

Feature Discovery in Gray Level Imagery for One-Class Object Recognition

M. W. Koch and M. M. Moya

Abstract— Feature extraction transforms an object's image representation to an alternate reduced representation. In one-class object recognition, we would like this alternate representation to give improved discrimination between the object and all possible non-objects and improved generalization between different object poses. Feature selection can be time-consuming and difficult to optimize so we have investigated unsupervised neural networks for feature discovery.

We first discuss an inherent limitation in competitive type neural networks for discovering features in gray level images. We then show how Sanger's Generalized Hebbian Algorithm (GHA) removes this limitation and describe a novel GHA application for learning object features that discriminate the object from clutter. Using a specific example, we show how these features are better at distinguishing the target object from other non-target object with Carpenter's ART 2-A as the pattern classifier.

I. INTRODUCTION

An imaging sensor can provide a tremendous amount of data. Taking advantage of this information requires correctly interpreting patterns corresponding to various objects in the image. To automate image interpretation, we first compress the information contained in the large images into a more compact representation. The compression must preserve not only characteristic object features but also their relative location. The most popular and most straightforward compression method has been extracting hand-crafted features from the image data. Selecting meaningful discriminating hand-crafted features frequently requires an empirical search process.

We propose an alternative to selecting hand-

crafted features. We propose using data-driven discovery algorithms, in a receptive field neural network architecture, to discover relevant discriminating features. These feature discovery networks rely less on human intuition and more on real data characteristics to find important data features. Linker has shown that adaptive networks can self-organize to discover features similar to those learned in the mammalian visual cortex during prenatal development [7]. Cooper et. al. [6] have developed a learning law that explains many visual cortical plasticity experiments such as monocular and binocular deprivation and reverse suture.

The self-organizing multiple layer receptive field network discovers complex features in raw image data. After completing self-organization, the network extracts the learned features from the raw data and passes the output to the classification stage, which must decide the class of the input image based on the feature location and strength information.

To recognize targets with good detection probability and acceptable false alarm rate, a classifier must recognize new target examples, and it must reject all possible non-targets, even those classes excluded from training. For example, if we train a person to recognize Greek symbols, we do not have to show him Chinese characters as negative examples. We call a classifier that possesses these capabilities a one-class classifier. This paper describes the proposed feature discovery network and reports the results of using it with the ART 2-A network [1] as a one-class classifier [9] for target recognition in gray scale imagery.

II. COMPETITIVE FEATURE DISCOVERY NETWORKS LIMITATIONS

Other investigators have used unsupervised competitive neural networks, such as the Neocognitron [4], to discover features in images. While competitive feature discovery networks can learn features in binary data, they have difficulty in learning subtle features present in gray level data. Applications that use the Neocognitron with gray level imagery usually preprocess the data to produce a binary image [3, 8]. Other work has focused on training with bi-

This work performed at Sandia National Laboratories supported by the U.S. Department of Energy under contract DE-AC04-94AL85000. The authors are with Department 9133: Image and Signal Processing, Sandia National Laboratories, Albuquerque, NM 87185-0355

MASTER

nary data, but recognizing analog images [12].

Competitive neural networks learn prototypes from the raw image data. If we analyze the range of the raw data feature types, we find features containing a constant bias produce stronger outputs than low frequency features, which produce stronger outputs than high frequency features. Self-organized competition on the raw receptive field data will learn prototypes for only those features that produce strong outputs: the constant bias and some low frequency features. The networks ignores the middle and high frequency features, which often provide fine discrimination ability.

We also desire a compact feature set that adequately spans the target receptive field vector space, T^n . The competitive process learns by finding cluster prototypes in T^n . To adequately span T^n , the competition process finds an excessive number of vectors, since there is no compactness constraint. If we limit the number of vectors, then the learned features may not adequately span T^n . To prevent redundancy in gray level feature discovery, we need the ability to break input signals down into their basic components.

III. THE GENERALIZED HEBBIAN ALGORITHM

Instead of competition, the Generalized Hebbian Algorithm (GHA) [11] uses deflation to force the output nodes to learn unique features. Deflation removes previously learned receptive field features from future receptive inputs. This procedure forces the network to learn new features and decomposes the receptive field vector spaces into basic components. The GHA has the following learning law:

$$\Delta \mathbf{w}_j = \beta (y_{jk} \mathbf{u}_{jk} - y_{jk}^2 \mathbf{w}_j) \quad (1)$$

$$\mathbf{u}_{jk} = (\mathbf{x}_k - \sum_{i < j} y_{ik} \mathbf{w}_i) \quad (2)$$

In these equations, k is the iteration number, $\Delta \mathbf{w}_j$ is change in weights for node j , \mathbf{x}_k is the input vector, \mathbf{w}_j is the weight vector for node j , $y_{ik} = \mathbf{w}_i \cdot \mathbf{x}_k$ is the output activation for node i , \mathbf{u}_{jk} is the deflated input for node j , and β is the learning rate. To form the deflated input, equation 2 removes previously learned features, $W_p = \{\mathbf{w}_i, i < j\}$, from the current input. The Oja term, $y_{jk}^2 \mathbf{w}_j$, forces the weights to have a magnitude of one [10]. The Hebbian term, $y_{jk} \mathbf{u}_{jk}$, increases the weight if the output, y_{jk} , and the deflated input, \mathbf{u}_{jk} are correlated [5]. The Hebbian term learns the major component present in the deflated input. The previously learned feature vectors, W_p , span a subspace V^n . Removing W_p from the input projects that input onto a subspace orthogonal to V^n . The Hebbian learning law pro-

duces a new weight vector consisting of a linear combination of the deflated input vectors. Thus, each new weight vector is also orthogonal to all previously learned vectors, and the entire set of learned weight vectors are mutually orthogonal. In addition to breaking receptive fields down into their components, the generalized Hebbian algorithm will converge to a unique solution independent of the initial weights [11].

IV. LEARNING FEATURES TO DISCRIMINATE TARGET OBJECTS FROM CLUTTER

To learn receptive fields that characterize target objects, we present the feature discovery network with image chips, which contain a gray level target image on background clutter. The background clutter might contain grass, bushes, trees, or rocks. When processing receptive fields from these image chips, a network can learn not only target features, but also some clutter features.

To learn features with better target discrimination capabilities, we have developed a novel GHA training method. First, we train the GHA with image chips containing only clutter. Here, the GHA will learn N_c features that decompose the clutter receptive field space into basic components. Then, we temporarily let the clutter features be the first N_c target features, and start learning with the $N_c + 1$ target feature. When learning the $N_c + 1$ feature, the GHA uses deflation to remove the previous N_c clutter features from the inputs that train the $N_c + 1$ target feature. When presented with a target chip receptive field containing mostly clutter, the deflation procedure removes most of the information. When the receptive field contains target pixels, the deflation procedure leaves the information that distinguishes the target receptive field from the clutter. The GHA then learns from these deflated inputs. After learning, we remove the clutter-trained features and use only the target-trained features for feature extraction.

V. RECEPTIVE FIELD ARCHITECTURE

Figure 1 shows a block diagram of the GHA feature discovery architecture. The receptive field extraction module takes an $N \times N$ image and extracts $(N - M + 1)^2 M \times M$ receptive fields at every input image location. A 2-D Gaussian attenuates the extracted receptive fields. The Gaussian peaks in the receptive field's middle. This attenuation makes the network more sensitive to the receptive field's middle, which discourages the network from learning shifted feature versions. Shifted versions do not add any new information, since the feature extraction process extracts image features at every position.

The GHA then processes the attenuated receptive fields.

Figure 2 shows the feature extraction architecture. The network correlates the learned weight vectors with the input image, scales the correlated outputs, and applies a nonlinear full wave rectifier to produce the S-layer outputs. The outputs require scaling because the output amplitudes vary greatly, a result of the frequency dependence described in Section II. To scale the outputs, the network computes each output's average and standard deviation over the training set. During feature extraction, the network scales the output by removing the average and dividing by the standard deviation. The full wave rectifier allows the local variance estimation by the second layer.

The C-Layer has similarities to the Neocognitron's C-Layer. For each spatial position, it computes a root mean square (RMS) average across all the planes. The C-Layer then compares each S-Layer output to its spatially corresponding RMS average. Only those values greater than the average survive the competition. The C-Layer then blurs and downsamples the result of this competition. The Gaussian blurring function provides a local variance estimate of the feature correlation output, and a reduced resolution output so that downsampling will not discard information. The downsampling produces a compressed representation of the input image.

VI. THE SAR AUTOMATIC TARGET RECOGNITION PROBLEM

We have applied our feature discovery work to the problem of automatic target recognition in synthetic aperture radar (SAR) imagery. Specifically, we have sought to recognize a vehicle, named V_1 , at arbitrary azimuthal angles while rejecting clutter and other vehicle returns. See Figure 3 for an example of simulated returns of two vehicles at various azimuthal angles.

Microwave radar has the advantage that it can penetrate cloud cover and darkness [2], whereas passive optical sensing cannot. Images formed using microwave radar look very different than optical images. For example, man-made surfaces often behave as microwave mirrors and produce large specular reflections. This and other phenomena make SAR images difficult to interpret.

VII. NEURAL ARCHITECTURE: FEATURE DISCOVERY AND ONE-CLASS CLASSIFICATION

Figure 4 gives a block diagram of the two stage self-organizing network architecture. We assume that

a separate region-of-interest preprocessor searches large images and finds candidate target image chips. The first stage extracts features from the image chip. We have chosen not to extract standard visual image features, such as edges and textures, because SAR images differ so significantly from visual images. Instead, we use the GHA and the architecture described in section V to discover features. For pattern classification, we use the ART 2-A neural network [1].

VIII. THE DATA FOR CLASSIFIER: TRAINING AND TESTING

Our training and test data comes from two sources. The SRIM SAR simulation tool¹ generated the simulated vehicle data. The SRIM data set contains seventy-two V_1 vehicles at 5° azimuthal increments and fifty V_2 vehicles at various depression and azimuthal angles. Figure 3 shows a few selected V_1 and V_2 vehicle image chips at the same orientations. The figure illustrates the difficulty of distinguishing the two vehicle types.

MIT Lincoln Laboratories collected the clutter data using the Advanced Detection Technology Sensor (ADTS). The clutter data set consists of one hundred image chips extracted at random from the ADTS images. The clutter data represents images of trees, grass, bushes, and roadways.

The target training set contains half of the seventy-two V_1 target class images. These thirty-six training images represent azimuth rotations at every 10° starting at a 0° azimuth. We test the classifier for the following two types of generalization: within-class and between-class [9]. For within-class generalization, the test set contains the thirty-six target class images not used for training. The fifty V_2 vehicle images comprise the between-class generalization test set.

IX. RESULTS

Using the training data described in the previous section, we trained two feature discovery networks to learn 24 11×11 feature detectors. A Gaussian with a $\sigma=4$, attenuated the training receptive fields, and both networks used a learning rate of $\beta=0.001$. The input image chips have dimensions of 46×46 or 2116 pixels. The S-Layer has 24 planes, each with 36×36 pixels. After applying a Gaussian blur ($\sigma=1$) and a 4:1 downsample, the output C-Layer contained 24 8×8 planes for a total of 1536 pixels. This represents a 25% compression. In the image chips, the target and non-target vehicles are centered within ± 1 pixel. When testing the classifier,

¹Developed by the Environmental Research Institute of Michigan in Ann Arbor, MI.

shifts of ± 1 . When a false alarm occurred on one of the nine shifted non-target versions, we counted it as a false alarm for all nine versions. When a correct detection occurred on one of the nine shifted target versions, we counted it as a correct detection for all nine shifted versions.

A. Independently Trained Features

We trained the first feature discovery network with receptive fields from target chips. Figure 5-A shows the features discovered by this network. The features represent a set of orthonormal vectors that describe the target's receptive field space. Figure 6-A shows the performance results for the ART 2-A classifier. The figure shows how a larger vigilance or smaller category size reduces the false alarms, but decreases the detection probability. The vigilance 0.778 produces the smallest false alarms and largest detection probability. At this vigilance, the network correctly classifies 100% of the V_1 targets and 85% of the V_2 non-targets. Without feature extraction, the ART 2-A would correctly classify 86% of the V_1 vehicles and 91% of the V_2 vehicles. We have also trained a two layer Neocognitron using the same training and testing sets. For the Neocognitron discovered features, the ART 2-A would correctly classify 89% of the V_1 vehicles and 92% of the V_2 vehicles.

B. Target and Clutter Features

We train the second feature discovery network with the procedure described in Section IV. We train the first eight receptive fields with only clutter receptive fields, and the remaining twenty-four with only targets. The latter twenty-four features discriminate between target receptive fields and clutter receptive fields. After training, we remove the first eight clutter trained features and use the remaining features for feature extraction. Figure 5-B shows the features discovered by this second network. Figure 6-B shows the ART 2-A classifier performance results. Notice the range of vigilances where the classifier has 100% performance on both vehicles. These new features greatly enhance the classifier performance.

X. CONCLUSION

We have discussed how the competitive neural networks cannot learn features that break receptive fields down into their basic components. This leads to redundant features not adequately covering the receptive field vector space. The Generalized Hebbian algorithm learns orthonormal vectors that break a receptive field space into basic components. We have developed a novel application of the Generalized Hebbian algorithm to learn features that discriminate target receptive fields from clutter

receptive fields. These new features greatly enhance the classifier performance in a SAR automatic target recognition problem.

REFERENCES

- [1] G. A. Carpenter, S. Grossberg, and D. B. Rosen. ART 2-A: an adaptive resonance algorithm for rapid category learning and recognition. *Neural Networks*, 4(4):493-504, 1991.
- [2] D. C. Munson, Jr. and R. L. Visentin. A signal processing view of strip-mapping synthetic aperture radar. *IEEE Transactions of Acoustics, Speech, and Signal Processing*, 4(12):2131-2147, December 1989.
- [3] C. E. Daniell, D. H. Kemsley, W. P. Lincoln, W. A. Tackett, and G. A. Baraghimian. Artificial neural networks for automatic target recognition. *Optical Engineering*, 31(12):2521-2531, 1992.
- [4] K. Fukushima and S. Miyake. Neocognitron: A new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognition*, 15(16):455-469, 1982.
- [5] D. O. Hebb. *The Organization of Behavior*. Wiley, New York, 1949.
- [6] N. Intrator and L. N. Cooper. Objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks*, 5(1):3-17, 1992.
- [7] R. Linsker. Self-organization in a perceptual network. *IEEE Computer*, 21:105-117, March 1988.
- [8] J. I. Minnix, E. S. McVey, and R. M. Inigo. A multilayer self-organizing artificial neural network for invariant pattern recognition. *IEEE Transactions on Knowledge and Data Engineering*, 4(2):162-172, April 1992.
- [9] M. M. Moya, M. W. Koch, and L. D. Hostetler. One-class classifier networks for target recognition applications. In *Proc. World Congress on Neural Networks*, pages 797-801, Portland, OR, July 1993. International Neural Network Society.
- [10] E. Oja. A simplified neuron model as a principal component analyzer. *J. Math. Biology*, 15:267-273, 1982.
- [11] T. D. Sanger. Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks*, 2(6):459-472, 1989.
- [12] C. Ting and K. Chuang. An adaptive algorithm for neocognitron to recognize analog images. *Neural Networks*, 6(2):285-299, 1993.

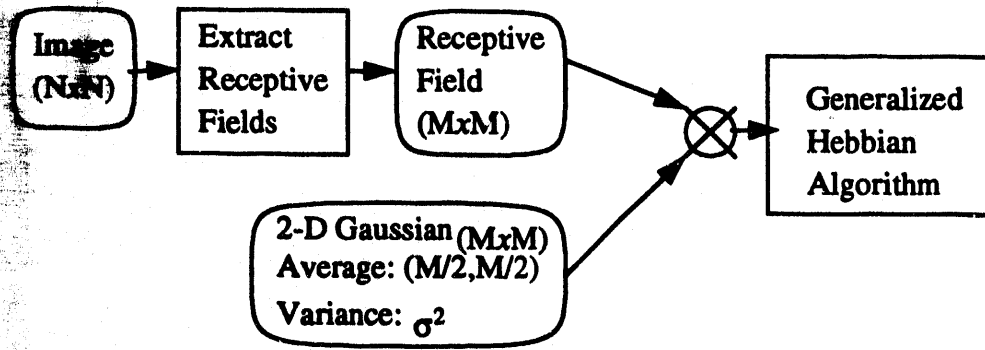


Figure 1. Block diagram of approach for learning receptive fields with the Generalized Hebbian Algorithm.

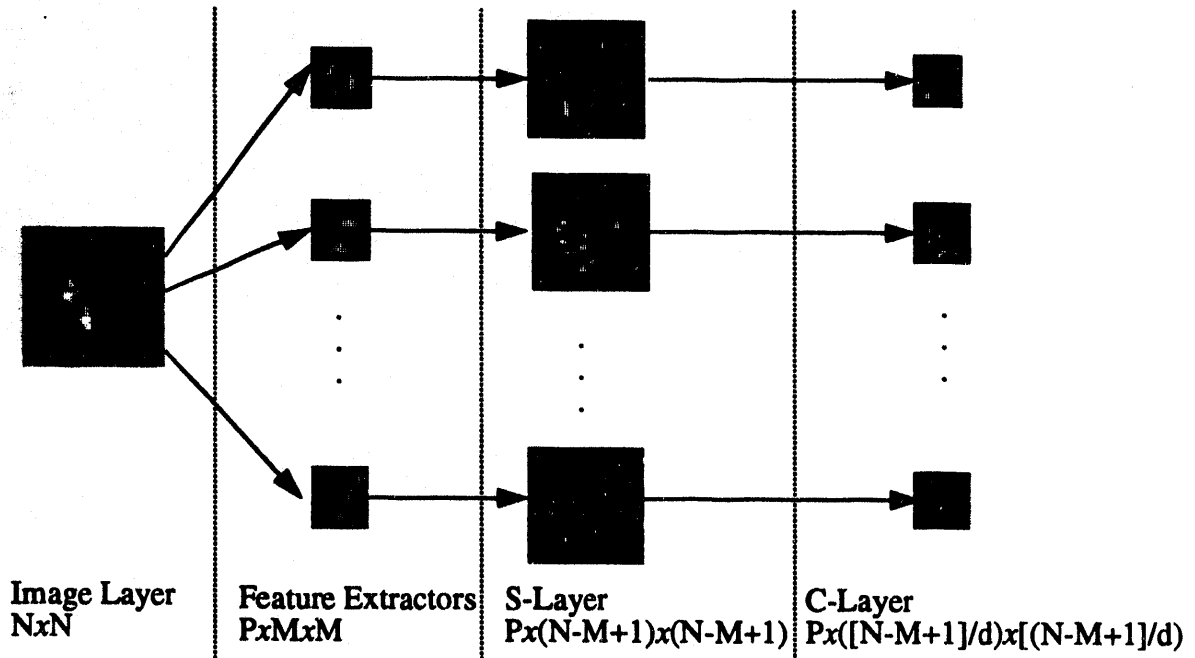


Figure 2. Feature extraction architecture.




Figure 3. Vehicle V_1 and Vehicle V_2 simulated returns





Figure 4. Overall SAR automatic target recognition architecture.




A) Features trained with targets only.

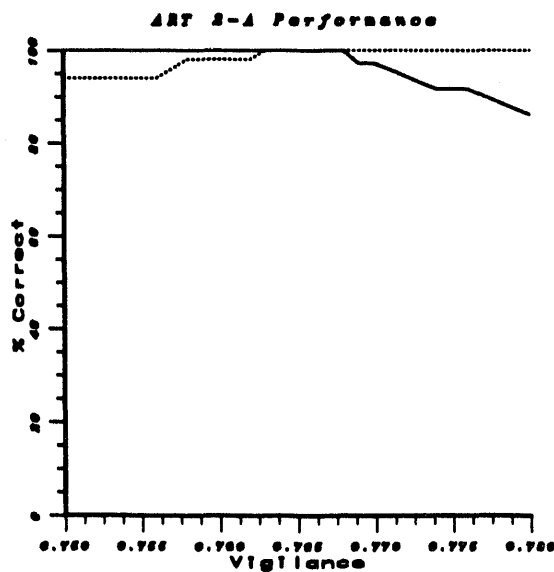


B) The features learned to discriminate targets from clutter.

Figure 5. Feature extractors learned by the Generalized Hebbian Algorithm.



A.



B.

Figure 6. Performance results. A) Performance results for features trained with targets only. B) Performance results for features trained to discriminate targets from clutter.

END

DATE

FILMED

3/24/94

