nature web matters

11 Marcl

Is There an Intelligent Agent in Your Future?

JAMES HENDLER

> *"The White House said the increased technology spending – mentioned briefly by President Clinton during last week's State of the Union address – could be used, for example, to create "intelligent agents" that roam the Internet collecting information" ... (AP News Service, January 24, 1999)*

Imagine that you are lucky enough to find some free time in your schedule, a you decide to take a trip. What do you do? If you're like most people, you co a travel agent to arrange the details. You describe your needs (where you wa go and when), the constraints that you need to impose (how much you are w to spend, the hotel must provide child care) and some personal preferences (: preferred airline, you'd like to sit in an aisle seat). The travel agent, using a combination of information sources (flight schedules, hotel guides) and guid past experiences, recommends where you might go. Once you confirm your the agent generates the itinerary, books the flights and generally does all the that you don't want to bother with.

Now think about your voyages on the internet. Wouldn't it be nice to have ar agent to help you in these travels? You could specify your needs (to find jou articles describing experiments in a particular field), your constraints (that a certain experimental system or a particular reagent was involved), and your preferences (that articles are written by research groups at major universities internet travel agent, would then find you some possibilities and ask your ap of its choices. Upon your agreement, the agent would download the articles : in any registration details needed for the web sites, arranging credit card pay to those that need it, any number of mundane things that you don't want to b with.

The vision of such *intelligent agents* is quite compelling and many people nc believe they will be necessary if we are ever to tame the increasing complex: caused by the accelerating and virtually uncontrolled growth of the World W Web. The present situation can only get worse as new web technologies (for example new mark-up languages such as XML) permit the further integratio more complex data sources (for example database searches, real-time simula of complex systems, multimedia presentations, *etc*.).

If you are unconvinced of the need for help try searching the Web using a fe scientific keywords in your favourite search engine. Look at the ten or twent item found. Now imagine the web was a hundred times it's current size (whi the predicted growth in the next 2-3 years). That item might be number 1,00 2,000, buried under the avalanche of false hits, misused terms, misinformatic all those other problems that already cause many of us to consider the Web a 'necessary evil' rather than a true boon to our scientific endeavours.

At this moment such agents do not exist, despite the great need for their serv
There are technical reasons why developing such agents is hard, but conside
work is under way to improve the current state of the art.

What makes a good "agent"?
Returning to the travel agent analogy, what makes a travel agent good as opp
to one you wouldn't want to use?

First, you must be able to communicate with the travel agent. Both you and y
agent must speak the same language. Any agent that cannot understand whe
want to go or what you want to do there will be of little help to you.

Second, the agent must be able to act as well as suggest. This is the differenc
between some sort of travel advisor and a real travel agent. An advisor migh
you decide where to go, but when you needed to book flights and hotels, ren
and all the other minutia that is involved in a trip, you have to do that yourse
real travel agent on the other hand can actually do these things for you.

Third, an agent can do things without supervision. Imagine trying to book a
complex trip, through a travel agent which required you to be on the phone v
all details were planned and finalized. You'd need to wait while they looked
and contacted hotels and airlines, checked with previous travellers whether t
enjoyed their trips to the same places, searched through price lists and day/da
lists, checked with the hotel to see if it was open that day, *etc*. Forget it! Goo
travel agents collect some information from you, go off on their own to find
options, call you back to suggest and discuss possibilities, take the itinerary y
prefer and then go book it.

Finally, a key ability of good travel agents is that they use their experience to
you. They don't send you to hotels which other customers have told them we
awful. Similarly, if you travel a lot your travel agent should learn about your
preferences: which airlines you like, what times you like to fly, what sort of
you prefer to rent, all those little things that can make the difference betweer
good and bad trip. But further, the best travel agents learn how much, or how
little, you want them to do: some people like to micromanage their entire itir
some prefer to let the agents make all the choices, and others like something
middle.

The ideal internet agent
A good internet agent needs these same capabilities. It must be *communicati*
able to understand your goals, preferences and constraints. It must be *capabl*
able to take options rather than simply provide advice. It must be *autonomou*
able to act without the user being in control the whole time. And it should be
adaptive; able to learn from experience about both its tasks and about its use
preferences.

Let's look at each of these in turn, reviewing current research in these areas.

- **Communicative**
  What does it mean to be able to communicate with someone (or somet
  in the computational case)? Greatly simplifying an issue which has be
  cornerstone of intellectual thought for many years, useful communicat
  requires shared knowledge. While this includes knowledge of a langua
  the words and syntactic structures – it is even more focused on know

about the problem being solved. To deal with a travel agent, you need able to talk about travelling, to interact with a florist you need some knowledge of flowers, and to deal with an internet agent you must sha vocabulary about the area of concern.

A key problem with current search engines is that although based in language, they have no knowledge of the domains of interest. Searchi physics papers would be enhanced if the search engine actually 'knew' something about physics (how experiments are performed, which wor ambiguous, whether papers are theorerical or empirical, etc.), rather th simply looking for the appearance of key words. Not all squids are superconducting quantum interference devices!

Solutions to this problem usually involve 'ontologies'. While this term part of the technical jargon of artificial intelligence researchers, the ba concept is simple – an ontology is a formal definition of a body of knowledge. The most typical type of ontology used in building agents involves a structural component. Essentially a taxonomy of class and subclass relations coupled with definitions of the relationships betwee these things.

In addition, the ontologies contain some kinds of 'inference' rules – th can be explicit rules about an item, or they can be 'structural' inference provided by a system. An example of the former might be a rule like " a car, then X has four tires." An example of the latter might be rule lik one thing is part of something else, and that latter thing is itself a com of an assembly, then the first item is a part of the assembly." (that is, t hubcap is part of a tire, a tire is part of a car, therefore the hubcap is pa a car).

As an example of an ontology, consider the one shown at http://www.cs.umd.edu/projects/plus/SHOE/tse/tseont.html. This onto was developed at the University of Maryland in a joint effort of compu scientists and FDA and veterinary biologists to help in the developme new tools to interact with information about transmissible spongiform encephalopathies (the most famous being BSE or 'mad-cow disease'). ' ontology links a number of concepts such as diseases to diseased anim symptoms to diseases, *etc*. It is worth noting that this ontology is far fr complete. It doesn't get down to prions or other deep molecular conce but it does provide key concepts that are being developed for an intern agent that will help users find internet-based information to aid in mak risk-assessment decisions.

If an ontology can be made machine readable, it allows a computer to manipulate the terms used in the ontology, terms that make sense to us who understand this information. The computer doesn't understand thi information, in any deep sense of the term, but it manipulates terms th user understands. This allows for a form of communication between u and computer, which in turn enables the creation of software products, the agents we've been discussing, which can represent the needs, preferences, and constraints of the user.[i]

- **Capable**
  For an agent to be capable, it must be able to take action in some sort

world. As mentioned before, the difference between an agent and an advisor, is that the agent not only provides advice, but also provides a service, the ability to do things on the Web without you needing to kn details. Unfortunately, the current state of the art is limited by the need know much about the specifics of the internet sources the agent will in with. For example, suppose you want a paper from some particular ph journal, but the journal charges for it. If you want an internet agent to able to download this paper for you, it would need to know where on t page the price is, where it communicates your credit card number, in v formats the various pieces of information about you need to be put, *etc* fact, this is exactly the information that you would need to use the pag yourself, and the variations from site to site is one of the reasons the c web is so hard to use.

This is a difficult problem, and many current internet communities are worrying about how to overcome the syntactic issues involved in findi appropriate items on web pages (for example XML). While agreement starting to emerge, a lot of engineering is still to be done to encode information about internet sources and about how to manipulate them. Intelligent agent researchers are watching these developments closely, as the web becomes increasingly 'agent-friendly' more capable agents being developed.

- **Autonomous**
  One of the more contentious issues in the design of human-computer interfaces arises from the contrast between 'direct manipulation' interf and autonomous agent-based systems (see http://www.acm.org/sigchi/chi97/proceedings/panel/jrm.htm). The proponents of direct manipulation argue that a human should always b control – steering an agent should be like steering a car – you're the you're active the whole time. However, if the software simply provide interface to, for example, an airlines booking facility, the user must ke needs, constraints and preferences in his or her own head.

  Software that can help make wiser decisions, but is not capable of doi anything for you or needs constant steering to be effective, may be a u advisor, but falls short of the time saving tool that a true agent would b truly effective internet agent needs to be able to work for the user whe user isn't directly in control. Of course, if the software were too auton this could cause problems. Thus, the key to autonomy is finding the ri level for the task at hand.

  Enabling autonomy is a difficult programming task, particularly becau is very dependent on features of the area in which a program is operati travel agent must have very different levels of autonomy than a real es agent. For human agents the level of autonomy is clearly defined, sometimes by law, sometimes by customs that have grown up over ma years. For internet agents it is much less clear (hence the debate descri above). It is even unclear whether the variation among tasks will be hi low for internet agents. It seems unlikely that an agent helping find, or and download physics papers would be much different from one in bic However, the agent helping find biological papers might have very dif constraints than one that visits protein databases to download data. Th paper agent might check with you about costs, but might be given auto to do all the downloading. The database agent might check with you a downloading (which could fill my entire disk if queried carelessly), bu be concerned about costs for public or low-cost databases. Unfortunat a technology this new, the agents' rules are not established by precede

One realm being explored is the so-called *mixed-initiative* approach. I systems the internet agent varies its autonomy based on factors like co the resources needed, or other variables the user might wish to control Sometimes, the system is in charge, pre-authorized to make decisions suggest alternatives. Other times, the user is in charge, taking control steering the decision making.

In another sort of mixed-initiative system the agent essentially 'looks the user's shoulder. Such a system might make suggestions as to whic pages to look at next, predict downloads to pre-process so that by the the user wants to see a particular page it has already been downloaded otherwise take actions directly based on the actions the user is taking interacting with a browser or other network aware system. Mixed-initi systems could also have the potential to learn a users preferences over and customize its actions accordingly.

- **Adaptive**
  The best way for a system to find the appropriate level between aiding user and overstepping the bounds of what the user would prefer, is for learn the users' preferences. Similarly, the internet agent with predeter pages to visit is quite limited. Thus, a truly useful agent should be abl adapt its behaviour based on a combination of user feedback and environmental factors. For example, if the agent visits a web site that longer works, it should learn to stop going there. If the site has change features, the agent should learn the new ones or should ask the user fo in reorganizing.

  Such adaptive behaviour can be achieved in a number of ways. The si way is to group users based on some set of features, and then to assum similarity between them. This can work fairly easily. A new user can a questionnaire that allows the system (using a statistical clustering algorithm) to figure out to which cluster of other users this user belon Preferences associated with that group are then assumed to work for t user.

  As an example, a now defunct movie finder site used to ask the user t on a scale of 1-10 each of ten movies. The users' results were then clu with others who answered the questions in a similar way. That inform was then used for new movie recommendations. Members would ente preferences on other movies, and that would be used to provide 'advic movie choices to others in the same group. A similar technique has be used at a number of sites that sell various commercial products[ii].

  Other forms of machine learning techniques are also being explored f Internet use. One particularly clever application attempts to learn to recognize web advertisements and to strip them from sites (http://www.cs.ucd.ie/staff/nick/research/ae/ for more information and relevant literature citations). Other applications range over a wide vari learning techniques and web behaviours, far too many to review here. Several interesting examples of learning agents have been developed

Carnegie-Mellon University's Text Learning Group (http://www.cs.cmu.edu/afs/cs/project/theo-4/text-learning/www/index.html; other pointers to research in this area can b found at a number of the "agents resource sites" linked to this article).

What is the state of the art in agent design?
To build agents with the desired capabilities is a challenge the computer scie research community is now attacking – but what is the state of the art today bad news is that the sort of internet agents I've been describing so far are few far between – usually running in research labs and not robust enough for ex use. The good news is that many of the components to build such agents are beginning to move beyond research exercises and it is not hard to imagine su agents coming into common use in the next few years. However, there are st number of limiting factors that must be overcome before scientists will be routinely using internet agents to augment their search, downloading and oth research needs.

Building autonomy into web agents is not the limiting factor. This capability even particularly difficult given that modern web-based computer languages Java™ now provide numerous 'class libraries' aimed at providing such applications. In fact, in my undergraduate computer science classes I often h students write such agents as a final project. Using simple tool sets, the stude are able to build agents that wander the web looking for particular concepts a suggesting to the user some web sites to look at. While nice demonstrations, however, these systems are still too limited for use by scientists in commerci applications.

The key limiting factor at present is the difficulty of building and maintainin ontologies for web use. The most basic need in interacting with an agent is a language in which to communicate. While it is possible to 'fake' these seman (with the program reacting appropriately to keywords, for example), an agen is truly useful must have a lot of knowledge about the problem being solved. travel agent doesn't know about geography (Where is the Caribbean?), transportation (What airlines go there?), lodging (Is that a good hotel?), econ (Can I afford to stay there?), *etc*. then we cannot easily communicate our nee the internet agent doesn't understand the area in which it must work (molecu biology, particle physics, *etc*.), it is not able to find appropriate resources any better than current keyword based approaches[iii].

Organizing ontologies
Unfortunately, building these ontologies is a daunting task, especially as extremely detailed knowledge is needed to provide truly useful searches. Eve when produced such ontologies must be brought to the web in a machine rea form. A further problem arises in a trade-off between the depth of the ontolo and the difficulty in encoding the knowledge. Thus, the SHOE system descri above (http://www.cs.umd.edu/projects/plus/SHOE), is aimed at being writte edited by scientific experts without specialized computer science training; bu currently provides a relatively shallow ontology. Other research groups, for example the Knowledge Sharing Laboratory at Stanford University (http://www.ksl.stanford.edu/currentproj.html), are developing more comple ontologies that encode deeper knowledge. However, these approaches requir knowledge of artificial intelligence knowledge representation techniques, something that isn't in the typical training of the average scientist.

It is worth noting that major efforts are underway to overcome these problen to develop new tools for creating ontologies and/or bringing them to the wet High Performance Knowledge Base Initiative, sponsored by the US Defense Advanced Research Projects Agency is one example (see http://www.teknowledge.com:80/HPKB/ for more info). Thus, there is reaso believe that as the current set of web tools (like SHOE) get more capable, an the high end tools (like those at KSL) become more accessible, the bottlenec developing ontologies will be overcome.

Improvement is also being seen in the effort to make agents more capable. M forces are now driving online journals and other scientific content providers explore the greater use of agent-based systems. Current search engines, usin; keyword based techniques, are inadequate for providing the detailed sort of searches needed by the scientific community. Further, XML and other advan web languages are being used to organize scientific material, making it easie web agents to find key aspects of scientific documents (these can be as simp author names and affiliations or as complex as identifying components in sequences described in figures). These languages also make it easier for agei become 'capable' as they can more clearly identify what payment is required information is needed for downloads, *etc*.

The area of agent-based systems is a hot one. We have the technology to bui software agents that are communicative, capable, autonomous and adaptive key behaviours needed to help make our internet journeys more fruitful. The limiting factors in building such systems are being overcome, and new apprc are emerging from information technology research laboratories around the ' In short, if you're not now using agent-based technology, don't worry, you sc will be.

Agent research sites
As with most hot topics on the web, there are many sites which describe woi agent-based systems. The following are some links that may be useful in leai more about research in this area:

- The UMBC Agent Web – http://www.cs.umbc.edu/agents/
- The National Research Council of Canada Agent resource list – http://ai.iit.nrc.ca/subjects/Agents.html
- The French "@gency" Frontpage – http://hoegaarden.iutc3.unicaen.fi bin/pywiki
- The European Community Agentlink project – http://www.agentlink.
- In addition, a number web sites also describe interesting research in ag based systems. As well as those directly cited in above, some others th usually up to date and include demonstrations are
- The University of Washington "Softbots" Project – http://www.cs.washington.edu/research/projects/softbots/www/softbot
- The MIT Media Laboratory Agents Research group – http://agents.www.media.mit.edu/groups/agents/
- The Ariadne Project at the Information Sciences Institute – www.isi.edu/ariadne

(Please note, as with all other web pages, the quality of the pages above is variable and char The above are starting pages that point to many useful web sites, but no endorsement of the:

pages or the information thereon is intended.)

James Hendler
Dept of Computer Science
University of Maryland
College Park, MD 20742
http://www.cs.umd.edu/~hendler

[i]The ontology shown previously is actually represented in a formal machine readable langua
called SHOE. The web-literate reader might wish to look at the document source for the
tseont.html page, and see both the HTML form, which is displayed, and the SHOE form, wl
a set of HTML extensions that allow ontology use on the World Wide Web. See
http://www.cs.umd.edu/projects/plus/SHOE for more details.

[ii] This technique is currently being explored by several popular and large e-commerce sites,
however I will not provide links to commercial sites from this article.

[iii] At the time of this writing, a common internet search engine given the string "Can you fir
papers discussing the instability of retroviral proviruses in chromosomes" returned more tha
1,000,000 hits, with no recent research papers in the top 100 – surely we can do better thar

---