



Association for the Advancement of Artificial Intelligence

445 Burgess Drive

Menlo Park, CA 94025

(650) 328-3123

www.aaai.org

For press inquiries only, contact:

Sara Hedberg

(425) 444-7272

hedberg@aaai.org

Summary of Applications: Backgrounder **Innovative Applications of Artificial Intelligence (IAAI-08) Conference**

For 20 years, the Innovative Applications of Artificial Intelligence (IAAI) Conference has awarded high-impact computer applications that utilize AI in new ways with a significant economic return on the investment (ROI) costs to develop and deliver the system. Typically the ROI is quantitatively measured in dollars saved, staff work hours saved, increased productivity by current staff, etc.

To study the past 20 years of applications is to see the evolution of these systems from rare, standalone applications running on specialized LISP processing machines (the early language of AI), to highly ubiquitous applications running on PCs and various handheld devices threaded through the fabric of our modern lives, corporations, and government agencies.

Several years ago, the IAAI organizers expanded the scope of the conference to include not only fully deployed applications, but also emerging innovative applications under development with preliminary promising results. This year's 22 IAAI papers include four award-winning deployed applications and 18 emerging applications. Together, they constitute a sampling of the myriad ways that AI techniques are being harnessed to solve a wide range of problems and issues around the globe.

This year, the number of award-winning and accepted papers from outside the U.S. has grown to reflect the increasing globalization of the AI community, and the widespread use of AI technology. While there have been a few papers through the years from China, this year finds 2 papers from China, and 2 from Hong Kong. There is an award-winning deployed application from Hong Kong that behind the scenes of the 2008 Olympics Equestrian Competition will be scheduling the requisite staff and volunteers to support the events. The three emerging applications from China include (1) an intelligent system to automatically teach English to middle school and university students, (2) a Semantic Web based system documenting and searching Traditional Chinese Medicine databases and for discovering and learning about Chinese herbal and Western pharmaceutical interactions, and (3) a vehicle delivery route planning system that optimizes the daily delivery schedule for one of the largest chain of restaurants in Hong Kong. Together, these four applications are an interesting sampling of the

increasing global adoption of AI research and technology through the fabric of modern living.

The 2008 IAAI papers fall into several distinct categories. Those that are deployed are *starred*, the others are all emerging, experimental applications still under development and evaluation.

- **Aerospace** (1 paper) (page 3) – this diagnoses faults in electrical power systems
- **Computing** (6 papers) (page 3)- three are network security/management systems, one enables interoperability of legacy systems for the US military, and two are intelligent interface systems
- **Education** (3papers) (page 5) – two automatically teach foreign language (*deployed*), and one helps select the appropriate U.S. public school for an individual student (*deployed*)
- **Environment/Earth Sciences** (4 papers) (page 7) – one predicts air turbulence for aircraft safety, one aids crop selection for sustainable agriculture, one plans flights for airborne Earth science observation missions, and one is in the process of generating maps of the Earth using Landsat sensor data that will be used by the US Geological Survey and most likely the next generation of Google Earth.
- **Medical/Health care** (5 papers) (page 9) – one detects heart disease, one treats epilepsy, a third is a Web-based knowledge repository for traditional Chinese medicine as well as a mechanism for discovering and documenting Chinese herbal/ Western medicine interactions, one monitors surgical activity in the operating room, and one detects handwriting deficiencies for therapeutic purposes.
- **Resources scheduling/ Allocation** (3 papers) (page 11) – one schedules staff and volunteers for the 2008 summer Olympics Equestrian competition (*deployed*), one performs large military convoy management in an urban setting, and the last schedules daily delivery vehicles for one of the largest chain restaurants in Hong Kong.

Aerospace

Diagnosing Faults in Electrical Power Systems for Spacecraft and Aircraft *NASA Ames Research Center (Moffett Field, CA), University of California Los Angeles*

Electrical power systems in aerospace vehicles support critical subsystems such as avionics, propulsion, life support, and thermal management systems. More broadly, electrical power systems play central roles throughout modern society and therefore proper management of their health is critical. This system provides a diagnostic capability for an electrical power system testbed, including a high-level specification language which supports auto-generation of Bayesian networks, and compiles the networks into arithmetic circuits.

Computing

Reinforcement Learning for Vulnerability Assessment in Peer-to-Peer Networks

Oregon State University (Corvallis, OR)

This paper describes the use of reinforcement learning for proactive assessment of computer network vulnerability and security in the context of denial-of-service attacks in peer-to-peer networks. Such a tool could help network administrators and designers assess and compare the vulnerability of various network configurations and security measures to optimize their choices for maximum security. Results of this investigative work show that this RL-based approach is able to significantly outperform a number of heuristic strategies in terms of the severity of the attacks discovered. Results also suggest some possible network design lessons for reducing the attack potential of an intelligent attacker.

Real-time Alert Correlation Using Stream Data Mining Techniques

University of New Brunswick (Fredericton, New Brunswick, Canada)

Network security is vital to our computer-dependent world. Intrusion detection is one of the major techniques for protecting computer networks. This system provides a more condensed, focused and useful view of the network taking real-time alerts, and aggregating and mining them into structure patterns. It provides the system with automatic analysis of alerts, resulting in significant savings of administrative cost, and identifying potential threats more efficiently.

In-the-Dark Network Traffic Classification Using Support Vector Machines

Wake Forest University (Winston-Salem, NC)

This system improved network management by classifying network packet traffic, an important feature for managing such functions as network bandwidth allocation and security policy enforcement. This work addresses the problem of single flow in-the-dark TCP traffic classification, where in-the-dark means that the only parameters available are those that describe the session flow, e.g., packet size, direction and inter-arrival time. Using Support Vector Machines in combination a spectrum representation of packet flows are providing a highly accurate, fast and robust way for classifying common application protocols.

Finding Ontological Correspondences for a Domain-Independent Natural Language Dialog Agent

University of Maryland (College Park, MD, USA)

This experimental system offers an intelligent speech-based dialog interface for various devices (e.g., PDA, GPS) to facilitate communication. For instance, consider an automobile that turns the stereo on and off, adjusts the temperature by receiving verbal commands instead of pushing buttons, or a GPS route planner that negotiates with users about different routes and priorities instead of

the user trying to find how to operate the device and configure his preferences. The agent gathers the necessary information, determines and sends the appropriate commands to the device for execution. Other applications of a dialog agent that converts user utterances into machine understandable commands extend to a number of realms such as robots that could provide services to patients in hospitals or perform routine tasks for the elderly at home that could interact through conversation, and a PDA and schedule planner that communicates with users through a built-in dialog agent rather than current modalities.

Personalization of Telecommunications Services as Combinatorial Optimisation

BT (UK), 4C, UCC (Ireland)

This emerging application allows users to select and compose telecommunications network applications and features (e.g., call-divert-on-busy, find-me, etc.), including finding optimal reconfigurations of network features when a user's preferences violate the technical constraints. Such service personalization refers to the problem of selecting which features should be active when, and how to resolve conflicts (e.g., common signals being used for different purposes, or "call-divert at lunch time" and "voice-mail on Fridays").

Enabling the Interoperability of Large-Scale Legacy Systems

Knexus Research Corp. (Springfield, VA), CDM Technologies (San Luis Obispo, CA), Naval Research Laboratory (Washington, D.C.)

This emerging system addresses the problem of achieving interoperability of large-scale legacy relational databases via an Intelligent Mapping Toolkit, which enables mapping of metadata and instances between relational data models. It uses a distributed "federation" of mapping agents case-based similarity assessment and learning to semi-automatically acquire domain specific lexicons and thesauri to improve mapping between different databases.

The approach was evaluated on the United States Transportation Command reference data, which are shared across client military organizations nationally and internationally. The USTRANSCOM is responsible for the movement of all military personnel and materials worldwide. It must match data entities such as airports, equipment and product codes across databases. Currently at USTRANSCOM, 25 full-time staff members maintain and distribute over 800 data entities to over 1,000 client applications, with four full-time analysts performing the mapping. This method explores automating all or part of the mapping task to significantly speed it or and reduce human errors. Initial testing indicates the IMT significantly reduces mapping errors and can reduce staffing requirements for this task.

Education

A Case Study of AI Application on Language Instruction: CSIEC (*Deployed*) *Peking University (Beijing, China)*

Computer Simulation in Educational Communication (CSIEC) is an intelligent web-based human-computer system for English instruction for students in China. Using natural language processing, it includes a dialogue system for instruction, and learning assessment for teachers and students. Functions include grammar gap filling exercises, talk shows, and free chats to accommodate students with different backgrounds and abilities.

English is considered an international language gaining increasing attention in developing countries. In China, English is required in school and higher education. CSIEC supports the learning theory that one of the best ways to learn a foreign language is frequent communication with a native speaker. It bridges the gap of there not being enough qualified speakers to provide one-to-one student/teacher interaction by providing an interactive web-based conversational partner, or intelligent “chatbot” accessible via text or voice.

There are more than 30,000 registered users. The system has been available free on the Internet for six months. It is also integrated into English classes in various universities and middle schools, where 60.5% of users rate it as “like” or “like very much”. In a junior middle school class that used the system, exam scores improved from 64.39 to 90.81 demonstrating higher English proficiency as a result of using the chatbot. The system is still being further refined and enhanced, and the development team is also exploring other applications of the underlying architecture for related fields, e.g., computer aided testing for language learning, computer assisted writing and translation, etc.

Tactical Language and Culture Training Systems: Using Artificial Intelligence to Teach Foreign Languages and Cultures (*Deployed*)

Alelo, Inc. (Los Angeles, CA)

This system helps people quickly acquire communicative skills in foreign languages and cultures through an interactive game format. Between 20,000 and 50,000 learners worldwide have used TLCTS courses, largely by US marines and soldiers, as well as other military service members in other countries. Over 1,000 people download copies of courses each month – in a typical month there are about 910 downloads of Tactical Iraqi, 115 of Pashto, and 146 of French. Depending on the type of exercise, the system can give feedback on pronunciation, morphological and grammatical forms, word choice, or cultural pragmatics.

Anyone with a .mil email account can register and download copies either for their own use or for installation in computer labs. Currently courses are offered in Tactical Iraqi, Tactical Pashto, and Tactical French. Additional courses are under development for businessmen, workers for nongovernmental organizations, and high school/college students.

Many learners who download copies for their own use study them to completion - requiring on the order of 100 or more hours of training. For military unit training, 20 to 40 hours of training are more the norm. In one study of marines trained with Tactical Iraqi who completed their tour of duty in Iraq in December 2007, independent evaluations found the marines trained with Tactical Iraqi were able to perform many communicative tasks on their own without interpreters. This enhanced the battalion's operational capability, enabled the battalion to operate more efficiently, and resulted in better relations with the local people.

The system utilizes automated speech recognition, virtual humans with 'believable' behavior, intelligent monitoring of users' progress to guide learners in their training, and intelligent systems content authoring tools. It is currently delivered on videogame-capable personal computers equipped with headset microphones. A version deliverable on iPods is under beta test, and a prototype has been developed for the Sony PlayStation Portable. A pilot Chinese course for college and high school Chinese programs is under development, as well as a pilot Cherokee game intended to help preserve Native American language and culture.

The Law of Choice and the Decision Not to Decide (*Deployed*)

University of North Carolina (Charlotte, NC)

The SmartChoice Program provides parents with school recommendations for individual students based on parents' preferences and students' needs, interest, abilities, and talents. It supports the public school choice at the primary and secondary level that is one of the key elements of the U.S. No Child Left Behind Act of 2001. If a school does not meet assessment goals for two consecutive years, by law the district must offer students the opportunity to transfer to a school that is meeting its goals.

The system employs content-based recommender systems techniques to generate a personalized list of schools that have a demonstrated track record with children who share a child's interest, skills, and needs. The beta version of the system is deployed and live for a focus group of 50 participants who have been using it for the January and March/April 2008 Charlotte-Mecklenburg school choice period.

Environment/Earth Sciences

Application of Artificial Intelligence to Operational Real-Time Clear-Air Turbulence Prediction

National Center for Atmospheric Research (Boulder, CO)

University of Colorado at Boulder (Boulder, CO)

Air turbulence prediction is critical to the safety of the millions of people who fly every year. Although fatalities are low, 65% of all weather-related commercial aircraft incidents can be attributed to turbulence encounters, and

major carriers estimate they receive hundreds of injury claims and pay out tens of millions of dollars annually.

The current automated turbulence forecasting system, which was funded by the Federal Aviation Administration and used by the National Oceanic and Atmospheric Administration's Aviation Weather Center, is based largely on subjective observations of 'moderate' turbulence reported by pilots and large-scale quantitative atmospheric data to identify regions where aircraft-scale eddies are likely to form using fuzzy logic.

Recently, a better source of turbulence observations, termed, in-situ data, has become available from the sensors actually on aircraft. The sensor data is so much more plentiful than pilot observations that researchers now have enough data to explore additional AI techniques for forecasting. This paper describes two such methods being explored - support vector machines and logistic regression - both with promising results.

Learning to Improve Earth Observation Flight Planning

NASA Ames Research Center (Moffett, CA)

This system integrates machine learning with planning and data visualization for the management of mobile sensors for Earth science investigations related to understanding changes to the Earth's ecosystem (e.g., climate change or pollution.) Typically a mission will last several weeks. Because the process under investigation (e.g., a hurricane or pollution plume) may be changing daily, a cycle of planning, observation, analysis of data acquired from observation, and model revision and prediction occurs throughout the mission. This system supports what are called mixed observation missions which combine observations from heterogeneous ground, airborne, or space-borne sensors which must be managed in concert.

Local Search for Optimal Global Map Generation Using Mid-Decadal Landsat Images

NASA Ames Research Center, Goddard Space Flight Center, Aerospace Corporation

This application is being developed to help automate the selection of images from the Landsat data from 2004-2006 to construct a data map that will be distributed to the public at no charge through a US Geological Survey website. In addition to providing benefits to researchers in the Earth Sciences, it will likely become the next generation backdrop for Google-Earth. Eventually, over 300,000 images must be evaluated to select the roughly 9,500 that will compose the map. A number of factors must be considered, such as selecting the most cloud-free image, and selecting image data that are seasonally consistent with neighboring scenes. In addition, since there was a malfunction in the image scanner on Landsat 7 since 2003, the image covers only 78% of the land area, so the system must compensate by combining two images of the same scene taken on different days to produce a composite image that partially or fully closes the gaps.

Future plans include using the system to sift through the data from the Landsat mission that is expected to continue until 2012, to create a state mosaic for all 50 states in the U.S., and to generate complete moon maps using Clementine image data. It should also assist Earth scientists by providing improved automated tools to study the Earth's changing ecosystem.

Crop Selection for Optimal Soil Planning using Multi-objective Evolutionary Algorithms

Universidad Nacional de Asuncion (San Lorenzo, Paraguay)

This experimental system aids farm managers develop optimized plans for which crops to cultivate based on soil characteristics. This minimizes the need for soil treatment, thereby reducing the costs and potential environmental damages. It is geared towards sustainable land use practices, with benefits that decrease rural poverty, protect watersheds, increase biodiversity, increase sustainable agricultural production, and increase food security. The system employs multi-objective evolutionary algorithms using 8 processors.

Medical/Health Care

Learning Sparse Kernels from 3D surfaces for Heart Wall Motion Abnormality Detection

Siemens & Motorola, USA

Coronary heart disease (CHD) is a global epidemic and the leading cause of death worldwide, killing 17 million annually according to the World Health Organization. CHD can be detected using an echocardiogram - an ultrasound of different 2-D cross-sections of the left ventricle of the heart. Echocardiograms are difficult to interpret, even for experts. This project is exploring a novel automatic technique using a machine learning classifier to detect the regional wall motion abnormalities of the LV that indicate coronary heart disease. Experimental results on a set of echocardiograms for 320 cases collected in routine clinical practice at one hospital demonstrate the potential of the proposed approach. The researchers plan experiments to confirm that this methodology is very general and can be applied to any classification problem involving 3D surfaces.

Adaptive Treatment of Epilepsy via Batch-mode Reinforcement Learning

McGill University (Quebec, Canada)

This system utilizes machine learning techniques to optimize the deep brain stimulation strategy for the treatment of epilepsy. It chooses the optimal sequence of stimulation treatments to apply to minimize the frequency and duration of seizures. Currently, this experimental system learns an optimal stimulation policy using labeled training data from in vitro animal brain tissues. Results indicate this is an effective means of reducing the incidence of seizures by 25% while minimizing the amount of stimulation applied by a factor of 10. The next tests will determine whether these results carry over to the human

model of epilepsy. If so, they will mean less risk of damage to brain tissues, and the battery onboard the neuro-stimulator has a much longer life (installing a new battery currently requires surgery).

Semantic Web Development for Traditional Chinese Medicine

Zhejiang University (Hangzhou, China)

This is the first use of Semantic Web technology to build, manage and use the largest computerized knowledge repository of Traditional Chinese Medicine (TCM). TCM is an ancient medical system that still accounts for around 40% of all health care delivered in China. The system bridges the gaps between a plurality of over legacy and heterogeneous relational databases from both TCM and Western Medicine. It enables ontology-based query and search across database boundaries. These databases provide knowledge and information services for TCM practitioners in drug usage, information retrieval, and safety analysis.

Using the Semantic Web, the paper describes a global Chinese herb-Western prescription drug interaction network the researchers have mapped. This interaction network can be mined to discover and interpret interesting patterns of TECM herbal/ Western drug interactions. This project illustrates how the semantic Web can connect data across domain and community boundaries to facilitate inter-disciplinary and cross-cultural studies. It also leverages the progress of biomedical informatics to address the preservation and modernization of the ancient Chinese cultural heritage of TCM.

On-line Recognition of Surgical Activity for Monitoring in the Operating Room

Technische Universitat Munchen (Munich, Germany), LORIA-INRAI (France), Klinikum Rechts der Isar (Munich, Germany)

This system aims to design support systems for the Operating Room ranging from simple tasks such as the activation of lights and calling the next patient, to more complex ones such as automatic reporting. It uses on-line recognition of what is happening inside the OR based on recorded signals. The paper describes using the system to recognize on-line the surgical steps performed by the surgeon during a laparoscopic surgery using 11 real surgeries performed by different surgeons in several ORs. In the future this capability could be used to develop systems to support surgeons, trainees, medical staff and administrative staff within hospitals.

COACH - Cumulative Online Algorithm for Classification of Handwriting Deficiencies

Bar-Ilan University (Ramat-Gan, Israel), University of Haifa (Haifa, Israel)

This innovative system is designed as an online handwriting evaluation tool for classifying and remediating handwriting deficiencies. Standard practice has been to perform such analysis manually by trained therapists at some expense,

and with built-in subjective evaluation. This system lowers the cost of evaluation, increases objectivity and enables repeated testing that can accompany therapy. It uses machine learning and data mining techniques and has achieved an 80% success rate on average, results considered to be very good given that diagnosing writing deficiencies is not an exact science and perfect classification is not expected.

Resource Scheduling/ Allocation

Using AI for Olympic Equestrian Event Preparation (*DEPLOYED*)

City University of Hong Kong (Hong Kong SAR)

This Web-based AI “Workforce Management System” will support scheduling and managing approximately 1,000 part-time volunteers and staff at the 2008 Beijing Olympic Equestrian Events which will be held in Hong Kong. It is delivered via the Web, accessed daily by all volunteers and staff as the main communications medium for rosters, duties, schedules, instructions, messages, notices, etc. The Equestrian Events will include two different competition venues – a main site for jumping and dressage, and a second for the cross country event.

This resource scheduling problem is quite complex requiring assigning each shift and job duty, ensuring the appropriate number of adequately skilled volunteers/staff, while optimizing availability and reducing the amount of travel required between venues. Further complexities include the fact that most of the work shifts are not fixed, vary daily depending on job nature and workload, and the event schedules can change for any number of unexpected event delays and cancellations. (Coincidentally, a typhoon passed through Hong Kong just before the test events and caused some minor delays and schedule changes.)

This system was tested in 2007 at a scaled down dress rehearsal event after which the International Olympic Committee determined that the system passed the readiness requirements for the 2008 Olympics. It will ensure the games run smoothly, can adjust quickly to unexpected events, and will keep everyone involved in sync.

Producing schedules manually would be a large-scale headache. The AI-based system centralizes data in one place, ensuring data consistency and avoiding potential confusion in schedules and assignments. It also represents a tremendous savings in manpower: both in terms of building the schedules and notifying staff/volunteers. Assignments are generated within a fraction of a second.

A Hybrid Approach to Convoy Movement Planning in an Urban City

National University of Singapore (Singapore)

This system dynamically routes convoys of vehicles within a road transportation network in an urban city. It addresses the need for a defensive strategy against when a military convoy transports large volumes of material

and manpower, with or without an armed escort. It plans the movement of up to a hundred convoys through a real road network, and dynamically re-plans when changes occur while the plan is being executed (e.g., road congestion, enemy attacks, etc.) In manual planning, this function took hours to generate and was usually far from optimal. This experimental system was evaluated by end users, and found to render an approximately 20% improvement in solution quality with vast improvements in computation time with respect to the existing non-AI method used. Consequently, the prototype was deployed for field testing.

A Vehicle Routing System to Solve a Periodic Vehicle Routing Problem for a Food Chain in Hong Kong

South China University of Technology (Guangzhou, China), Red Jasper Ltd. (Hong Kong), Hong Kong University of Science and Technology (Hong Kong)

This system has been developed for one of the largest food and restaurant chains in Hong Kong. Its aim is to optimize the supply delivery schedule given a set of constraints such as daily changes in quantities ordered by each of the few hundred stores located all over Hong Kong, required delivery time, and drivers must have substantially the same route each day. Using this experimental system, results shows that at least 14 vehicles could be saved, where each vehicle is valued at one million Hong Kong dollars with 0.3 million Hong Kong dollars of running cost for a total of 18.2 million annually (\$2.34 million U.S.) Using the system shows a possible savings of up to 23% in asset costs and annual operating costs for the food chain.

About AAI

Founded in 1979, the Association for the Advancement of Artificial Intelligence (formerly the American Association for Artificial Intelligence) (www.aaai.org) is a nonprofit scientific membership society devoted to advancing the science and practice of AI. Its mission is to: (1) advance the scientific understanding of the mechanisms underlying intelligent thought and behavior, (2) facilitate their embodiment in machines, (3) serve as an information resource for research planners and the general public concerning trends in AI, and (4) offer training for the current and coming generations of AI researchers and practitioners. The Association has sponsored the annual AAI/IAAI conferences, highly regarded in the AI field, since 1980 and 1989 respectively.