HRL Laboratories: Thinking Outside the Box

APPLIED RESEARCH IN AN AGE OF BOTTOM LINES

I. INTRODUCTION

HRL makes advances in electronics, information & systems sciences, materials, sensors, and photonics: from basic research to product delivery. We are producing pioneering work in high performance integrated circuits, high power lasers, antennas, networking, and smart materials. HRL technologies fly in satellites and fighter jets, ride on diesel locomotives, and support the systems of the future. Each year, HRL's intellectual property base grows with patents and trade secrets in key technology areas.

HRL has a rich history of discoveries and innovations dating back more than 60 years to the days when Howard Hughes first created Hughes Research Laboratories to address the most challenging technical problems of the day. Under that name, and now as HRL, this organization has a long-standing reputation of serving the national interest through contract and internal R&D. We continue to work with government agencies and laboratories, and also collaborate with universities and academic institutions.

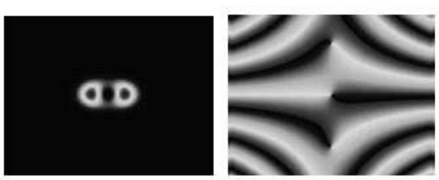

Fig. 1. Amplitude and phase of the wavefunction of two electrons in an anisotropic quantum dot as computed by a few-body code designed by HRL

Over 95% of our our energetic 300-member technical staff have advanced degrees - more than 70% have Ph.D. degrees. We focus on high performance game-changing technologies where we bring unique perspectives and capabilities. Our multi-disciplinary workforce lends itself to development of creative and innovative solutions that cross conventional technology boundaries to produce breakthrough solutions.

This article focuses on the Information and System Sciences Laboratory (ISSL), one of the four technical laboratories at HRL. We will briefly describe the types of research going on there, and then present two representative projects.

II. INFORMATION SCIENCES RESEARCH AGENDA

The ISSL is developing technology to enable smart networks and systems. These are systems that can reason about and adapt to changes in the environment, goals, or their own capabilities, can learn from experience to improve their performance, and can intuitively interact with and respond to their users. This requires broad-based, multi-disciplinary activities in adaptive filtering and learning, human-computer interaction, large-scale networking systems, and computational sciences.

We are combining strengths in mathematics, theoretical physics, computational science and physics-based modeling tools to accurately simulate a variety of important physical phenomena relevant to various experimental groups within HRL (see Figure 1). These models permit realistic analysis of the properties of electronic materials and devices and the phenomena of electromagnetic scattering and propagation.

We apply cognitive science theories to real-world problems, including reasoning by analogy, learning via mental models, and perceiving occluded objects. HRL is actively involved in research on 3D visual and auditory environments, ubiquitous geo-spatial tracking for applications in augmented and virtual reality, and multimodal interaction using dialog and gestures. Applications include command and control, soldier-centric warfare, driver-centric transportation, and remote presence.

In communications and networks, we produced a state-of-the-art wireless platform to analyze and evaluate connectivity, latency, interference, security, quality of services, and congestion issues for a wide variety of application and data networks. Applications include satellite networks, airborne communication networks, vehicular networks, large-scale battlefield networks, and embedded networked sensing systems.

We are developing a single comprehensive architecture to seamlessly integrate perception, memory, planning, decision-making, action, self-learning and affect to address the full range of human cognition. The work focuses on goal-driven scene understanding, language communication, and learning sequentially planned behaviors, as well as on the comprehensive brain-like cognitive architecture.

Fig. 2. A team of pherobots built for Darpa Software for Distributed Robotics program.

We are interested in the dynamics of organization, communication, and control in living organisms, biological systems, and social networks. This is helping us produce high-value systems that exhibit the next-generation capabilities of self-optimization, self-awareness, self-diagnosis, self-regulation, self-healing, self-generation, and reflection.

We are applying evolutionary and neuromorphic techniques to systems where both the software and hardware learn and adapt to their environment. In Fig. 2 we show a swarm of simple robots that coordinate by means of a communications analogue to insect pheromones, to perform mapping of a building and to detect hidden targets.

III. Project Focus on SwarmVision: Advanced Classifiers for Object Recognition and Cognitive Swarms for Fast Search

Objects in a visual scene must be located and classified before they can be combined into events. Typically, classification of objects in an image is performed using features extracted from an analysis window that is scanned across the image. This sequential deterministic search can be very computationally intensive, especially if a small window is used, since a classification must be performed at each window position.

Conventional approaches have utilized motion-based segmentation using background estimation methods to reduce the search space by generating areas of interest around moving objects.This approach fails if the object is motionless or if significant background motion is present, as is the case for motion imagery.

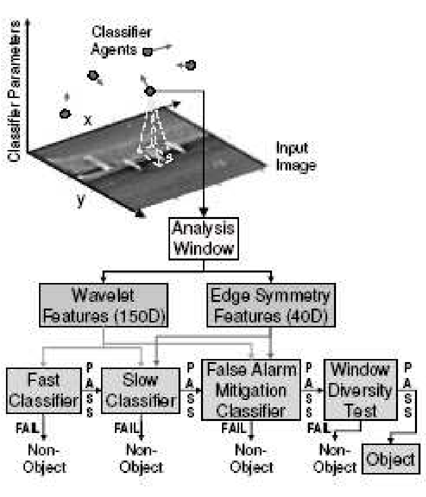

Fig. 3. Cognitive swarms and advanced object classifiers for fast search and object detection in video streams.

HRL's unique cognitive swarm approach to searching for objects combines feature-based object classification with efficient search mechanisms based on the Particle Swarm Optimization (PSO) dynamics developed by Kennedy and Eberhart (1995). Inspired by the flocking behaviors of animals and insects, the PSO algorithm is effective for optimization of a wide range of functions. The algorithm explores a multi-dimensional solution space using a cooperating swarm of search entities or "particles" where the degree of success of each particle in maximizing the objective attracts other members of the swarm. PSO is similar in its generality to genetic algorithms in that it can be used for discontinuous and noisy solution spaces since it only requires an evaluation of the objective function at each particle position; no gradient information or assumptions such as convexity are needed. However, unlike genes in genetic algorithms that compete with each to *win* in a competition for good solutions, in PSO the particles cooperate to explore the solution space and find good solutions. This results in highly efficient search properties. In addition, the evolution of good solutions is stable in PSO (e.g., small changes in the representation result in small changes in the solution), which results in improved convergence compared to GA.

The basic cognitive swarm concept is illustrated in Fig. 3. The objective is to find multiple instances of an object class in an input image. The "cognitive" PSO particles move in a solution space where two of the dimensions represent the x and y coordinates in the video frame. The key concept in our approach is that each particle in the swarm evaluates an objective function value consisting of the classification confidence that the particle's receptive field matches a targeted object in the frame. All cognitive particles in the swarm implement the same classifier, only the classifier parameters vary as the particle visits different positions in the solution space. This recasts the object detection problem as an optimization problem. The solution space dimensions represent location and size of the analysis window and may also include other parameters like rotation.

Cognitive swarms offer a much more efficient method for finding objects in an image compared to searching based on scanning the image, pyramidal approaches, or using gradient information, especially if the scale of the object is not known beforehand. Our experimental results show large speedups over exhaustive search; for example, over 70x speedup to locate and classify one pedestrian of known height (80 pixels) in a 480x700 pixel image.

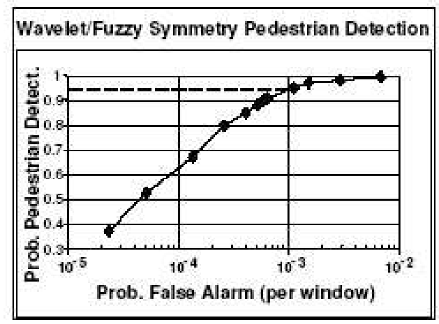

Fig. 4. Window-level probability of detection vs. false alarm rate for human pedestrian classifier that achieves a detection rate of 95% at a very low false alarm rate of 0.1%, and 98% detection for a false alarm rate of 0.3%.

The number of false alarms per image is greatly reduced because the focus of attention of the swarm is quickly directed towards likely objects, which is very important for practical applications (see Fig. 4). The results shown in the figure were obtained on videotaped humans in urban and rural environments under various illumination conditions. The analysis window classification time for our pedestrian classifier is 0.3 msec on a 3 GHz PC. This combination of accuracy and speed is superior to any published results known to us. The framework also provides a natural way to incorporate expectations based on previous recognition results, moving object cues, or externally-supplied rules. For example, if a vehicle has been detected, a human-detection cognitive swarm can be made to focus its attention near the vehicle to "catch" people exiting or entering.

Fig. 3 illustrated some of the object classifiers HRL has developed. This novel approach for object classification utilizes a combination of Haar wavelet and fuzzy edge symmetry features and a cascade of neural network subclassifiers. The features can be calculated quickly using high speed integer arithmetic. A subwindow must be classified as an object by a subclassifier in the cascade in order to proceed to the next (higher complexity) subclassifier. Non-object subwindows are usually rejected early in the cascade, resulting in high speed without sacrificing accuracy.

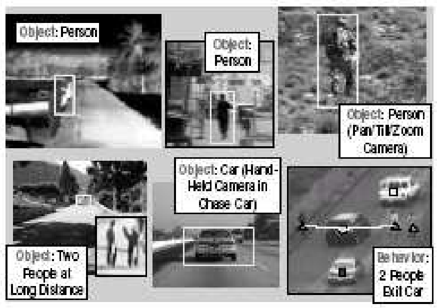

Fig. 5. Detection Examples.

We have used our classifier methodology to create classifiers for other objects as well, such as vehicles and boats. Some example cognitive swarm detections using HRL's advanced object classifiers and cognitive swarms are shown in Fig. 5. HRL has used cognitive swarms successfully in applications for our LLC members, and we are currently adapting them for weapons detection.

IV. Project Focus on System Health Prognosis

Diagnosis of a system determines what failed in the system. It uses observations of the failure such as symptoms of failure, or failed tests. In contrast, prognosis asks what is likely to fail in the near future. It requires not only evidence about present system health as measured by sensors, but also data on health trends, the extent of past system use (e.g., miles, hours, or cycles of operation), and expected future use (possibly focused on a particular mission for which we make prognosis). These multiple pieces of evidence are combined to arrive at system health prognosis.

We have developed a novel framework for prognosis, Fig. 6. The heart of the framework is a probabilistic reasoning engine that produces probability of failure of system components at the end of the mission. It employs a Bayesian network model of the system and multiple sources of evidence for prognosis. The evidence about the previous usage and expected usage for the mission, i.e. future usage, is derived from maintenance and usage data bases and from mission specification. The evidence about present health of components is obtained by applying signal processing and feature extraction algorithms on sensor measurements. Health history of components is used to project health into the future i.e. to the end of the mission. Here trending algorithms are applied to produce the evidence. All elements of the evidence are fused in the reasoner.

Bayesian networks were first proposed as a tool for reasoning in the presence of uncertainty nearly twenty years ago. Many diagnostic systems based on Bayesian networks have been described in literature and some of them have been implemented and deployed in the field. But application of Bayesian networks to prognosis requires a reasoner that is different from those used for diagnosis. An example of a Bayesian network graph developed for a flight actuator is shown in Fig. 7. The graph constitutes a structure of the model. The nodes of the graph are annotated with model parameters, which are conditional probability tables. In Fig. 7 they are shown as histograms.

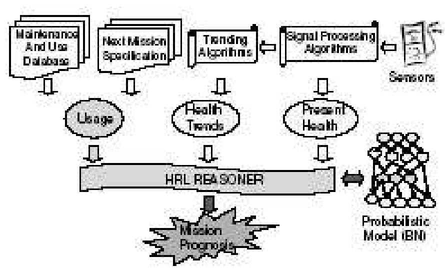

Fig. 6. Prognosis Framework Based on Bayesian Network Models and Probabilistic Reasoner.

In two phases our novel reasoner supports both diagnosis and prognosis. In phase one, a diagnostic phase, the inputs are the evidence on the present system usage and the present health. Given this evidence and system model, the reasoner produces a list of component failures ranked by probability of occurrence. In phase two - prognosis phase - the reasoner takes in evidence on usage for the future mission and evidence on health trends at the end of the mission. The output is a ranked list of probabilities of component failures at the end of the mission.

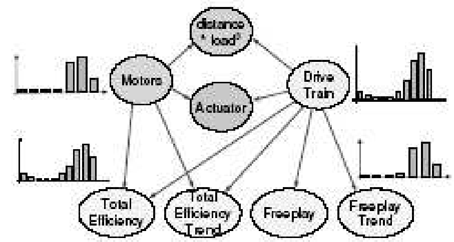

Fig. 7. Bayesian Network Model for Flight Actuator - Structure and Distributions. Motors and Drive Train represent components, Actuator is a subsystem, the node at the top stands for evidence of usage and the four nodes along the bottom represent present and future health evidence.

In addition to the unique reasoner we have also designed a special layered form of Bayesian network. The structure and parameters of the network are customized to diagnosis and prognosis. The layered Bayesian model is much easier to create and requires fewer parameters. Moreover it reduces the computational burden during reasoning. We have developed an editor for the layered Bayesian models, which uses simple tabular representation of the model information. It is intended for experts familiar with the system and does not require knowledge of Bayesian networks. We have also developed a family of software tools for diagnostic/prognostic model evaluation and debugging.

We have used our methodology and tools in development of diagnosis and prognosis solutions for many real-life complex systems including diesel locomotives, automobiles, and aircraft. Our solutions became a part of commercial software provided for some of the systems. We were also successful in extending our methodology and tools to other problems such as decision support for law enforcement and data analysis for homeland security related purposes.

Contact Information
President: Dr. M.W. (Matt) Ganz
310-317-5200
mganz@hrl.com

VP: Dr. C.G. (Conilee) Kirkpatrick
310-317-5374
ckirkpatrick@hrl.com

HRL Laboratories, LLC
3011 Malibu Canyon Rd.
Malibu, CA 90265
http://www.hrl.com