A Novel Process for Littoral Target Recognition: Preliminary Analysis Results[1]

Robert J. (Bob) Jannarone * John T. (Tyler) Tatum *
Brainlike, Inc.
April, 2005

Background

This report gives results based on a novel, automated process for identifying littoral targets in real time. The results include a preliminary analysis of a sonar image dataset that was kindly provided by the Naval Surface Warfare Center at Panama City, Florida, which in turn was funded by the Office of Naval Research. The dataset included several hundred images that contained zero to three littoral targets per image. The results came from a concurrent learning process developed by Brainlike, Inc.

The results in this report were produced to evaluate how the Brainlike process offers the following features and benefits.

- Fully automated deployment, allowing targets to be identified against noisy backgrounds without needing trained operator evaluation.
- Auto-adaptive learning, allowing targets to be identified accurately and robustly when background conditions vary substantially between or within images.
- Rapid, compact, and energy efficient operation, allowing deployment on remote sensor platforms such as sonar buoys and unmanned undersea vehicles (UUVs); remote deployment, in turn, reduces transmission costs and leaves time for other methods to be used for increased accuracy.
- Identification of simple features, allowing robust performance over varying target configurations, simpler understanding of target recognition results, and fast development cycles without the need for highly trained data analysis experts.
- Reasonably high target hit rates and low false alarm rates, offering complementary value to established target recognition methods, as well as stand-alone accuracy levels that would allow remote sensor platforms to identify multiple targets in real time.

The results in this report add to a large body of existing reports [1-8], showing that the Brainlike process adds significant value for detecting unforeseen anomalies under changing conditions. These new results show that the Brainlike process also adds value in cases where targets can be identified from training data prior to field deployment. These results also show that the Brainlike process adds complementary rather than

---

[1] The data and background information for this analysis were kindly provided by Dr. Michael Traweek from the Office of Naval Research, along with Dr. Gerald L. Dobeck and James T. Cobb of the Naval Surface Warfare Center, Panama City, Florida.

alternative or replacement value to established target. While the results from this preliminary analysis are surprisingly strong in terms of speed and simplicity, they are not as accurate as they will be when established feature extraction methods are used as well.

Key Results

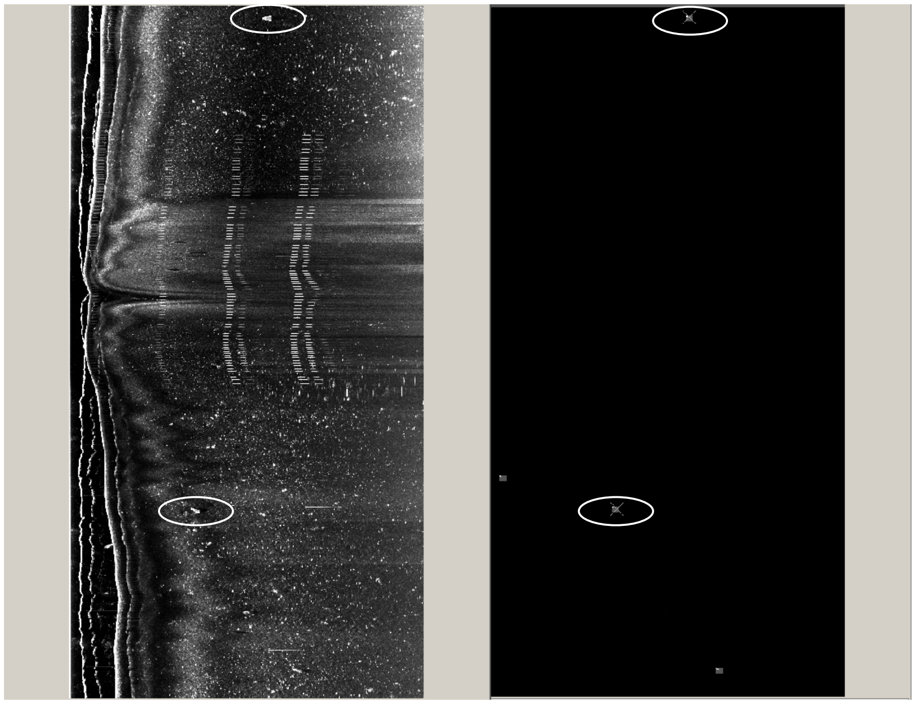

Figure 1. A Typical Image and Brainlike Result

Figure 1 shows a typical image from the dataset for this study on the left, along with a corresponding Brainlike target recognition result on the right. The image in the left frame was produced by a UUV over a sixty second period. The image includes 1,000 rows, with the bottom row corresponding to the least recent 60 millisecond time slice and the top row corresponding to the most recent 60 millisecond time slice. Each row represents reflections from an active sonar ping during its 60 millisecond time slice. Each such row includes 512 pixels, with short range reflections appearing as white pixels on the left and long range reflections appearing as white pixels on the right. Each pixel value is on a one byte gray scale, ranging from black = 0 to white = 255.

The two white ellipses in the left frame of Figure 1 each contain a target, and the frame contains only those two targets. For each target a white patch appears on the left indicating a reflected signal, and a darker patch appears on the right indicating a shadow. This target-shadow profile is typical of all other targets that were observed within the dataset.

The right frame in Figure 1 shows how the Brainlike process identified corresponding regions in the left image as targets. For clarity, each of the two actual targets in the left image is represented by a corresponding x in the right frame. The four gray patches within the right frame show that the Brainlike process identified both targets in the left image but also produced two false alarms.

Based on a training data subset, the Brainlike process established a target recognition objective function and identified 9 out of 12 targets within 20 training images, while producing a total of five false alarms. When exactly the same target objective function was used on a testing data subset, which had not been examined at all prior to testing, the Brainlike process produced a variety of results, based on different cutoff values. For one such cutoff value the Brainlike process identified 60 out of 77, while producing fewer than two false alarms per image.

Most notably, the Brainlike process operated as follows:

- It was automated, updating learned parameters and producing alarms automatically, *during* each time slice rather than at the end of each image.
- It was fast, computing feature values, evaluating targets, and updating learned parameters at a rate of two milliseconds per time slice on an office personal computer; if necessary, the process could have run a thousand times faster when used in conjunction with a field programmable gate array (FPGA).
- It was compact, requiring only a small portion of available memory and allowing all computations to be performed on a single FPGA chip if necessary.
- It was simple to use, allowing key features to be identified in less than 40 hours of analysis time and without the need for advanced target recognition training.
- It left ample room for accuracy improvement, by requiring only a fraction of the available computing power, using only crude feature extraction methods and producing feature values that would be compatible with alternative methods such as neural network and nonlinear discriminant analysis.

Analysis Details

The dataset included 514 images, obtained during operational testing of a UUV over a 15 month period. The dataset was separated into (a) a training subset, made up of the first 117 images that had been obtained during the period, (b) a preliminary testing subset, made up of the next 139 images obtained during the period, and (c) the remaining 258 images, which were reserved for future analysis.

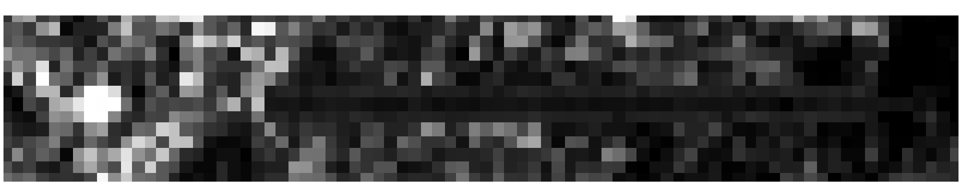

Figure 2. A Typical Type I Target

Targets in the training subset were examined for the existence of features that could be used efficiently in conjunction with the Brainlike process. Preliminary analysis identified two target types. The first type, shown in Figure 2, produced bright reflections that range from 3 to 5 pixels high and long shadows that are 3 pixels high, with shadows having lower within-row variation than background rows. The second type, shown in Figure 2, produced bright reflections on the left that range from 4 to 8 rows high shadows on the right that range from 5 to 8 rows high, with shadows being darker than background rows.

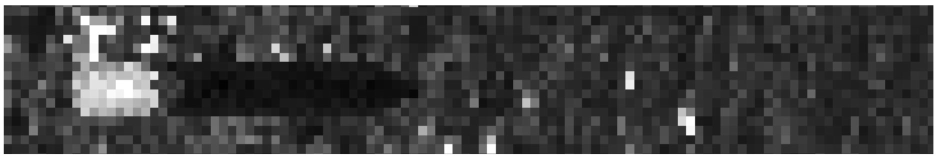

Figure 3. A Typical Type II Target

After identifying the above two target types, two features for each type were constructed. A value for each such feature was computed within each of several windows for any given image and used to determine whether or not the window contained a target. Each such window was 12 rows high and 25 columns wide, so that any given image contained 481,156 windows (481,156 = [1,000–12]×[512–25]).

In keeping with the Brainlike process, each of the four feature values for each window was compared to four corresponding expected values for each window. These four expected values, in turn, were based on a continuous learning process, which updated learned parameters automatically on a window by window basis. The result of this comparison produced four deviance values for each window, with two such values corresponding to Type I and the other two corresponding to Type II. Each such deviance value was automatically standardized by the Brainlike process to have an expected value of 0 and an expected standard error of 1.

With feature values obtained and standardized in this way, establishing a reasonably accurate and robust target classification rule was relatively simple. On a standardized scale, target feature values were found to be consistently different from non-target feature values, with difference criteria being about the same from one target to another and from one image to another. Based on these differences, a simple linear combination of the two features for Type 1 was identified and a target cutoff value based on that combination was set, and likewise for Type 2. To set a global objective function, a simple sum of the two type values was computed and a cutoff criterion was obtained in a similar way. When training dataset results were obtained, targets were usually found to produce three or more alarms in a window, and a classification rule was created accordingly, identifying a window as containing a target if three or more pixels in the window generated an alarm. The only refinement to this simple criterion was an added set of simple rules, which limited the objective function value if target and shadow values did not occur within the same window.

After feature functions were created and combined in this way, correct classification results were obtained and examined using the first 20 training set images, which contained 12 targets. For the pre-determined cutoff value, the Brainlike process identified 9 out of the 12 targets, along with five false alarms. For lower cutoff values, the process identified more targets and false alarms, and for larger cutoff values, the process identified fewer targets and false alarms, as is the usual case with target classification analysis.

Correct classification results were then obtained for the testing data subset. In order to reflect actual field conditions where no tuning of the objective function would be possible after deployment, exactly the same objective function that had been established for the training dataset was used for the test dataset, which had not been examined at all prior to testing.

Cutoff Criterion Factor	**False Alarms**	**Hits**	**False Alarms per Image**	**Proportion of targets hit**	**Hit to False Alarm Ratio**
1.6	1	7	0.01	0.09	7.00
1.5	1	7	0.01	0.09	7.00
1.4	1	11	0.01	0.14	11.00
1.3	2	17	0.01	0.22	8.50
1.2	3	29	0.02	0.38	9.67
1.1	4	37	0.03	0.48	9.25
1.0	6	44	0.04	0.57	7.33
0.9	14	52	0.10	0.68	3.71
0.8	68	56	0.50	0.73	0.82
0.7	259	60	1.90	0.78	0.23
0.6	882	65	6.49	0.84	0.07
0.5	4,706	72	34.60	0.94	0.02

Figure 4. Test Dataset Alarm Statistics

The Brainlike process produced a variety of results, based on different cutoff values, as shown in Figure 4. The first column gives a cutoff value factor that was used to produce the other results in the table. The row with a cutoff value of 1.0 used the same cutoff criterion that was predetermined in the training dataset; the row with a cutoff value of 0.8 used 0.8 times that cutoff criterion, and so on. As the table shows, for a cutoff value of 0.7, the Brainlike process identified 60 out of the 77 total targets in the test dataset, while producing an average of 1.90 false alarms per image. Other rows have a similar interpretation.

Figure 5 shows the test dataset results in the form of a hits versus false alarm curve (also called a Receiver Operating Characteristic or ROC curve). Only the gray cells from Figure 4 are plotted in Figure 5 for clarity. As the figure shows, the Brainlike process produced very few false alarms per image until target hit rates approached 0.8, at which point the false alarm rates increased substantially. Interestingly, for these and higher hit rates studied, most of the false alarms occurred within the left-most 100 image columns, where no actual targets were found but where substantial background noise existed along the lines shown in the left side of Figure 1.

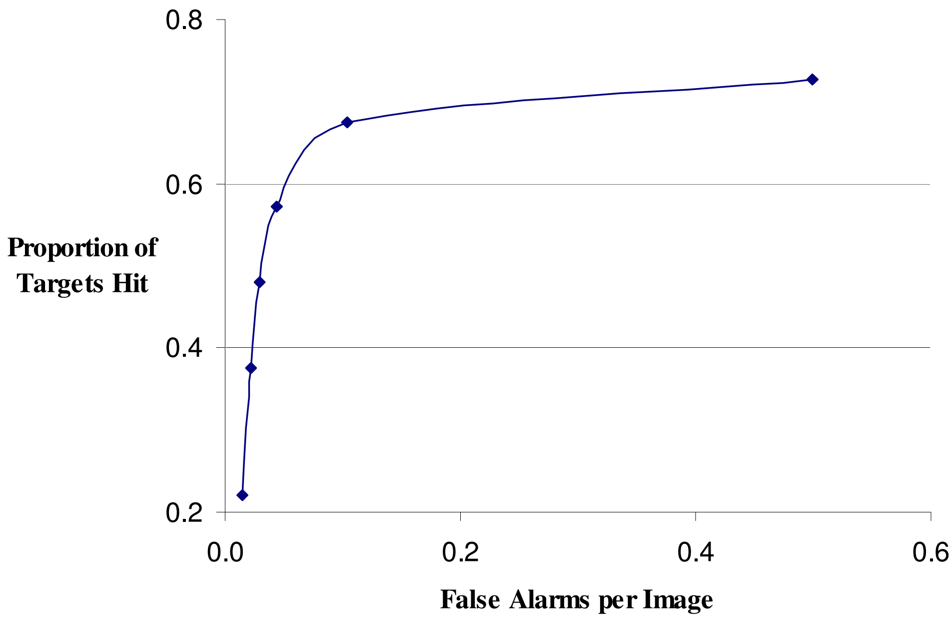

Figure 5. Test Dataset Hits *versus* False Alarms Plot

Regarding computing efficiency, Brainlike algorithms computed feature values, evaluated targets, and updated learned parameters at a combined rate of less than two milliseconds per time slice, running on a personal computer with a 927 megahertz processor. As with other Brainlike processes, all results were obtained using core algorithms that can be broken down into parallel processes and implemented on an FPGA. Implemented in that way, results could have been obtained about a thousand times faster, that is, at processing rates of about one two microsecond per row.

Conclusions

To show how Brainlike processes could be used with limited domain knowledge, the results in this report were obtained with little knowledge of how the dataset was obtained, how decisions based on target recognition results might be made, and what computing resources might be available. As a result, just how Brainlike processes will add the most littoral image processing value awaits further study. However, the following conclusions seem clear.

As with other monitoring, surveillance, and target recognition applications [3-7], optimal alarm cutoff values will depend on operational factors such as costs of false alarms and missed targets, as well as on how decisions are to be made once alarms have been produced [8]. For example, if optimal decisions required distinguishing targets from false alarms with very high precision, operating with low cutoff values as shown in the bottom rows in Figure 4 might be best, and accuracy improvements from complementary target recognition methods might be required. If instead, decisions during routine UUV fly-bys involved identifying regions like minefields with many targets present, operating with cutoff values near the top rows in Figure 4 might be better. Reason: as shown in the upper row entries in the right-most column in Figure 4, alarm rates could be expected to jump by a factor of 7 or higher if UUVs were to enter such regions, and such jumps would be easily detectable by a Brainlike change detector.

If more accuracy is needed, the results support using the Brainlike process as a useful complement to established target recognition procedures. Reason: the Brainlike process costs little in terms of computing time and resources, leaving extra resources available for complementary processing.

Since the results in this report were reasonably accurate and they were generated very quickly, they may also point toward real-time Brainlike process deployments at remote sensor arrays. When deployed in this way, the Brainlike process could serve as a real-time controller, marking entries into minefield regions, triggering anomaly alerts, and starting more accurate target recognition processes, while otherwise minimizing costly transmission and energy use.

References

1. R.J. Jannarone, *Concurrent Learning and Information Processing: A Neurocomputing System that Learns during Monitoring, Forecasting, and Control.* Chapman & Hall, New York, 1997.
2. R.J. Jannarone, Concurrent Learning and Performance Information Processing System, U.S. patent [#]5,835,902, 1998.
3. R.L. Smith and R.J. Jannarone, "An Auto-Adaptive Statistical Procedure for Tracking Structural Health Monitoring Data," *Proc. SPIE*, **5391**, 166-176, 2004.
4. **www.Brainlike.com**: "Brainlike Shallow Water Intrusion Detection."
5. **www.Brainlike.com**: "Brainlike Monitoring Improvement Illustration."
6. **www.Brainlike.com**: "Brainlike Homeland Security Relevance."
7. **www.Brainlike.com**: "Brainlike Anti-Bio-terrorism Relevance."
8. **www.Brainlike.com**: "Brainlike Return on Investment Analysis."

*

Robert J. (Bob) Jannarone is President of Brainlike Surveillance, Inc., a San Diego company that delivers and evaluates monitoring value for defense and homeland security. Before Brainlike, Bob founded, capitalized, and ran Netuitive, Inc., a company that delivers computer server monitoring software. Before that he was a tenured (engineering, statistics, and psychology) professor and an active (ONR- and DOE-funded) researcher. Bob also served in the U.S. Navy as a submarine nuclear engineer. He has authored many patents, articles, and a book that describe concurrent learning and

information processing. Bob has degrees in computer science, statistics, and psychometrics, including a Ph.D. from the University of California at Berkeley.

John T. (Tyler) Tatum is Executive Vice President of Brainlike Surveillance, Inc. His career began as a programmer with Netuitive, Inc. (then known as Rapid Clip Neural Systems, Inc.), in 1997. After leaving Netuitive in 2001, Tyler ran Project Planet, a leading hotel environmental services company. At Brainlike, Tyler is focusing on delivering valuable monitoring solutions to satisfied customers very quickly, and ensuring Brainlike's commercial success in the defense and homeland security sector. Tyler holds a degree in Computer Science from Cornell University.