systems that are too large to be simulated with a discrete element approach, PowderSim incorporates a continuum-based SPH module, which when considering the addition of a calibrated, cohesive, constitutive model (Lunar Regolith Constitutive Model (LRCM)), is a novel use of mesh-free methods. Because of the discrete and continuum methods implemented in the same framework, the software can capture dynamic particulate material behavior at a variety of spatial scales from the coarse-grain scale (DEM) to the bulk scale (SPH). The DEM capability also supports clustering, which allows it to capture a rich variety of shape detail. Advanced contact models and charge spots capture many effects of contact plasticity and hysteresis, roughness, adhesion, and electrostatic interaction of particles. The SPH capability for bulk material behavior uses the LRCM to capture the critical-state behavior of cohesive lunar regolith.

Multiple-Frame Detection of Subpixel Targets in Thermal Image Sequences
This technique has applicability in fire detection, and tracking ships, ground vehicles, and aircraft.

NASA's Jet Propulsion Laboratory, Pasadena, California

The new technology in this approach combines the subpixel detection information from multiple frames of a sequence to achieve a more sensitive detection result, using only the information found in the images themselves. It is taken as a constraint that the method is automated, robust, and computationally feasible for field networks with constrained computation and data rates. This precludes simply downloading a video stream for pixel-wise co-registration on the ground. It is also important that this method not require precise knowledge of sensor position or direction, because such information is often not available. It is also assumed that the scene in question is approximately planar, which is appropriate for a high-altitude airborne or orbital view.

This approach tracks scene content to estimate camera motion and finds geometric relationships between the images. An initial stage identifies stable image features, or interest points, in consecutive frames, and uses geometric relationships to estimate a "homography" — a transformation mapping between frames. Interest points generally correspond to regions of high information or contrast. Previous work provides a wide range of interest point detectors. In this innovation, SIFT (Scale Invariant Feature Transform) keypoints recovered by a difference of Gaussians (DoG) operator applied at multiple scales are used. A nearest-neighbor matching procedure identifies candidate matches between frames. The end result of this first step is a list of candidate interest points and descriptors in each frame.

An important benefit of SIFT detection is that the system permits absolute georeferencing based on image contents alone. The SIFT features alone provide sufficient information to geolocate a hot pixel. This suggests an initial characterization phase where the remote observer transmits high-contrast, SIFT descriptors along with images of the (fire-free) surface. The ground system, with possible human assistance, would determine the SIFT features' geographic locations.

During regular operations, the system can query the database to find geographic locations of new observations. Any preferred single- or multiple-channel detection rule is applied independently in each frame with a very lenient threshold. Then, the algorithm matches consecutive detections across potentially large displacements, and associates them into tracks, i.e., unique physical events with a precise geographic location, that may appear in multiple frames. Finally, the system considers the entire sequence history of each track to make the final detection decision.

Metric Learning to Enhance Hyperspectral Image Segmentation
NASA's Jet Propulsion Laboratory, Pasadena, California

Unsupervised hyperspectral image segmentation can reveal spatial trends that show the physical structure of the scene to an analyst. They highlight borders and reveal areas of homogeneity and change. Segmentations are independently helpful for object recognition, and assist with automated production of symbolic maps. Additionally, a good segmentation can dramatically reduce the number of effective spectra in an image, enabling analyses that would otherwise be computationally prohibitive. Specifically, using an over-segmentation of the image instead of individual pixels can reduce noise and potentially improve the results of statistical post-analysis.

In this innovation, a metric learning approach is presented to improve the performance of unsupervised hyperspectral image segmentation. The prototype demonstrations attempt a superpixel segmentation in which the image is conservatively over-segmented; that is, the single surface features may be split into multiple segments, but each individual segment, or superpixel, is ensured to have homogenous mineralogy.

A segmentation strategy was tested based on the "Felzenszwalb" algorithm for its simplicity and computational efficiency. This approach represents the hyperspectral image as an 8-connected grid of pixels that can begin as independent segments. Edges between nodes represent the distance between neighboring spectra, and each is weighted according to a measure of distance between pixels. The algorithm iteratively joins neighboring pixels together into larger segments, and describes each segment by the minimum spanning tree of edges that joins all segments in the cluster.

Hyperspectral segmentation algorithms partition images into spectrally homogenous regions. However, the exact definition of homogeneity is dependent on the chosen similarity metric. The segmentation algorithm is augmented with a task-specific distance metric. Here, a Mahalanobis distance metric is used, learned from training data. By leveraging a (small) set of labeled pixels with known mineralogical interpretations, the metric suppresses uninformative spectral content. Multiclass linear discriminant analysis (LDA) is used to maximize the ratio of between-class vs. within-class separation, defined by the Rayleigh quotient computed over labeled training data. Other distance metrics and segmentation strategies are possible, and can be substituted for these choices in modular fashion as different applications demand.

❷ Basic Operational Robotics Instructional System

Lyndon B. Johnson Space Center, Houston, Texas

The Basic Operational Robotics Instructional System (BORIS) is a six-degree-of-freedom rotational robotic manipulator system simulation used for training of fundamental robotics concepts, with in-line shoulder, offset elbow, and offset wrist. BORIS is used to provide generic robotics training to aerospace professionals including flight crews, flight controllers, and robotics instructors. It uses forward kinematic and inverse kinematic algorithms to simulate joint and end-effector motion, combined with a multibody dynamics model, moving-object contact model, and X-Windows based graphical user interfaces, coordinated in the Trick Simulation modeling environment.

The motivation for development of BORIS was the need for a generic system for basic robotics training. Before BORIS, introductory robotics training was done with either the SRMS (Shuttle Remote Manipulator System) or SSRMS (Space Station Remote Manipulator System) simulations. The unique construction of each of these systems required some specialized training that distracted students from the ideas and goals of the basic robotics instruction.