An Application of Integrated Sensing and Processing Decision Trees for Target Detection and Localization on Digital Mirror Array Imagery

Carey E. Priebe†, David J. Marchette††, Youngser Park†, Robert R. Muise†††

† Center for Imaging Science, Johns Hopkins University, Baltimore, MD
†† Naval Surface Warfare Center, B10, Dahlgren, VA
††† Lockheed Martin, Orlando, FL

1. Introduction

We demonstrate the applicability of the Integrated Sensing and Processing Decision Trees [1] (ISPDTs) iterative denoising methodology to a set of digital mirror array [2] (DMA) imagery. The sensor sequentially collects 256 Hadamard [3] frames in order to produce a hyperspectral data cube that measures 256x256 pixels spatially and consists of 266 spectral bands. However, the sensor is (dynamically) programmable; given a pattern recognition task of interest it is desirable to learn a sequence of just a few Hadamard frames that will accomplish the task sans the requirement to collect (and process) the entire data cube. We demonstrate the construction, on one training hyperspectral data cube, of an iterative denoising tree that allows for detection and localization of a target using just two Hadamard frames; the entire hyperspectral data cube need not be collected in order to successfully perform the task.

2. Methods

We apply the ISPDT iterative denoising methodology to the hyperspectral vehicle imagery, I_t, t = 0, 1, 2, 3, 4. We use the hyperspectral data cube collected at time t = 2, I_2, as our training data. We focus on 64x210 image swaths, so as to obviate issues of sensor skewness and edge effects. Figure 1 depicts the full image (left) and the swath (right) for Hadamard frame 110 of the training data I_2. The swath shows the region delineated by the small box representing the training information regarding the target of interest. (No additional information regarding vehicle/non-vehicle pixels within the target box is utilized.)

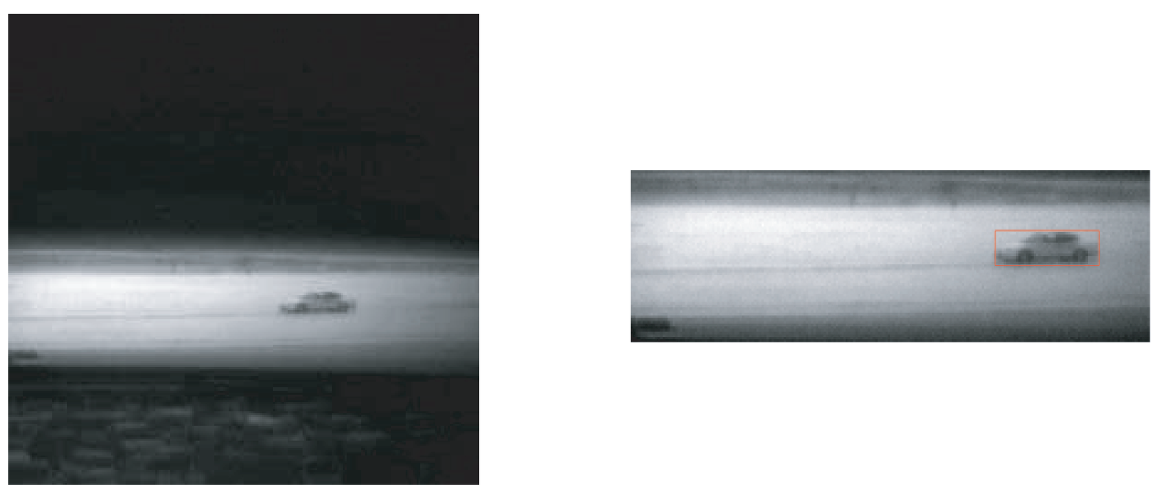

Figure 1. Full 256x256 image of hadamard frame 110 (left) and 64x210 swath (right) for I_2

No single Hadamard frame provides adequate performance in segregating target from non-target pixels, and so Hadamard frame 110 is selected for use at the root by virtue of its performance in providing the most clustering clarity -- a new Hadamard frame at each node is determined based on the best separation between the target box and the non-target pixels by using the adjusted Rand index criterion [4]. Model based clustering [5] is employed, in which a Bayesian Information Criterion (BIC) is used to determine the complexity and type of a Gaussian mixture fit to the data. The clusters are then defined in terms of the posterior likelihoods of the individual components. The Hadamard frame 110 pixels (gray scale intensity) cluster into three clusters; the spatial location of the pixels for these three clusters are represented by blue, green and yellow in the root node of the tree depicted in Figure 2.

In building an ISPDT, after clustering at a node, each cluster is processed in a branch of the tree. This subsequent processing proceeds in an analogous fashion to that of the root: if the pattern recognition task at hand can be adequately addressed (for the data falling to that branch) then tree growth (along that branch) is halted. It is necessary, of course, to perform the search for the best Hadamard frame once again in each branch, conditionally upon the results of all previous clusterings. In this application, the left-most (green) and right-most (yellow) clusters at the root of the ISPDT depicted in Figure 2 yield essentially pure non-target branches and no further tree growth is necessary in these branches, while the middle (blue) root cluster contains many of the target pixels as well as a significant number of the non-target pixels. This middle cluster must be processed further. In this middle cluster, we nd that Hadamard frame 151 provides clustering into two clusters; furthermore, one of these clusters (the left-most node at level three of the ISPDT depicted in Figure 2) contains a significant number of target pixels and no non-target pixels. Ergo, this leaf is labeled as the “target leaf”.

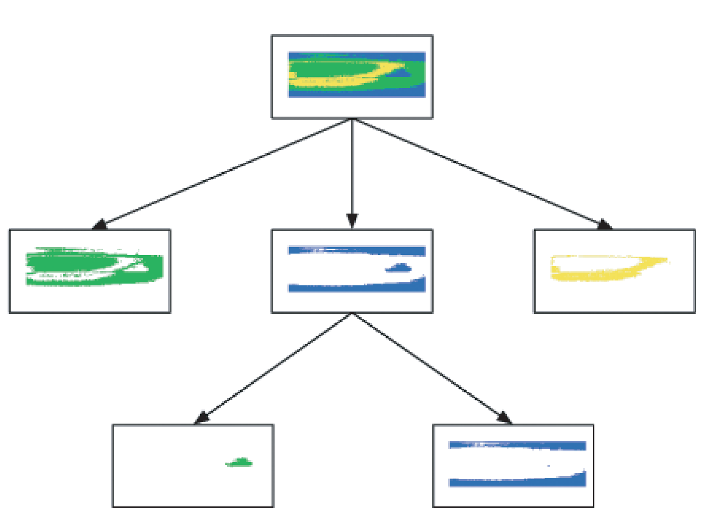

Figure 2. ISPDT constructed on I2 (training). The left-most node at level three of the tree is the “target leaf”

References:
[1] C.E. Priebe, D.J. Marchette, and D.M. Healy, “Integrated Sensing and Processing Decision Trees,” IEEE Trans. Pattern Analysis and Machine Intelligence, 26 (6), 699-708 (2004).
[2] R.A. DeVerse, R.R. Coifman, A.C. Coppi, W.G. Fateley, F. Geshwind, R.M. Hammaker, S. Valenti, F.J. Warner, and G.L. Davis, “Application of spatial light modulators for new modalities in spectrometry and imaging,” Proc. SPIE 4959, 12-22 (2003).
[3] M. Harwit and N.J.A. Sloane, Hadamard Transform Optics (Academic Press, New York, 1979).
[4]L. Hubbert and P. Arabie, “Comparing partitions,” Journal of Classification, 2 193-218 (1985).
[5] C. Fraley and A.E. Raftery, “Model-based clustering, discriminant analysis, and density estimation,” Journal of the American Statistical Association, 97 611-631 (2002).

http://www.cis.jhu.edu/~parky/isp

3. Results

We present successful detection/localization results on the vehicle imagery test data I_t, t = 0, 3, 4, described above, using the ISPDT trained on I_2 and depicted in Figure 2. There is no vehicle present at time t = 0, and (see Figure 3a) no pixels fall into the “target leaf”. This implies no detection, as desired. For t = 3, 4 there is a vehicle present, and (see Figure 3b and 3c) in each case the existence of pixels falling into the “target leaf” implies detection, as desired. Furthermore, the spatial location of these pixels indicates that the detection is indeed “on target”.

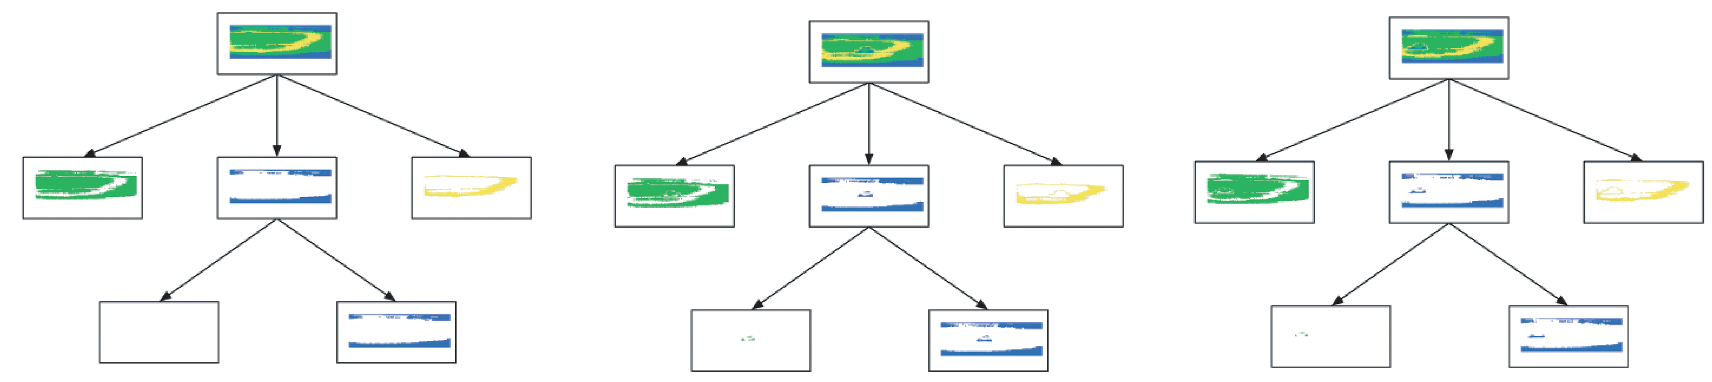

(a) (b) (c)

Figure 3. ISPDT detection/localization results for I_0 (a), I_3 (b), and I_4 (c) for testing. No vehicle is present at time t = 0, and no pixels falling into the “target leaf” implies no detection, as desired. Meanwhile, there is a vehicle present at time t = 3 and 4, and the existence of pixels falling into the “target leaf” implies detection, as desired. The spatial location of these pixels indicates that the detection is indeed “on target”.

Figure 4 depicts a more elaborate ISPDT, requiring non-trivial conditioning/adaptation, generated for I_3. In this tree, as before, Hadamard frame 110 is used at the root, and in the middle cluster at level two Hadamard frame 151 provides clustering into two clusters. While the “target leaf” -- the left-most node at level three -- provides target detection, there are a significant number of target pixels that fall to the right-most node at level three -- a non-target node.

When training on I_2, tree-building stops at this point. If training on I_3 (presented in Figure 4) it is preferable to continue the tree-building, using Hadamard frame 145 for additional accuracy. Thus the left-most node at level four is an additional “target leaf” for this illustration, and the two target leaves together allow more target pixels to be identified (while still yielding no false detections). This tree would be realized by dynamically programming the sensor to collect Hadamard frame 145 only for cases in which frames 110 and 151 leave detection unresolved.

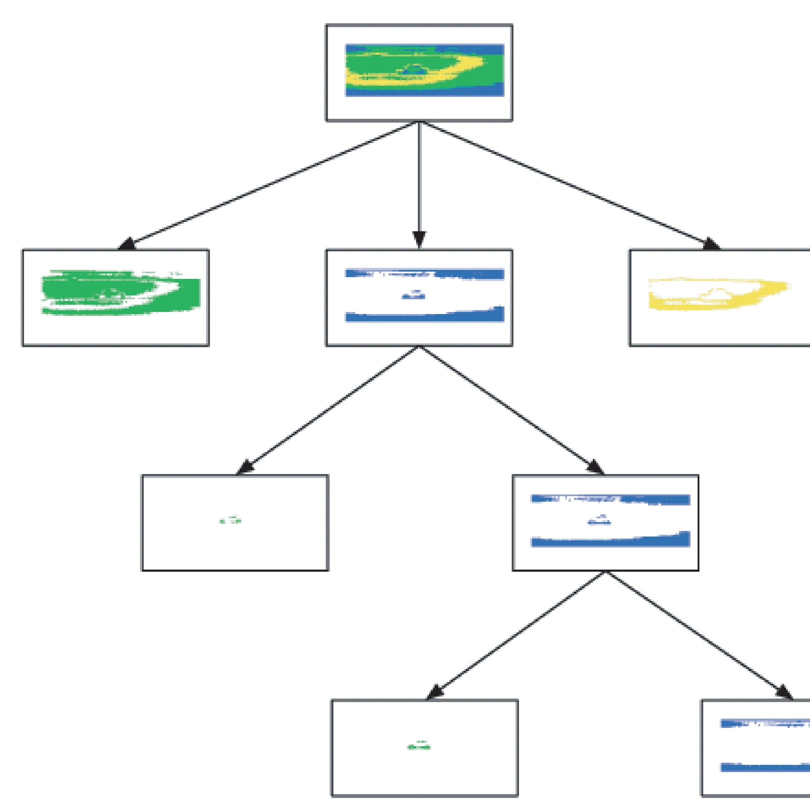

Figure 4. Illustration of a more elaborate ISPDT, requiring non-trivial conditioning/adaptation, for I_3.

4. Conclusions

We have successfully applied ISPDT to DMA hyperspectral imagery for detection/localization, demonstrating the potential of an integrated sensing/processing suite consisting of a dynamically programmable DMA sensor and ISPDT processing.

This example demonstrates detection and localization. When multiple target types are possible, a subsequent classification step -- possibly involving additional Hadamard frames, is employed. If, as is likely in practice, detection is not so “perfect” and some few off-target pixels are identified as falling into the “target leaf” (but still many on-target pixels are so identified) then post-processing under some spatial dependence scheme, such as maximum a posteriori (MAP) spatial filtering, can be used to perform the ultimate detection/localization.

Finally, while it may seem (and is indeed the case) that processing more elaborate than our individual pixel-based approach (such as the use of an edge detector) would make the detection and localization task trivial, the extremely time-sensitive nature of the pattern recognition applications envisioned here preclude elaborate processing schemes.

Work supported in part by the Applied and Computational Mathematics Program (ACMP) of the Defense Advanced Research Projects Agency (DARPA).