

# Persistent hyperspectral adaptive multi-modal feature-aided tracking

Andrew C. Rice<sup>a</sup>, Juan R. Vasquez<sup>a</sup>, John Kerekes<sup>b</sup>, Michael J. Mendenhall<sup>a</sup>

<sup>a</sup>Numerica Corporation 2661 Commons Blvd. Ste. 210 Beavercreek, OH 454341

<sup>b</sup>Chester F. Carlson Center for Imaging Science Rochester Institute of Technology  
54 Lomb Memorial Drive Rochester, NY 14623-5604

## ABSTRACT

An architecture and implementation is presented regarding persistent, hyperspectral, adaptive, multi-modal, feature-aided tracking within the urban context. A novel remote-sensing imager has been designed which employs a micro-mirror array at the focal plane for per-pixel adaptation. A suite of end-to-end synthetic experiments have been conducted, which include high-fidelity moving-target urban vignettes, DIRSIG hyperspectral rendering, and full image-chain treatment of the prototype adaptive sensor. Corresponding algorithm development has focused on: motion segmentation, spectral feature modeling, classification, fused kinematic/spectral association, and adaptive sensor feedback/control.

**Keywords:** Hyperspectral, Feature aided tracking, Performance driven sensing, Adaptive sensors

## 1. INTRODUCTION

### 1.1 Background and Objectives

While current intelligence, surveillance, and reconnaissance (ISR) assets are very effective in certain situations, it is recognized that new ISR challenges require the collection and processing of all available information to achieve high detection and tracking performance. To provide complementary data and enable the synergistic gain from the fusion of multiple sensing modalities, there have been recent efforts to incorporate multiple sensors on the same platform. This approach can be effective and is likely to be the only feasible way to collect simultaneous data from widely disparate systems such as visible electro-optical imagers, microwave synthetic aperture radars and low frequency signals intelligence systems. That is because these systems require vastly different detector material and apertures (optical lenses and antennae) due to the physics involved.

However, within a limited spectral range (e.g., visible through infrared electro-optical wavelengths) one can envision an integrated multi-modality sensor that does not violate the laws of physics and takes advantage of the ongoing advances in microelectromechanical systems (MEMS) and nanotechnology. The miniature nature of these devices allow concepts that can be integrated at or near the detector focal plane, thus leading to compact size and much simplified alignment techniques. On-device co-alignment addresses one of the limitations of multiple single-mode sensors, which is the inability to collect perfectly aligned imagery across modalities thereby limiting the possible fusion gain. In addition to gains possible from integrated multi-modality sensors, there is an opportunity for improved ISR performance from sensors designed to optimally adapt based on feedback from integrated real-time exploitation algorithms in a *performance driven sensing* paradigm.

An example of such a sensor – the motivation for this research – is the Rochester Institute of Technology Multi-Object Spectrometer (RITMOS),<sup>1</sup> shown in Figure 1. This sensor utilizes a digital micro-mirror array at the focal plane of its fore-optics. Nominally, the micro-mirrors reflect light along an imaging path which produces a fully framed image. Individual micro-mirrors can be commanded to flip, directing that portion of the scene along a hyperspectral imaging path. A spectrometer and high resolution focal plane array measure the per-pixel hyperspectral signature of these pixels. Ideally, the dispersed pixels should not be allowed to overlap along their axis of dispersion; this manifests as a constrained resource optimization problem via micro-mirror tasking.

---

Further author information: (Send correspondence to A.C.R.)

A.C.R.: Email: Andy.Rice@Numerica.us, Telephone: 1 937 427 9703

J.R.V.: Email: Juan.Vasquez@Numerica.us, Telephone: 1 937 427-9725

J.P.K.: Email: kerekes@cis.rit.edu, Telephone: 1 585 475-6996




Figure 1. The Rochester Institute of Technology Multi-Object Spectrometer (RITMOS)

## 1.2 Approach

This effort is focused on the development and evaluation of hyperspectral feature-aided tracking algorithms in a persistent, remote-sensing employment; adaptive sensor control being the key enabler to this approach. A tightly-coupled effort is researching adaptive, multi-modal sensing at the device level.<sup>2</sup> Together, a performance-driven-sensing research framework is being formed. Simulated data has been used to facilitate early algorithm integration and testing.

Figure 2 illustrates a fully automated feature-aided target tracking system based on the desire to detect and track moving targets. The system supports on-the-fly target modeling (no *a priori* model dependence), the use of features to aid tracking, target classification and background modeling. Targets are initialized primarily with motion segmentation algorithms, but additional feature-based segmentation is used to provide detections on targets that have stopped (e.g., “sitters”). The data association component uses detections that have been classified to determine the appropriate pairings of these detections to existing targets (or as new targets). A multiple hypothesis tracking (MHT) approach provides robustness for ambiguous situations such as closely spaced targets, short-term occlusion, and dynamic target behavior.<sup>3</sup>

## 1.3 Overview

A description of the data synthesis technique is given in Sec. 2. Particular attention is given to employment of the RIT Digital Imaging and Remote Sensing Image Generation (DIRSIG) model, as well as radiometric sampling of the resulting radiance fields. Methods for feature aided tracking are discussed in Sec. 3, including measurement-to-track data association and track-to-track stitching. The detector-level sensor resource management (SRM) approach is given in Sec. 4. Finally, examples of tracker results are given in Sec. 5.

## 2. DATA SYNTHESIS

A significant focus for this effort was to establish and employ a high-fidelity remote-sensing hyperspectral simulation environment to support multiple target feature-aided tracking research and development. The criteria for this simulation required reasonable duplication of real-world challenges such as atmospheric attenuation, multiple scattering, attenuated transmittance, shadowing, spectral mixing, and the stochastic spectral/spatial

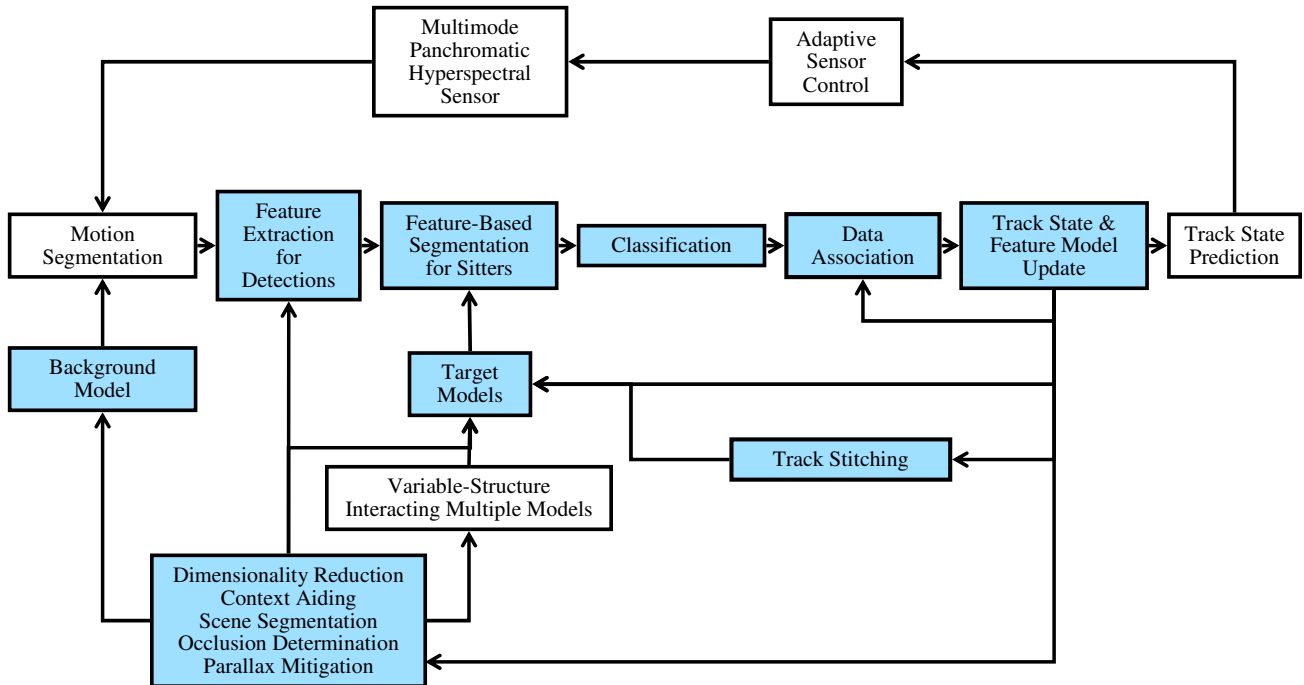


Figure 2. Multiple target feature-aided tracking architecture with adaptive sensor control. Shaded boxes indicate the elements most impacted by the availability of HSI feature data.

nature (i.e., texture) of materials. Additionally, the ability to rapidly synthesize complex target kinematics in a clutter and obscuration dense environment was desirable. Finally, a precise treatment of coordinates and terrain was necessary in order to generate subpixel level truth data for performance evaluation.

The DIRSIG model is a first principles based synthetic image generation model developed by the Digital Imaging and Remote Sensing Laboratory at Rochester Institute of Technology.<sup>4</sup> Here, DIRSIG is used to render the panchromatic imaging-side data, as well as the per-pixel hyperspectral data, as directed by the SRM. The simulation incorporated a large urban area with 211 moving vehicles and random permutations of 75 spectral signatures and 5 geometrical models. Vehicle motion was generated in conjunction with the SUMO traffic simulation tool, as described in a companion paper.<sup>2</sup> The scenario represents 210 seconds of capture at 10 fps, resulting in 1200 total frames after imposing dropouts. The sensor array size is 880 x 560 pixels at 17  $\mu\text{m}$  with 3 times spatial oversampling. The hyperspectral response is 0.4 - 1.0  $\mu\text{m}$  with a bin spacing of 0.01  $\mu\text{m}$ . The sensor platform operated at 3000m altitude, stationary, and nadir looking to provide an overall FOV of 660m x 420m with a GSD of 0.75m.

Hyperspectral data synthesis is a computationally intense endeavor, and is commonly performed in a batched computation prior to exploitation. Due to the adaptive feedback nature of this system, coupled with the SRM constraints permitting only a fraction of the total pixels to be hyperspectral at a given instant, an alternate, hybrid rendering approach was used. The panchromatic images were rendered with DIRSIG in a traditional fully-framed batch mode. Next, DIRSIG was integrated into the tracking testbed via its “interactive-mode,” whereby single rays of path-radiance can be rendered individually. Thus, the SRM’s decision at each frame caused only those pixels called for by the SRM to be rendered with full hyperspectral resolution. This method also presented a solution to the single-threaded architecture of DIRSIG. Since the tracking testbed is hosted by Matlab, which is also single-threaded, a Java rendering manager was developed. This architecture is shown in Fig. 3. In this way, the tracking testbed uses many DIRSIG instantiations distributed across many local or remote machines to render one frame’s worth of hyperspectral pixels in parallel. As a means of comparison, it takes 2 hours to render a full hyperspectral frame (all pixels) for this simulation, producing 2GB of data. The same frame, where the SRM selected 600 pixels to be hyperspectral and on-demand rendering was distributed across seven DIRSIG nodes, took 45 seconds and generated 3MB of data. Furthermore, since Monte Carlo trials

of the same tracking scenario frequently demand the same pixels, a persistent cache is used to store path-radiance values for a further decrease in rendering time.

While simulated data is extremely useful in the development process, it is important to ensure that real-world phenomenology is properly represented. The following discussion presents a key aspect of the simulation software that focuses on meeting this very need. Simulations yield synthesized sensor-reaching-radiance images relative to the front aperture of a nominal sensor. Ultimately, the exploitation algorithms are intended to process sampled data representative of a real world imaging spectrometer. Radiometric sampling is the process which converts the spectral sensor-reaching-radiance values ( $\frac{W}{cm^2 \cdot sr \cdot m}$ ) into digitally sampled voltage values. This process accounts for the first-order effects of any filters, dispersive elements, refractive optics, detector elements, integration time, shot noise, and the A/D converters with readout circuitry, and is given extensive treatment in Schott.<sup>5</sup> First, the front aperture spectral sensor-reaching radiance  $L$  must be converted to irradiance  $E$  on the focal plane, given by the well known *camera equation*:

$$E = \frac{L}{G\#} \quad (1)$$

where  $G\#$  is a measure of optical throughput. For simple lens models, the following relationship is fairly accurate:

$$G\# = \frac{1 + 4(f\#)^2}{\tau_\lambda \pi} \quad (2)$$

Notably  $f\# = \frac{f}{d}$ , and the focal length of the lens  $f$  is also an input parameter (solely for the purpose of FOV determination). We then hold  $f$  constant during the post-processing of a simulation and leave the aperture diameter  $d$  as a free parameter. The wavelength dependent transmittance of the entire optical train is given by  $\tau_\lambda$ .

Now, the spectral irradiance at the focal-plane ( $\frac{W}{cm^2 \cdot m}$ ) must be converted to a voltage signal  $S$  in a manner consistent with a real-world EO detector element:

$$S = \frac{Et_{int}A_d \frac{dV}{dN} QE \lambda^2}{hc} [V] \quad (3)$$

Where  $t_{int}$  is the integration time (seconds), and  $A_d$  is the effective detector area ( $cm^2$ ).  $\frac{dV}{dN}$  is the output sensitivity or charge-to-voltage conversion factor (volts per electron).  $QE$  is the quantum efficiency of the detector (electrons per photon).  $hc$  is Planck's constant by the speed of light in a vacuum.

Dark current can be a significant source of additive noise, but is generally the most straightforward to account for. In practice, dark frames are measured with a real-world instrument in-situ with all irradiance blocked, at the same ambient temperature and integration time as the imagery. Frequently an imaging spectrometer may employ agile integration times; dark frames should be measured for each distinct timing. The appropriate dark frame is then subtracted from collected images. Assuming this procedure is always performed, the residual dark current noise effects depend on how well the dark frame models the random quantum radiometric events during the imagery integration. There is also a minor effect on quantization noise and saturation due to dark frame subtraction. Residual dark current noise is not considered a primary noise source, and has not yet been added to the system simulation.

Focal-plane read-out circuitry is complex, involving many charge hand-off operations and analog to digital conversions. The noise introduced by this circuitry is frequently characterized by a simple metric, the RMS electron readout error,  $N_{\mathbf{RMS}e^-}$ . This is the expected value of the number of electrons lost or gained during a read-out, and tends to be independent of integration time. The readout noise  $N_{RO}$  is modeled as a voltage equivalent

$$N_{RO} = N(0, N_{\mathbf{RMS}e^-}^2) \frac{dV}{dN} [V] \quad (4)$$

where  $N(\cdot, \cdot)$  is a normal distribution.

Shot noise arises from the quantum nature of light and electrical charge. Particularly for high photon count signals, the random arrival nature of photons can become the most significant source of noise in the system. In

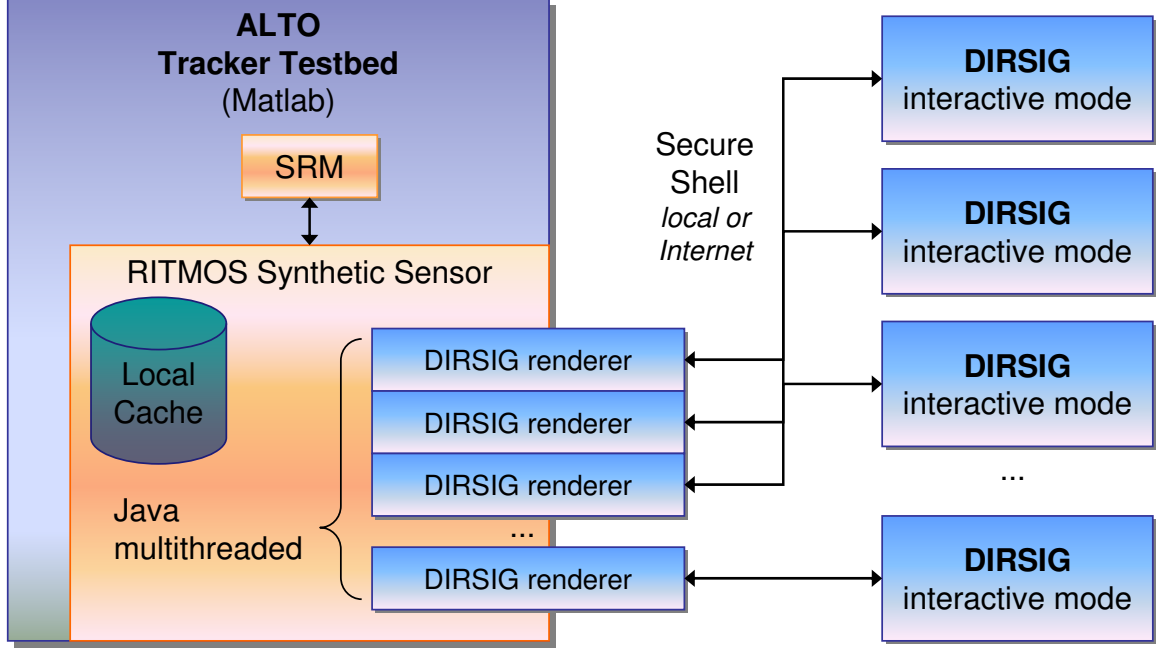


Figure 3. Distributed, on-demand DIRSIG rendering architecture.

the system simulation, shot noise  $N_{\Phi}$  is modeled as a voltage equivalent

$$N_{\Phi} = N(0, S) [V] \quad (5)$$

Saturation occurs when a maximum signal level  $V_{SAT}$  (V) is reached during the read-out procedure for a detector element.

Finally, the signal+noise voltage at each detector element is converted into a quantized digital number  $S_{DN}$  with  $b$  bits of resolution.

$$S_{DN} = \left\lfloor 2^b \frac{\max \{ \min \{ S + N_{RO} + N_{\Phi}, V_{SAT} \}, 0 \}}{V_{SAT}} + 0.5 \right\rfloor \quad (6)$$

This inherently introduces the appropriate quantization noise. The resulting synthetic signal can be treated as if it came from a real imaging spectrometer: exploited while in memory, saved to storage as a 16-bit tiff file, down-sampled to an 8-bit imagery file, or perhaps compressed for storage or transmission. Radiometric sampling is ultimately a crucial part of the sensor simulation; if done poorly, it is almost certain that the exploitation algorithm performance results will be vastly optimistic or pessimistic and will not match experiments on real data.

### 3. FEATURE-AIDED TRACKING

In general, the family of remote spectral sensors being studied are capable of providing both a feature measurement as part of its intrinsic signal (e.g., spectral radiance), and a temporally derived kinematic (e.g., motion) observation. It is of great interest to optimally utilize feature information in the framework of the kinematic tracking system. Incorporating features into the measurement-to-track-association cost is a good way to do so. This class of methodologies has the potential to resolve kinematic track ambiguities before computationally expensive downstream processing, such as MHT maintenance, is wasted.

Let  $C_{i,j}$  be the incremental cost of associating track  $\mathbf{x}_i$  with observation  $\mathbf{z}_j$ . This cost will be the sum of several component costs. First, the kinematic cost is determined from the logarithm of a likelihood  $C_{i,j}^K = -\log(p(\mathbf{z}_j|\mathbf{x}_i))$ , where the probability is determined from the Mahalanobis distance between  $\mathbf{x}_i$  and  $\mathbf{z}_j$ . Next, one or more feature

costs are similarly computed.  $C_{i,j}^{F_n}$  is the log-likelihood based cost of associating the track with the measurement in some dimension  $n$  of feature space. The total feature-aided association cost is then

$$C_{i,j} = \mu_K \tilde{C}_{i,j}^K + \sum_n \mu_{F_n} \tilde{C}_{i,j}^{F_n}, \quad (7)$$

where  $\mu$  is a weighting term and  $\tilde{C}$  is a normalized cost. The constant weighting terms are empirically or analytically determined *a priori*, and indicate the relative importance of each component cost. For example, setting  $\mu_K$  to 1 and all  $\mu_{F_n}$  to 0 is equivalent to a kinematic-only tracker. In general, the total cost should be stable and reasonable from one association to another within the same frame of data, as well as from one frame to the next. This is especially true since the windowed MHT builds rolled-up hypothesis costs from histories of cost data.

Several approaches exist to arrive at the normalized costs. In an empirical approach, constant offsets  $C_0$  can be applied. This essentially sets a boundary between “good” and “bad” costs by setting the origin of the corresponding cost axis:

$$\tilde{C}_{i,j}^K = C_{i,j}^K - C_0^K \quad (8)$$

$$\tilde{C}_{i,j}^{F_n} = C_{i,j}^{F_n} - C_0^{F_n} \quad (9)$$

This normalization approach relies on establishing reasonable offsets, which may be brittle across complex operating conditions. Another approach “learns” costs boundaries based on all measurements  $g$  which fall within the same track gate  $G$ :

$$\tilde{C}_{i,j}^K = \frac{C_{i,j}^K}{\sum_{g \in G} C_{i,g}^K} \quad (10)$$

$$\tilde{C}_{i,j}^{F_n} = \frac{C_{i,j}^{F_n}}{\sum_{g \in G} C_{i,g}^{F_n}} \quad (11)$$

This approach is less ad-hoc than empirical offsets, but becomes less valid when the true measurement is *not* within the gate, and degenerates when only one measurement is within the gate.

As part of the feature-aided association scoring research, a track stitching capability has been implemented. Much like the measurement-to-track associator, the track-to-track stitcher uses both kinematic and feature-based cost to test the likelihood of a prior terminated track associating (stitching) with a new track. A single-hypothesis, recursive, promotion-time track stitcher has been implemented. It is deemed a single-hypothesis approach as it evaluates the potential to stitch a new track *once* at time of promotion to all previously terminated tracks under the assumption that a good match would represent a target that was temporarily occluded or dropped for some reason. This has proven to be very effective at tracking targets through medium-length ambiguities which are typically beyond the reach of the MHT. A more sophisticated multiple hypothesis track stitcher has been conceived for future development which will provide more robust stitching for the multi-target case and with deferred decision-making. In addition, it enforces a test for kinematic feasibility when considering a track stitch versus attempting to stitch with all previously terminated tracks.

A simple feature model, the spectral angle, has been implemented as an initial capability. The spectral angle is exponentially age weighted, such that a target’s spectral feature model adapts over time. Another benefit of age weighting is that misassociations leading to corrupting feature updates will tend to be “forgotten”.

#### 4. SENSOR RESOURCE MANAGEMENT

An emerging class of hyperspectral imagers allow for increasingly fine-granularity sensor resource management (SRM), extending to per-pixel adaptation via MEMS technology. These imagers are ideally suited for performance driven sensing. Generally, limitations are imposed on how the pixels may be tasked at any given time; these arise due to spectral dispersion, optical geometry, bandwidth, and the like. For instance, an instrument may nominally collect a single broadband irradiance value at each pixel, but allow a subset of pixels to collect full

hyperspectral measurements. Pixel adaptation then becomes constrained optimization problem which attempts to maximize the utility of the sensor, employed here in a target tracking sense. Previous research<sup>6</sup> has shown that track-centric heuristics are a reasonable manner in which to form a utility function  $U_{rc}(t)$  in pixel row/column (RC) space. It is nontrivial, however, to formulate  $U_{rc}(t)$  in an operational context: hypothesis-dense and deeply-windowed MHT, moving platform with position/orientation uncertainty, and causal-time/latency constraints. In order to capture these issues, the notion of the predicted utility function is introduced:  $U_{rc}(t_{k+1|k})$ , which is the predicted utility distribution one time-step into the future given only the current knowledge of the system. Here, we present such a formulation developed within the tracker testbed.

The SRM module executes once per frame after that frame's exploitation, incorporation and MHT maintenance is complete. Thus, when the SRM executes at the end of frame  $t_k$ , all track state estimates are *posterior*, denoted  $\hat{\mathbf{x}}_i(t_k|k)$ , and incorporate the kinematic position and velocity of the track. The SRM must then determine how to optimize the next frame of data, namely  $U_{rc}(t_{k+1|k})$ ; for this reason all tracks are propagated via an extended Kalman filter to their subsequent *prior* state estimates  $\hat{\mathbf{x}}_i(t_{k+1|k})$  and associated covariances  $P_i(t_{k+1|k})$ .




Figure 4. Utility distribution function in pixel row/column space.

Intrinsic to the observation measurement processing are several state transformations. In the case of a moving platform, these transformations vary as a function of time. Measurements originate in pixel row/column (RC) space, and are immediately transformed into a position-azimuth-elevation (PAE) space via the transformation  $H_{rc \rightarrow pae}(t_k)$ , which is derived from the camera model and platform position/orientation. The MHT natively utilizes an East-North-Up (ENU) coordinate system, which is locally fixed and tangent to the Earth, and is

Cartesian. The transformation from PAE to ENU is nontrivial, and for ground targets is facilitated by ray-casting towards a digital elevation model, such as from the 2000 Shuttle Radar Topography Mission. Hence it is given as a linearized approximate per-track  $i$ , valid only in its vicinity:  $H_{pae \rightarrow enu}^i(t_k)$ . Notably, in the moving-platform case,  $H_{rc \rightarrow pae}(t_{k+1|k})$  and  $H_{pae \rightarrow enu}^i(t_{k+1|k})$  necessarily rely on a predicted platform state. Based on these state transformations, the position with uncertainty of each track can be modeled as a scaled and rotated 2D Gaussian probability distribution function in the row/column space of the next frame, denoted  $N_{rc}^i(t_{k+1|k})$ .

Consider  $\mathbb{X}(t)$ , the union of track state estimates  $\hat{\mathbf{x}}_i(t)$  within the entire MHT at time  $t$ .  $\mathbb{X}(t)$  will contain tracks which are partially redundant (have similar positions, but perhaps different covariances), and mutually exclusive (exist in different hypotheses and share historic measurements). Each track  $i$  carries in its state a hyperspectral, age-weighted feature model  $\mathbf{f}_i(t)$ . A “new track” flag  $A_i^N(t)$  is 1 when a track has never received a HSI measurement, and is otherwise 0. An age coefficient  $A_i^F(t) = 1 - \exp^{-(t-\tau)}$  is based on the time  $\tau$  of the last hyperspectral update for the track. Should a track lose simple panchromatic-based motion measurements for the  $m$  frames prior to time  $t$ , its “missed motion measurement” coefficient  $A_i^M(t) = 1 - \exp^{-m}$  will reflect this loss. Lastly, an ambiguity metric  $A_i^A(t)$  is based on the density and statistical closeness of each track and its in-hypothesis neighbors.

The utility function is then:

$$U_{rc}(t_{k+1|k}) = \max_{\hat{\mathbf{x}}_i(t_{k+1|k}) \in \mathbb{X}(t_{k+1|k})} N_{rc}^i(t_{k+1|k}) [C^N A_i^N(t_{k+1|k}) + C^F A_i^F(t_{k+1|k}) + C^M A_i^M(t_{k+1|k}) + C^A A_i^A(t_{k+1|k})] \quad (12)$$

This corresponds to the combination of many potentially overlapping 2D Gaussians, shown in Fig. 4. Values for the constants  $C$  are tuning parameters which adjust the relative importance of the track maintenance functions. Prior research<sup>6</sup> has explored selecting  $C$  based on multi-objective evolutionary algorithms in order to optimize  $U_{rc}$ . Finally, the inherent adaptive sensor constraints are applied to  $U_{rc}$  in order to task the micro-mirrors.

## 5. RESULTS

Figure 5 illustrates the challenging environment for this tracking scenario along with a snapshot of the tracking performance. All moving targets are automatically tracked without the need for user interaction. This includes track initiation, deletion, stitching of tracks that become occluded for both short and long periods of time. Crossing targets and dynamic behaviors are also tracked. All targets are tracked at all times with the exception of a few frames required for initiation. Figure 6 shows intermediate results for one iteration of the SRM control loop. Figures 7, 8, and 9 demonstrate the types of tracking challenges which this system architecture attempts to resolve. Future efforts will focus on improving the purity of track ID’s assigned to each track, with the anticipation that HSI features will provide a significant improvement over the kinematic only results.






Figure 5. Multiple target feature-aided tracking results.



(a) Three targets under track at time  $t_{k|k}$ . The small green ellipse indicates track position with uncertainty, whereas the larger blue ellipse indicates the 1-second position prediction with uncertainty.

(b) Utility map for the tracks predicted to time  $t_{k+1|k}$ , which manifests as a multi-modal mixture of 2D Gaussians which are scaled and rotated to match the track uncertainty. Occasionally, as with track #25, they grow so large that they are clipped. Visible in track #6, a different track from an inferior hypothesis has changed the utility distribution. White x's indicated the SRM's hyperspectral tasking decision.



(c) HSI Pixel hits (red) overlaid on pan image at time  $t_{k+1|k+1}$

Figure 6. Example results of SRM Utility computation and constrained optimization for a single iteration of the control loop. In this case, two of three tracks were successfully updated with new hyperspectral signatures. The remaining track, #25, was just missed due to undergoing a high-dynamics turn and competing column-wise for pixels against track #13. Notably, there is no penalty for collecting surplus hyperspectral pixels, as the measurement-to-track association algorithm is able to ignore them.



(a) A track enters occlusion.

(b) The track misses motion measurements, association with motion measurements, indicated by the red position estimate.

(c) The track resumes motion measurements, association with motion measurements.

Figure 7. Example results of the feature-aided MHT.






Figure 8. Three targets are successfully tracked through kinematic ambiguity, high target dynamics, merged measurements, and tree canopy occlusion. The high degree of ambiguity drives the SRM to allocate many HSI pixels on these tracks, thus assisting the measurement-to-track association challenge.



(a) Track #135 enters occlusion at time  $t_k$ .



(b) A newly initialized track emerges from occlusion at time  $t_{k+11s}$ , and is successfully stitched to track #135.

Figure 9. Track #135 undergoes an extended occlusion under tree canopy for approximately 11 seconds. An intersecting street with traffic meets mid-occlusion. This occlusion duration may be beyond the reach of a tractable, real-time, video-rate MHT. Feature-aided track stitching is useful in this type of highly ambiguous situation.

## ACKNOWLEDGMENTS

This effort has been sponsored by the Air Force Office of Scientific Research (AFOSR) under agreement number FA9550-08-1-0028 (AFOSR-BAA-2007-08) “Integrated Multi-modal Sensing, Processing, & Exploitation”

Many contributions for this paper have been made by additional colleagues at the Rochester Institute of Technology. A companion paper describes their efforts and findings.<sup>2</sup>

## REFERENCES

- [1] Meyer, R. D., Kearney, K. J., Ninkov, Z., Cotton, C. T., Hammond, P., and Statt, B. D., “Ritmos: a micromirror-based multi-object spectrometer,” *Proc. SPIE 5492*, 200–219 (2004).
- [2] Kerekes, J. P., Presnar, M. D., Fourspring, K. D., Ninkov, Z., Pogorzala, D. R., Raisanen, A. D., Rice, A. C., Vasquez, J. R., Patel, J. P., MacIntyre, R. T., and Brown, S. D., “Sensor modeling and demonstration of a multi-object spectrometer for performance-driven sensing,” *Proc. SPIE 7334* (2009).
- [3] Blackman, S., [*Multiple-Target tracking with radar applications*], Artech House, Norwood, MA (1986).
- [4] Schott, J. R., Brown, S. D., Raqueño, R. V., Gross, H. N., and Robinson, G., “An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development,” *Canadian Journal of Remote Sensing* **25**(2), 99–111 (1999).
- [5] Schott, J. R., [*Remote Sensing: the image chain approach, 2nd ed.*], Oxford University Press, Oxford, NY (2007).
- [6] Secrest, B. and Vasquez, J., “Optimal spatial sampling of hyperspectral imagery for fusion with panchromatic video in multitarget tracking,” Proceedings of the IEEE, Sensors Applications Symposium (2009).