

behavior as their predictability grows during the course of learning. The difference between the actual occurrence and the prediction of the reinforcer is denoted as the reinforcement prediction error signal. In what is known in the art as the temporal-difference model (TD model) of Pavlovian learning, the reinforcement prediction error signal is used to learn a reinforcement prediction signal. The desired reinforcement prediction signal is defined as a weighted sum of future reinforcement signals wherein reinforcements are progressively discounted with increasing time into the future. The reinforcement prediction error signal progressively decreases during learning as the reinforcement prediction signal becomes more similar to the desired reinforcement prediction signal.

- Algorithms based on the TD model ("TD algorithms") have been analyzed and shown to be implementations of a variant of dynamic programming. Machine-learning studies have shown that TD algorithms are powerful means of solving reinforcement learning problems that involve delayed reinforcement.

ment. Examples of such problems include board games, which involve delayed rewards (winning) or punishments (losing).

- Mid-brain dopamine neurons are so named because they modulate the levels of activity of other neurons by means of the neurotransmitter chemical dopamine. The cell bodies of dopamine neurons are located in the brain stem and their axons project to many brain areas. Activities of dopamine neurons have been found to be strikingly similar to the prediction error signals of the TD model.
- Real neural signals include spikelike pulses, the times of occurrence of which are significant. The term "biological spike coding" denotes, essentially, temporal labeling of such pulses in real neurons or in a neural-network model.

This concludes the background information.

The present neural-network model incorporates biological spike coding along with some basic principles of the learning by synapses in the cortex of the human brain. According to the learning rule of the model, synaptic weights are adapted when pre- and

postsynaptic spikes occur within short time windows. In simplified terms, for a given synapse and time window, the synaptic strength is increased in the long term if the presynaptic spike precedes the postsynaptic spike or is decreased in the long term if the presynaptic spike follows the postsynaptic spike. This learning rule has been shown to minimize prediction errors, indicating that the neural network learns an optimal dynamic model of an external process.

This work was done by Tuan Duong, Vu Duong, and Roland Suri of Caltech for NASA's Jet Propulsion Laboratory.

In accordance with Public Law 96-517, the contractor has elected to retain title to this invention. Inquiries concerning rights for its commercial use should be addressed to:

*Innovative Technology Assets Management
JPL*

*Mail Stop 202-233
4800 Oak Grove Drive
Pasadena, CA 91109-8099
(818) 354-2240*

E-mail: iaoffice@jpl.nasa.gov

Refer to NPO-41691, volume and number of this NASA Tech Briefs issue, and the page number.

◉ Evolutionary Computing Methods for Spectral Retrieval

Solutions to the inverse problem of spectral retrieval are found in a computationally efficient process.

NASA's Jet Propulsion Laboratory, Pasadena, California

A methodology for processing spectral images to retrieve information on underlying physical, chemical, and/or biological phenomena is based on evolutionary and related computational methods implemented in software. In a typical case, the solution (the information that one seeks to retrieve) consists of parameters of a mathematical model that represents one or more of the phenomena of interest.

The methodology was developed for the initial purpose of retrieving the desired information from spectral image data acquired by remote-sensing instruments aimed at planets (including the Earth). Examples of information desired in such applications include trace gas concentrations, temperature profiles, surface types, day/night fractions, cloud/aerosol fractions, seasons, and viewing angles. The methodology is also potentially useful for retrieving information on

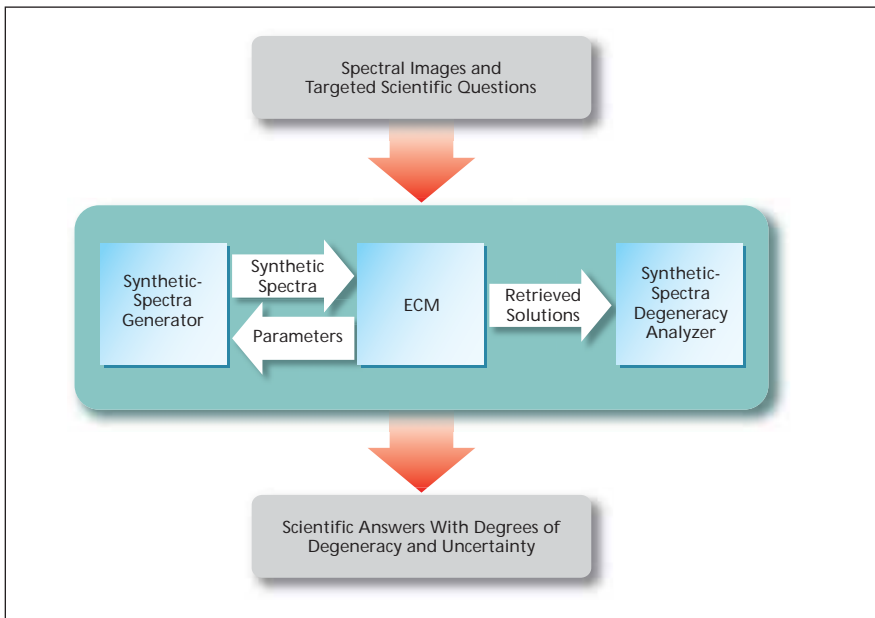
chemical and/or biological hazards in terrestrial settings.

In this methodology, one utilizes an iterative process that minimizes a fitness function indicative of the degree of dissimilarity between observed and synthetic spectral and angular data. The evolutionary computing methods that lie at the heart of this process yield a population of solutions (sets of the desired parameters) within an accuracy represented by a fitness-function value specified by the user. The evolutionary computing methods (ECM) used in this methodology are Genetic Algorithms and Simulated Annealing, both of which are well-established optimization techniques and have also been described in previous *NASA Tech Briefs* articles. These are embedded in a conceptual framework, represented in the architecture of the implementing software, that enables automatic retrieval of spectral and angular data and analysis of the retrieved solu-

tions for uniqueness. This framework is composed of three modules (see figure):

1. The central core, which consists of the aforementioned ECM;
2. The synthetic-spectra generator, which, coupled with the ECM, generates a population of automatically retrieved spectral solutions; and
3. The synthetic-spectra degeneracy analyzer ("degeneracy" is used here in the mathematical sense of signifying the existence of multiple equally valid solutions for a given set of data and user-defined accuracy), which applies several well-established mathematical methods to characterize the uniqueness (or the degeneracy, which is essentially the lack of uniqueness) of the solutions within the population.

One advantage afforded by this ECM-based methodology over traditional spectral retrieval methods is the ability to perform an automatic, unbiased search for all solutions within the entire parameter



This Conceptual Framework for Spectral Retrieval, and the software that implements it, comprises three modules that interact in an iterative process.

space, using criteria that make searching computationally far more economical than in complete-enumeration (“brute force”), Monte Carlo, or random

searches. (As used here, “unbiased” characterizes a search that does not depend on initial *ad hoc* guesses by experts.) Other advantages include the following:

- Optimal solutions are found;
- Better interpretations of planetary spectral and angular data are possible, and initial tests have shown these interpretations to be consistent with ground truth; and
- The methodology is not limited to specific problems, and can be extended to solve problems of greater complexity.

This work was done by Richard Terrile, Wolfgang Fink, Terrance Huntsberger, Seungwon Lee, Edwin Tisdale, Paul Von Allmen, and Giovanna Tinetti of Caltech for NASA’s Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).

In accordance with Public Law 96-517, the contractor has elected to retain title to this invention. Inquiries concerning rights for its commercial use should be addressed to:

*Innovative Technology Assets Management
JPL*

*Mail Stop 202-233
4800 Oak Grove Drive
Pasadena, CA 91109-8099
(818) 354-2240*

E-mail: iaoffice@jpl.nasa.gov

Refer to NPO-42564, volume and number of this NASA Tech Briefs issue, and the page number.

Monitoring Disasters by Use of Instrumented Robotic Aircraft

Real-time synoptic data would help in coordinating and planning responses.

Ames Research Center, Moffett Field, California

Efforts are under way to develop data-acquisition, data-processing, and data-communication systems for monitoring disasters over large geographic areas by use of uninhabited aerial systems (UAS) — robotic aircraft that are typically piloted by remote control. As integral parts of advanced, comprehensive disaster-management programs, these systems would provide (1) real-time data that would be used to coordinate responses to current disasters and (2) recorded data that would be used to model disasters for the purpose of mitigating the effects of future disasters and planning responses to them.

The basic idea is to equip UAS with sensors (e.g., conventional video cameras and/or multispectral imaging instruments) and to fly them over disaster areas, where they could transmit data by radio to command centers. Transmission could occur along direct line-of-sight paths and/or along over-the-horizon paths by relay via spacecraft in orbit around the Earth. The initial focus is on

demonstrating systems for monitoring wildfires; other disasters to which these developments are expected to be applicable include floods, hurricanes, tornadoes, earthquakes, volcanic eruptions, leaks of toxic chemicals, and military attacks.

The figure depicts a typical system for monitoring a wildfire. In this case, instruments aboard a UAS would generate calibrated thermal-infrared digital image data of terrain affected by a wildfire. The data would be sent by radio via satellite to a data-archive server and image-processing computers. In the image-processing computers, the data would be rapidly geo-rectified for processing by one or more of a large variety of geographic-information-system (GIS) and/or image-analysis software packages. After processing by this software, the data would be both stored in the archive and distributed through standard Internet connections to a disaster-mitigation center, an investigator, and/or command center at the scene of the fire.

Ground assets (in this case, firefighters

and/or firefighting equipment) would also be monitored in real time by use of Global Positioning System (GPS) units and radio communication links between the assets and the UAS. In this scenario, the UAS would serve as a data-relay station in the sky, sending packets of information concerning the locations of assets to the image-processing computer, wherein this information would be incorporated into the geo-rectified images and maps. Hence, the images and maps would enable command-center personnel to monitor locations of assets in real time and in relation to locations affected by the disaster. Optionally, in case of a disaster that disrupted communications, the UAS could be used as an airborne communication relay station to partly restore communications to the affected area.

A prototype of a system of this type was demonstrated in a project denoted the First Response Experiment (Project FiRE). In this project, a controlled outdoor fire was observed by use of a thermal multispectral scanning imager on a