

Intelligent Information Extraction to Aid Science Decision Making in Autonomous Space Exploration *

Erzsébet Merényi^a, Kadim Tasdemir^a, and William H. Farrand^b

^aRice University, Department of Electrical and Computer Engineering, Houston, TX, USA;

^bSpace Science Institute, Boulder, CO, USA

ABSTRACT

Effective scientific exploration of remote targets such as solar system objects increasingly calls for autonomous data analysis and decision making on-board. Today, robots in space missions are programmed to traverse from one location to another without regard to what they might be passing by. By not processing data as they travel, they can miss important discoveries, or will need to travel back if scientists on Earth find the data warrant backtracking. This is a suboptimal use of resources even on relatively close targets such as the Moon or Mars. The farther mankind ventures into space, the longer the delay in communication, due to which interesting findings from data sent back to Earth are made too late to command a (roving, floating, or orbiting) robot to further examine a given location. However, autonomous commanding of robots in scientific exploration can only be as reliable as the scientific information extracted from the data that is collected and provided for decision making. In this paper, we focus on the discovery scenario, where information extraction is accomplished with unsupervised clustering. For high-dimensional data with complicated structure, detailed segmentation that identifies all significant groups and discovers the small, surprising anomalies in the data, is a challenging task at which conventional algorithms often fail. We approach the problem with precision manifold learning using self-organizing neural maps with non-standard features developed in the course of our research. We demonstrate the effectiveness and robustness of this approach on multi-spectral imagery from the Mars Exploration Rovers Pancam, and on synthetic hyperspectral imagery.

Keywords: space exploration, autonomous science, knowledge discovery, neural computation, self-organized learning

1. BACKGROUND AND MOTIVATION FOR NEURAL ON-BOARD COMPUTATION

Autonomous operation of unmanned spacecraft, aircraft and other mobile robotic vehicles has long been a goal for application to exploration of remote frontiers of space and hard-to-access or dangerous environments of Earth. Robotic capabilities already exist that allow hazard avoidance by smart navigation, using fast, fault tolerant, reliable on-board computing devices withstanding harsh environments.^{1,2} However, systems do not yet exist that are able to develop and employ sophisticated understanding of science or intelligence data to enable trustworthy autonomous decision making based on information learned in situ. As an example, the Mars Exploration Rovers have excellent hazard avoidance capabilities based on perceived terrain properties, but do not have sufficient on-board understanding of science data to recognize rare minerals or other scientifically interesting surface features or objects. Thus the rovers could not make autonomous decisions to turn attention to interesting science opportunities instead of passing them by according to preprogrammed navigation commands.

*Copyright 2008 Society of Photo-Optical Instrumentation Engineers. This paper was published in the Proc. DSS08 SPIE (Defense and Security Symposium, Space Exploration Technologies, March 17–19, Orlando, FL) Vol. 6960 and is made available as an electronic reprint with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

Further author information: (Send correspondence to Erzsébet Merényi)

Erzsébet Merényi: E-mail: erzsebet@rice.edu, Telephone: 1 713 348 3593

Kadim Tasdemir: E-mail: tasdemir@rice.edu, Telephone: 1 713 348 2285

William H. Farrand: E-mail: farrand@spacescience.org, Telephone: 1 720 974 5825

Similarly for orbiters around planets, the capability of real-time preferential employment of the most appropriate instruments and actions, triggered by on-board alerts from science data, does not yet exist. The longer the delay in communication between Earth and spacecraft, the more scientific opportunities may be lost without such capabilities.

Autonomous commanding of robots, however, can only be as reliable and as sophisticated as the information extracted from the data collected on-board and provided for decision making. In the case of scientific (as well as surveillance and other) applications this often means teasing out, in sufficient detail, relevant information from a mass of complicated high-dimensional (multivariate) data, or from a combination of those. Information extraction and knowledge discovery for complex autonomous science-driven decision making involves two levels:

1. Detailed and accurate automated analysis (deep information mining) of individual data sets collected by various on-board instruments;
2. Automated interpretation (automatic labeling) of objects in scenes by combining the information obtained from the analysis of the individual data sets.

For example, on level 1, class maps of the upper few centimeters of surface covers can be produced from spectral imagery, soil moisture maps from microwave or radar data, maps of the upper half meter of soils from gamma spectroscopy, and elevation maps from laser altimetry, separately. On level 2, these derived products can then be jointly evaluated, along with knowledge from other sources (such as previous analyses, existing geologic maps, spectral libraries, etc.), for higher-level scientific interpretation. An attractive idea on how such higher-level interpretation can be done in an autonomous scenario with fuzzy logic was recently put forward by Furfaro et al.³ The quality of information extraction on level 1 is critical for the level 2 interpretation task.

This paper concentrates on the level 1 task, which involves both

- Finding precisely what we are looking for: targets of known character;
- Discovering what we do not know we should be looking for.

The first item can be done by supervised classification into predefined classes, or by predictive models, which were developed or trained with the known target characteristics. The second item can be accomplished through unsupervised approaches such as Principal Components Analysis, or various clustering methods.

Deep information mining of scientific and other similarly complex data, collected in today's missions, has proved difficult, however, with conventional approaches. A prime example in space and Earth applications is spectral and hyperspectral imagery, acquired in most missions for the wealth of information it contains. Several factors make these data challenging: high dimensionality, many clusters of extremely varying sizes, shapes, densities, and sometimes subtle differences in the spectral signatures of species that are significantly different from an application standpoint. The most interesting scientific phenomena are often spatially small, comprising a statistically insignificant number of pixels, yet their discovery is of utmost importance.

We approach these challenges with neural computation. Neural computing architectures strive to mimic the intelligent information processing of nervous systems, enabled by massive parallelism, dense interconnectivity of many very simple processors (neurons), and other observed properties of brains. The massive parallelism makes neural architectures also suitable for implementation in compact, high-speed hardware, which can be embedded in on-board processing and decision making systems. We employ a specific neural architecture called the self-organizing map (SOM), to capture some ways the cerebral cortex is believed to organize and summarize sensory data, and derive detailed knowledge of the environment.

In this article we focus on the discovery potential of neural maps. We use two data sets to demonstrate the detail that unsupervised neural SOM clustering can produce. One data set is a multispectral image collected by the Mars Exploration Rover Spirit in Gusev Crater. The other data set is a synthetically generated hyperspectral image of an urban scene, with over 70 different materials in it. In what follows we first give a short background on self-organized learning and on some advanced features that we developed in our research, then we present the two case studies. We offer future directions in closing.

2. DISCOVERY THROUGH SELF-ORGANIZED LEARNING

Self-organizing maps (SOMs) were invented more than 25 years ago, and have been used extensively for many applications. For reasons of space constraints we assume that the reader has generally familiarity with SOMs. For a comprehensive review, see ref.⁴ or ref.⁵ for a short summary. Here we only review the main salient points of this paradigm, and some advanced features that we contributed through our research, and which are not available in standard packages.

2.1 Background on Self-Organizing Maps

The neurons in self-organizing neural maps (SOMs) learn to collectively represent the a priori unknown structure of a data set through simultaneous competition and collaboration among the locally acting neural units, in an unsupervised iterative learning procedure. This involves a) finding an optimal distribution of prototype vectors — the neural weights — in the data set (which is an adaptive vector quantization process), and b) simultaneously organizing the prototypes on a rigid lattice, according to their similarity relations. During learning, the prototypes mold themselves to look like the data vectors and similar ones (those that represent a cluster of similar data vectors) occupy contiguous regions in the SOM, after sufficient learning. These regions can be identified by various methods, most commonly by computing the vector distances of the prototypes that are adjacent in the SOM lattice, thereby demarcating the dissimilar groups. Visualization of this information is then used in semi-manual extraction of the clusters. Figure 1, left, shows a typical visualization of a learned SOM, where differences (vector distances) of the adjacent prototypes are visualized as “fences” on a black-to-white gray scale. Black means no difference, white means large difference. This SOM learned a simple synthetic data set containing 20 classes 16 of which are roughly equal in size, approximately 1024 pixels each. Four of the classes are rare, with 1, 16, 64, and 128 pixels, respectively. The single-pixel class is mapped to the center prototype in the bottom row. Immediately to its right the 128-pixel class is represented in the 4-prototype area. The 16-pixel class occupies the diagonal neighbor at the upper right of the latter group. Black prototypes have no data points mapped to them thus represent no clusters, typically between the strong corridors of double fences. The groups of prototypes which represent similar data vectors are well separated, including the representatives of the rare classes.

For noisy real data with many clusters, the representation of rare species (small clusters) can be suppressed in a quantization (SOM or other), or not resolved at all, with the given number of prototypes. Nature’s way to ensure that important rare signals are noticed is through a “perceptual magnet” effect,⁶ which preferentially magnifies the area of the cortex that represents the rare stimuli. Such magnification can be induced in an SOM based on the theory by Bauer et al.,⁷ without prior knowledge of the data distribution. We showed in systematic studies that this theory — formally proved only for 1-D data and 2-D data with uncorrelated dimensions — can be applied to broader classes of data.⁵ When forced “negative magnification” is applied, which essentially commands the SOM learning to be more alert to rare events, small classes are represented by more than their fair share of prototypes, which increases their detectability and thus the chance of their discovery.

We demonstrated this approach to be effective in detailed information extraction from intricate voluminous real data. Real examples are the detection of rare mineralogy in Mars Pathfinder multispectral imagery,⁸ or finding very small spectrally distinct, 3–10 pixel spatial objects from an urban multi- and hyperspectral images.⁹ Importantly, one does not need to know the data distribution, not even the fact that rare clusters exist in the data, for the magnification to take place; thus SOM magnification is truly a discovery tool.

Depending on the degree of magnification some clusters will be enlarged in the SOM at the expense of shrinking the representation area of other clusters, since the SOM has a finite size. Ideally, for general mapping, each data cluster should get a representation (the number of SOM prototypes) proportional to its size. This is effected by a magnification factor of 1, which corresponds to a mapping with maximum information theoretical entropy. This in turn means that the probability density function of the data is approximated by the SOM prototypes in the most faithful way possible. (The original Kohonen SOM is suboptimal in this sense, as discussed in ref.⁵ and references therein.) In the case studies here, we applied this regime. In an autonomous environment one could have two SOMs, one with negative magnification, to work in an aggressive discovery mode for rare species, and another one with maximum entropy mapping, for general mapping.

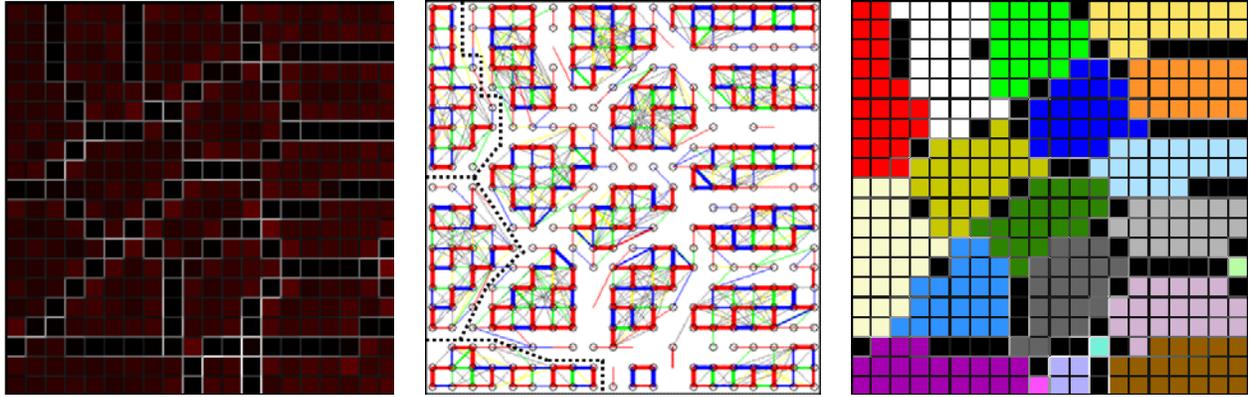


Figure 1. **Left:** Example of visualization of the SOM’s knowledge, for a simple synthetic data set with 20 clusters, through layering the sizes of the prototypes’ receptive fields, and the relative dissimilarities of the prototypes that are neighbors in the SOM. Each grid cell in this 20 x 20 SOM lattice describes attributes of one prototype vector. The intensity of the red monochrome color of the cell is proportional to the size of the receptive field (number of data points mapped to the corresponding prototype). The similarities (Euclidean or other distances of prototypes in data space) are encoded as black-to-white “fences” between the grid cells (modified U-matrix⁵). Black means high degree of similarity, white indicates high degree of dissimilarity, and thus delineates cluster boundaries. In a semi-manual procedure, a user can outline the perceived boundaries and thereby extract the clusters. **Center:** A novel SOM visualization, ConnVis, proposed in ref.¹⁰ Details are explained in the text. While this visualization looks more difficult to decipher (especially for complicated, real data) than the visualization at left, it in fact provides a cleaner guidance for locating cluster boundaries. The dotted lines indicate a few of the boundaries. **Right:** Clusters extracted from the visualization shown in the center.

2.2 Toward Automation of SOM Clustering

Clustering is done automatically with many methods, including the extraction of data clusters from an SOM by clustering the SOM prototypes.¹¹ The critical issue is what level of detail existing automatic clusterings can achieve. Are they capable of separating many clusters of varying sizes, shapes, densities, capable of discovering small clusters? For example, in a recent study⁹ K-means did relatively well on an 8-band spectral image of an urban area, identifying clusters that were clean-cut superclusters of clusters produced by an SOM. However, in a 200-band hyperspectral image of an urban area (with more surface material groups than in the 8-band image), K-means not only delineated less detail, but severely confused clusters compared to the SOM. The SOM discovered many spectrally unique (known, meaningful) material groups, some of which were as small as 3 – 6 pixels in the image. However, capable as the SOM may be in detecting the structure of a complicated data manifold, one has to extract that knowledge by determining which groups of prototypes represent similar data items (data clusters). For complex cases this can be daunting with semi-manual approaches.

Presently available automated SOM-segmentation methods are working with simplifying assumptions that result in less than desired accuracy in the characterization of the data structure. For example, in ref.¹¹ small clusters were missed. In the popular SOMPAK (implemented as a MATLAB toolbox), it is assumed that every SOM prototype is the center of a separate cluster, regardless of how similar some prototypes may be to one another. Thus the SOMPAK may return many clusters of the same character instead of detecting them, correctly, as one large cluster. Existing semi-manual methods can do a much better job from various visualizations (such as in Figure 1), but automation of that human-machine interaction has eluded the SOM community so far. A more extended review of this subject and references are given in ref.¹⁰

A recently proposed novel representation¹⁰ of the SOM’s knowledge is through a weighted graph of the quantization prototypes, which describes the structure of a data manifold (on the prototype level) in more detail than other known representations. This graph, which is a *weighted* equivalent of the Delaunay triangulation, is obtained by connecting the prototypes that are centroids of neighboring Voronoi cells, by weighted edges where the weights are proportional to the number of data points that select the two neighbor centroids as their closest and second closest prototypes. The weights thus signify the connectivity strengths between each pair of prototypes —how many data points connect them — indicating the local density distribution within Voronoi

cells with respect to their neighbors. Consequently, the weighted graph provides a description of the generally unisotropic “connectedness” of a prototype to its Voronoi neighbors, which separates dense parts of the data manifold from poorly connected or discontinuous regions, thereby delineating cluster boundaries. This graph can be expressed by a connectivity matrix that contains the pairwise connectivity strengths of the prototypes, and which we denote by CONN. We define the degree of similarity of two prototypes by their connectivity strength. Weakly connected or unconnected prototypes have low similarity (low CONN value), high CONN value means strongly connected — very similar — prototypes. This similarity measure enables more detailed cluster extraction than other existing methods can do, because density representation is finer than on the Voronoi cell level. Other methods that take density into account represent the density on the level of the Voronoi cells.

When quantization prototypes are obtained by SOMs, CONN visualized over the 2-dimensional SOM lattice reveals the topological structure of high(er)-dimensional data manifolds in 2 dimensions, regardless of the data dimensionality. This is illustrated in Figure 1, center. The connectivity strengths are first binned, with thresholds *automatically* determined solely from data statistics. Then each two SOM prototypes (represented by small circles centered in the grid cells) are connected with a line whose width indicates the (binned) connectivity strength. The line widths thus express the local density relative to the global data distribution. The connectivity strengths of each prototype are also ranked and color coded, which reveals the most-to-least dense regions, local to the corresponding prototype, in data space. Thus the colors, from red, to blue, yellow, green, and grey shades beyond that, indicate the local importance of the connections (the most-to-least important Voronoi neighbors) of this prototype. We refer the reader to ref¹⁰ for a more detailed explanation. Figure 1, center, shows the CONN visualization of the same SOM as in Figure 1, left. We extract clusters (Figure 1, right) from this visualization by detecting poorly connected or unconnected prototypes. At first sight it may be harder to see the cluster boundaries from Figure 1, center, than from the visualization on the left, because of the somewhat unusual look and the highly detailed nature of this representation. In order to guide the reader’s eye, we marked some of the paths along which prototypes are poorly connected or disconnected, outlining several of the clusters. As seen from Figure 1, right, all 20 clusters, including the rare ones, were found. Some prototypes were not assigned to clusters. Those prototypes represent transitions between clusters, or they are “empty” (no data point is mapped to them). Ref¹⁰ gives more details and instructive examples on cluster extraction from CONN visualization. For simple data, the two knowledge representations in the example in Figure 1 can serve equally well for cluster identification. For complicated data, however, the CONN visualization makes it easier to find the boundaries.

Currently we extract clusters semi-automatically from CONN visualization, which results in detailed structure identification (as the Mars example below, and previous analyses⁹ show). Owing to the nature of the CONN representation, and to the experience gathered from semi-automatic exercises, we expect to be able to fully automate the capture of clusters from non-visualized CONN representation. As mentioned above computing the thresholds for determining low and high connectivity strengths is already done automatically, derived purely from the data statistics. Adding connectivity based similarity criteria to the grouping of prototypes will complete the automation of cluster extraction from a learned SOM, on the level of sophistication we present in this paper.

3. DATA ANALYSIS EXAMPLES

3.1 Unsupervised Mapping of Materials in a Mars Exploration Rover Scene

From sol 159 of its mission to the present day, the Mars Exploration Rover Spirit has been exploring a region dubbed the Columbia Hills (in honor of the space shuttle Columbia astronauts) in Gusev Crater (at approximately 14.6 S latitude, 175.5 E longitude). Each of the hills was named in honor of one of the Columbia astronauts. Spirit is equipped with a Panoramic Camera (Pancam) with stereoscopic and multispectral imaging capabilities. The Pancam scene analyzed here was collected on sol 608 of Spirit’s mission near the summit of Husband Hill (named after shuttle commander Rick Husband). Rocks found in the Columbia Hills are generally clastic in nature and several chemical classes have been identified.¹³ These rock types have also been analyzed in terms of their visible and near infrared (VNIR) multispectral properties^{12,14} and thermal infrared emittance.¹⁵ The scene examined is a 700 x 450 pixel 7-band image from the Pancam’s right “eye”, data which has one band centered at 432 nm and six more bands with central wavelengths from 754 to 1009 nm. It includes the rock outcrop named “Tenzing” and the rock target “Whittaker” as well as other unnamed rocks.

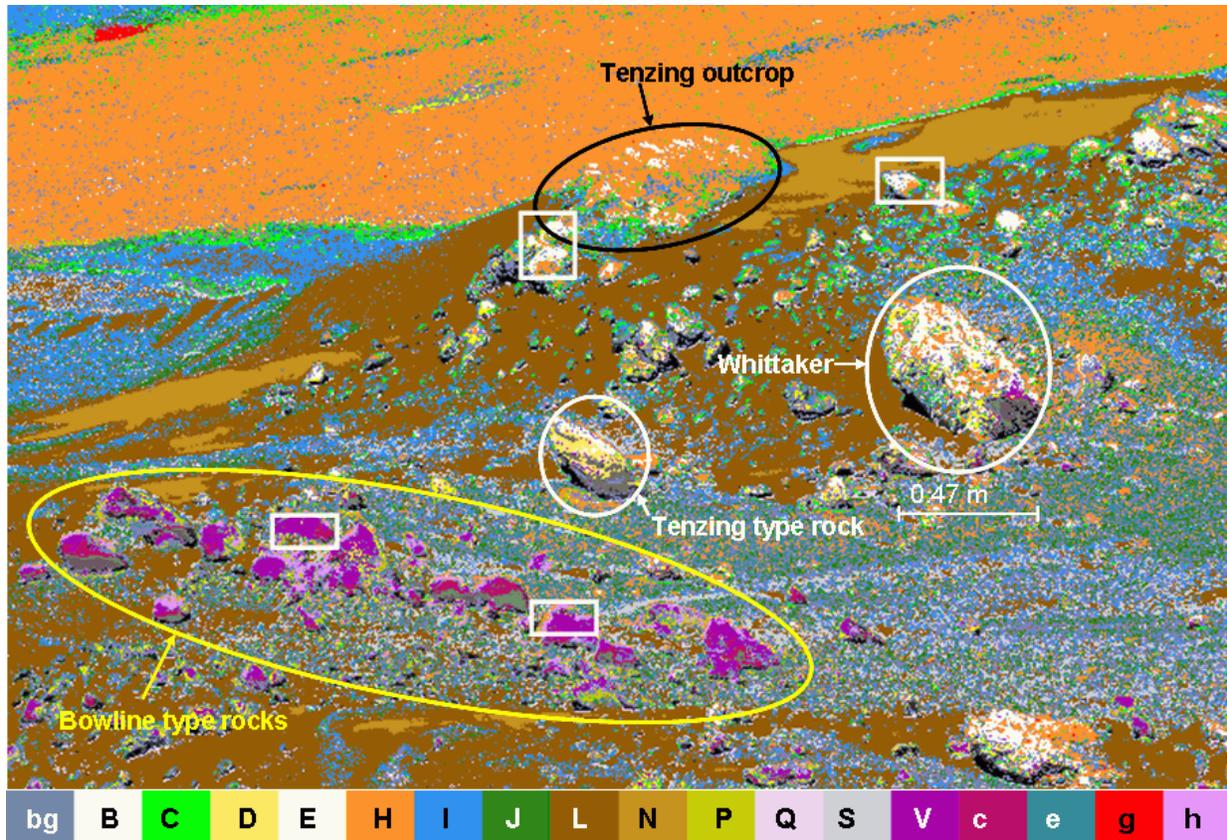


Figure 2. Cluster map of the scene near the summit of Husband Hill (MER Pancam image taken on Sol 608) with clusters extracted from CONN visualization (ConnVis¹⁰). The spatial resolution in the near field is indicated by the scale bar at the rock Whittaker. Clusters are color coded according to the color wedge. These 17 clusters coincide with previously identified spectral classes. For example, classes B, D, and E over the rock targets Whittaker and Tenzing correspond to the Jibsheet spectral class identified by Farrand et al.¹² The classes V, P, and h correspond to the Methuselah spectral class, which in this area is best represented by the rock target Bowline. The rocks within white rectangles are example areas which were less clearly resolved in a previous analysis, shown in Figure 3. The cluster map also indicates different mixtures of soil: H is a mixture with heavy influence from rock and an airfall dust component; N is bright drift; L is an intermediate albedo soil; J and e are coarse-grained soils with a presumed basaltic composition; and I, L and S are finer-grained soils also likely of basaltic composition.

A cluster map of the scene with 17 clusters extracted from CONN visualization is shown in Figure 2. Some clusters correspond to well known surface spectral classes identified in previous analyses, whereas some are related and could be deemed subtypes of the more representative spectral classes. A product of one of those earlier analyses, a composite of fraction images derived from constrained energy minimization (CEM) analysis,¹⁶ is shown in Figure 3 which has rocks most closely resembling the Tenzing outcrop in green, another spectral class of rocks with a deeper 904 nm absorption band in red, and bright drift material in blue. Mean spectra of the 17 clusters are shown in Figure 4. We describe the geologic meaning of the clusters below.

The far field of the scene (upper left part of the image, most extensively covered by the orange class, H) is a mixed class, most resembling the spectral signature of soil with substantial influence from rock fragments and also with an airfall dust component. The scene includes soils of various compositions including darker-toned, coarser-grained soils with a presumed basaltic composition (J, e) and lighter-toned, finer-grained soils which are also presumed to be of bulk basaltic composition, but much more oxidized (I, L, S). Some of these lighter-toned soils (bright drift, N) are higher in albedo and are likely mobile drifts. Intermediate albedo soils (L) are likely immobile. The dune near the summit of Husband Hill, which is mostly covered with bright drift and intermediate

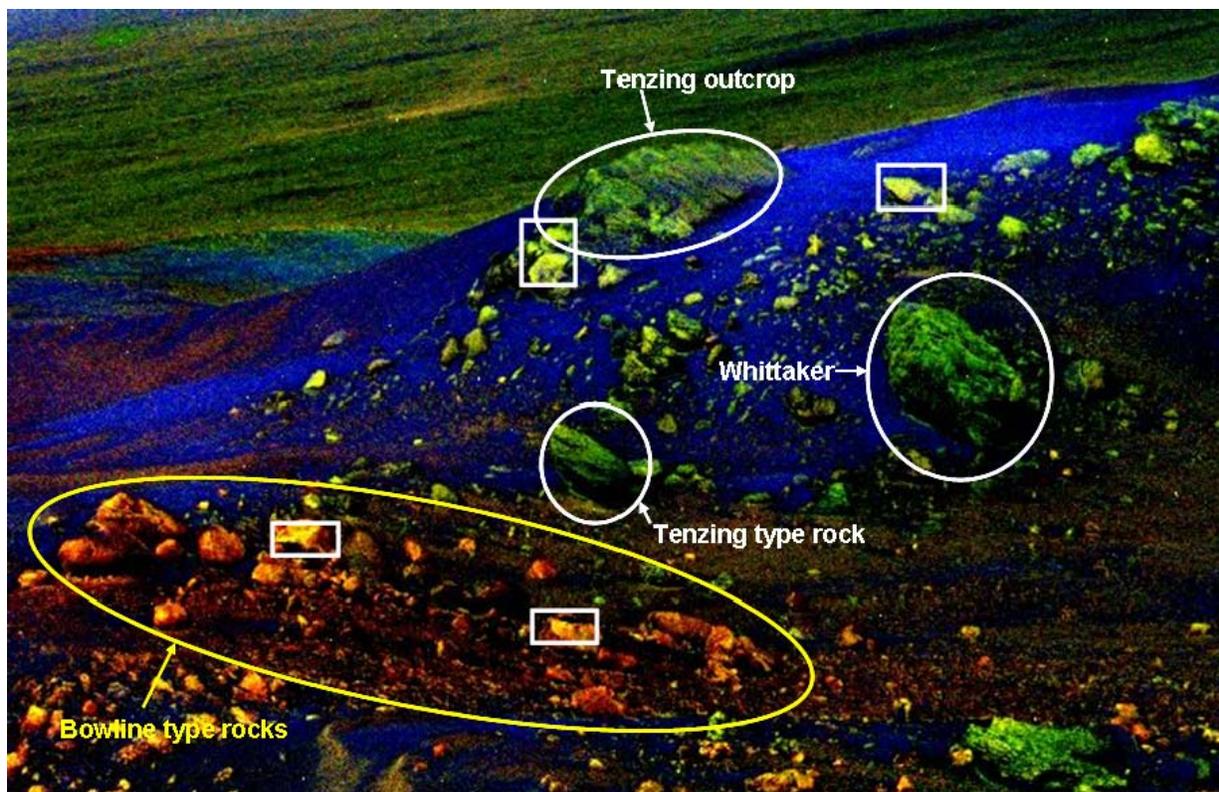


Figure 3. Composite of CEM¹² fraction images of rocks with strong 900 nm absorption band (red, Bowline type rocks), shallow 900 nm band (green, Tenzing type rocks) and bright drift (blue). Some Bowline type rocks (two are shown by rectangles in the reddish region) have yellow spots which indicate they have similar high values for red and green endmembers. Similarly, some Tenzing type rocks on the dune have yellow spots (two are shown in rectangles) due to some degree of confusion in the CEM map between the Tenzing class and rocks/soils in the area close to Tenzing outcrop.

albedo soil, is separated from the far field with a border (C, Q) that is possibly a mixture of sand grains. The big rock on the dune (mostly covered by the orange cluster, H, due to airfall dust) is the “Tenzing” outcrop. Tenzing is a subtype of the Jibsheets spectral class, characterized by a shallow infrared absorption feature centered between 900 and 934 nm (Farrand et al. 2007; in preparation). The white rocks (B) in the cluster map are Jibsheets type rocks whereas yellowish (D, E) rocks are rocks with a variant of the typical Jibsheets spectral signature. The MER team has observed that some rocks near the summit of Husband Hill had rougher textures and some had smoother textures.¹⁷ The changes in the scattering properties due to those different textural properties might cause some spectral differences. Indeed the two subclusters of Tenzing (D, E) which show slight but consistent spectral differences (Figure 4) may reflect this. In the CEM map, there is confusion between the Tenzing class and rocks/soils in the area close to the Tenzing outcrop (yellow Tenzing rocks within the rectangles in Figure 3), which is better distinguished in the cluster map. Figure 5, left, shows the slight differences among the spectra of subclusters B, D and E and compares them to the Jibsheets superclass in ref.¹²

Another spectral type of rock in the scene is close to that of the nearby rock target Bowline. Spectra of these materials indicate that Bowline, and rocks with similar spectra have a deeper near infrared absorption feature than a similar, but shallower absorption feature in the Tenzing outcrop and related rocks (Figure 4). The rocks in purple (V) and in brownish green (P) are Bowline types where P has a shallower near infrared band compared to V. Dark cherry (c) areas are partly shaded Bowline rocks whereas light pink (h) areas are a mixture of small Bowline rocks and soil. In the CEM image, some part of the “Bowline” rocks, which should be classified as the “Bowline class” rocks, are colored yellow since those rock spectral classes received similar data numbers for the red and green endmembers. This is another case where our approach yields more clarity than the CEM analysis.

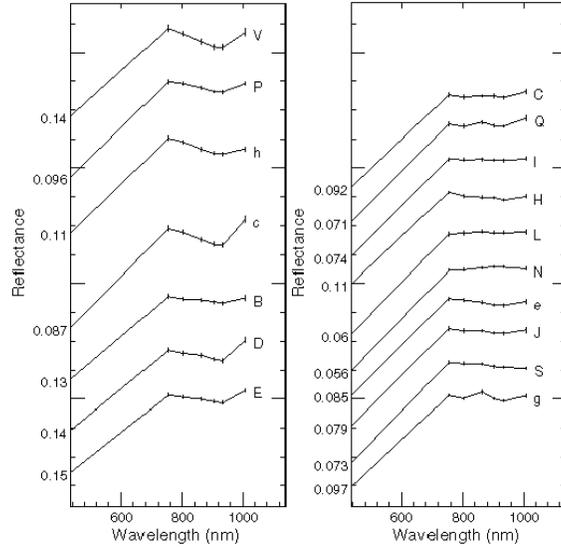


Figure 4. Mean spectra of the 17 clusters in Figure 2, vertically offset for viewing convenience, and organized into groups by major types, as applicable: Bowline (V, P, h), Jibsheet and Tenzing (B, D, E), soil and soil mixtures (I, H, L, N, e, J, S) and g. The spectra of subclusters have slight but consistent differences whereas the different major types have greater spectral dissimilarities. The spectrum of g distinctively marks the red cluster near the top left corner of the image. The spectral shape suggests relation to clusters Q and C (which are likely sorted sand grains at the border of dunes), yet it is different enough to describe a separate phenomenon, perhaps a dust devil. A detailed comparison of the Jibsheet and Bowline subclusters are given in Figure 5.

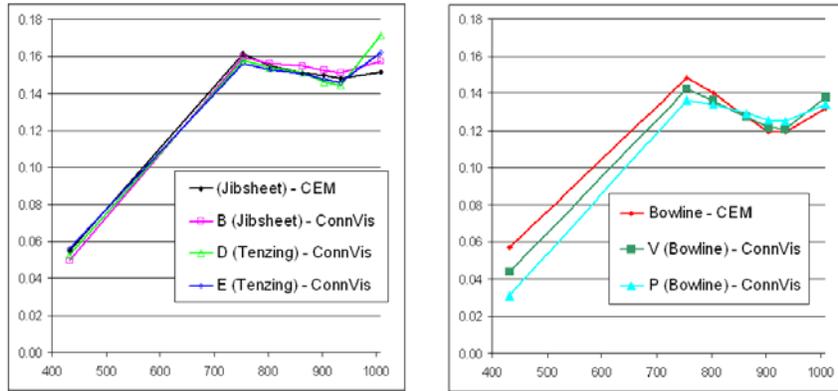


Figure 5. Comparison of spectral characteristics of subclusters with their parent types. **Left:** Jibsheet and Tenzing. The Tenzing subcluster has a deeper absorption band than Jibsheet in the near infrared with D showing the deepest band. **Right:** Bowline subclusters. P is a variant of V with a shallower band.

A number of these 17 clusters are not new but rather indicate variations within clusters such as I, S, J and e, which are soil mixtures. At this point we do not know whether these variations are geologically meaningful yet they are interesting from a clustering point of view because of their spatial coherence and spectral distinction from the parent clusters. One cluster we discovered (g, red spot) near the top left corner, has a distinct spectrum, appears spatially very coherent and occurs only at this location in the scene. This cluster might be a dust devil, however the spatial resolution of the Pancam image is too poor to unambiguously identify a dust devil shape.

Our analysis of this MER Pancam image has shown that a number of geological units identified in previous research were found through unsupervised self-organized (SOM) learning, and using the CONN visualization for cluster extraction. Some units such as Jibsheet and Bowline were also segmented into subtypes which may be geologically meaningful. Our analysis has provided a comprehensive mapping of the entire scene. Spectral evi-

dence in Figure 5 shows that clusters extracted from this comprehensive mapping reliably match the prototypical clusters described by one of us (Farrand) in previous work from smaller sampling areas. This gives us confidence that this clustering can be a viable candidate for autonomous on-board information extraction. As we discussed in section 2.2 we expect to fully automate this clustering process.

3.2 Discoveries in an Urban Hyperspectral Image

While the MER scene clustering in section 3.1 seems to match well the knowledge previously derived by geologists and spectroscopists, the lack of detailed validation data limits objective evaluation of the analysis method. In this section we present clustering of a hyperspectral urban image, which was synthetically generated through the rigorous DIRSIG modeling procedure at the Rochester Institute of Technology.^{18,19} Owing to its simulated nature, this image has ground truth for every pixel, allowing objective evaluation of analysis results on the 1-pixel scale. The characteristics of this data set are close to that of a low-altitude AVIRIS image: it comprises 400 x 400 pixels in 210 image bands in the 0.38 to 2.4 μm visible-near-infrared spectral window, with a spatial resolution of 2 m/pixel. The image is very realistic as seen from a color composite in Figure 6, upper left. It contains over 70 different surface materials, including vegetation (trees and grass), about two dozens of various roof shingles, a similar number of sidings and various paving and building materials (bricks of different brands and colors, stained woods, vinyl and several types of painted metals), and car paints. This multitude of clusters, widely varying in their statistical properties in 210-dimensional data space, presents a larger challenge than the multispectral image from the Spirit rover.

Preprocessing of this image consisted of empirical line correction to remove the (simulated) atmospheric effects and to convert radiances to reflectances, and a brightness normalization to remove linear illumination effects. Clusters extracted from the SOM and mapped back to the spatial image are shown in Figure 6, upper right. The SOM lattice that produced this clustering consists of 40 x 40 neurons, shown in Figure 6, lower left, with the groups of prototypes that represent the various data clusters color coded. The correctness of the clusters can be checked against a truth map that accompanies this synthetically generated scene.

Not all identified clusters are color-coded because it is difficult to visually distinguish many colors. For the sake of clarity we chose to limit this presentation to 38 clusters. One can view more clusters interactively, with our software tools. In order to save colors for interesting, smaller units such as some man-made materials which have the greatest variety in this image, we lumped all types of trees into three clusters (J, medium blue, Z very dark blue, and g, purple-blue). Besides these tree units, the SOM is dominated by three large clusters: K (dark green), healthy grass, T (flesh color), grass mixed with dirt, and V (light green), older asphalt paving. One interesting small vegetation unit we kept is Y (warm reddish brown), wedged between J and T in the SOM, which is a somewhat more distressed grass than the one in cluster T, and shows in the two playing fields in spatially distinct patterns, top center, and lower right of the scene. Clear separation of the large clusters is indicated by well formed “fences”, and empty prototypes (those cells with the gray background color “bg”). In the lower right of Figure 6 the same cluster map is presented as in the upper right, but with the large background clusters removed, to highlight the small units. Many of these are different types of roofs and show up in the columns of houses among the trees. A large variety of small clusters occupy highly structured areas in the upper left and lower left corners of the SOM. A number of these map obvious units: k (mauve) corresponds to the large building in the center, with gravel roof. B (orchid) outlines the L-shaped building in the lower center. Its roof is tan colored asphalt shingle, which also appears on houses elsewhere. Cluster U (lilac, upper left in the SOM lattice) identifies two larger buildings which seem to be spectrally unique in the scene, with mixed brown asphalt shingle roof. O (split pea color, represented by a single prototype along the fence between V and K) and P (brown, two prototypes between Z and T) perfectly match the green and red surfaces of the tennis court, just right of the center. Toward the upper right, cluster E (light blue) is a building covered with glass. Glass is not unique to this building, however. The two large squares in the track field, lower center of the scene, are tarps, one grey and one black, placed there for calibration purposes.

Including the larger buildings we listed above we mapped 20 different roof materials, which are mostly various asphalt shingles, some with subtle spectral differences. An enlarged cut-out from the upper left corner of the cluster map in Figure 7, left, contains 15 roof types. At the bottom two tiny objects are circled, which were mapped, along with several others, into clusters “c” (best spectral match to “Saturn hood, white”) and E (glass).

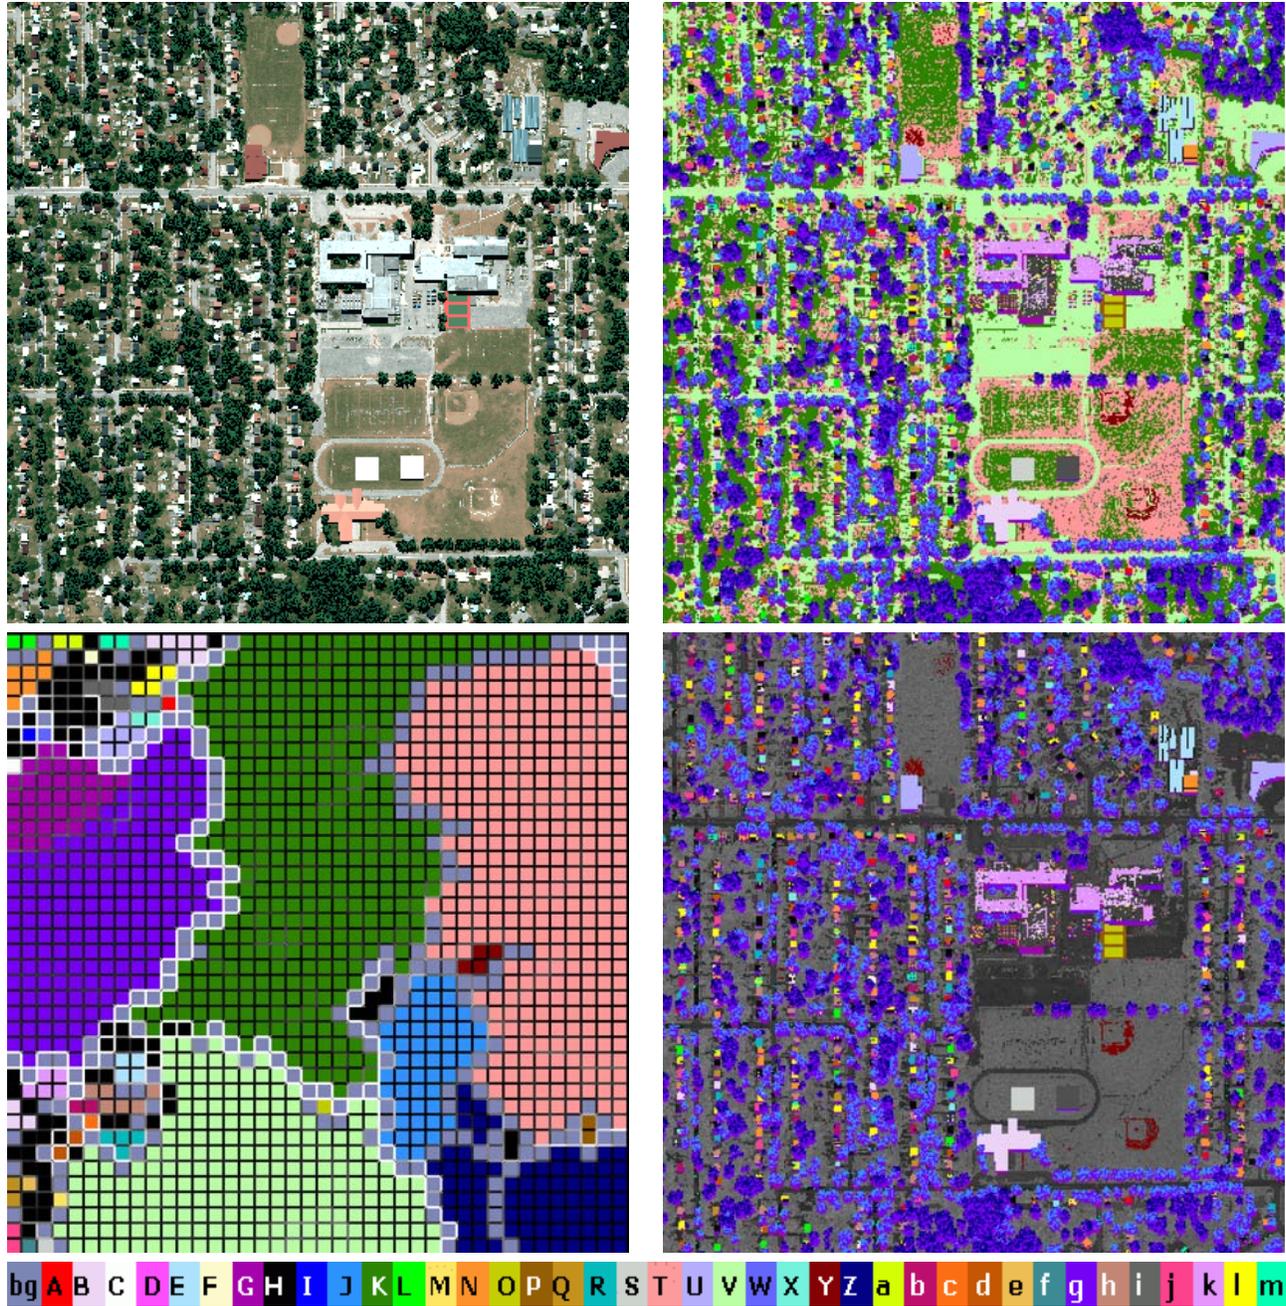


Figure 6. **Upper Left:** A color composite made from 3 selected spectral bands of the 400 x 400 pixel, 210-band synthetic hyperspectral urban image. The hyperspectral image was generated by the DIRSIG procedure at the Rochester Institute of Technology.^{18,19} The scene covers an area of 800 x 800 square meters (2 meters/pixel spatial resolution). Besides the obvious vegetation (trees and grass), the approximately 70 different surface cover types contained in this image include a large number of various roof materials, pavings, and several types of car paints. **Upper Right:** Partial cluster map, extracted from the self-organizing map in the lower left frame, using the modified U-matrix visualization shown in Figure 1, left. Each material type is coded by a unique color keyed in the color wedge at the bottom. **Lower Left:** The SOM with discovered clusters color-coded. The cells with medium grey color, appearing mostly along cluster boundaries are SOM prototypes with no data points mapped to them. Some areas — shown as black cells, which have data mapped to them — were left unclustered, for reasons of colors limitations. **Lower Right:** The same cluster map as in the upper right, with the large background clusters (grass, paved roads and lots) removed to provide better contrast for the many different roof materials, and other unique spectral types such as tennis courts. Twenty of these clusters are roof types. Spectral plots comparing the characteristics of extracted clusters with true classes are in Figures 8 and 10.

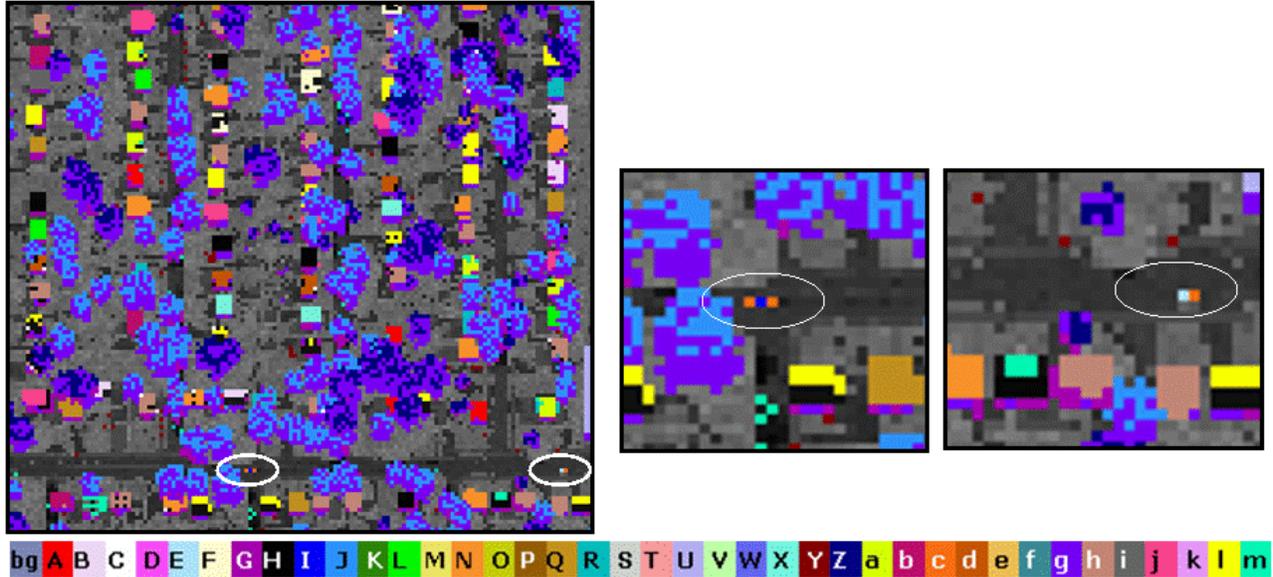


Figure 7. **Left:** The upper left corner of the cluster map enlarged, to show the variety of roof types. Based on the match of their spectral characteristics with ground truth classes, these clusters are interpreted as follows. B (orchid): tan asphalt shingle; F (pale yellow): black/grey asphalt shingle; L (true green): forrest green asphalt shingle; N (orange): twilight grey asphalt shingle; Q (ocher): wood stained red; R (teal): brown/black new asphalt shingle; X (turquoise): asphalt shingle, dark and light; “a” (neon greenish yellow): weathered black asphalt shingle; “b” (maroon): old white asphalt shingle; “h” (mauve): old weathered white asphalt shingle; “j” (cherry): fair clored brick; “l” (intense yellow): asphalt shingle, Sequoia tile. Spectral comparison of selected roof clusters with true classes is given in Figure 8. **Right:** Two of the cars found in this scene from their spectral signatures, discussed below. The left one is showing two pixels in cluster “c” (dark orange), consistent with a white Saturn hood, the middle pixel is in cluster I, distinguishing the roof of the car. The car in the right inset has a light blue pixel in cluster E, consistent with glass (windshield), and another pixel in cluster “c”, again a white Saturn hood. More details on cars are below.

These are cars, enlarged on the right in Figure 7. We further discuss the identification of cars below. The clusters discovered by the SOM show remarkable spectral match with the ground truth classes. Figure 8 samples a few of the roof materials, where in the upper row we plotted the mean, and also the envelope (the extremes in each spectral band) of the SOM-discovered clusters to give a sense of how clean these clusters are in comparison to the ground truth (bottom row).

Perhaps the most interesting part of this analysis is the discovery of cars in the scene, since they each occupy 2–3 pixels only, and are varied in spectral character. Figure 7 pointed out two occurrences. There are more cars scattered in the image. For a focussed discussion we zoom in on the parking lot in the center, in Figure 9, left. When grass, paving, and grass/dirt ground is excluded from the cluster map (center inset), the remaining colored pixels are easier to see (right inset). The magenta pixels (cluster G) correspond to shadows cast by cars, the other colored pixels within the white box in this right inset identify several materials. Light blue pixels (cluster E) signify glass, those locate the windshields of four cars: one in the left column of the left half of the lot, and three in the left column of the right half of the lot, toward the upper part. Of these four cars, each of which shows in two pixels, two have body parts in cluster “c” (dark orange), which stands for white Saturn hood paint. The truth map labels them both as white Saturn, but the lower one is given as the roof of a white Saturn. The car between the two Saturns in the same column is a Ford Focus, and indeed it was assigned to a unique cluster, “f” (teal). To the right of the Ford Focus stands another white Saturn, correctly assigned into cluster “c”. To the right of the Saturn below the Focus, (third from top, in the right column) a dark VW was left unclustered. In the left half of the lot, left column, a 3-pixel “c” object correctly matches yet another white Saturn (hood), and a few spaces below the car whose windshield shows in light blue (E) has its body assigned to cluster W. This is a dark blue BMW, which has a single-pixel occurrence in the scene, according to the truth map. Cluster W is

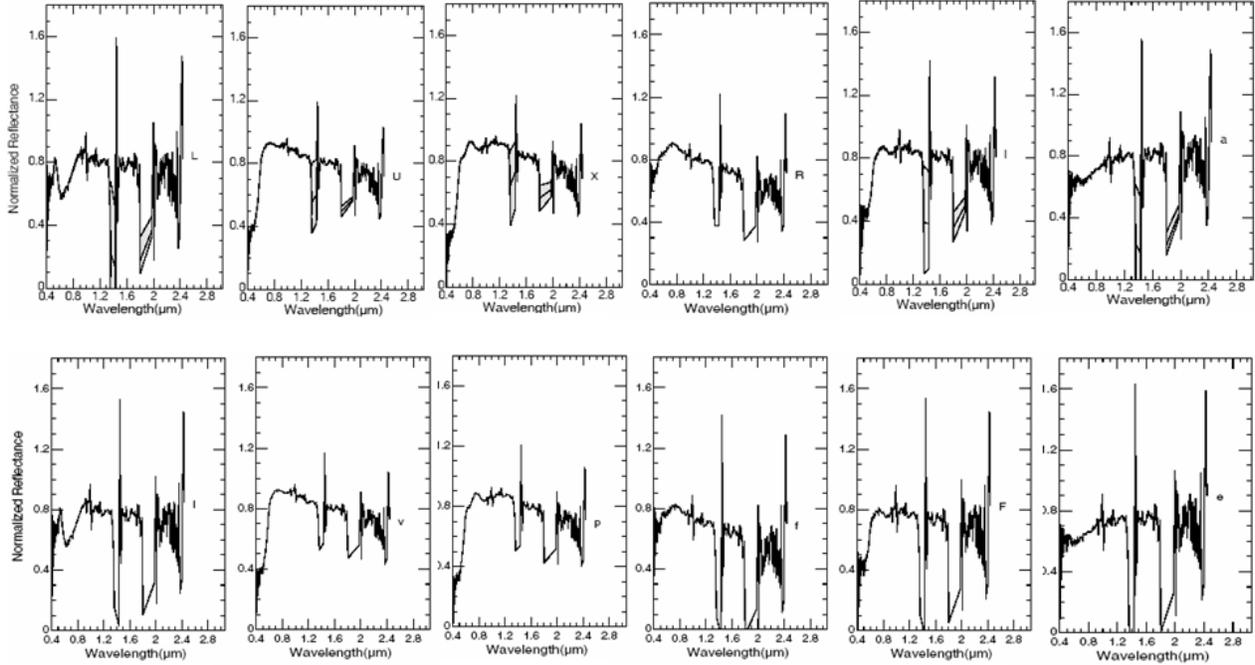


Figure 8. Spectral characteristics of selected roof clusters. Upper row: the mean spectra of SOM clusters. Lower row: the mean of the true classes. For the SOM clusters we also plotted the envelopes (the extremes in each spectral band) to show how tight they are. These and other roof clusters comprise a few hundred pixels each, distributed over 10 – 20 houses across the scene. The labels in the lower row are not the same due to different assignment scheme than for the SOM clusters. The materials represented are, from left to right, using their SOM cluster labels, L: asphalt shingle, forest green; U: asphalt shingle, mixed brown; X: asphalt shingle, dark; R: asphalt shingle, brown/black, new; l: asphalt shingle, Sequoia tile; a: asphalt shingle, black, weathered. Noise in bands at the edges of the 1.2 – 1.35 and 1.8 – 2.0 μm windows is remnant of imperfect atmospheric correction.

represented uniquely by one prototype in the upper left of the SOM lattice (this prototype only maps this single pixel). The 3-pixel Saturn on the road in Figure 7, and one car (another white Saturn) at the bottom left of the boxed area here have their center (presumably the roof) assigned into cluster I. Cluster I has 12 pixels in it, and while its mean signature does not exactly match that of the true spectrum of a white Saturn roof, it has a shape closer to Saturn roof than to Saturn hood (shown in Figure 10), thus it separates the roofs from the hoods of the white Saturs. (A difference can also be seen in the color composite inset.) There are 81 pixels in cluster “c”, most of which are not identifying white Saturs, but white metallic objects such as the ones regularly spaced left of the boxed area. Their spectral signatures are almost perfect match to the white Saturn hood. At the top of the right column of the right half of the lot a grey BMW occupies 3 pixels, all three assigned into different spectral clusters (M, yellow, C, white, and “e”, sand). The spectral plots in Figure 10 give some insight to this confusion. Apparently some other species in the scene, such as wood, gravel roof, have sufficient gross similarity and a much larger number of pixels than cars occupy, to overwhelm and attract those car signatures to the same SOM prototype. Similarly, the signature of the single-pixel Ford Focus was grouped with 6 other pixels, whose spectra align with the main shape of the Focus spectrum. Two cars in the parking lot, a VW and a Volvo, do not show in this clustering. These are cases where the SOM may need longer learning or larger real estate to separate the confused species; or negative SOM magnification may help. These possibilities will be considered in follow-up research. Future work on investigating this level of detail for information extraction will also include the use of more advanced versions of simulated hyperspectral imagery that the DIRSIG group^{18,19} is capable of producing.

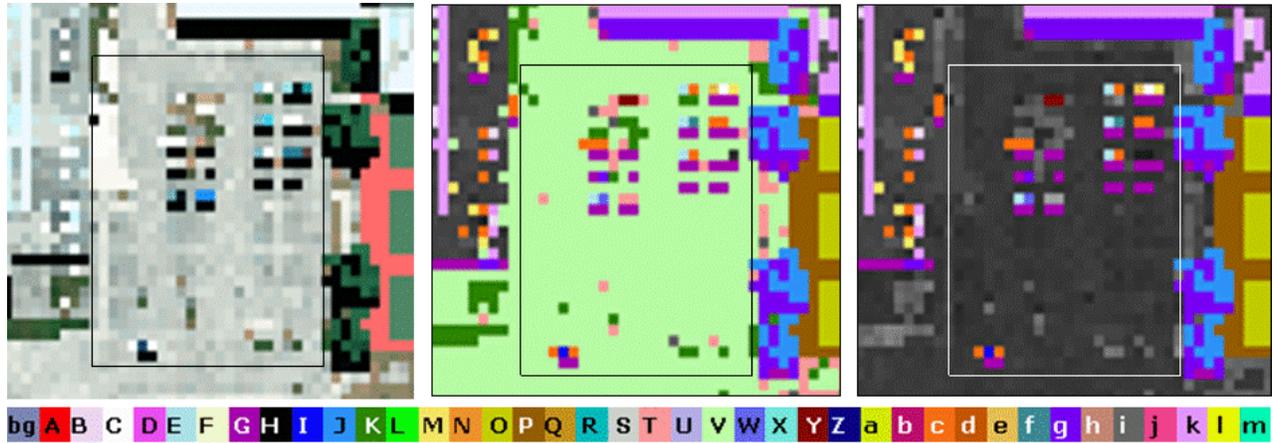


Figure 9. Enlarged detail of the parking lot near the center, left of the tennis court, showing several different makes of cars, **Left:** from the color composite in Figure 6, upper left; **Center:** from the cluster map in Figure 6, upper right; **Right:** from Figure 6, lower right. The cars, discussed in the text, are within the box shown in each inset.

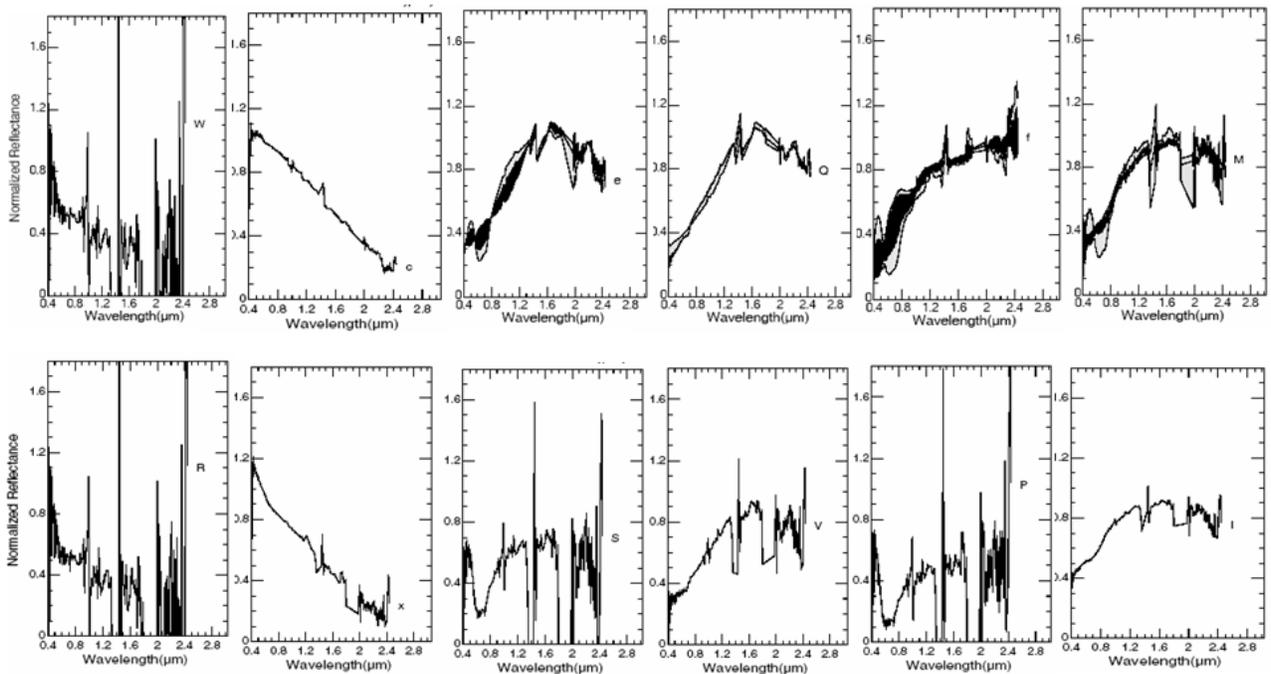


Figure 10. Spectra of cars and related materials. Upper row: mean spectra of SOM clusters. Bottom row: mean spectra of the true classes. For the SOM clusters we also plotted the envelopes (the extremes in each spectral band) to show how clean or mixed the clusters are. The labels in the lower row are not the same due to different assignment scheme than for the SOM clusters. The materials represented are, from left to right, using their SOM cluster labels, W (violet blue): dark blue BMW; “c”: white Saturn hood; in the bottom row the spectrum of the truth class “white Saturn roof” is shown, to give an idea of the difference between the two body parts of the same car; the SOM cluster “c” matches the “white Saturn hood” truth spectrum perfectly. “e” (sand): the mean of this cluster is very close to “wood, stained red” but apparently other spectra (including the grey BMW signature, shown at the bottom), were mixed in. “Q” (ocher): wood stained red; “f” (teal): Ford Focus hood; this cluster attracted 6 other pixels which show general similarity to the overall shape of the Focus spectrum; M (yellow): closest to “gravel roof” but obviously modulated by the spectrum of the grey BMW, among other signatures. Noise in bands at the edges of the 1.2 – 1.35 and 1.8 – 2.0 μm windows is remnant of imperfect atmospheric correction.

3.3 Significance For On-Board Decision Making

We showed in the previous sections that many details can be mapped from complex data sets such as hyperspectral imagery, with self-organized learning. This can outline spectrally distinguishable spatial objects in a precise way even if they are very small, and occupy statistically insignificant portions of the image. The spectral clusters can be labeled by finding matches in a comprehensive spectral library, or by human-in-the-loop if unknown spectra are discovered. The detail and precision are critical for the success of further interpretation beyond this level. For example, identifying single pixels as glass, or "white Saturn hood" is within the power of our approach. This enables a higher-level inferencing system that uses spatial context, to recognize them together as a car (as opposed to the glass being part of a building); or whether the "white Saturn paint" locates a car or a metallic structure. Another example, showing the importance of being able to find units with subtle but consistent differences, is the SOM discovery of the red cluster close to the upper left corner of the image from the MER Spirit rover in section 3.1. Its spectral characteristic suggests relationship to loose sand grains seen as thin borders of dunes elsewhere in the image, yet it is consistently different, and shows in a lenticular spatial shape, very different from the shape of dune borders. Could it be a dust devil? Such question may be answered by an inferencing scheme that has an on-board knowledge-base of other existing analyses and data of the same area, or rules derived from human experience (such as in ref³). The work we presented here strives to provide accurate and detailed information for input to higher-level interpretations.

4. SUMMARY AND FUTURE DIRECTIONS

We demonstrated an approach to processing massive and complex scientific data with brain-like neural computation, for a) sophisticated understanding of data, and b) for the potential speed that massively parallel neural learning machines can achieve in appropriate hardware. Both are critical for on-board decision making in terrestrial and planetary missions.

Unsupervised discovery, presented in this article through two complex data sets, is one of the most promising applications of self-organized learning by neural computing systems. We also use self-organized neural learning to aid in precise supervised classification of a complex data set into many predefined classes. These capabilities, augmented by neural feature extraction (data compression)²⁰ can be packaged together to produce systems that will be able to facilitate highly intelligent data understanding,²¹ with the appropriate level of sophistication to support higher-level science interpretations and science-driven decisions. This neural computing approach, when implemented in massively parallel hardware on board autonomous vehicles, will enable real-time discoveries of unexpected differentiated classes, as well as detection of targets with known signatures, represented within massive, complex data sets. While fabrication of the needed highly capable neural chips with appropriate scale-up properties is still a challenge, nanotechnology is expected to provide that capability soon. This approach promises to combine the intelligence of neural computing algorithms with the speed needed for real-time exploration, decision-making, and operations.

ACKNOWLEDGMENTS

We thank Prof. John Kerekes of the Rochester Institute of Technology for providing the synthetic hyperspectral image used for this work, and Maj. Michael Mendenhall for his help with preprocessing the hyperspectral image. Partial support by the Applied Information Systems Research Program (grant NNG05GA94G) and Mars Data Analysis Program (grant NAG5-13294) of NASA's Science Mission Directorate is greatly appreciated.

REFERENCES

- [1] Tunstel, E., Howard, A., and Huntsberger, T., "Robotics challenges for space and planetary robot systems," in [*Intelligence for Space Robotics*], Howard, A. and Tunstel, E., eds., *TSI Press Series*, 3–20, TSI Press (2006).
- [2] Some, R., "Space computing challenges and future directions," in [*Intelligence for Space Robotics*], Howard, A. and Tunstel, E., eds., *TSI Press Series*, 21–42, TSI Press (2006).

- [3] Furfaro, R., Dohm, J. M., Fink, W., Kargel, J., Schulze-Makuch, D., Fairén, A., Palmero-Rodriguez, A., Baker, V., Ferré, P., Hare, T., Tarbell, M., Miyamoto, H., and Komatsu, G., “The search for life beyond earth through fuzzy expert systems,” *Planetary and Space Science* **56**, 448–472 (2008).
- [4] Kohonen, T., [*Self-Organizing Maps*], Springer-Verlag, Berlin Heidelberg New York (1997).
- [5] Merényi, E., Jain, A., and Villmann, T., “Explicit magnification control of self-organizing maps for “forbidden” data,” *IEEE Trans. on Neural Networks* **18**, 786–797 (May 2007).
- [6] Kuhl, P., “Human adults and human infants show a “perceptual magnet” effect for the prototypes of speech categories, monkeys do not,” *Perception & Psychophysics* **50**(2), 93–107 (1991).
- [7] Bauer, H.-U., Der, R., and Herrmann, M., “Controlling the magnification factor of self-organizing feature maps,” *Neural Computation* **8**(4), 757–771 (1996).
- [8] Merényi, E., Farrand, W., and Tracadas, P., “Mapping surface materials on Mars from Mars Pathfinder spectral images with HYPEREYE.,” in [*Proc. International Conference on Information Technology (ITCC 2004)*], 607–614, IEEE, Las Vegas, Nevada (2004).
- [9] Merényi, E., Csató, B., and Taşdemir, K., “Knowledge discovery in urban environments from fused multi-dimensional imagery,” in [*Proc. IEEE GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas (URBAN 2007)*], Gamba, P. and Crawford, M., eds., IEEE Catalog number 07EX1577, Paris, France (11–13 April 2007).
- [10] Tasdemir, K. and Merényi, E., “Data topology visualization for the Self-Organizing Map,” in [*Proc. 14th European Symposium on Artificial Neural Networks, ESANN’2006, Bruges, Belgium*], 125–130 (26–28 April 2006).
- [11] Vesanto, J. and Alhoniemi, E., “Clustering of the self-organizing map,” *IEEE Transactions on Neural Networks* **11**, 586–600 (May 2000).
- [12] Farrand, W. H., III, J. F. B., Johnson, J. R., and Blaney, D. L., “Multispectral reflectance of rocks in the columbia hills examined by the mars exploration rover spirit: Cumberland ridge to home plate,” *Lunar and Planetary. Science XXXVIII* (1957) (2007).
- [13] Squyres, S. W., Arvidson, R. E., Blaney, D. L., Clark, B. C., Crumpler, L., Farrand, W. H., Gorevan, S., Herkenhoff, K. E., Hurowitz, J., Kusack, A., McSween, H. Y., Ming, D. W., Morris, R. V., Ruff, S. W., Wang, A., and Yen, A., “The rocks of the columbia hills,” *Journal of Geophys. Res.: Planets*, **111**, E02S11, 10.1029/2005JE002562 (2005).
- [14] Farrand, W. H., III, J. F. B., Johnson, J. R., Squyres, S. W., Soderblom, J., and Ming, D. W., “Spectral variability among rocks in visible and near infrared multispectral pancam data collected at gusev crater: Examinations using spectral mixture analysis and related techniques,” *Journal of Geophys. Res.: Planets*, **111**, E02S15, 10.1029/2005JE002495 (2006).
- [15] Ruff, S. W., Christensen, P. R., Blaney, D. L., Farrand, W. H., Johnson, J. R., Moersch, J. E., Wright, S. P., and Squyres, S. W., “The rocks of guser crater as viewed by the mini-tes instrument,” *Journal of Geophys. Res.: Planets*, **111**, E12S18, 10.1029/2006JE002747 (2006).
- [16] Farrand, W. and Harsanyi, J., “Mapping the distribution of mine tailings in the coeur dalene river valley, idaho through the use of a constrained energy minimization technique,” *Remote Sensing of Environment* **59**, 64–76 (1997).
- [17] Herkenhoff, K. E., Squyres, S., Arvidson, R., and the Athena Science Team, “Overview of recent athena microscopic imager results,” in [*Lunar and Planetary Science XXXVIII, abstract #1421*], (2007).
- [18] Schott, J., Brown, S., Raqueo, R., Gross, H., and Robinson, G., “An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development,” *Canadian Journal of Remote Sensing* **25** (June 1999).
- [19] Ientilucci, E. and Brown, S., “Advances in wide-area hyperspectral image simulation,” in [*Proceedings of SPIE*], **5075**, 110–121 (May 5–8 2003).
- [20] Mendenhall, M. and Merényi, E., “Relevance-based feature extraction for hyperspectral images,” *IEEE Trans. on Neural Networks, in press* (May 2008).
- [21] Merényi, E., Farrand, W. H., Brown, R. H., Villmann, T., and Fyfe, C., “Information extraction and knowledge discovery from high-dimensional and high-volume complex data sets through precision manifold learning,” in [*Proc. NASA Science Technology Conference (NSTC2007)*], ISBN 0-9785223-2-X, 11 (June 19 – 21 2007).