CS 395-22

Conversational Techniques
Currently Used
in Computer Games

Spring 2005

Overview

- What are conversations for?
- Dialogue Trees
- Restricted languages
- Keyword spotting
- Chatbots
  - Basic patterns
  - Statistical techniques

We are magicians

- Goal: Enhance the player's experience via illusion

- As engineers, we are trying to deliver the intended experience, within the budgetary limitations of the design and construction process

Factors working against us

- Conversational NPCs lack world knowledge, experience, vocabulary, communication skills that we expect in other humans

- Game development takes place within very tight budget restrictions
 - Money, time, %CPU, memory

Factors working for us

- Players want to be deceived
- People tend to enter social relationships with things
  - cf. Reeves & Nass, The Media Equation: How people treat computers, television, and new media like real people and places
  - cf. Lester *et al* The Persona effect: Affective impact of animated pedagogical agents *Chi97*
- Games can provide tightly constrained contexts
  - Limits what makes sense to say, and how
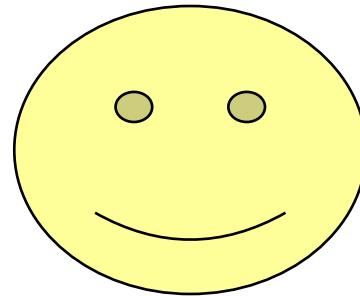
Why have conversations in games?

A prototypical situation

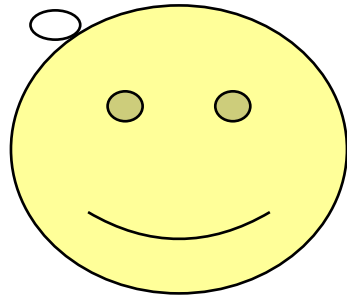- Player encounters NPC in game, and wants to communicate. How do we handle this?

???
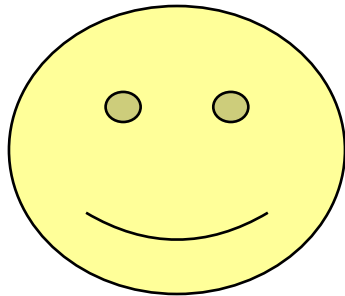
Player Character   Non-Player Character

The Basic Problem
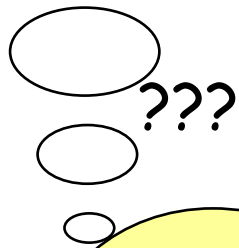
Player Character
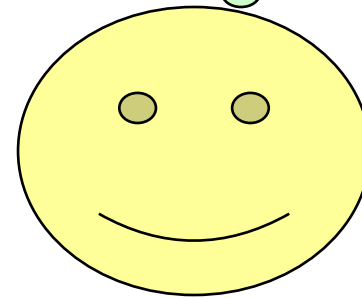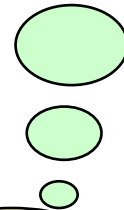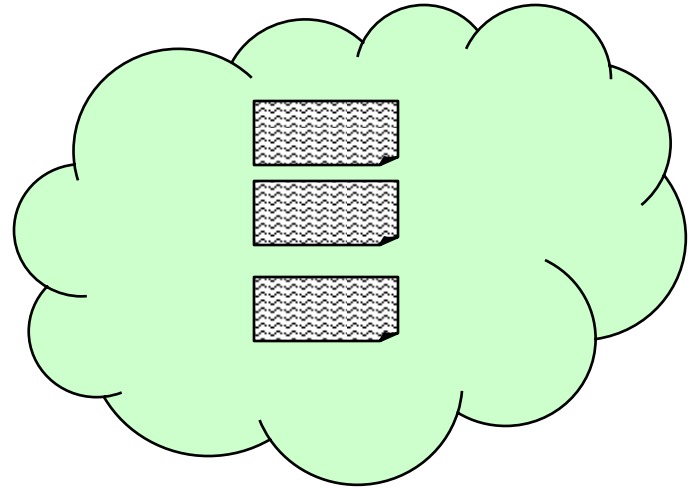
Non-Player Character

Habitability in Natural Language Interfaces

- Key principle: System should not produce language that it cannot in turn understand
  - People naturally tune their language to match their communication partner
  - If system cannot understand things that it has produced, it is considered stupid

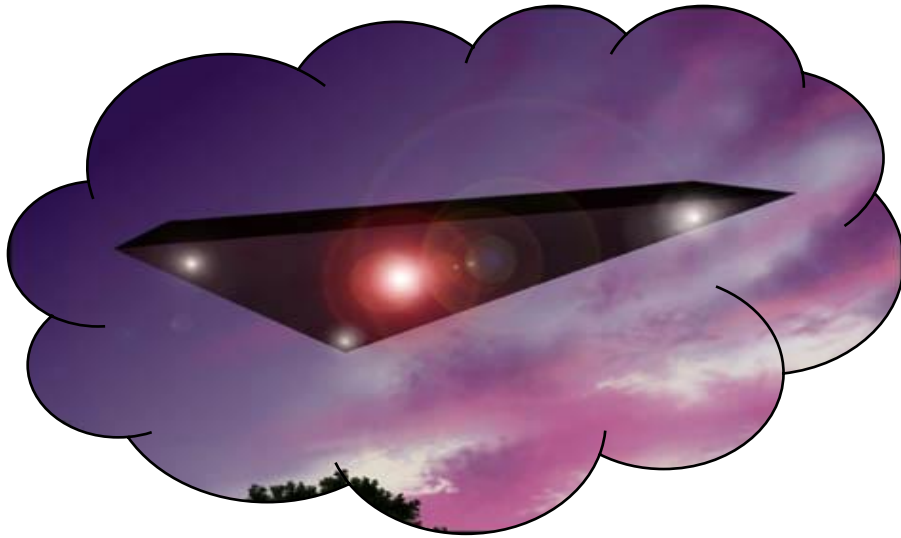# Problem: How do players know what they can say?

- Solution 1: Tell them a restricted set of choices

- Solution 2: Give them a restricted language to express themselves in

- Solution 3: Work hard to interpret what they say

Dialogue Trees

- Forget language processing, just use menus
  - Player gets list of choices
  - Response player gets depends on what they picked + game state
- Many embellishments
  - Dialogue can be re-stated in player character's voice
  - Not really trees, but directed graphs: Can include joins, loops
  - Menu options conditional on game state, including player attributes
- By far the most common solution

Advantages of dialogue trees

- Easy to construct
  - Writers can create them with minimal to zero programming skills
- Designer has full control over the experience
- Low runtime budget
  - Just a small set of datastructures

Problems with dialogue trees

- Player loses sense of control
- Hand-authoring does not scale as conversations grow in length
  - Assume n branches per turn
  - For t turns, n^t-1 nodes to write, n^t distinct end states (if really a tree)

Restricted Languages

- Classic example: Interactive fiction
  - Ask Merlin about sword
  - Ask about babelfish
  - Tell zaphod about towel
- Query space limited by context of game
- Only need to find the right nouns

Keyword spotting

- Let user type free text
- Look for words that matter in the context of what the NPC has to communicate

could you tell me where the toilet is?
There aren't any. The designers forgot. You'll just have to cross your legs.

Chat-O-Mat

Pattern matching

- Use more of the context of the player's utterance
- An extremely crude, zeroth-order version of parsing

ELIZA: The progenitor

Men are all alike.

IN WHAT WAY?

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE?

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed much of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

It's true. I am unhappy

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

I need some help, that much seems certain.

WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

Perhaps I could learn to get along with my mother.

TELL ME MORE ABOUT YOUR FAMILY

My mother takes care of me.

FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU

My father.

YOUR FATHER

You are like my father in some ways.

WHAT RESEMBLANCE DO YOU SEE

You are not very aggressive but I think you don't want me to notice that.

WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE

You don't argue with me.

WHY DO YOU THINK I DON'T ARGUE WITH YOU

You are afraid of me.

DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU

My father is afraid of everybody.

WHAT ELSE COMES TO YOUR MIND WHEN YOU THINK OF YOUR FATHER

Bullies.

DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

How Eliza Works

- Keyword spotting + transformation rules
- Example: (??b YOU ?v ME) →
  (WHAT MAKES YOU THINK I ?v YOU)
  - ??b matches a sequence of words
  - ?v matches a single word
- If no keywords spotted,
  - Respond with one of a set of leading utterances
  - Return to some previous utterance
  - (Also known as *confusion scripts*)

Chatbots: The modern versions

- Mostly based on structure of Eliza, just bigger
    - Example: ALICE
        - http://www.alicebot.org/
        - Tons more transformation rules
        - Coverage updated by analysis of logs kept by the system
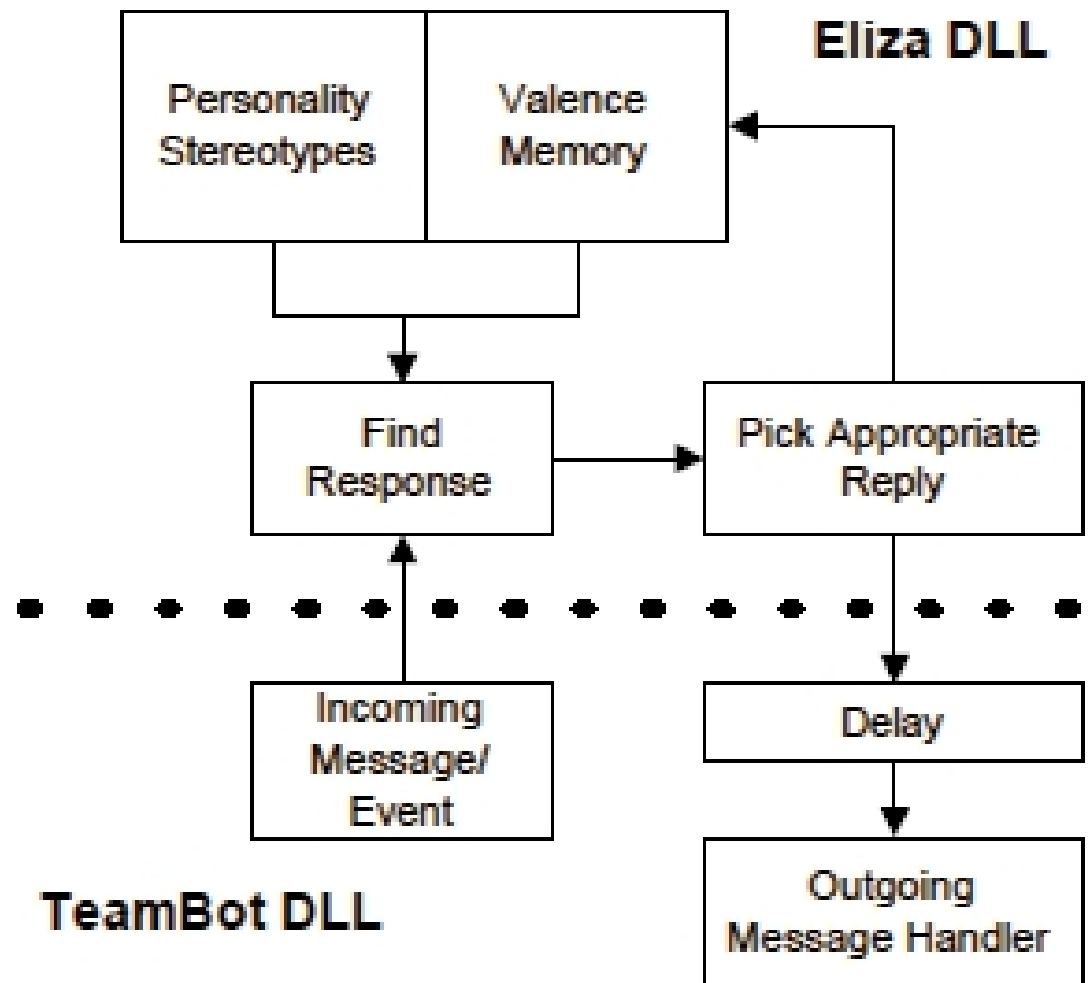   - Used in a variety of computer games

-get out of my way sonny, I've a ship to save
If I were you, I'd pack it in and go to bed.
-after I've saved this ship
I think you need to chill out. Why not go up to the music room
and relax to the cool sounds of Bobby Headcase and his
Laid-Back Loafers? Mmm, nice.

Chat-O-Mat

Example: TT13M

- Provide personality for Half-life bots
- Chat with fellow players
- http://www.cs.north estern.edu/~rob/pub ations/making-the-human-care.pdf

Eliza DLL

```
┌─────────────────────────────┐
│ Personality  │  Valence     │
│ Stereotypes  │  Memory      │
└─────────────────────────────┘
        │                ▲
        ▼                │
┌──────────────┐   ┌──────────────────┐
│ Find         │──▶│ Pick Appropriate │
│ Response     │   │ Reply            │
└──────────────┘   └──────────────────┘
        ▲                  │
- - - - - - - - - - - - - -│- - - -
        │                  ▼
┌──────────────┐   ┌──────────────────┐
│ Incoming     │   │ Delay            │
│ Message/     │   └──────────────────┘
│ Event        │           │
└──────────────┘           ▼
                   ┌──────────────────┐
                   │ Outgoing         │
                   │ Message Handler  │
                   └──────────────────┘
```

TeamBot DLL

Example: Cobot in LambdaMOO

- Project goals: Understand how to build software that can function socially with people
  - http://cobot.research.att.com/
- While this was for research, it provides some interesting ideas for making social NPCs

Cobot's life in LambdaMOO

- Over first three months, 550,000 separate events
  - One event every 11 seconds
  - Operating mostly continuously since 1999
- Didn't say much in the beginning
  - 3 events/1000 directed at it
- Once a month of statistics were gathered, it would offer them up
  - 50 events/1000, more than double that of next most popular denizen of the Living Room

Data that Cobot collects

- For each user encountered:
  - Histogram of verbs used by them
  - Histogram of verbs used on them
- For each verb invoked in the room:
  - A histogram of users invoking it
  - A histogram of users that have been its target
- For each pair of users it sees interact:
  - A histogram of the verbs they have used on each other

Providing statistics

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%."

Unintended consequences

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.

...

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!

Another Cobot Example

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.

Making Cobot plausible

- Cobot generates emotive responses
  - Classic chatbot strategy: Randomly select from several plausible replies
- Extended text chat abilities
  - Use word matching from input to sentences in large corpus of text
    - Corpus includes Unabomber Manifesto, scripts for "The Big Lebowski" and "Planet of the Apes", articles about the JonBenet murder, alien abduction, and paranormal experience.
    - Change contents of corpus slowly, to mimic changing themes or concerns
  - Still does not maintain state across utterances

Reinforcement learning

- Experiment: Let Cobot initiate actions
  - Learn via feedback from users when actions are appropriate
    - Verbs like "hug", "spank"
  - Actions include trying to start conversations, social commentary (state of room), introducing users to each other

- Some results:
  - Cobot learned appropriate behaviors, responding in patterns based on the reinforcement received
  - A small set of users became "parents", interacting an order of magnitude more with it than other users

Information Retrieval

- The classic Information Retrieval problem:
  - Inputs: A query Q (i.e., the player's input) and a corpus of possible responses
  - Output: The most appropriate choice from the corpus
- Here, could divide each corpus element into two parts:
  - \<Possible player utterances, Response to be made\>
  - Only use possible player utterances for matching

Simple Model: TF-IDF

- Think of each word appearing in a corpus of documents as a dimension in a vector space
  - Many distinct words → lots of dimensions
- Not all words count the same
  - Throw out some words (*stop words*)
    - Ex: the, a, an, …
  - Weight by *term frequency* (i.e., how often the word appears in the document)
  - Weight by *inverse document frequency* (i.e., how often the word appears in every document in the collection)
- Scaling word counts in documents by TF*IDF leads to vector representations useful for retrieval
  - Pick maximum of cosine(query, document)

More complex model: Latent Semantic Analysis

- Start with large corpus, compute reduced dimensionality space
- Again, use cosine metric to pick closest response
- Turns out to be useful for many purposes
  - Dialogue-based tutoring systems (cf. Autotutor)
  - Grading essay tests (cf. ETS)

Markov Models

- Markov model = Finite state machine with probabilities for each state transition
- Suppose state is not directly observable
  - Each state randomly produces some output
  - The same output might be produced by multiple states
  - Example of a *Hidden Markov Model*
    - Primary method for speech recognition used today
    - We'll see them used later in the quarter...

MegaHal (Hutchens)

- Parsed utterances into sequence of word/non-words
- Learned two Markov models from text corpus
  - Given any sequence of four symbols, predict what will follow next
  - Given any sequence of four symbols, predict what preceded them.
- These two models used to generate responses from keywords
- Trained over a large corpus:
  - Hand-crafted sentences designed in order to create a personality for MegaHAL, including sentences containing a false name, age and occupation.
  - Encyclopaedic information taken from the Web, on topics such as geography, music, sports, movies and history.
  - A selection of sentences picked from transcripts of previous Loebner contests.
  - Lines of dialogue taken from scripts for movies and television shows.
  - Lists of popular quotations.
  - A small amount of text in languages other than English.

MegaHal Algorithm

1. Segment user's input into word/non-word sequence
2. Extract keywords by throwing away stop words
3. Generate candidate replies from each keyword (hundreds), using Markov models plus transformations (i.e.,
4. Choose the most surprising reply
   1. $I(w|s) = -\log_2 P(w|s)$, for w = keywords in reply $P(w|s)$ is computed by Markov models
5. Update Markov models based on user's input

MegaHal Example

- Appearance of learning:

  User: Roentgen discovered the X-rays.

  MegaHAL: THE NOBEL PRIZE FOR PHYSICS IN 1900 WAS WON BY ROENTGEN.

  User: Who discovered the X-rays?

  MegaHAL: ROENTGEN DISCOVERED THE X-RAYS.

Intriguing, but still quite limited

- "In general MegaHAL's conversations are not as successful as this. Most users are satisfied with typing in rude words to see how the program responds." Hutchens, http://megahal.alioth.debian.org/How.html
- Might be a useful confusion script generator
  - How to avoid generating red herrings for the player?