

**JOIGNEZ
L'UTILISABILITÉ
À L'AGRÉABLE**
CESART



CESART est fière de commanditer
les soirées conférences
UPA-Montréal et de vous offrir
rafraîchissements et café.


Bonne soirée!

Intelligent User Interfaces

Kristina Pitula

-
- Traditional user interface
 - Why add intelligence?
 - How do you add intelligence?
 - Multimodal interfaces
 - Why use agents?
 - Issues and open questions
 - Agents in the news

Traditional UI



Traditional UI (c.2)

Traditional UI has 3 main components:

- **Presentation**
 - Input/output mechanisms
- **Dialogue Control**
 - Defines the set of user input triggers and the corresponding actions and/or output by the system
- **Application Interface**
 - APIs between the UI and application(s) which effectuate the requested action

Traditional UI (c.3)

- Metaphor of *direct manipulation*
 - Visualize the objects
 - Actions on objects in the interface correspond to actions on real objects in the system
 - Nothing happens unless the user makes it happen
- Knowledge about the user, application, and interaction is implicit
 - Intelligence may be “hard coded” within the input-action triggers
 - No explicit representation nor reasoning about what is occurring
- Designed for a closed, static, relatively small and well structured information world

Basic usability principles

- UI is effective
- UI is efficient
- UI is easy to learn & “natural” to use
 - What is “natural” depends on the culture, knowledge, and physical abilities of the individual in the target environment
- User is in control of the system
- System is predictable
 - System always gives the same response given the same input
- System is transparent
 - User can understand how the system works

Why add intelligence?

- UIs are getting too complex
 - Within a given system
 - Across the multiple systems a user is exposed to
 - Across the range of access devices available
- UIs are too inflexible
 - Predefined text, graphics, images and sound
- UIs don't change as the users' needs change
 - Different user populations
 - Individual user over time
 - More untrained users using more complex systems
- Need to reduce cost, time and expertise required to develop good UIs

Why add intelligence? (c.2)

- How can the interaction be made clearer and more efficient?
- How can the interface provide better support for the users' tasks, plans and goals?
- How can information be presented more effectively?
- How can the design and implementation of “good” interfaces be made easier?

Why add intelligence? (c.3)

- Adapt to the needs of different users or different situations
- Anticipate the needs of users
- Choose what information to present to users
- Provide explanations of actions or results
- Support a range of media & modes
- Take initiative and make suggestions to users
- Take on tasks for users

How do you add intelligence?

- Make knowledge explicit within the system so that the system can reason and act on what it knows
- Incorporate concepts & techniques drawn from *AI*
 - Knowledge representation techniques
 - Inference mechanisms
 - Multimodal interaction
 - Machine learning
 - Agents
- Increase efficiency, effectiveness and generality
 - Knowledge & reasoning strategies to interpret inputs, determine actions, design responses
 - Symbolic techniques increase understanding, maintainability and modularity
 - Knowledge can be factored out for use in other applications

How do you add intelligence? (c.2)

1. Presentation

- + Input processing augmented by intelligent media analysis
- + Output rendering augmented by intelligent media design

2. Dialogue Control

- + Media analysis & interpretation
- + Media design
- + Interaction management
- + Knowledge representation & inference

3. Application Interface

How do you add intelligence? (c.3)

+ Media analysis & interpretation

- Interpret ambiguous and/or incomplete user directives ... according to context, focus, etc.
- Parse mixed and asynchronous multimedia input

+ Media design

- Coordinate multimedia output
- Generate output on the fly from common representation ... according to characteristics of device, user preferences, etc.

How do you add intelligence? (c.4)

+ Interaction management

- Multimedia fusion
- Discourse modeling
- Plan recognition & generation
- User modeling
- Presentation design

+ Knowledge representation & inference

- User models
- Discourse models
- Domain models
- Task or context models
- Media models

“Data Mountain”

- 3D document or web page management system
 - [George Robertson et al, MicroSoft Research, 1998]
- Uses 3D visual and audio cues combined with 2D interaction techniques to sort and arrange documents
- Exploits human perceptual skills and spatial memory to improve users’ ability to find things back
 - Spatial relationships, cognitive chunking, pattern recognition, object constancy, attention & curiosity...

Multimodal Input integration

“Put that there” [Bolt, 1980]

- The first “Spatial Data-Management System”
- User in a mixed virtual/real space manipulates block shapes on a large-screen graphic display using speech and manual pointing
- Connected speech recognition limited to a maximum of 5 words per sentence
- Gestures captured by a “sensor cube” attached to the participant’s watchband

“Put that there” (c.2)




Figure 1

Sketch of Media Room

“Put that there” (c.3)

- User can create, modify, move and delete objects on the graphical display
 - Speech and gestures used simultaneously
 - User can use pronouns to refer to objects or locations
 - e.g. “*Put that there*” rather than “*Put the blue triangle to the right of the green square*”
 - Gestures used to resolve what the pronouns refer to
- Input integrated by searching for the synchronized gestural acts corresponding to the spoken referents

Multimodal interfaces

- Combine variety of UI interaction styles
 - Graphic-based: GUI, WIMP, animation, visualization, 3D, VR
 - Linguistic-based: command, speech, natural language
 - Gesture based: pointing, gestures, eye-tracking, motion, etc.
- Combination of modes leads to new interaction styles
 - Users can perceive multiple things at once
e.g. vision and sound, tactile feedback, sense of equilibrium, possibly taste and smell
 - Users can do multiple things at once
e.g. speak and use hand gestures, look at things, use body language (position, orientation)

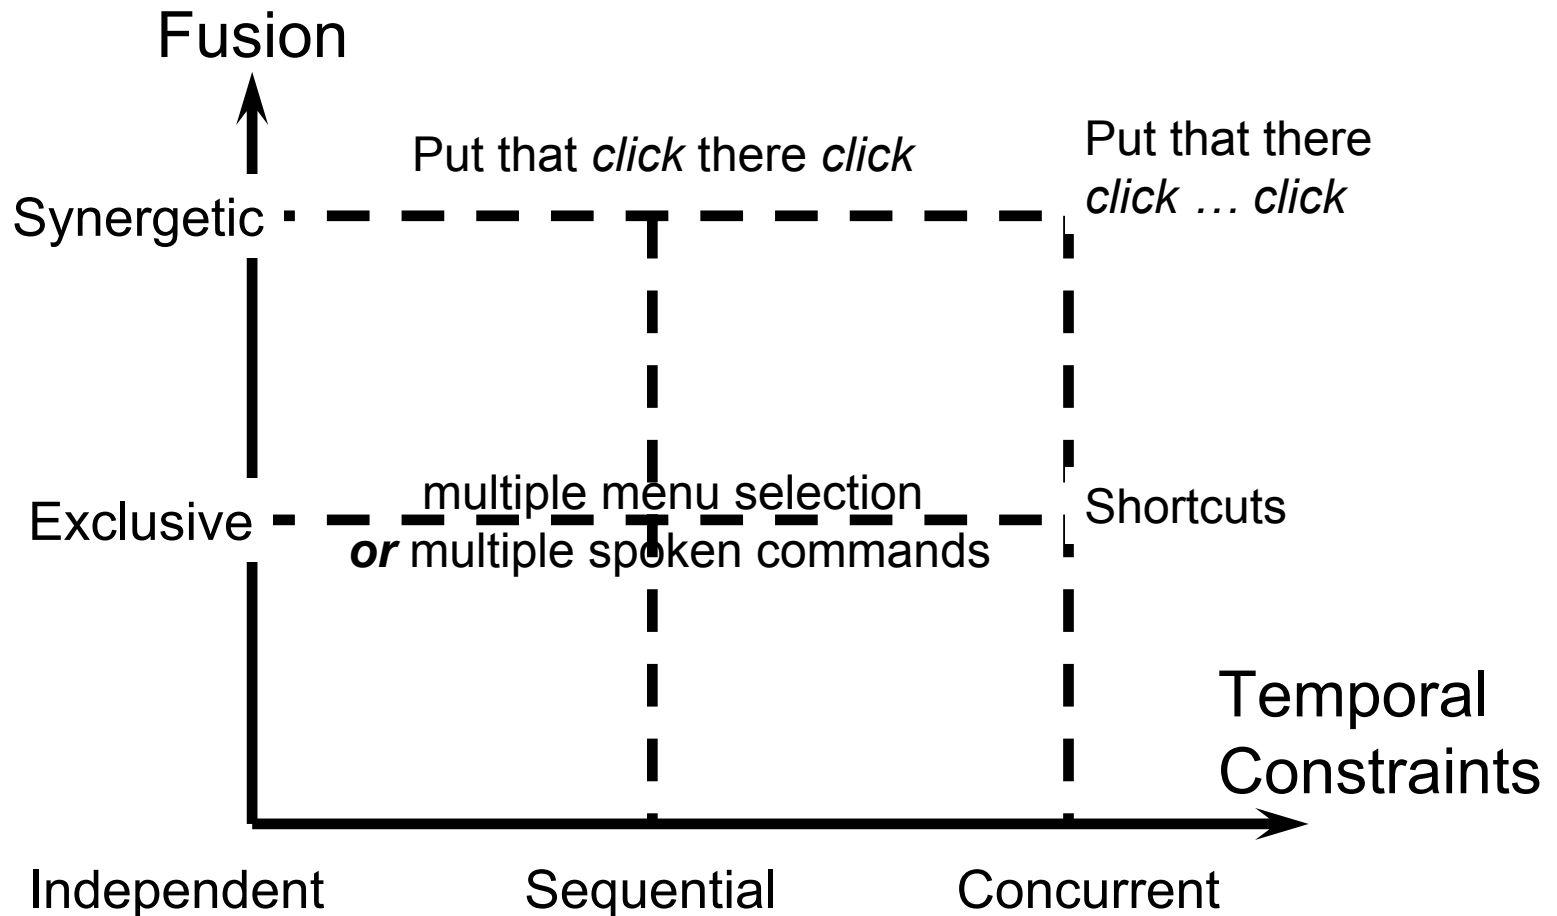
Multimodal interfaces (c.2)

- Different modalities are better for different tasks ... in different environments ... using different devices
 - e.g. point to location on map, gesture to indicate location in space, speech to refer to objects...
- Multiple modes improve interaction
 - Provide a higher bandwidth i.e allows more information to be communicated within the same time frame
 - User can choose most appropriate mode
 - Reduces errors
 - More natural way of interacting
 - More engaging to the user

Multimodal interfaces (c.3)

- Input from different modes has different characteristics
 - Modalities may be simultaneous or not
 - Input may be complimentary, redundant, ambiguous, incomplete
 - Inputs may overlap in time or be delayed in respect to each other
- Input from different modalities has to be integrated
 - Interpretation must rely on various models (*User, Discourse, Domain, Task or Context, Media models, etc.*)
 - *Fusion* according to temporal constraints that reflect the intertwining of modes

Taxonomy [Balbo, 1993]



Boeing's Virtual Reality Aircraft Maintenance Training Prototype

[Duncan et al., 1999]

- 3D multimodal system that uses an avatar in a VR environment to assess the maintainability of a new aircraft design and train mechanics in maintenance procedures
- It combines speech and gesture and uses an avatar in a VR environment
- The avatar is driven by magnetic trackers attached to a human actor
- The avatar's motion in the VR environment reflects the human's in the real world
- The user can give commands such as "*hand me the socket wrench*", "*give me that*", "*take me to the E4 table rack*"...

Machine learning

- Ability to learn provides considerable value
 - Some tasks cannot be defined well except by example (e.g. driving)
 - Important relationships may be hidden within vast amounts of data
 - Certain characteristics may not be completely known at design time
 - Environments change over time
 - New knowledge is constantly becoming available
- Wide range of *Machine Learning* approaches available
 - By observation, by asking questions or taking advice, by generalizing from a set of examples (induction), case based reasoning, reasoning by analogy...

Machine learning (c.2)

1. Data mining and knowledge discovery
 - use historical data to improve decisions
 - e.g. medical records to improve medical diagnosis, credit risk analysis, detecting credit card fraud, analysing shopper behaviour
2. Software applications that are difficult or impossible to program by hand
 - autonomous driving (Alvin, Pomerleau)
 - speech recognition, learning strategies for game playing
3. Self customising programs
 - a news reader that learns user interests

Why use agents?

- An *agent* is a concept from AI
 - Persistent
 - Has goals, sensors and effectors
 - Is autonomous, learns, and cooperates
 - ... with users &/or other agents &/or applications
 - This does not prevent it from competing.
- Can autonomously decide what actions to take in the current situation to maximize progress toward its goals

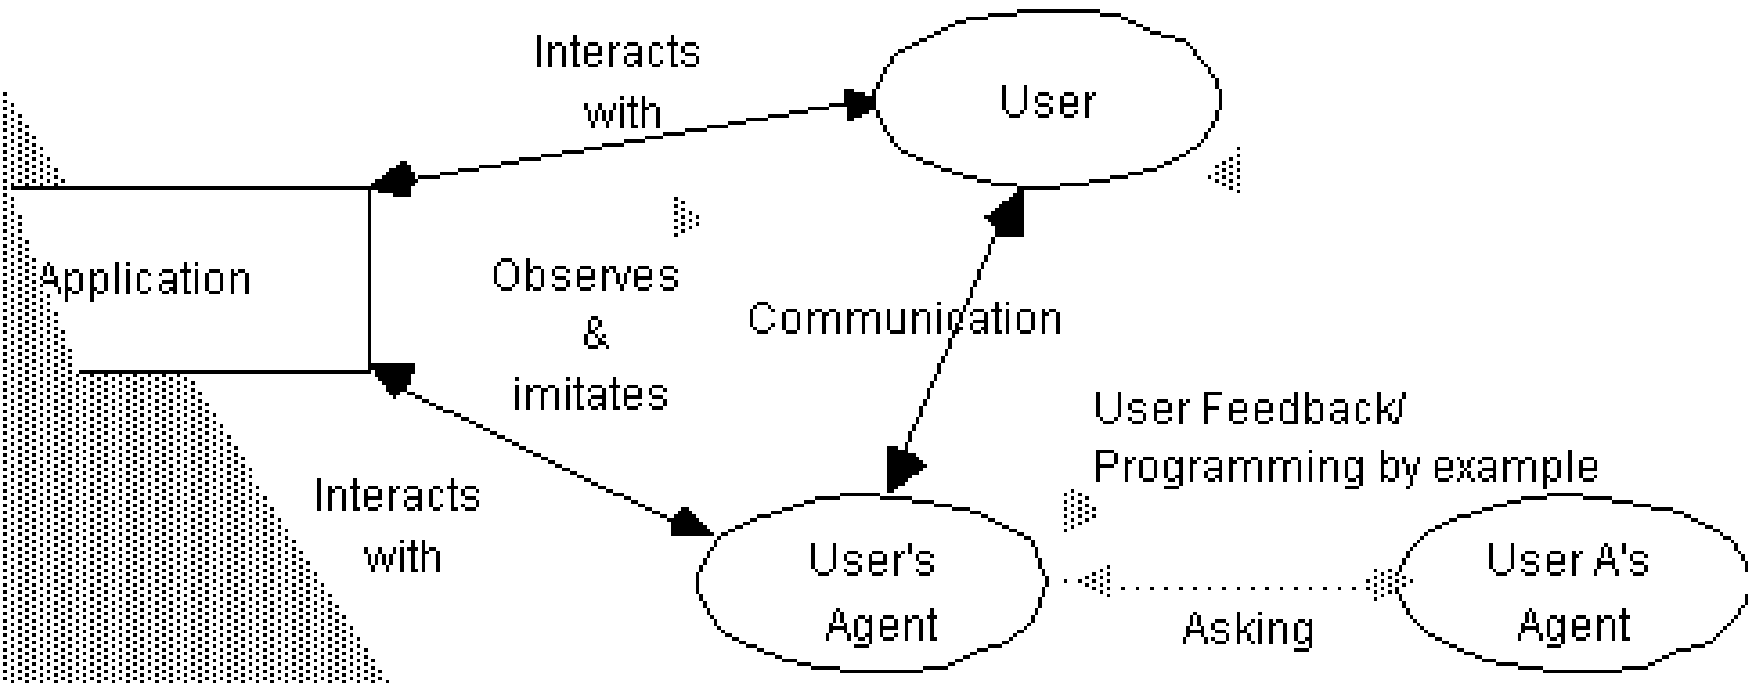
Why use agents? (c.2)

- Agents can perform a range of tasks
- Can interact with user or work in the background
 - Personal assistants e.g. personal news editor, e-shopper, web guide, tutor, etc.
 - Service agents work in the background e.g. web indexing, information retrieval, phone network load balancing, etc.
- Can reside locally, remotely, or move around
- Can be programmed by the user, a knowledge engineer, or be capable of learning on their own

Why use agents? (c.3)

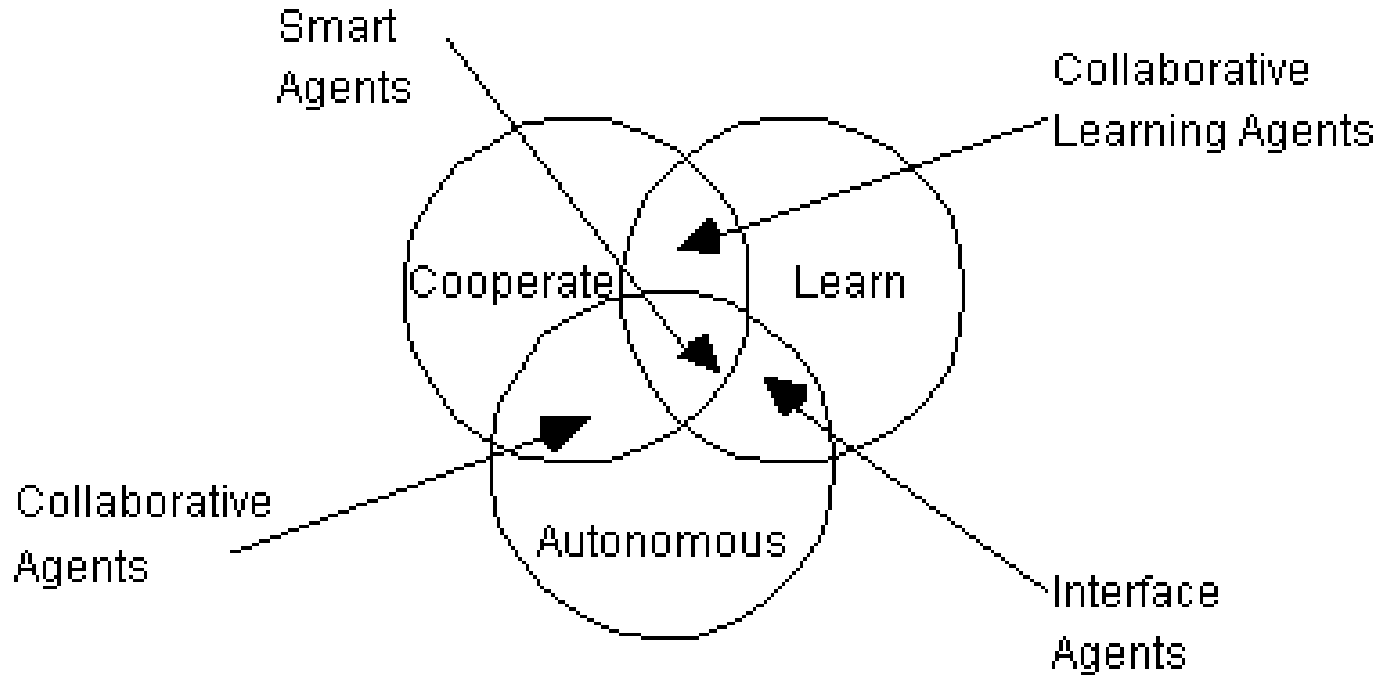
- Metaphor becomes one of *Indirect Management*
 - User can delegate to agents that know the user's interests, habits, preferences
 - Agents can make suggestions and/or act on behalf of the user
 - Agents can continue working even when the user is not active
- A user interface agent can fulfill a number of roles
 - Assist the user in communicating with the system
 - Learn the user's profile
 - Perform routine chores
- Agents must display competence to gain the user's trust

Why use agents? (c.4)



Taken from "Software Agents: An Overview", Hyacinth Nwana, Knowledge Engineering Review, Vol. 11, No. 3, pp.1-40, Sept. 1996

Why use agents? (c.5)



Taken from “Software Agents: An Overview”, Hyacinth Nwana, Knowledge Engineering Review, Vol. 11, No. 3, pp.1-40, Sept. 1996

New interaction styles

- Different modes provide multiple ways for a system to be aware of the user
 - Language
 - Gesture
 - Eye tracking
 - Motion detection
- Awareness allows us to capitalize on basic human behaviour
 - Attention
 - Emotion
 - Curiosity

Attentive UI

- Attentive UI is aware of the user's attention space
 - What the user is currently paying attention to
- Can reason about the user's focus of attention
 - Uses an attention model to estimate task priority in relation to other tasks and events, and to predict changes in the user's focus of attention
- Attempts to minimise user disruption by negotiating turns gracefully
 - Uses i.e. peripheral channels to get the user's attention
 - Waits for the user's acknowledgement before moving to the foreground
- Augments attentive resources
 - Similar to similar to the Cocktail Party Effect, it magnifies information in the user's estimated focus of activity while attenuating peripheral detail

Affective UI

- Affective UI is aware of emotion
 - Can respond accordingly
- Emotion plays an essential role in cognitive functions
 - Impacts rational decision making, perception, learning and more
 - Too little or too much emotion can seriously impact decisions
- Emotion essential in establishing relationships
 - Ability to recognise and express emotions essential to natural communication
 - Emotion essential in acquiring social competence & establishing trust

Affective UI (c.2)

- How do you model emotions?
- What range of emotions to express?
 - Happiness, sadness, compassion, grumpiness, puzzlement, curiosity, frustration, fear, anger ...
- How to express emotions?
 - Certain portion expressed through content
 - Significant cues provided by facial expressions, body language, tone of voice
- Which emotion to display in response to a specific situation?
 - If emotions are learned, can agents become irrational or emotionally unstable? ...
- The *uncanny valley* thesis [Masahiro Mori, 82]
 - Non-human entity with nearly human behaviour but slightly off-kilter seems eerie, disquieting, and makes people uncomfortable

Application domains

- Intelligent tutoring
- Intelligent help
- Information filtering
- Natural language interaction
- Multimedia interaction
 - Media abstraction
- Personal assistants
- Embodied agents
- Robots
- Control systems
- E-commerce
- Manufacturing
- Telecommunications
- Traffic control
- Military applications
- Games
- ... etc...

What is IUI?

- Not every “intelligent” system has an “intelligent” interface
 - A knowledge based system such as an expert system may have a command line interface
- Not every “good” interface is an “intelligent” interface
 - Does not necessarily incorporate AI techniques
 - ... nor does it need them*
- An “intelligent” interface incorporates AI techniques to support user interaction
- An “intelligent interface agent” differs in that it is proactive and autonomous ... to the desired degree

Many open issues & questions

- Usability principles for intelligent interfaces
 - User remains in control
 - System remains predictable & transparent
- Better understanding of how intelligence *can improve* the interaction ... rather than get in the user's way
 - Adding intelligence does *not* replace good UI design
- How do agents gain a user's trust?
- Who is responsible if an agent makes an inappropriate decision?
- Methods and tools to facilitate design, development and maintenance of IUI
- Methods and tools to evaluate IUI

Agents in the news

Driving a Smart Vision

- Wireless Week (01/28/02) Vol. 8, No. 4, P. 21; Albright, Peggy

The Intelligent Transportation Society of America (ITSA) proposes to build intelligent transportation systems.

Their report proposes an integrated network of transportation information, and the use of the Internet, telematics, and cellular and wireless technologies to allow for real-time reporting of infrastructure conditions, maintenance activities, weather conditions, and traffic congestion.

- Automatic crash detection systems a key element.
- Hands-free driving also considered

Agents in the news

Erasing the Blind Spot: Driver's Aid Averts Traffic Jams

- New York Times (12/26/02) P. E5; Dizikes, Peter

Driver-assistance systems to help drivers respond faster to traffic conditions

Equipping just 10 percent to 20 percent of all cars would dramatically reduce the frequency that isolated incidents slow down the entire freeway (Dirk Helbing et al., Dresden University of Technology)

Radar tracks speed and distance from the vehicle ahead; Drivers alerted of dangerous situations by i.e. audio signals, dashboard displays, or vibrations on the steering wheel... Sensors track vehicles moving into the driver's "blind spot"... Other mechanisms to correct lane-drift.

Manufacturer TRW Automotive is working on a system linked to control that keeps the car at the same speed as the vehicle in front.

Agents in the news

21st century cars hit the road

– Business Week, September 4, 2003:

... "Safety systems will move from passive protection, such as airbags, to active systems that use radar and cameras to watch for danger. 'We've put airbags just about every place you can,' notes John Weiner, a U.S. product-planning manager at Toyota (TM). 'Within the next five years, the car will use algorithms to anticipate hazards and intervene or warn the driver.' ... Voice commands, already used in some BMWs, Jaguars, and Lexuses, will become far more common to help drivers juggle the proliferating functions in their cars. ... Privacy issues will arise, too. If cars start to record and transmit data on things like speed, where the driver has been, and whether the wipers were on in a storm, then the family car could become a snitch on its driver."

Agents in the news

Bots Battle, Breed in A.I. Test

- Wired News (02/07/02); Delio, Michelle

Noel Sharkey at Magna Scientific Adventure Center ... to stage largest artificial intelligence experiment ever when several dozen robots are released into a prepared habitat to see if they can evolve using basic survival instincts.

Robots will be programmed into "predator" and "prey" modes:

- Prey robots will feed (recharge their batteries) using solar power
- Predator robots will feed on the prey by leeching their battery power via an "energy sucking fang."
- All robots will have sensors to track down prey, avoid predators
- Robots that perform well will be allowed to breed by combining artificial genes into new neural networks.

Experiment may yield research that could be applied to machines for use in undersea exploration and outer space.

Agents in the news

Robotic Heroes

- Laptop (12/02) Vol. 21, No. 23, P. 108; Hendrickson, Nancy

Advanced robots that can function in areas or situations too dangerous for humans are finding use in the military, law enforcement, and even the space program.

Battlefield robots under development or in use include iRobot's Bloodhound, a machine that can autonomously navigate its way to wounded soldiers and administer medical treatment; and Packbot, another iRobot product that was used in Afghanistan to search caves for al Qaeda or Taliban fighters

Agents in the news

Govt to help develop robot mine hunters

– Yomiuri Shimbun, Daily Yomiuri On-Line Jan. 2003

Japanese government to cooperate with domestic heavy machinery makers to develop high-technology devices for removing land mines in Afghanistan

Land mine detectors that can withstand sudden changes in temperature in the mountains and high-tech robot arms capable of digging into tough terrain to where land mines are buried will be developed

Agents in the news

UW professors develop terrorist defence robots

– USA Today June, 2003

David Thayer et al. at University of Wyoming are creating small robots to detect and disable chemical targets.

Each robot is autonomous; can move, reason, react, and communicate, and has multi-modal sensor arrays to detect chemical, biological or even radiological plumes..

100 robots are organized in a "swarm intelligence" network. The swarm converges on a chemical target — and can even eliminate a chemical spill — without exposing people to the contaminants. It can also pinpoint the spot from where a contaminant was released.

Agents in the news

Smart steps for eliminating spam

– By Bob Parsons. Entrepreneur.com, September 22, 2003:

"Spam, or unsolicited bulk mail, is an inexpensive and easy way to advertise to a vast audience. Today, an estimated 50 percent of all e-mails handled by Internet service providers (ISPs) are spam. It's also estimated that the average U.S.-based office worker sifts through as many as 100 pieces of spam every day.

... The best way to clean up your inbox is to download an intelligent spam filter that recognizes and separates the good from the bad. To date, the best on the market is the 'Bayesian' spam filter, which uses artificial intelligence to detect spam according to your personal e-mail preferences. This self-learning filter uses a probability-based mathematical theory ... Bayes' theory is based on the number of times an event has or has not occurred, and the likelihood that it will occur in the future."

Agents in the news

Chatting with on-line characters

– By Sebastian Ruple. PC Magazine News, July 24, 2003

"While today's intelligent online characters, or bots, have disappointed some people, two prominent partners have launched a new effort to find useful e-learning and customer service applications for virtual people.

Oddcast, a company that makes conversational characters, and the ALICE AI Foundation, a non-profit research organization focused on advancing AIML (Artificial Intelligence Markup Language) have announced a partnership to create smarter intelligent online characters. The technology allows for personal interaction with online agents that can function as customer service agents, tutors, and the like."

Agents in the news

Robot Tarzan' helps forest work

– By Jo Twist. BBC. December 29, 2003

"The hi-tech Tarzan of the robot world, nicknamed Treebot, is the first of its kind to combine networked sensors, a webcam, and a wireless net link. It is solar-powered and moves up and down special cables to take samples and measurements for vital analysis. Treebot has been developed by scientists at the US Centre for Embedded Network Sensing in California. ... Eighteen months in development, the main difference between Treebot and other fixed sensors is its autonomous nature and its ability to communicate with other devices and sensors."

Agents in the news

Software finds tunes you want to hear

– By Matthew Daneman. Democrat and Chronicle, **July 7, 2003:**

"Mitsunori Ogihara, chairman of the UR computer science department, has applied for a patent on a piece of software that would, in essence, 'listen' to songs and categorize them into specific genres and by emotional content. The result could be a future where a listener in the mood for happy jazz pieces or maudlin country tunes could send the program filtering through radio stations or among digital music files stored in a computer's memory to find what the listener wants. ... The software categorizes musical genres and emotional content by analyzing signals and patterns in songs. The ultimate goal, Ogihara said, is to create personalized software that recognizes signals and learns its owner's musical tastes. 'You have only to tell the software, 'This is what I think of as jazz; this is what I think of as rock,'"he said."

Agents in the news

Virtual helper makes independence a reality

– By Deborah L. Shelton. St. Louis Post-Dispatch, July 12, 2003:

"Open the drapes. Brew the coffee. Prepare the shower. Ralph's routine is ordinary enough. But Ralph has an impressive work ethic, laboring 24 hours a day, seven days a week - year in, year out. Ralph moved in with Don Holbert, of Sedalia, Mo., over a year ago. Holbert, 59, contracted polio when he was 5. Though paralyzed below the waist, Holbert was able to manage for himself until his wife, Barbara, died in May 2001. Without her, even some of the simplest tasks around the home, like opening the blinds, became impossible. That's where Ralph comes in. Ralph now adjusts the thermostat, turns lights on and off and reads stories from the newspaper. Ralph is a helper. A housemate. A talkative companion. Ralph is a computer. To be more precise, Ralph is a voice-operated computer and home automation system, programmed to function using artificial intelligence."

Agents in the news

Computer games learn new tricks

– By Mark Ward. BBC, 12.30.03

Yves Guillemot, boss of French game studio Ubisoft ... gives his views on the way that the industry has to change to keep players interested. 'You should at least have the same game experience even if you are not as good,' he says.

Ubisoft is working on ambient AI that watches what a player does and adapts the game and how the plot plays out to their skill levels, he says. 'We should be able to improve the game for particular kinds of people,' he says. 'It is about making sure you can understand the reactions of the players to give them the things that will really work for them. It is about AI reacting to your abilities. If you cannot do something after 20 tries it makes sure you still progress.'