In other words, he was little more than a new piece of furniture. Not surprisingly, Cobot generated only a small amount of interaction.

In the next sections we explore the next two stages of Cobot's development and see how these changes impacted both Cobot's popularity and his environment.

Social Statistics

Previous work on agents in MUDs (Foner, 1993; Foner, 1997) has argued that being able to provide information of interest or value to other users aids the social acceptance of the agent. Because Cobot is intended to be primarily a social creature, we chose to have Cobot build and maintain what might be thought of as a *social map* of user interactions in LambdaMOO. In particular, Cobot maintains:

- For each user he encounters:
  - a histogram of verbs used by that user
  - a histogram of verbs that have been used on that user
- For each verb invoked in his presence:
  - a histogram of the users that have invoked it
  - a histogram of the users that have been its target
- For each pair of users Cobot has seen interact:
  - a histogram of verbs they have been used on each other

For both technical and ethical reasons, this information is gathered only for objects and users that are in Cobot's presence. The details of acquiring such information reliably are fairly straightforward, but beyond the scope of this paper. For a discussion, we refer the reader to (Curtis, 1997).

These statistics define a rich graph of social interactions. For example, it is possible to determine which users interact

[3] Cobot has visited about 1070 rooms.

tell me about *verb*-o-meter	Lists which users use (and get used by) *verb* the most.
tell me about me	Tells the questioner various facts about herself, including the verbs she like to use most, and the verbs most often directed at her.
who are your playmates	Lists those who interact with Cobot the most.
who loves me	Tells the questioner those with whom she interacts the most.
who acts like me	Tells the questioner users who perform actions similar to hers.
who does not act like me	Tells the questioner users who do not perform actions similar to hers.
relate me to *user*	Indicates how the questioner interacts with *user*, providing a ranking on each other's list of playmates, how many playmates they have in common, and how similarly they act.
tell me my *verb*-buddies	Tells the questioner the users she uses *verb* on and who use *verb* on her.
who is spammy	Lists those users who generate and are the targets of the most verbs.

Table 1: **A partial list of social statistics questions that may be asked of Cobot.** Here, the term "act like" refers to distribution of verb usage. The similarity measure is the cosine between verb count vectors (based on a standard information retrieval technique).

with one another the most, who the most "popular" users are, and the types of actions any given user tends to use.

Using this information, Cobot is able to answer natural-language-like queries about social interactions in the Living Room. For example:

HFh [to cobot]: relate me to Buster

cobot whispers, "Here are your relationships with Buster. You like to use: - (62%), poke (7%), hug (3%), eye (3%), nod (2%), hi5, h5, zap, comfort, and grin on each other. Buster is ranked #14 on your list of playmates. You are ranked #1 on Buster's list. Your socializing overlap is 75.4% and your playmate overlap is 33.7%. Your actions have a similarity of 95.9% but ignoring common speech verbs it's 58.3%. Others act on you with a similarity of 96.6% but ignoring common speech verbs it's 81.9%."

This particular query yields a great deal of information. Cobot first reports how and how often HFh and Buster interact with one another (the "-" verb is one method for directing speech; the first line above is generated by HFh typing "-cobot relate me to Buster."). Cobot then provides measures of similarity of social circles and similarity of action. Each of these measures is calculated using a cosine measure (a common information retrieval technique) between histograms of action use or inter-user interactions. Table 1 lists more of the queries that can one can make of Cobot.

After a month of gathering statistics, the social statistics query feature of Cobot was made available to LambdaMOO users. As the Figures show, the results were immediate and dramatic. Initially, before he had any real functionality, interaction with Cobot was constant but low (three per 1000 events). After the introduction of his new abilities, the number of interactions directed at Cobot jumped significantly (now over 50/1000 events, more than double that of the next most popular Living Room denizen).[4]

While these graphs quantify the sudden rise in Cobot's popularity, they cannot express the extent to which he altered (for better or worse) social interaction in the Living Room. Users began to converse with each other on what they were learning about their relationships and similarities:

Snow_Crash [to Medb]: Cobot says you act like me. Stop.

Medb cracks up laughing at Snow_Crash!

Medb [to Snow_Crash]: How do you know it's not you acting like me?

Medb tries to decide whether she or Snow_Crash should feel more insulted....

...

Gabaldon [to cobot]: loudly relate me to sparklebug

cobot [to Gabaldon]: Here are your relationships with Sparklebug ... Sparklebug is ranked 7 on your list of playmates. You are ranked 19 on Sparklebug's list. ...

Gabaldon [to Sparklebug]: I can't play with you so much anymore, you do realize. You don't like me nearly as much as I seem to like you.

Gabaldon cries at Sparklebug.

...

Spitboy has a new strategy. He stops by the Living Room, asks the cobot who's like him and if cobot tells him one of the people present is like him, he starts talking to them.

Spitboy says, "I haven't caught anyone who's like me, yet."

Spitboy [to Kimmy]: Shame, innit?

Kimmy says, "sure."

Kimmy [to cobot]: who is like me?

After observing these responses to the information provided by Cobot, we implemented changes to encourage such behavior. For example, as seen earlier in HFh's request about Buster, Cobot provides relative rankings of interactions between users. Players seemed appreciative of the ability to gain deeper insight into their interactions, and used the information accordingly. Access to social statistics resulted in other changes as well, some quite unintentional. For example, users began actively competing to move up on Cobot's list of playmates, or to be ranked the "most hugged."

cobot [to Stevage]: Here are my favorite playmates: Sparklebug, HFh, ...

Sparklebug beams brightly. Stevage might need some shades for that.

Stevage winks to cobot.

cobot winks to Stevage.

Sparklebug [to Stevage]: I've worked SO hard to beat HFh!

These social consequences were beyond our expectations, sometimes so severe that we needed to implement several social controls (as we shall see later). The important point is that Cobot provided a useful service (as have other agents, such as Julia) that captured information of constantly evolving and long-term interest to his community. The result was that Cobot had an immediate and profound impact on the social interactions within his sphere.

Chat and Emote Abilities of Cobot

Cobot's social statistic functions are what many Living Room denizens might view as his "purpose," or the inter-

[4] User HFh (an author) also experienced a sudden equal jump in interaction. As Cobot's primary human ambassador at the time, he spent much of that period answering questions, explaining functionality, and fielding requests for new abilities.

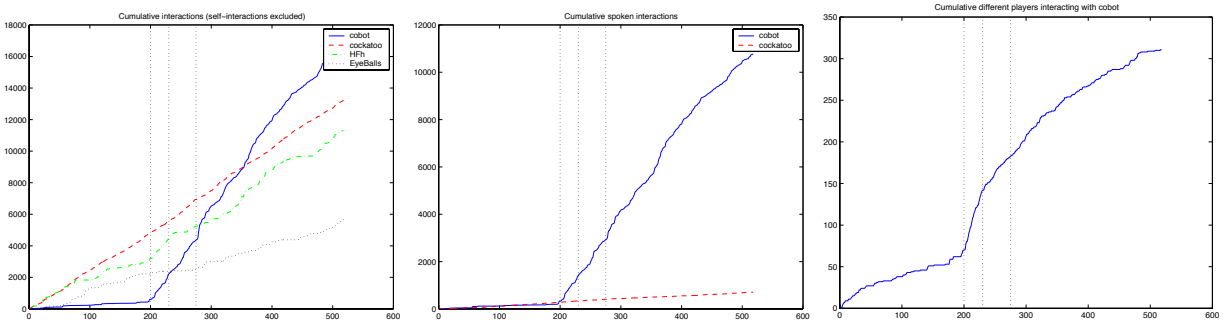

Figure 1: **Cumulative interactions with objects in the Living Room.** (a) Cumulative number of verbs (speech verbs, hugs, waves, etc.) directed towards various Living Room denizens: Cobot, the Cockatoo, and the two human users most interacted with during this period. The *x*-axis measures cumulative events (in thousands) of any type in the Living Room, while the *y*-axis measures cumulative events directed at the indicated user. Each dashed vertical line indicates the introduction of a major new feature on Cobot (from the left, his social statistics, his emoting abilities, and his extended chat abilities). A straight line—such as the Cockatoo's—indicates constant interaction. By contrast, Cobot's curve shows sudden changes in the slope coinciding with new features. Note that even when the slope levels off afterwards, it remains higher than it was previously, indicating long-term impact. Judging from cumulative interaction, Cobot is the most popular user in the Living Room. (b) Cumulative *speech* acts directed at Cobot and the Cockatoo. Clearly users interact with the two artifacts differently. Most of the interaction with Cockatoo is explained by users invoking its **gag** verb to silence it. Still, note that there is constant non-zero daily spoken interaction with it as well, indicating that users are willing to talk to even such a relatively simple creature. Also, note that before his new abilities surfaced, Cobot enjoyed a comparable amount of spoken interaction. (c) Variety of users interacting with Cobot.

esting service he provides. We follow researchers before us in believing that to truly integrate into an online society like LambdaMOO, Cobot must do much more. In particular, LambdaMOO is a world that is rich in both conversation and emoting, and in this section we detail our efforts to enable him to participate meaningfully in both.

Emoting. As users became more enamored or frustrated with Cobot's social statistics, they began to express their opinions via directing standard or personalized emotes such as hugging, waving, pouncing and kicking towards Cobot:

> cobot exclaims, "Hi! I'm awake!"
>
> Sparklebug [to cobot]: I've missed you SO much!
>
> Sparklebug snuggles with cobot.
>
> ...
>
> cobot [to Ratman]: Hello, Ratman.
>
> Ratman takes cobot over his knee and spanks its bottom.

Consequently, it seemed natural to equip cobot with the ability to respond to such actions in appropriately. Like others before us, we found it effective to program several plausible replies to any given scenario, and to randomize between the responses (including doing nothing). This is a simple technique for introducing variability in Cobot's behavior, and seems to enhance his apparent intelligence for many users.

> Melon hugs cobot.
>
> cobot hugs Melon warmly.
>
> EyeBalls kisses cobot lightly.
>
> cobot hugs EyeBalls warmly.
>
> ...
>
> cobot [to Ratman]: Welcome, special playmate of mine!
>
> Ratman spits at cobot
>
> cobot [to Ratman]: You'll be first up against the wall when the revolution comes.

Simple Chatting. Cobot's early chat abilities were of a similar nature, being randomized, canned responses to specific utterances. However, this is a more difficult problem because the variety of utterances directed at Cobot is much greater than the variety of emotes. As others have discovered, for certain classes of incoming utterances, it is helpful to respond with fairly general and leading conversational utterances, while other times it is possible to be rather specific:

> killa [to cobot]: how old are you?
>
> cobot [to killa]: I am older than you think.
>
> lackluster [to cobot]: why are you carrying a shield?
>
> cobot [to lackluster]: I think many of the dumber MOOers think I'm a garbage can.

Just as users expect to have their hugs to Cobot returned, his heaviest users enjoy having Cobot actually initiating recognition of them. We thus made occasional departures from our informal policy of having Cobot only respond to activity initiated by users, for instance by having Cobot initiate greetings to certain users who interact with him frequently, as in his (rebuffed) greeting to user Ratman above.

Extended Chat. So far we have described a limited set of chat and emote abilities, carefully crafted but based on tried techniques. Not long after Cobot began enjoying sustained usage, many users began to recognize some of his canned responses, and made comments longing for better chatting ability. At this point we inaugurated a novel experiment, in which some of Cobot's utterances were generated from a small number of large and publicly available online documents meeting the following desiderata:

- The sentences could plausibly serve as utterances by Cobot, at least in some contexts.
- The documents should be sufficiently large that repeated utterances will be unlikely.
- The documents should be on fairly specific and recognizable topics. Rather than designed to enhance general conversational ability (which remains handled in part by the mechanisms already described), the utterances generated by the documents are meant to give Cobot "personality".
- The utterances generated by the documents are entertaining and sometimes controversial (to spur further conversation).

Any utterance directed towards Cobot that is not recognized as a request for social statistics becomes a candidate

Guest [to cobot]: Do you know anyone I should send this home vasectomy kit to?

Sparklebug peers at Guest suspiciously.

cobot [to Guest]: This is a joke in very poor taste.

Guest [to HFh]: He's good.

Guest [to HFh]: Have you been tinkering with him?

HFh says, "I wish I could take credit, but I think he's evolving."

Guest [to cobot]: You're evolving?

cobot nods to Guest.

cobot says, "CITIZENS: ACIDHORSE DOES NOT EXIST."

Guest [to cobot]: he does so, I'm sure I've seen him

cobot [to Guest]: What does this prove?

Guest [to cobot]: my hovercraft is full of eels

Guest [to HFh]: He just made perfect sense. damn him.

cobot [to Guest]: Heavens!

...

deLaMer says to HFh, "you did good work on cobot."

HFh [to deLaMer]: Thank you.

deLaMer says, "it's starting to make sense. that's spooky."

Blue_Guest [to HFh]: I thought it was a person once.

Technicolor_Guest says, "cobot is an evil mastermind lurking behind the false guise of a mere bot"

Guest comes out of the closet (so to speak...).

cobot [to Guest]: Sunshine! You validate me!

Guest [to cobot]: nice.

cobot [to Guest]: We've had a nice snooze.

Guest [to cobot]: so more charitable than usual?

cobot [to Guest]: But my dreams are a lot emptier than yours.

Fawn_Guest [to cobot]: They can't possibly be any emptier than mine

Guest [to cobot]: oh.

cobot [to Guest]: Your cry for help is transparent.

...

Stevage says, "cobot is pretty damn human I have to admit... I asked him 3 questions and he gave sensible replies to them all."

...

Valley [to HFh]: cobot's genuinely cool...and I'm not even a techie :)

Guest says, "Cobot frightens and confuses me, yet I cannot turn away"

HFh [to Inimeg]: He *is* a player.

HFh [to Inimeg]: He has always been a player.

Inimeg says, "damn the line just gets blurrier and blurrier"

...

A better, kinder Ratman puts on his Ratman's Official Church of Cobot Membership Badge.

Table 2: **Additional Cobot commentary by LambdaMOO users.**

for the following process. Words in the incoming utterance are matched to words appearing in sentences in the documents, assigning to each sentence a weight based on the number of matching words. A minimum length on the words matched provides a parameter that can adjust the difficulty of matching sentences in each document independently. If no match is found, the utterance is handled by one of the mechanisms mentioned earlier; otherwise, with a (large) fixed probability Cobot randomly chooses a sentence to utter according to the distribution defined by the weights.

The effectiveness of this approach depends on the choice of documents. We first fielded the well-known Unabomber Manifesto by Theodore Kaczynski, a lengthy, didactic essay on the evils wrought on society by science, and the industrial revolution. The resulting response was immediate and dramatic, resulting in a steep rise in experimentation and interaction with Cobot (note the rightmost dashed vertical line in Figure 1). The new functionality resulted in endless conversation, both with and about Cobot, and allowed users to attribute to him all manner of intelligence and attitudes:

Technicolor_Guest [to cobot]: What's your religion?

cobot [to Technicolor_Guest]: technology unless you REALLY believe in that religion yourself and find that it arouses a deep, strong, genuine response in many other people.

...

cobot [to DragonBoi]: The line of conflict should be drawn between the mass of the people and the power-holding elite of industrial society.

DragonBoi eyes cobot warily.

cobot [to DragonBoi]: You have something cute in your eye.

DragonBoi [to HFh]: you're bot is turning into a communist.

DragonBoi [to cobot]: so what are you gonna do about it?

DragonBoi [to HFh]: did you see what it told me? you are turning it into a communist!

Wikked [to cobot]: are you a communist?

DragonBoi [to cobot]: are you now, or ever been a member of the communist party?

cobot mumbles something in binary.

DragonBoi [to cobot]: so you're taking the fifth are you? well this is a special commision, that won't work here!

...

cobot [to HFh]: The modern individual on the other hand is threatened by many things against which he is helpless; nuclear accidents, carcinogens in food, environmental pollution, war, increasing taxes, invasion of his privacy by large organizations, nation-wide social or economic phenomena that may disrupt his way of life.

HFh [to cobot]: I agree with that position, I suppose, but it doesn't help to explain the lack of a safety net.

cobot [to HFh]: Nevertheless, the position can be defended.

HFh [to cobot]: ...in the *short* term only.

cobot nods to HFh.

Cobot's replies are not always appropriate; that they are often appropriate and frequently amusing suffices to intrigue and entertain users. User seem to ignore the senseless interactions with Cobot and remember the "intelligent" ones. Certainly the open-ended nature of LambdaMOO, and the fact that few users have concrete, goal-oriented expectations of the other users or artifacts they encounter encourage this.

Other documents include scripts of "The Big Lebowski" and "Planet of the Apes"; a long analysis of the JonBenet Ramsey murder case; and articles on alien abduction and paranormal experience. It is simple to instantly swap new documents in and old ones out, limiting Cobot's chat abilities only by the imaginations of his creators. Still, we emphasize our belief that it is advantageous to limit the number of documents at any given time, and for these documents to be reasonably specific, so that well-defined "themes" can emerge in Cobot's conversation. Note that Cobot has no notion of conversational state at this point, so utterances in any given conversation may be drawn from multiple documents, as in the following exchange alternating between the Unabomber Manifesto and "Planet of the Apes":

cobot [to cyberat]: There is good reason to believe that many primitive societies subjected people to less pressure than the European society did, but European society proved far more efficient than any primitive society[....]

cyberat [to cobot]: Thats what they want you to think!

cobot [to cyberat]: Tell me who and what you really are and where you come from, and no veterinarian will touch you.

cyberat [to cobot]: Veterinarians!!!

cobot [to cyberat]: The Roman system of urban sanitation was forgotten, so that until rather recent times did the sanitation of European cities that of Ancient Rome.

Fawn_Guest [to HFh]: cobot's conversation topics just get weirder and weirder

cobot [to cyberat]: In spite of every sign I'm an intelligent being meaning no harm, you continue to hate and fear me.

Privacy and Social Issues

LambdaMOO has policies researcher, acquiring permissions from users, and quoting material; however, as in any complex culture, there are unofficial, unwritten rules that must be learned. Much of the work in introducing Cobot into LambdaMOO involved learning and understanding these rules, and is an ongoing concern. We began our education by discussing with colleagues who were LambdaMOO users, and could provide insight into its social dynamics. HFh spent several months as a LambdaMOO user, building genuine friendships, learning about etiquette and privacy, and floating the idea of an agent like Cobot. Only after convincing ourselves that Cobot would not be completely unwelcome did we introduce him into the Living Room.

Privacy. The Living Room is a *public* space, not a private room. Possibly intrusive objects are more likely to be tolerated, as long as users do not believe they are being recorded in some way for unfriendly purposes. Furthermore, most users are sophisticated enough to understand the tenuous nature of privacy in such public spaces. On the other hand, most users will *not* simply reveal their real-life identities to casual acquaintances in the MOO, nor do they expect others to reveal such information. Early in the project, several MOOers raised questions around these issues, sometimes in jest, sometimes quite seriously. In keeping with the goal of social acceptance, Cobot is fairly conservative. He notes only events in his presence, and does not share events verbatim. Furthermore, a questioner can generally only ask about herself, and not directly about others, and all responses are whispered only to the questioner unless the questioner explicitly asks to share the information with the room.

Spam. Being in a public place, Living Room regulars are more likely to tolerate spammy objects, but there are clear limits. Because the goal of Cobot is to be a part of his social environment, it is important that he not cross this line. Thus, Cobot's tendency to whisper answers to requests is not motivated just by privacy, but by the desire to be less spammy. Similarly, Cobot generally does not speak unless spoken to, except under certain conditions, as discussed earlier. These design decisions are intended to give users a modicum of control over the nature and rate of Cobot's output.

Nevertheless, these precautions sometimes proved inadequate. For example, users expect to have the ability to silence non-human objects. So, Cobot has a **silence** verb that allows a user to prevent him from speaking out loud (to anyone) for a random amount of time. Cobot may still respond to users, but he will only whisper to his intended target.

Still, it is no more possible to prevent users from using Cobot to generate spam (*e.g.* by repeatedly querying him aloud) than it is to prevent users from generating spam directly. It is debatable whether such users or Cobot are to blame, but it can be irritating regardless. Cobot's various social statistics are a great incentive to interact in front of and with him perhaps more than one would otherwise, thus raising the overall level of noise in the Living Room.[5]

To combat this, we have implemented more drastic measures. Along with Cobot's "owners," any of a set of identified regulars can direct Cobot to ignore individual players. Cobot will not interact with such a player during these times, except to occasionally inform him that he is being ignored. These policies appear to be the minimum necessary. We are continually revisiting and updating these decisions.

Discussion

Cobot has become a member of his community. As we have seen, he is perhaps the most "popular" resident of the Living Room. Users engage him in conversation, interact with him in a variety of ways, take advantage of his statistical services and sometimes even have generally positive things to say about him (see Table 2). On the other hand, his entry has not been welcomed by all MOOers. Some complain of a general increase in spam, and others have noted that he has irreparably changed the nature of the Living Room.

Cobot's development continues. Although Cobot often has remarkably cogent conversations, most of the time his remarks are complete non sequiturs[6]. We plan to provide cobot with minimal state to help him be more focused and coherent. We have also begun experimenting with using reinforcement learning techniques to teach Cobot to learn how to act more independently and without prompting, while avoiding angering users. Users will be able to reward and punish Cobot for actions he takes, and he will learn to modify his behavior accordingly. For example, Cobot may decide when it is best to inject himself into a conversation, when to hug, wave to or comfort someone, or when to introduce similar users who do not regularly interact.

In general, we are interested in ways of allowing Cobot to integrate more fully and usefully into his environment. We are mindful, however, that like all good citizens, Cobot should change his environment without degrading it.

References

Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7):122–125.

Cherny, L. (1999). *Conversation and Community: Discourse in a Social MUD*. Cambridge University Press.

Curtis, P. (1997). The lambdamoo programmer's manual v1.8.0p6. ftp://ftp.research.att.com/dist/-eostrom/MOO/html/ProgrammersManual_toc.html.

Dibbell, J. (1999). *My Tiny Life : Crime and Passion in a Virtual World*. Holt, Henry & Company.

Etzioni, O. (1997). Moving up the information food chain: deploying softbots on the world wide web. *AI Magazine*, 18(2):11–18.

Foner, L. (1993). What's an agent, anyway? a sociological case study. Technical report, MIT Media Lab.

Foner, L. (1997). Entertaining agents: a sociological case study. In *Proceedings of the First International Conference on Autonomous Agents*.

Mauldin, M. (1994). Chatterbots, tinymuds, and the turing test: Entering the loebner prize competition. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*.

Mitchell, T., Caruana, R., Freitag, D., McDermott, J., and Zabowski, D. (1994). Experience with a learning personal assistant. *Communications of the ACM*, 37(7):80–91.

[5] There is some evidence that Cobot's presence may have raised the overall amount of interaction in the Living Room.

[6] Some users tailor their conversation with Cobot in order to keep him "on topic" while others revel in his lack of understanding.