Information Sciences

Adaptive System Modeling for Spacecraft Simulation

Lyndon B. Johnson Space Center, Houston, Texas

This invention introduces a methodology and associated software tools for automatically learning spacecraft system models without any assumptions regarding system behavior. Data stream mining techniques were used to learn models for critical portions of the International Space Station (ISS) Electrical Power System (EPS). Evaluation on historical ISS telemetry data shows that adaptive system modeling reduces simulation error anywhere from 50 to 90 percent over existing approaches.

The purpose of the methodology is to outline how someone can create accurate system models from sensor (telemetry) data. The purpose of the software is to support the methodology. The software provides analysis tools to design the adaptive models. The software also provides the algorithms to initially build system models and continuously update them from the latest streaming sensor data. The main strengths are as follows:

- Creates accurate spacecraft system models without in-depth system knowledge or any assumptions about system behavior.
- Automatically updates/calibrates system models using the latest streaming sensor data.
- Creates device specific models that capture the exact behavior of devices of the same type.
- Adapts to evolving systems.
- Can reduce computational complexity (faster simulations).

This work was done by Justin Thomas of Johnson Space Center. For further information, contact the JSC Innovation Partnerships Office at (281) 483-3809. MSC-24419-1

Lidar-Based Navigation Algorithm for Safe Lunar Landing
This algorithm could be used as a sensor approach for navigation of autonomous air vehicles for military surveillance.

NASA's Jet Propulsion Laboratory, Pasadena, California

The purpose of Hazard Relative Navigation (HRN) is to provide measurements to the Navigation Filter so that it can limit errors on the position estimate after hazards have been detected. The hazards are detected by processing a hazard digital elevation map (HDEM). The HRN process takes lidar images as the spacecraft descends to the surface and matches these to the HDEM to compute relative position measurements. Since the HDEM has the hazards embedded in it, the position measurements are relative to the hazards, hence the name Hazard Relative Navigation.

HRN processing starts with an initial elevation map from the Hazard Detection and Avoidance (HDA) phase. This map is generated by mosaicking the lidar over the Hazard Map Area (HMA). A feature selector is applied to the map to find a reference surface point that is surrounded by significant terrain relief and is therefore easier to identify in subsequent lidar images. This reference point does not have to be the landing site, and it probably won't be because the landing site should be free of terrain relief.

Next, the gimbal points the lidar sensor at the reference point and a lidar image is taken. The lidar image is converted to 3D points and these points are transformed into the local level coordinate frame using the current knowledge of the spacecraft position and attitude. These points are re-gridded into an elevation map. This elevation map is spatially correlated with the HDEM to determine the position change of the reference point in the local level frame between where it was predicted to be given the current state and its observed position when the HDEM was constructed.

The reference point is not actually moving in the local level frame, so this change in position is actually a measurement of current navigation state error growth from the time the HDEM was created. Since attitude errors are expected to be very small, the change in position of the reference point is most likely due to errors in the position of the spacecraft. This process is repeated with multiple new lidar images as the spacecraft descends.

During descent, the correlation performance degrades due to the shrinking field of view, increasing resolution and changing in view angle. The ground sample distance (GSD) of the basemap should be no more than twice the GSD of the current lidar map. To prevent the correlation from failing, resulting in a loss of knowledge of the position error on the reference point, a new base map is generated for correlation. This new base map is created by mosaicking the lidar around the landing site. A new, higher-resolution elevation map is created from the lidar mosaic. The feature selector is applied to the new base map to generate a new reference point. Lidar images are then taken of this new reference point and correlated with the new base map.

The process of generating a new base map, then correlating lidar images to it, is repeated until the beginning of vertical descent (30 m). Each time the base-map changes, it is correlated with the previous base map to tie its position to the original HDEM. This correlation introduces a fixed error to the estimate of the change in position of the original reference point. Fortunately, this fixed error is a function of the resolution of the corresponding base map, so the fixed error contribution is decreasing.

The algorithm is related to other motion and velocity estimation algorithms, but is different because the data processed is 3D points, not camera images. This difference in input data makes a large difference in how feature selection and correlation are implemented. The algorithm also must handle oblique viewing angles and relative high sensor noise; both of these make HRN challenging. Finally the HRN algorithm actually commands the lidar to collect data during descent that is the best for HRN. This "Active Vision" approach was not used in previous work.

This work was done by David M. Myers, Andrew E. Johnson, and Robert A. Werner of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).

The software used in this innovation is available for commercial licensing. Please contact Daniel Broderick of the California Institute of Technology at danielb@caltech.edu. Refer to NPO-47115.

⟩ Tracking Object Existence From an Autonomous Patrol Vehicle
These techniques could be part of a mobile surveillance system attached to a ground vehicle, boat, or airplane.

NASA's Jet Propulsion Laboratory, Pasadena, California

An autonomous vehicle patrols a large region, during which an algorithm receives measurements of detected potential objects within its sensor range. The goal of the algorithm is to track all objects in the region over time. This problem differs from traditional multi-target tracking scenarios because the region of interest is much larger than the sensor range and relies on the movement of the sensor through this region for coverage. The goal is to know whether anything has changed between visits to the same location. In particular, two kinds of "alert" conditions must be detected: (1) a previously detected object has disappeared and (2) a new object has appeared in a location already checked.

For the time an object is within sensor range, the object can be assumed to remain stationary, changing position only between visits. The problem is difficult because the upstream object detection processing is likely to make many errors, resulting in heavy clutter (false positives) and missed detections (false negatives), and because only noisy, bearings-only measurements are available. This work has three main goals:

(1) Associate incoming measurements with known objects or mark them as new objects or false positives, as appropriate. For this, a multiple hypothesis tracker was adapted to this scenario.

(2) Localize the objects using multiple bearings-only measurements to provide estimates of global position (e.g., latitude and longitude). A nonlinear Kalman filter extension provides these 2D position estimates using the 1D measurements.

(3) Calculate the probability that a suspected object truly exists (in the estimated position), and determine whether alert conditions have been triggered (for new objects or disappeared objects). The concept of a "probability of existence" was created, and a new Bayesian method for updating this probability at each time step was developed.

A probabilistic multiple hypothesis approach is chosen because of its superiority in handling the uncertainty arising from errors in sensors and upstream processes. However, traditional target tracking methods typically assume a stationary detection volume of interest, whereas in this case, one must make adjustments for being able to see only a small portion of the region of interest and understand when an "alert" situation has occurred. To track object existence inside and outside the vehicle's sensor range, a probability of existence was defined for each hypothesized object, and this value was updated at every time step in a Bayesian manner based on expected characteristics of the sensor and object and whether that object has been detected in the most recent time step. Then, this value feeds into a sequential probability ratio test (SPRT) to determine the "status" of the object (suspected, confirmed, or deleted). Alerts are sent upon selected status transitions. Additionally, in order to track objects that move in and out of sensor range — and update the probability of existence appropriately — a variable "probability detection" has been defined and the hypothesis probability equations have been re-derived to accommodate this change.

Unsupervised object tracking is a pervasive issue in automated perception systems. This work could apply to any mobile platform (ground vehicle, sea vessel, air vehicle, or orbiter) that intermittently revisits regions of interest and needs to determine whether anything interesting has changed.

This work was done by Michael Wolf and Lucas Scharenbroich of Caltech for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-47274