Evolutionary Computational Methods for the Design of Spectral Instruments

Richard J. Terrile, Seungwon Lee, Giovanna Tinetti*, Wolfgang Fink, Paul von Allmen and Terrance L. Huntsberger

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109
818 354-6158
rich.terrile@jpl.nasa.gov

*- European Space Agency
University College London
Gower Street, London, WC1E6BT, UK
+44 (0)7912509617
g.tinetti@ucl.ac.uk

Abstract—We have developed a technique based on Evolutionary Computational Methods (ECM) that allows for the automated optimization of complex computationally modeled systems. We have demonstrated that complex engineering and science models can be automatically inverted by incorporating them into evolutionary frameworks and that these inversions have advantages over conventional searches by not requiring expert starting guesses (designs) and by running on large cluster computers with less overall computational time than conventional approaches. We have applied these techniques to the automated retrieval of atmospheric and surface spectral signatures from Earthshine observational data. We have demonstrated that in addition to automated spectral retrieval, ECM can also be used to evaluate the discriminability of scientific results as a function of requirements placed on the spectral model. An important application of this technique is for the optimization of design parameters for spectral instruments. [1] [2]

TABLE OF CONTENTS

1. INTRODUCTION

We have developed a technique based on Evolutionary Computational Methods (ECM) that allows for the automated optimization of complex computationally modeled systems. An important application of this technique is for the optimization of design parameters for spectral instruments. Evolutionary computation is a method that operates on a population of existing computational-based engineering models (or simulators) and competes them using biologically inspired genetic operators on large parallel cluster computers. We have demonstrated that complex engineering and science models can be automatically inverted by incorporating them into evolutionary frameworks and that these inversions have advantages over conventional searches by not requiring expert starting guesses (designs) and by running on large cluster computers with less overall computational time than conventional approaches [1,2,3]. The result is the ability to automatically find design optimizations and trades, and thereby greatly amplify the role of the system engineer. ECM was originally developed for the automated retrieval of spectral data [4,5]. In this application we randomly vary spectral input components and solve for a synthetic spectral fit to real data. We have found that when actual observational parameters are used as inputs, we get a very accurate fit to the data. However, the ECM technique also allows us to determine the range of input conditions that will also produce good fits to the data. Plotting this range of solutions on a principal components diagram allows us to determine the degeneracy of solutions that have non-ground truth inputs yet still fit the data. This is a way of demonstrating the robustness of discriminability of the spectral technique. We are now examining a secondary application of ECM where we can co-evolve the weighting of components of the spectral fitting (e.g. line centers, continuum values, spectral range, signal to noise, etc.). These parameters are tuned to maximize the spectral degeneracy in the principal components diagrams and thereby can be implemented to specify instrumental parameters and design requirements that optimize spectral discrimination.

In this paper we will describe the ECM application to the automated retrieval of atmospheric and surface spectral signatures from Earthshine observational data. We will also show how the synthetic spectral model can be used to define instrument requirements to optimize the discrimination of spectral features.

2. EARTHSHINE SPECTRAL DATA

Earthshine as an Analogue for Extrasolar Planets

The next generation of terrestrial planetary exploration missions (NASA Terrestrial Planet Finder-Coronagraph

[1] 1-4244-1488-1/08/$25.00 ©2008 IEEE.
[2] IEEEAC paper #1194, Version 6, Updated October 23, 2007

(TPF-C), ~0.5-1.6 microns, and NASA Terrestrial Planet Finder-Interferometer (TPF-I)/ESA-Darwin, ~6.5-18 microns [6,7] has been designed to provide spectral information about terrestrial planets outside our solar system. Daunting technological challenges prohibit the ability to obtain spatial resolution on the extrasolar planets, and severely limit both the spectral and temporal sampling of even the most interesting discovery. The spectral information provided would be therefore averaged over the visible planetary disk and the exposure time, which may be hours, days, or weeks, depending on the target. Due to these reasons, the interpretation of the observed spectrum may not be unique and instead a family of solutions (i.e., degeneracy) will provide an equally good explanation of the spectral features within a given accuracy.

Previous research in this area has focused on measurements and preliminary interpretation of the Earthshine on the non-illuminated side of the Moon [8,9], models of diurnal photometric variability on an Earth-like planet and models that simulate disk-averaged spectra from a description of a spatially resolved terrestrial planet on diurnal or seasonal time scales [10,11]. These 3D models are *direct models,* i.e., they compute observable quantities starting from a set of model parameters. Analysis of Earthshine reflected off of the Moon (Figure 1) provides an excellent analogue of an extrasolar planet spectra integrated into one spatial resolution element. Additionally, knowledge of the ground-truth observational conditions of the Earth at the time of the observations allows comparison of spectrally retrieved results.

Figure 1: Earthshine illuminated "dark" side of Moon. Copyright © 2005 by Jerry Lodriguss

Classical Retrieval Methods

Retrieval in general is the process that solves the *inverse problem*, which is the determination of the model parameters for a given set of observed quantities. The *inverse problem* in spectral retrieval consists of determining, given an observed spectrum, the combination (*solution*) of planetary and atmospheric conditions at the time of the observation. The existence of multiple equally valid solutions (i.e., degeneracy), and the lack of sensitivity both affect the inversion process often to the point of preventing a successful retrieval. This problem has been studied extensively using a number of different approaches. A good theoretical understanding of the inversion process has been achieved in studies of dynamical systems, information theory and complexity theory, but the underlying techniques have not yet been systematically applied within the atmospheric remote sensing community.

The traditional retrieval techniques developed to study the environments of the Earth and other planets in the solar system, are inadequate to analyze disk-time averaged spectra because they assume spatially homogeneous environments and short observational time scales [12]. Moreover, traditional techniques benefit from a fairly good knowledge of the environment under investigation. This information is used to constrain the initial parameter set, and the retrieval method will search for the best possible solution that reproduces the observed spectrum in the local domain of the initial condition. The assumption that the solution needs to belong to the local domain of the initial condition is an unavoidable limitation of traditional gradient-based methods [13].

We have employed Evolutionary Computational Methods (ECM), developed by the Center for Evolutionary Computation and Automated Design (CECAD) at the Jet Propulsion Laboratory, Caltech, to automatically retrieve planetary and atmospheric information from disk-averaged planetary spectra [4,5]. ECM has been coupled to the "direct" model developed by Tinetti et al. [10,11] that simulates the 3-D spectral response of the planet. ECM is used to optimize planetary parameters of interest (here surface type and cloud fraction) in an iterative process that minimizes a fitness function measuring the degree of similarity between observed and synthetic spectra. The specific functional formula of the fitness function depends on spectral resolution, spectral integral, and signal-to-noise ratio of the observed data. Repeated application of ECM automatically yields a population of solutions (parameter sets) within the user-defined accuracy (fitness).

Using this 3-D retrieval method the final fit between the observed spectrum and the synthetic one is no longer heuristic with respect to the space inhomogeneities, but is the result of an accurate screening of possibilities such as surface types, viewing geometries, phases, and illumination. The advantage of ECM over traditional retrieval methods is the ability to perform an automatic unbiased search (not dependent on initial, ad-hoc expert guesses) for all solutions within the entire model-defined parameter space, using search criteria that are computationally far more economical than complete enumeration (brute force), Monte Carlo, or random searches. The only a priori information used is what is built into the synthetic spectral models employed.

Figure 2: Phase and viewing geometry observed by Woolf et al. (2002) during their Earthshine measurement.

We present here first results of this method applied to the retrieval of surface and atmospheric parameters from an Earth disk-averaged spectrum in the optical, observed by Woolf et al., [8]. Further, we estimate and analyze the degree of degeneracy of the retrieved solutions.

3. Technical Approach and Methodology

The developed spectral retrieval framework is composed of three modules: the central core is ECM; on the front end is the Synthetic Spectra Generator, which, coupled with the ECM, generates a population of automatically retrieved spectral solutions; on the back-end is the Synthetic Spectra Degeneracy Analyzer, which allows an analysis of the uniqueness of individual solutions within the population. Figure 3 shows a schematic diagram of the ECM-driven spectral retrieval framework.

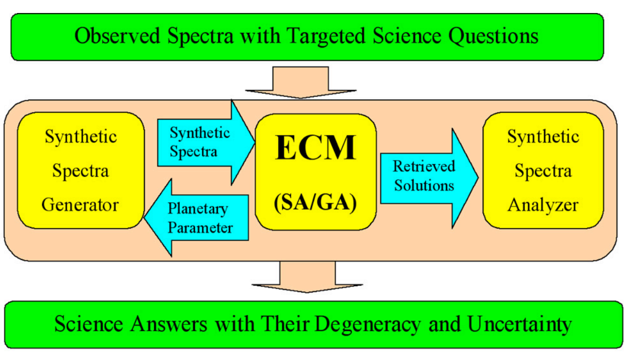

Figure 3: Schematic diagram of the spectral retrieval framework.

Evolutionary Computational Methods (ECM)

ECM is comprised of two well-proven multi-dimensional stochastic and evolutionary optimization techniques: Genetic Algorithms (GA) and Simulated Annealing (SA). Unlike other traditional optimization algorithms, both SA and GA are not fundamentally limited by restrictive assumptions about the search space such as continuity and existence of derivatives. GA and SA have been successfully used for a variety of high-dimensional optimization problems in space systems [1,14] and engineering and biomedical applications [15,16]. SA has also been applied to the problem of Rutherford backscattering spectral retrieval [17].

Genetic Algorithms (GA) —Genetic algorithms [18,19] rely on biologically inspired computational techniques that utilize evolutionary operators on a population of individuals. The process starts with an initial population of individuals (potential solutions), which then undergoes a sequence of unary (mutation) and higher order (crossover) transformations. The individuals strive for survival: a selection (reproduction) scheme, biased towards selecting fitter individuals, produces the individuals for the next generation. After several generations, the population converges to a set of optimal solutions, which represent the intrinsic degeneracy in the search space.

Simulated Annealing-related Algorithms (SA)—The objective of SA [20,21] is to minimize an energy function E, which is a function of N variables (the observational and compositional parameters for the synthetic spectral data). The minimization is performed by randomly changing the value of one or more of the N variables and reevaluating the energy function E. Two cases can occur: 1) the change in the variable results in a new, lower energy function value; or 2) the energy function value is higher or unchanged. In the first scenario, the new set of variables is stored and the change accepted. In the second scenario, the new set of variables is only stored with a certain likelihood (Boltzmann probability, including an annealing temperature). This process reduces the probability that the optimization algorithm becomes "trapped" in local minima as can be the case with "greedy" downhill optimization techniques (e.g., gradient-descent). The annealing temperature directly influences the Boltzmann probability by determining the likelihood of accepting an energetically unfavorable step. The temperature is gradually decreased (cooling schedule) and the overall procedure is repeated until the annealing temperature has reached its end value, a preset number of iterations has been exceeded, or the energy function has reached an acceptable user-defined level.

Synthetic Spectra Generator

The Synthetic Spectra Generator generates a library of synthetic spectra using radiative transfer and planetary models as well as planetary and atmospheric parameters as input (e.g. the size of the planet, gas-mixing ratio of the atmospheric components, temperature and pressure profiles, surface albedo, cloud/aerosol optical properties, stellar spectrum, etc).

In this work, the optical radiances were generated using the Spectral Mapping Atmospheric Radiative Transfer (SMART) model [22,23] for the clear sky and cloudy cases. The inputs to this model were the vertical profiles of temperature and gas-mixing ratios extracted from the Atmospheric Infrared Sounder (AIRS) Level 2 Simulation System [24] simulations. Estimates of surface reflectance and emissivity are based on ASTER data.

The planetary 3D geometry and disk-averaging technique we chose for our calculations, is described in Tinetti et al., [10,11]. The model uses a partition of the sphere, Healpix (Hierarchical Equal Area and Iso-Latitude Pixelization) that was originally implemented for the NASA-WMAP mission [25]. The planetary sphere can be resolved with an arbitrary number of pixels. A library of spectra can be built, (Figure 4) by running the radiative transfer codes for each pixel, for a variety of situations, including temperature profile, gas mixing ratios, surface type (ocean, vegetation, desert, ice etc.) or the cloud/aerosol type (Cirrus, Alto-Stratus etc.), viewing and stellar angles. By specifying the positions of the observer and the star, which determine the phase and the viewing geometry, a disk-averaged spectrum is computed by integrating the area-weighted pixels over the visible disk. A time-averaged spectrum can be obtained by integrating over time the contributions of the disk-averaged spectra for a rotating planet seen from a specified viewing point.

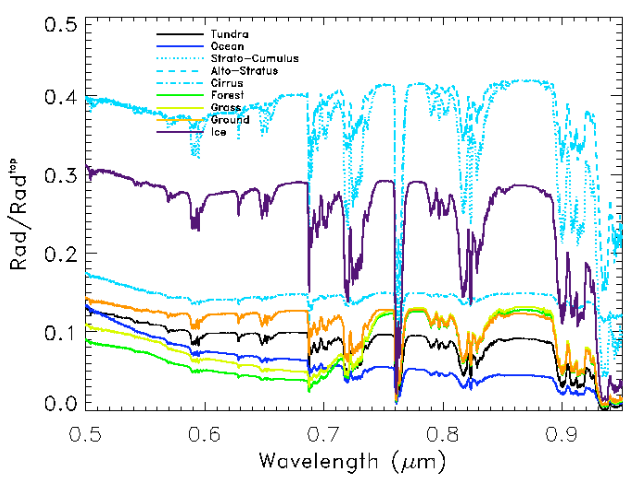

Figure 4: Disk-averaged spectra of the Earth in the optical, showing a planet covered by one surface or cloud type at a time. The phase selected for these simulations is the one shown in Figure 2.

Synthetic Spectra Degeneracy Analyzer

The solutions obtained with ECM are points in the high-dimensional solution (parameter) space. To characterize the degeneracy of the solutions found by ECM, the Synthetic Spectra Degeneracy Analyzer relies on several well-established mathematical techniques such as Level Set Analysis (LSA) and Principal Component Analysis (PCA).

Level Set Analysis—LSA groups ECM results into a set of clusters. If more than one cluster can be identified, the conclusion is that the retrieval process has produced a set of equivalent or degenerate solutions, which are precious input for the subsequent science-based analysis. The LSA process starts with assigning a "membership value" to each solution for each cluster. The number of clusters is a fixed parameter and the centers of the clusters in the parameter space are themselves obtained from an optimization process. The membership value is for example a normalized function of the Euclidean distance between solutions in the parameter space. A solution is said to belong to a particular cluster if the membership value is larger than a set threshold. For each fixed number of clusters, the percentage of solutions belonging to a cluster gives a measure of the quality of this particular clustering configuration. The best clustering configuration is the one with the largest percentage of solutions belonging to a cluster. Details about the LSA procedure are given in Huntsberger et al, [26].

Principal Components Analysis (PCA)—PCA is employed to visualize and characterize the solutions found by ECM. A Singular Value Decomposition of the matrix of column vectors representing the entire set of solutions found by a Genetic Algorithm or Simulated Annealing optimization is performed to find the set of deviations in the solutions (principal values) along a corresponding set of orthogonal directions (principal component vectors) which are sorted in descending order so that the first principal component is the direction of maximum deviation and the last principal component is the direction of minimum deviation. A projection of the solutions onto the plane defined by the first two principal component vectors tends to exhibit the largest separation between solutions. The solutions tend to clump together in clusters about an exemplar positioned at the center of each cluster, enabling an improved visualization of potential degeneracy in the solution set.

4. ECM Spectral Retrieval Studies

We have performed a set of ECM-driven retrieval studies in retrieving an Earthshine spectrum, which serves as an analogue for Terrestrial Planet Finder-Coronagraph (TPF-C) data. The Earthshine spectrum was measured by Woolf et al. with the Steward Observatory 2.3m telescope [8].

For the retrieval studies presented here, several approximations are made in order to confine the retrieval parameter space. First, a minimal set of component synthetic spectra is prepared by varying albedo types and stellar and viewing angles. The trace gas and temperature profiles of the atmosphere remain fixed. Nine different albedo types are considered and are characterized as ocean, forest, grass, ground, tundra, ice, high cloud, medium cloud, and low cloud. A component spectrum at a general angle is obtained through linear and bicubic spline interpolations. The resulting library of the component spectra allows us to explore the various configurations of albedo types within a fixed atmosphere profile. Second, the Earth's surface area is divided into 48 equal-area pixels in order to spatially resolve the visible planetary disk, and each of the pixels is assumed to consist of the same albedo properties. This approximation captures the effect of the overall averaged configuration of the albedo types, but eliminates the effect of the spatial distribution of the albedo-type configuration. Note that although the albedo properties are uniform through pixels, the contribution of each pixel to the disk-averaged spectrum is still different due to its relative location to the sun and the viewer.

The observer and solar positions were known from the observation and kept fixed, thereby constraining both the phase and the viewing geometry (Fig. 3) and reducing the number of illuminated pixels to 22. Using the selected observer-stellar positions, a disk-averaged spectrum was generated, for each of the cloud/surface type configuration prescribed by ECM. ECM is used to optimize these different cloud/surface fractions in an iterative process that minimizes a fitness function measuring the degree of similarity between the observed and the synthetic spectrum. Repeated application of ECM automatically yields a population of solutions (parameter sets) within the user-defined accuracy (fitness).

Two distinct retrieval studies were designed using ECM: i) evolution of one large population with 1000 individuals and ii) evolution of multiple (19) islands with 100 individuals in each island. The two retrieval studies are prepared to compare their effectiveness in terms of the fitting quality and degeneracy degree of retrieved solutions. As a first guess, it is expected that the first study has an advantage of a sufficiently large population size over the second study, while the second study has an advantage of diversity promotion over the first study. Besides the population size and the number of populations, the two studies use the same gene representations and algorithmic procedures for selection, reproduction, and replacement steps. A gene is represented by a real-valued parameter, which corresponds to the weight/percentage of one albedo type. A binary tournament is used for selection. Mutation with probability 0.1 per gene and Crossover with rate 0.8 per individual are used for reproduction. The population of the next generation is composed of the top 15% of individuals from the current population and the offspring generated from selection and reproduction procedures.

5. RESULTS

These two retrieval studies returned over 2800 automatically generated retrievals satisfying the error criteria (fitness) of 10% least squares match to the observed spectra. The first study resulted in 990 eligible solutions out of 1000 candidates in the final population. The second study led to 1835 eligible solutions out of 1900. This shows that the resulting fitting quality and success rate are comparable between the two studies. The total computational time for these studies is 20 hours with 16 processors working in parallel on a Linux-based cluster computer with 3.06 GHz Pentium IV CPUs. Figure 5 shows the synthetic spectra generated by several retrieved solutions in comparison with the observed spectrum.

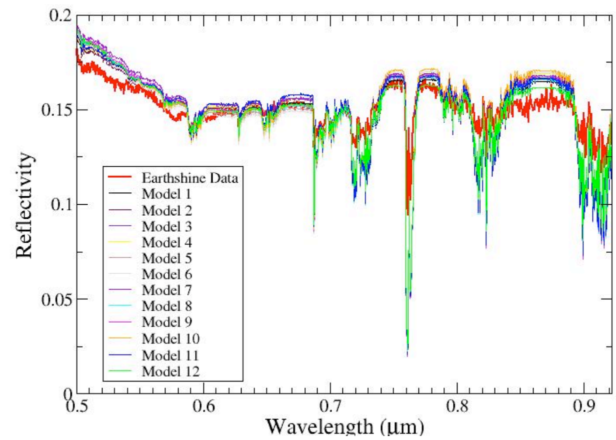

Figure 5: Synthetic spectra fitted to Earthshine data.

The retrieved solutions are processed to determine the degeneracy degree. They are classified into clusters using LSA. Figure 6 displays the albedo component configuration of the 17 identified cluster centers. The LSA-selected cluster centers illustrate the representative picture of the variations of the albedo configurations among all the 2825 retrieved solutions. According to the LSA results, the second study led to more distinct solutions (cluster centers) than the first study, indicating that the second study is more efficient in controlling the diversity of the solutions.

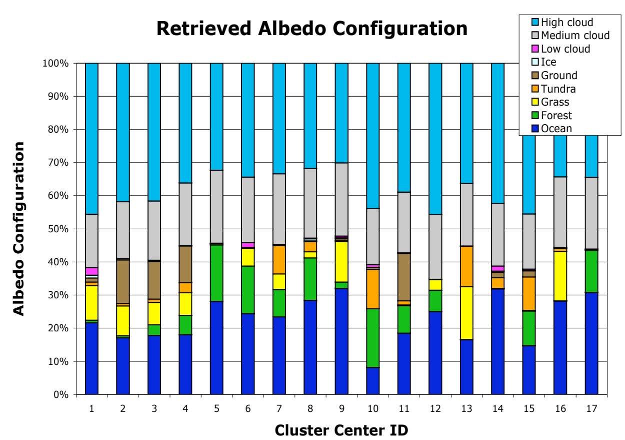

Figure 6: Albedo configurations of the synthetic spectra found by the *Synthetic Spectra Degeneracy Analyzer*. Each configuration leads to a synthetic spectrum that matches the Earthshine data within a 10% error bar.

6. DISCUSSION

The representative solutions of the retrieved solutions, shown in Figure 6, enable a science-based discussion of the results. Some of the cluster centers almost exclusively differ by the ratios of the grass to forest coverage (e.g., cluster 9 has a high ratio, whereas cluster 8 has a low ratio). This result suggests that grass and forest coverage are equivalent with respect to their spectral contribution within error-bars (Figure 7).

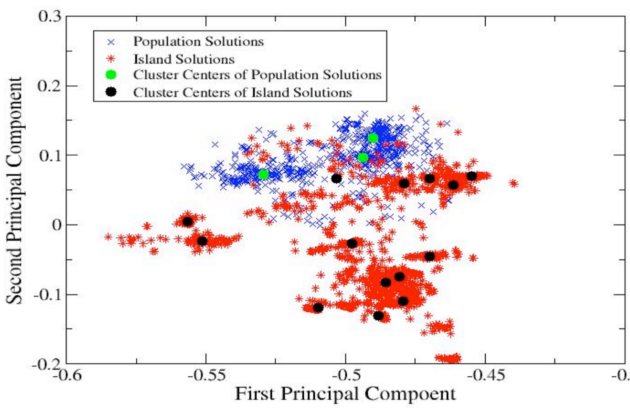

Figure 7: Principal Component Analysis results obtained with the Synthetic Spectra Degeneracy Analyzer.

The cluster center solutions also suggest that low clouds and ice cannot be retrieved from the Earthshine spectrum [8], in agreement with known observational parameters and satellite data (i.e., ground truth). Due to the viewing geometry and season of the year (N.H. summer), in fact, ice was not visible. As concerns low clouds, they are easily masked by medium and high clouds, which reside at higher altitude. It is important to note that these conclusions were reached without introducing any prior knowledge of the system. This demonstrates that the ECM-based retrieval framework introduced here seems to be capable of producing meaningful scientific results with a minimal set of *a priori* assumptions. ECM-based spectral retrieval will become particularly relevant in the context of extra-solar planetary studies, where close to nothing is known about the observed objects.

Outlook

The ECM-based spectral retrieval process can be refined by increasing the parameter space to a large number of components. Further, one can also iterate the manner in which fitnesses are computed by examining large numbers of solutions thereby optimizing fitness functions, i.e., for maximizing discriminability of degeneracies among solutions. In forthcoming work, we will explore/expand the search space as permitted by the scalability of the available computer resources. We will also run numerical experiments to assess the gain/loss in information content induced by an increased spectral resolution, which in turn requires an increased integration time necessary to complete the observations.

We have conducted preliminary experiments using a synthetic target spectrum instead of the observational Earthshine spectrum as a standard reference for automated spectral retrieval using ECM. Retrievals were run for various modifications of the fitness functions. Our standard method uses minimization of the area between the reference and retrieved spectra. Additional methods preferentially weight special regions of the target spectrum (e.g. continuum values, line centers, steep slopes, etc.). The results of PCA were compared for varying constraints on the fitness function to determine if clustering (i.e., spectral degeneracy) can be modulated. Initial results indicate that constraints on instrument requirements (i.e. wavelength range, signal to noise, resolution, etc.) can be determined from these ECM-based retrievals.

7. CONCLUSIONS

We have demonstrated that Evolutionary Computational Methods (ECM) can be used for automatic spectral retrieval and that the results are scientifically consistent with ground truth. We have further demonstrated that we can use clustering tools to discriminate classes of spectral fits and identify degeneracy (non-uniqueness) in solutions. The computational time used in these experiments indicates that full parameter retrievals are achievable with available computational resources in reasonable run-times.

Preliminary experiments using synthetic target spectra indicate that spectral instrument design parameters can be derived from scientific forward models of the observational environment. Evolutionary Computational Methods are applied to the forward models to retrieve large populations of synthetic spectra that can be evaluated under varying fitness conditions to maximize spectral discriminability.

ACKNOWLEDGEMENTS

The work described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

References herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

[1] Terrile, R. J., Adami, C., Aghazarian, H., Chau, S. N., Dang, V. T., Ferguson, M. I., Fink, W., Huntsberger, T. L., Klimeck, G., Kordon, M. A., Lee, S., von Allmen, P. A. and Xu, J. (2005) "Evolutionary Computation Technologies for Space Systems" IEEE Aerospace Conference Proceedings, Big Sky, MT, March 2005.

[2] Terrile, R. J., Kordon, M., Mandutianu, D., Salcedo, J., Wood, E. and Hashemi, M. (2006) "Automated Design of Spacecraft Power Subsystems." IEEE Aerospace Conference Proceedings, Big Sky, MT, March 2006.

[3] Terrile, R. J., Kordon, M., Postma, M., Salcedo, J., Hanks, D. and Wood, E. (2007) "Automated Design of Spacecraft Telecommunication Subsystems Using Evolutionary Computational Techniques." IEEE Aerospace Conference Proceedings, Big Sky, MT March 2007.

[4] Terrile, R. J., Fink, W., Huntsberger, T., Lee, S., Tisdale, E. R., Tinetti, G. and von Allmen, P. (2005) "Retrieval of Extra-Solar Planetary Spectra Using Evolutionary Computational Methods." Bull. Amer. Astron. Soc., 37.

[5] Terrile, R., Lee, S., Tinetti, G., Fink, W., Tisdale, E. R., Huntsberger, T. and von Allmen, P. (2006) "Spectral Retrieval of Exoplanetary Signatures Using Evolutionary Computational Methods." Proceeding of the EGU General Assembly, Vienna, Austria.

[6] Beichman, C.A., N. J. Woolf, and C.A. Lindensmith, *The Terrestrial Planet Finder (TPF),* JPL 99-3,1999.

[7] Lawson, P.R., Unwin, S.C., & Beichman, C.A., 2004, ``Precursor Science for the Terrestrial Planet Finder," JPL Publication 04-014.

[8] Woolf, N.J., P.S. Smith, W.A. Traub, and K.W. Jucks, The spectrum of earthshine: A pale blue dot observed from the ground, *Astrophysical Journal*, *574* (1), 430-433, 2002.

[9] Arnold, L., S. Gillet, O. Lardiere, P. Riaud, and J. Schneider, A test for the search for life on extrasolar planets - Looking for the terrestrial vegetation signature in the Earthshine spectrum, *Astronomy & Astrophysics*, *392* (1), 231-237, 2002.

[10] Tinetti, G., V.S. Meadows, D. Crisp,W. Fong, T. Velusamy, H. Snively, Disk-averaged synthetic spectra of Mars, Astrobiology, 5 (4), 2005.

[11] Tinetti, G., V.S. Meadows, D. Crisp, E. Fishbein, M. Turnbull, and J.P. Bibring, Detectability of Planetary Characteristics in Disk-Averaged Spectra I: the Earth Model, Astrobiology, 2006, Vol. 6, n. 1.

[12] Hanel, R.A., Conrath, B.J., Jennings, D.E., Samuelson, R.E., Exploration of the Solar System by infrared remote sensing (Cambridge Planetary Science Series, Cambridge University Press, 1992).

[13] Rodgers, C.D., Inverse Methods for Atmospheric Sounding: Theory and Practice (Series on Atmospheric, Oceanic and Planetary Physics - Vol. 2, World Scientific Publishing Co., 2004)

[14] Terrile RJ, Aghazarian H, Ferguson MI, Fink W, Huntsberger TL, Keymeulen D, Klimeck G, Kordon MA, Lee S, von Allmen P (2005b) Evolutionary Computation Technologies for the Automated Design of Space Systems; accepted for publication in 2005 NASA/DoD Conference on Evolvable Hardware, IEEE Computer Society

[15] Mitchell, M., "An Introduction to Genetic Algorithms", The MIT Press, 1996.

[16] Aarts, E., J. Korst, Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing (John Wiley & Sons, Chichester, 1989).

[17] Barradas, N.P., C. Jeynes, R.P. Webb, Simulated annealing analysis of Rutherford backscattering data, Appl.Phys.Lett. 71 (1997) 291-3.

[18] Holland, J.H., "Adaptation in Natural and Artificial Systems", The University of Michigan Press, Ann Arbor, Michigan, 1975.

[19] Goldberg, D.E., "Genetic Algorithms in Search, Optimization and Machine Learning", Addison-Wesley, 1989.

[20] Metropolis, N., A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, "Equation of State Calculation by Fast Computing Machines," J. of Chem. Phys., 21, 1087--1091, 1953.

[21] Kirkpatrick, S., C.D. Gelat, M.P. Vecchi,, "Optimization by Simulated Annealing," Science, 220, 671--680, 1983.

[22] Meadows, V.S., and D. Crisp, Ground-based near-infrared observations of the Venus nightside: The thermal structure and water abundance near the surface, *Journal of Geophysical Research-Planets*, *101* (E2), 4595-4622, 1996.

[23] Crisp, D., Absorption of sunlight by water vapor in cloudy conditions: A partial explanation for the cloud absorption anomaly, *Geophysical Research Letters*, *24* (5), 571-574, 1997.

[24] Fishbein, E., C.B. Farmer, S.L. Granger, D.T. Gregorich, M.R. Gunson, S.E. Hannon, M.D. Hofstadter, S.-Y. Lee, S.S Leroy; Formulation and Validation of Simulated Data for the Atmospheric Infrared Sounder (AIRS)}, IEEE Transactions on Geoscience & Remote Sensing, Vol. 41, N. 2, Feb 2003 1.

[25] Gorski, K.M., Hivon, E., and Wandelt, B.D., Analysis issues for large CMB data sets. Proceedings: Evolution of Large Scale Structure-Garching, 1998.

[26] Huntsberger, T.L., C. L. Jacobs, and R.L. Cannon, "Iterative fuzzy image segmentation", Pattern Recognition, Vol. 18, pp. 131-138 (1985).

BIOGRAPHY

Richard J. Terrile *created and directs the Center for Evolutionary Computation and Automated Design at NASA's Jet Propulsion Laboratory. His group has developed genetic algorithm based tools to improve on human design of space systems and has demonstrated that computer aided design tools can also be used for automated innovation and design of complex systems. He is a planetary astronomer and the co-discoverer of the Beta Pictoris circumstellar disk. Dr. Terrile has B.S. degrees in Physics and Astronomy from the State University of New York at Stony Brook and an M.S. and a Ph.D. in Planetary Science from the California Institute of Technology in 1978.*

Seungwon Lee *is a member of the technical staff at the Jet Propulsion Laboratory. Her research interest includes genetic algorithms, low-thrust trajectory design, nanoelectronics, quantum computation, parallel cluster computation and advanced scientific software modernization techniques. Seungwon received her Ph.D. in physics from the Ohio State University in 2002. Her work is documented in numerous journals and conference proceedings.*

Giovanna Tinetti *is a researcher at the Dep. of Physics and Astronomy of the University College London since September 2007. She is author of more than twenty peer-reviewed publications on planetary science, spectroscopy and exoplanet characterization, among which the most recent discovery of water vapor in the atmosphere of an extrasolar planet. After her PhD in theoretical physics from the University of Torino, Italy, she worked for four years at the California Institute of Technology and two years at the Institut d'Astrophysique de Paris.*

Wolfgang Fink *is a Senior Researcher at NASA's Jet Propulsion Laboratory in Pasadena, CA, Research Associate Professor of both Ophthalmology and Neurological Surgery at the University of Southern California, Los Angeles, CA, and Visiting Associate in Physics at the California Institute of Technology, Pasadena, CA. He is the founder and head of the Visual and Autonomous Exploration Systems Research Laboratory at Caltech (http://autonomy.caltech.edu). His research interests include autonomous planetary and space exploration, computational field geology, computer optimization, image processing and analysis, sensor data fusion, astrobiology, and biomedicine. Dr. Fink obtained a B.S. and M.S. degree in Physics and Physical Chemistry from the University of Göttingen and a Ph.D. in Theoretical Physics from the University of Tübingen in 1997. His work is documented in numerous publications and patents.*

Paul von Allmen *is the supervisor of the Applied Cluster Computing Technologies group and a senior researcher at the Jet Propulsion Laboratory. Paul is currently leading research in quantum computing, thermoelectric and nonlinear optics material and device design, nano-scale chemical sensors, semiconductor optical detectors, and low-thrust trajectory optimization. Prior to joining JPL in 2002, Paul worked at the IBM Zurich Research Lab on semiconductor diode lasers, at the University of Illinois on the silicon device sintering process, and at the Motorola Flat Panel Display Division as manager of the theory and simulation group and as principal scientist on micro-scale gas discharge UV lasers. Paul received his Ph.D. in physics from the Swiss Federal Institute of Technology in Lausanne, Switzerland in 1990. His work is documented in numerous publications and patents.*

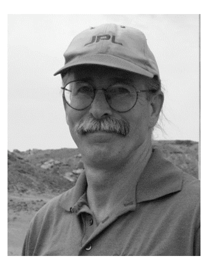

Terrance L. Huntsberger *is a Principal Member of the Technical Staff in the Mobility and Robotics Section at the Jet Propulsion Laboratory in Pasadena, CA, where he is the Manager for numerous tasks in the areas of multi-robot control systems, rover systems for access to high risk terrain, and long range traverse algorithms for rovers. He is an Adjunct Professor and former Director of the Intelligent Systems Laboratory in the Department of Computer Science at the University of South Carolina. His research interests include behavior-based control, computer* *vision, neural networks, wavelets, and biologically inspired system design. Dr. Huntsberger has published over 120 technical articles in these and associated areas. He received his PhD in Physics in 1978 from the University of South Carolina. He is a member of SPIE, ACM, IEEE Computer Society, and INNS.*