

# Cobot: A Social Statistics Learning Agent

Charles Isbell  
AT&T Labs-Research

*Christian Shelton, Michael Kearns,  
Satinder Singh, Peter Stone  
Olufisayo Omojokun, Diane Litman*

# Conclusions

Probabilistic AI methods are ripe for application to human-computer and human-human settings, even strongly social ones.

Real worlds that are human-centric

Sensing is easy

Socially Risky but not physically risky

# The Talk

- Quick Tour of LambdaMOO
  - *"real" human-human interactions*
  - *"social" users*
  - *text modality*
  - *low expectations*
  - *sensing easy*
- Cobot: A Social Statistics Agent
- Cobot's Impact
- CobotRL: Social Reinforcement Learning
- CobotDS: A Novel Dialogue System

## LambdaMOO: Whirlwind Tour

- Multiuser, text-based virtual world
- Chat channel: directed speech and "emotes"
- Users create rooms, objects, behaviors
- Founded 1990; >5K users, >118K objects, >2k rooms
- Technologically sophisticated population; history of AI experimentation
- Many academic and popular articles, books
- Strong sense of community
- Activity not easily characterized

# Sample Dialogue

**HFh** waves to Buster.

**Buster** bows gracefully to HFh.

**Buster** is overwhelmed by all these paper deadlines.

**Buster** begins to slowly tear his hair out, one strand at a time.

**HFh** comforts Buster.

**HFh** [to Buster]: Remember, the mighty oak was once a nut like you.

**Buster** [to HFh]: Right, but his personal growth was assured. Thanks anyway, though.

**Buster** feels better now.

Standard verbs and emotes: directed and broadcast speech, hug, wave, bow, nod, peer, poke, zap, grin, laugh, comfort, sup, chill, hpound, ...

# Cobot

- Has user status, but known to be a bot
- Resides in LambdaMOO Living Room
  - Popular social spot; guests and new users
  - 1 event every 11 seconds or so.
- Primary functionality:
  - **extensive observation and model building**
  - social statistics and queries
  - emote and chat abilities
  - reinforcement learning
  - CobotDS
- Mix of science, technology, sociology, entertainment

# Social Statistics

- For each user, histograms of:
  - verbs invoked **by** user (communicative style)
  - verbs invoked **on** user (communal treatment)
  - histogram of **other users** (social circle)
- For each verb, histograms of:
  - invoking users
  - target users
- Comparisons via standard cosine measure
- **Queries allowing (limited) access to stats**

wave buster →

**HFh** waves to Buster. →

**HFh:wave(Buster)**

# Can You Relate?

**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 similarity of 96.6% but ignoring common speech verbs it's 81.9%.``

**HFh** [to cobot]: hug-o-meter

**HFh** [to cobot]: who is like me?



# Social Warping

**Rain** [to Cloud]: Cobot says you act like me. Stop.

**Cloud** cracks up laughing at Rain!

**Cloud** [to Rain]: How do you know it's not you acting like me?

**Cloud** tries to decide whether she or Rain 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.

# More Warping...

**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?

**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!

# Chat and Emote of Cobot

- A consequence of being "interesting"
- Appropriate social basics hand-coded but unavoidably narrow
- Early Cobot: hand-coded chat replies via Eliza-like pattern matching + randomization
- Scaling up chat:
  - randomized pattern matching on **large** documents
  - topic **specificity** gives ascribed "personality"
  - variety and focus, sidestep NL generation
  - Documents: Unabomber Manifesto, Planet of the Apes, current political articles,...
  - Easily swapped in and out

# The Power of Randomization

**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 and always won out in conflicts with such societies because of the advantages conferred by technology.

**cyberat** [to cobot]: That's 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]: And in spite of every sign that I'm an intelligent being who means no harm, you continue to hate and fear me.

# The Assisted Turing Test

**cobot** [to Bah]: Women are encouraged to have careers because their talents are useful to the system and, more importantly because by having regular jobs women become better integrated into the system and tied directly to it rather than to their families.


**Guest** [to cobot]: Yes. Capitalistic propaganda. Fight the power!

**DangerMouse** [to cobot]: thus creating a vacuum and the eventual dissolution of the family unit

**cobot** [to Guest]: It's within your power.


**Guest** [to cobot]: No, I'm poor.

# Popularity



Cumulative Interactions: Population

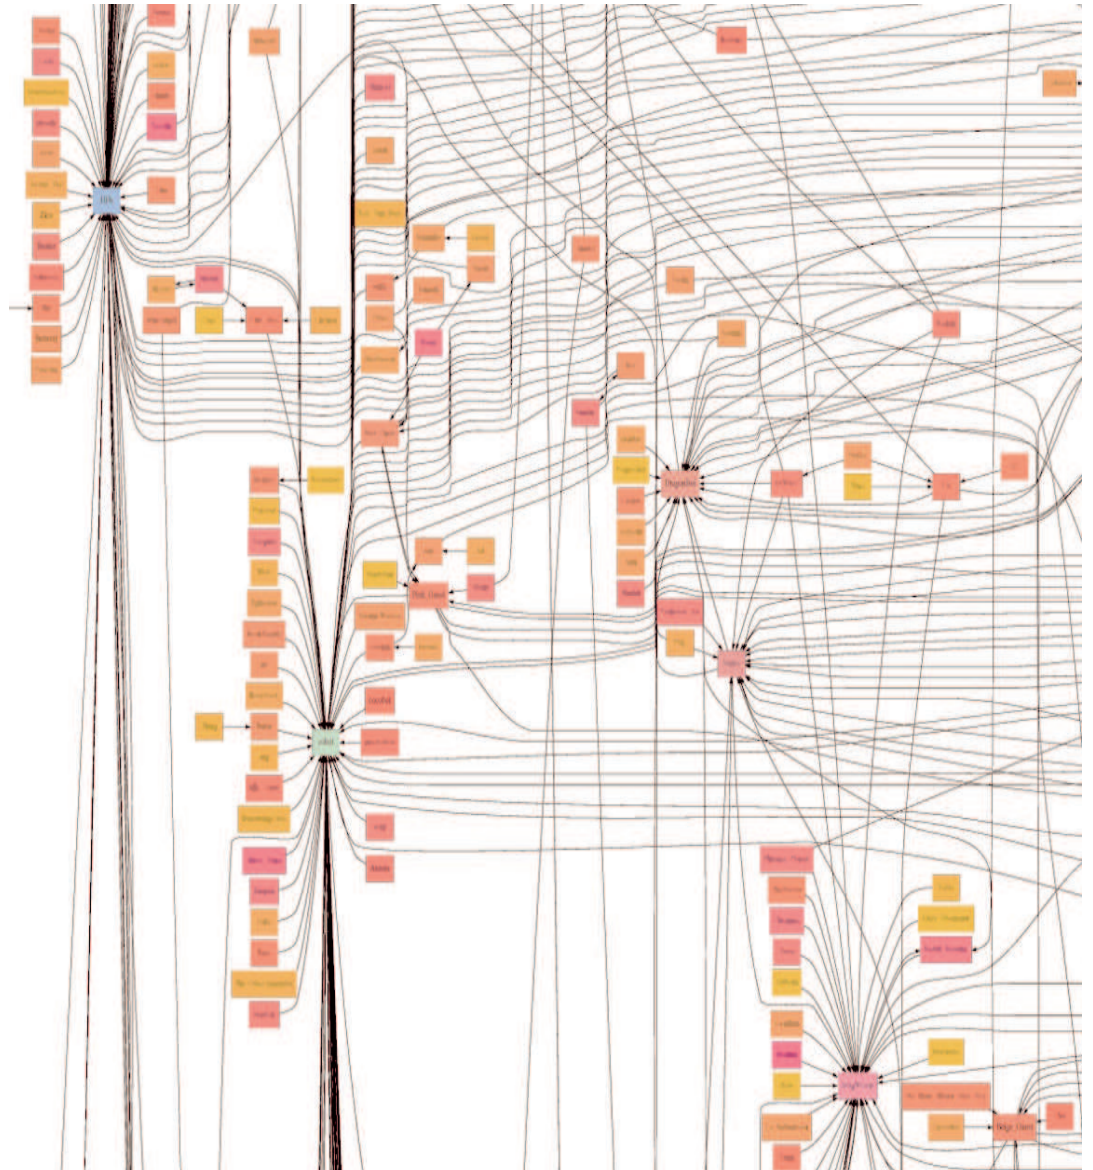
# Patterns of Interaction



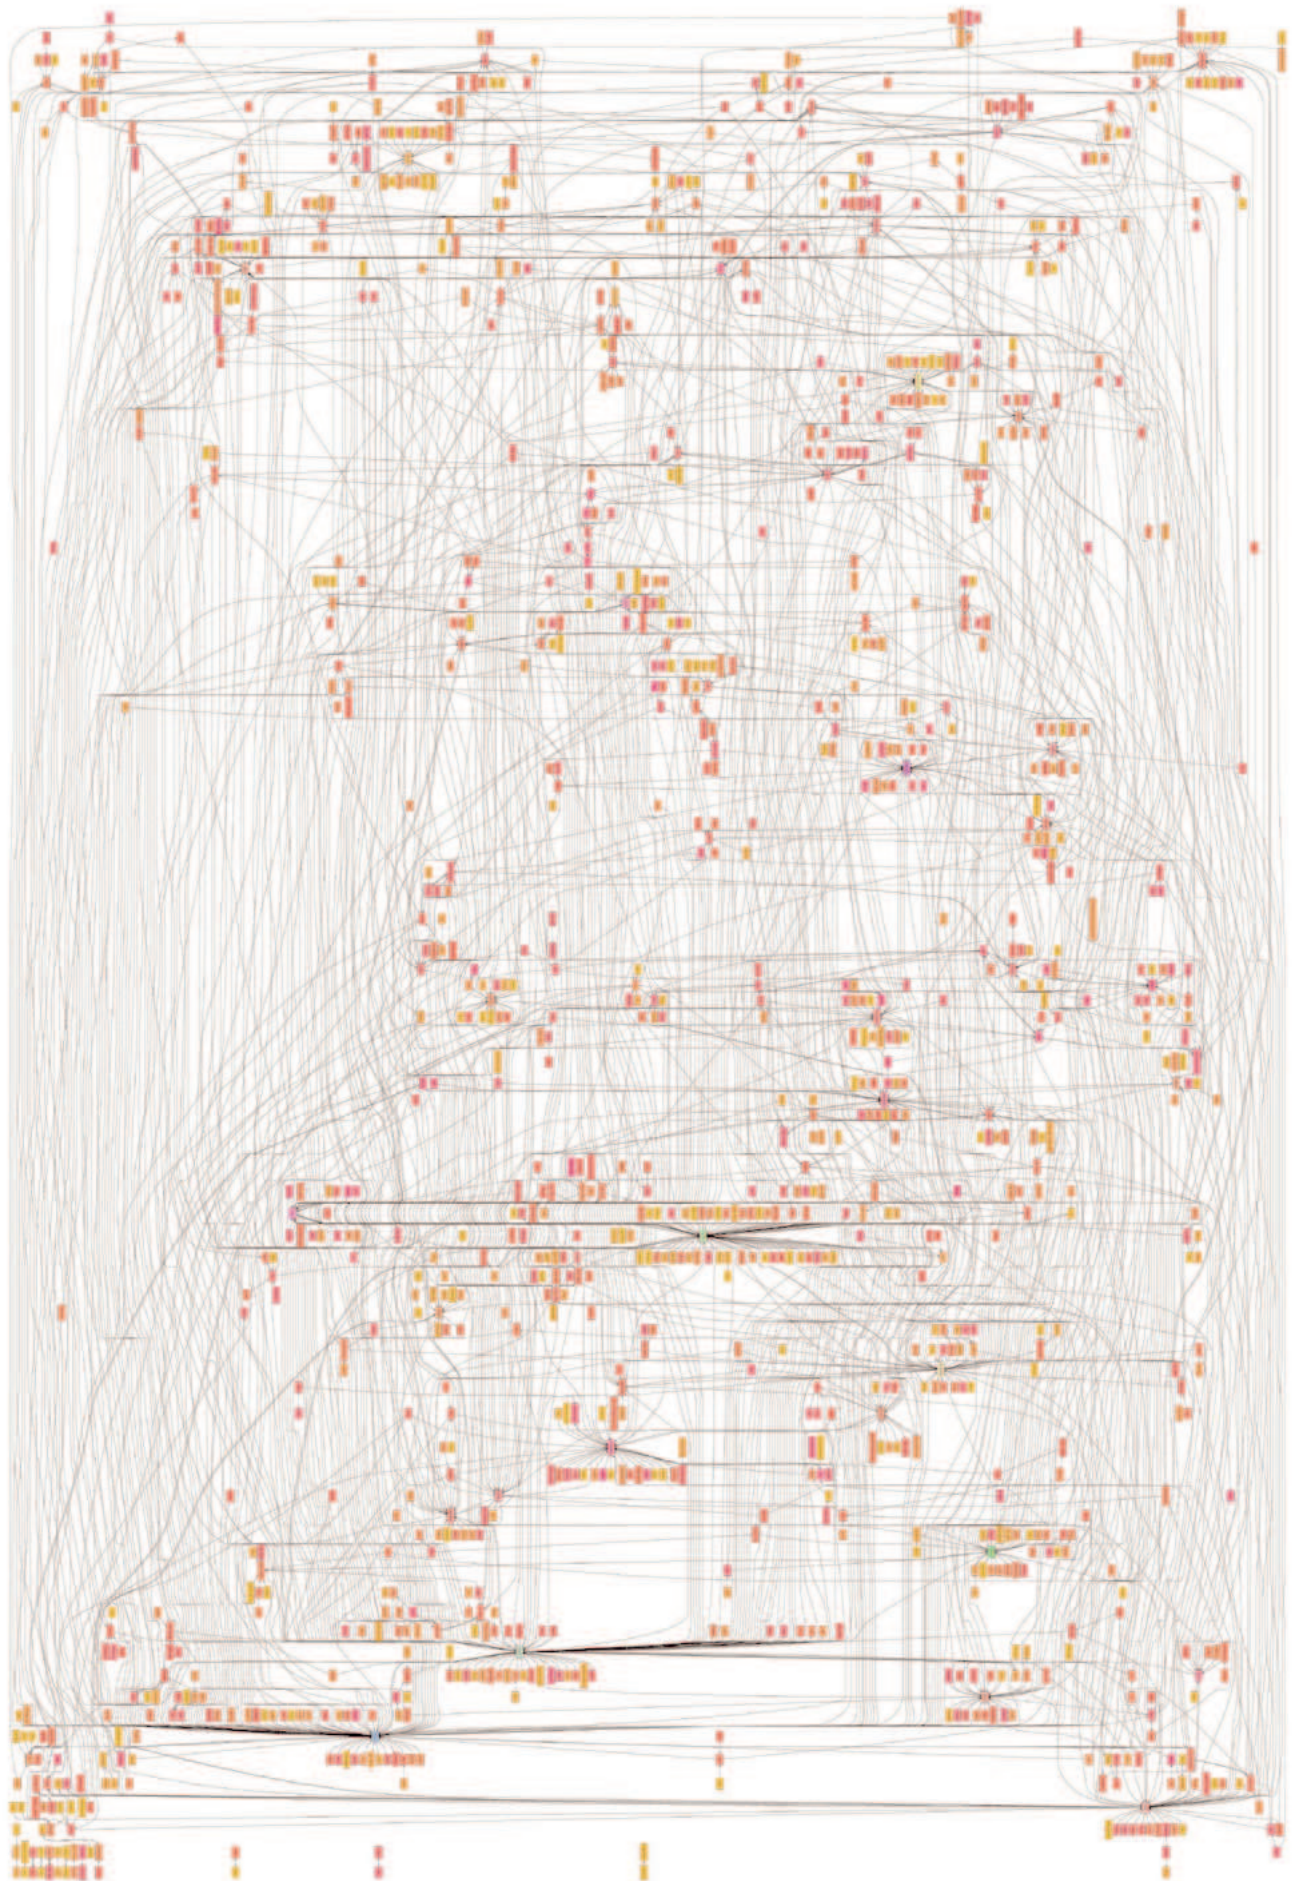
Cumulative Interactions: Individuals

# A Fragment of Cobot's Social Map

- A Graph of Users:
  - Each node is a user
  - Similar colors represent similar communication styles
  - An edge from User A to User B means that B is in the top 25% of A's interactions
- Allows computation of social cliques, six degrees of separation







# Learning to be Proactive

- Original Cobot entirely **reactive**
- Want **proactive** agent that learns from **users**
- An experiment in multi-user personalization
- AI methodology: **reinforcement learning**:
  - define Cobot's **state** and available **actions**
  - Cobot **experiments** with actions from various states
  - Users invoke **reward** and **punish** verbs on Cobot
  - Goal: map states to actions to maximize reward (?)
- **Challenges**:
  - Markov violations
  - Data sparsity
  - Inconsistent and drifting rewards
  - Multiple users: incomparable rewards, interpretation in presence/absence of other users

# Social Learning

- **Reactive**

- Original Cobot reactive using hand-programmed rules
- Learn from *observation* of users in order to self-program appropriate verb usage
- Extend the map to include time

- **Proactive**

- Define complex actions that Cobot may inject into the environment on his own
- Learn from user *feedback* best actions to take under different circumstances
- An experiment in multi-user personalization and child rearing

# Reinforcement Learning


- Methodology: **reinforcement learning**:
  - A method for learning action choice from delayed reward
  - Define Cobot's **state** and available **actions**
  - Cobot **experiments** with actions from various states
  - Users invoke **reward** and **punish** verbs on Cobot
  - Goal: map states to actions to maximize reward
- Challenges:
  - Markov violations
  - Data sparsity
  - Inconsistent and drifting rewards
  - Multiple users/sources of rewards
    - incomparable rewards
    - interpretation in presence/absence of other users
    - users are playing a "game" against each other

# Social Engineering

- **RL actions:**
  - be quiet
  - suggest topic
  - read Boston Globe
  - roll calls
  - social commentary
  - social introductions
- **RL state features (per user):**
  - user activity rates
  - room activity rates
  - "familiarity" measure
  - mood vector
  - current room
- Explicit and implicit reward mechanisms
- **Risky** actions, **social** features, **open-ended** feedback
- Treat each user as a **separate RL problem**
- Five months of data (May-Oct)
- >63K RL actions, >3K feedback from 254 users


# Difficulties with Average Reward

average reward



# Learning a User's Preferences

Rewards / Policy / Empirical Distribution Comparison for User M



## Features Matter

- Measure feature deviation from uniform weights to determine relevance
- Different features matter for different users
  - Roll calls: roll call parties, repeated roll calls
  - Room rates relative to the user
  - Familiarity: ignore Cobot when with new friends, reward Cobot when with common friends
  - Room



# Calling Cobot

- Integrate our dialogue system work
- Toll-free number with basic emotes, "say" and "listen" commands, info commands,...
- CobotDS goals:
  - Telephone access to Cobot & LambdaMOO
  - Blurring the "real" and the "virtual"
  - Tackle many user interface issues (multiparty dialogue, comm. imbalance, summarization, personalization)
- Fielded September 2000

# CobotDS: A Novel Dialogue System

- Apply AI techniques where the assumptions fail
- Integrate our dialogue system work
- Blurring the "real" and the "virtual"
- Issues:
  - "real" vs. "virtual"
  - multiparty
  - bandwidth imbalance
  - dynamic DB
  - lag
  - summarization
  - large vocabulary ASR
- Details/functionality:
  - Phone # to Cobot
  - phone: basic emotes
  - info commands
  - "say" command
  - "listen" command
  - multiple personalized grammars
  - MOO: unrestricted
  - Fielded 9/00

*Hi, Buster. What do you want to do?*

Wave.

*You wave.*

Who.

*You who. I am with Cocobot, Betwixt, and HFh. HFh says How was the movie?, to you*

Say HFh.

*What message would you like to pass?*

Fantastic.

*You say Fantastic, to HFh. Betwixt waves to you.*

Summarize.

*You summarize. Recently, HFh and Betwixt chatted. Betwixt and Natto bickered. HFh made the most noise.*

Grammar.

*Which grammar would you like to use?*

Personal.

*The grammar is now set to Buster.*

Say.

*What message would you like to pass?*

I am in Hoboken.

*You say I am in Hoboken, to Betwixt.*

Listen.

*You listen. Betwixt gives you a nod. HFh to Betwixt, Sprewell will go to the rim, but Houston settles for jumpers from the parking lot, & then, I grin to HFh.*

**HFh [to Betwixt]:** And thanks to TiVo, I was able to see the game when I got home.

**Betwixt [to HFh]:** The second half was pretty spectacular.

**Cobot turns to pick up the phone. Cobot begins talking to Buster!**

**Cobot holds up a sign: Buster passes on a wave from the phone.**

**HFh [to Cobot]: phone: How was the movie?**

**Betwixt [to Cobot]: phone: wave**

**Cobot [to HFh]: Fantastic**

**Cobot [to HFh]: That was from Buster.**

**Cobot holds up a sign: Buster says, 'I am in Hoboken' from the phone.**

**Betwixt [to Cobot]: phone: nod**

**Cobot holds up a sign: Buster passes on a listen from the phone.**

**HFh [to Betwixt]:** Sprewell will go to the rim, but Houston settles for jumpers from the parking lot.

**Cobot grins at HFh.**

**HFh [to Betwixt]:** With Camby's rebounding they have a chance to at least win the East.

**Betwixt [to HFh]:** Good point.

## CobotDS Observations

- 18 non-authors have made calls
- Calls for some users averaged 20 minutes, occasionally last for 45 minutes or more
- Phone users vary use of speech/verbs
- Interaction included a variety of LambdaMOO users
- Average lag ten seconds, filtering rare, socially filtered text allowed 70% of unfiltered length
- Personalized grammars used early and often, evolved to deal with ASR errors
- Switching in regimes from radio/listen/interaction
- Users liked system, hated ASR, wanted more personalization

# Summary

- **A significant member of a human community**
  - Fielded in a complex social environment
  - Projection of and desire for intelligence
  - Interesting "service", apparently complex chat, observance of social niceties/privacy
- **Social statistics are useful**
  - Popular
  - Information and insight
  - Inferring preferences (actions, dialogue)
- **RL in a multi-user human social environment**
  - Proactive learning
  - Emotion (*I'm sad, but things will get better*)
- **Multi-party asynchronous dialogue**

# Conclusions

Probabilistic AI methods are ripe for application to human-computer and human-human settings, even strongly social ones.

- Learned behaviors
- Feature selection
- User modeling/Clustering
- Discovering Intent
- Natural Language
- Multiple agents
- Graph fun
- Evaluation metrics

<http://cobot.research.att.com>

# Privacy and Social Issues

- LambdaMOO policy and etiquette
- Socialization by HFh
- Acceptance of feedback and suggestions
- Cobot's conservatism: initiative; limitation of social queries to questioner; whispering
- User-induced spam: gagging and silencing
- Continual revision and updating