THE APPLICATION AND DEMONSTRATION OF CLASSIFICATION TECHNIQUES FOR TORPEDOES

S C Thorpe BEng (Hons) CEng MIET,
QinetiQ,
Underwater Systems,
Winfrith Newburgh,
Dorset, UK.

1 INTRODUCTION

1.1 Overview of Torpedo Engagements

In a typical torpedo engagement, as shown in *Figure 1*, the launch platform will detect a threat target, confirm it and achieve lock-on. The intelligence regarding the target and environment, for example target location and type, are programmed into the weapon's fire control system, and the weapon launched. In its simplest form, this is an autonomous fire-and-forget process where the weapon transits at low speed, to maintain covertness, to the target uncertainty area (TUA) where it begins searching for the target using its sonar. Once the weapon locates and locks onto the target it uses its self-guidance system to home to either the point of target impact, or close proximity where the warhead is detonated.

Often, torpedoes are launched kilometres away from the target and typical engagements, unlike those using missiles, can last a significant amount of time (fifteen to thirty minutes) at a maximum speed of approximately 50km/h[1].

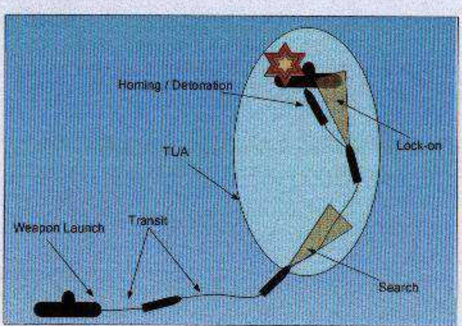

Figure 1 – Typical Torpedo Engagement

1.2 Change in Strategic Focus

With the end of the Cold War and the demise of the Soviet Union there has been a large export market for relatively inexpensive, yet highly capable, diesel electric submarines: Germany is advertising its highly advanced type 212 and 214 submarines, while Russia is still manufacturing for export to countries including the People's Republic of China[2]. As a result there has been a significant shift in strategic focus of naval operations from global confrontations to more complex regional conflicts in shallow water and littoral zones.

1.3 New Challenges

Using active sonar in these littoral environments against capable threat targets can present new challenges for torpedo operations because of:

- A significant number of returns from a diverse range of persistent false targets (rock outcrops, ship wrecks, etc.) that may appear target-like;
- High levels of ambient noise due to local shipping traffic and industry;
- Complex and poor sound propagation including refraction in the water column, caused by thermoclines, and multiple interactions with the sea surface and seabed.

These challenges can affect torpedo operations in a number of ways:

- More weapons wasted by engaging false targets, increasing stowage requirements together with procurement and maintenance costs;
- Longer engagements increasing the likelihood of the threat target manoeuvring and counterattacking, increasing the probability of sustaining battle-damage, casualties and fatalities.

1.4 Potential Solution

In order to reduce the impact of active sonar returns from false targets on torpedo operations, single-ping classification techniques have been applied and developed. These techniques have the ability to discriminate and reject returns attributed to false targets, whilst still retaining a significant proportion of returns from the threat target to allow effective prosecution.

This research and development has involved an approach which has:

- Proved that classification techniques can be used successfully on real data acquired through sea trials using representative sonar;
- Quantified the likely performance improvements using torpedo engagement modelling;
- Demonstrated the techniques in water using a closed-loop torpedo-like system architecture.

2 TORPEDO PROCESSING SYSTEM

A torpedo can be thought of as an intelligent, autonomous unmanned underwater vehicle (UUV) that after launch operates as a closed-loop system within predefined operating characteristics. These characteristics encompass a range of pulse transmissions (e.g. broadband and narrowband) and signal processing parameters.

A typical processing chain for a broadband transmission is shown in *Figure 2.* Sound received by the elements of the sonar array is digitised, beam-formed, and correlated using a replica of the transmitted wave form. A threshold is then applied to identify detections which are discrete crossings of an estimated background noise threshold, and these detections are grouped by range and bearing to form clusters. The tracker then associates clusters over time to form tracks, selects the most consistent track, and the torpedo uses its tactics and guidance to head towards the track.

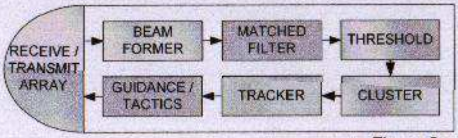

Figure 2 – Broadband Sonar Processing

3 CLASSIFIER SYSTEM DESIGN

3.1 Constraints

When designing a classification system for a torpedo there are a number of important constraints that need to be understood and addressed to ensure that the final system is accepted by torpedo upgrade programs. These constraints include:

- **Processing Load** – Due to the constraints of legacy torpedo design there are limited processing resources. It is therefore vital that the feature extraction and classification systems must be able to operate in real-time;
- **System Integration** – The classification system must be able to be integrated as seamlessly as possible within existing torpedo processing architectures. This is vital as it reduces risk and therefore the cost of technology insertion. Ideally it should be plug-and-play and fit between existing processing subsystems;
- **Feature Extraction** – In order to facilitate discrimination it is essential that a diverse range of features can be extracted;
- **Robustness** – Naval operations using torpedoes can encompass a diverse range of scenarios including different acoustic environments and targets. It is therefore essential that the classification system can function robustly in such a context;

3.2 System Architecture

When using the aforementioned design constraints it becomes apparent that certain tradeoffs must be made. For example, to reduce processing load the classification system must be located along the processing chain where information becomes reduced / summarised making it less processor intensive to extract features from. Although there is a dilution of information that could facilitate discrimination this can be offset by the grouping of information during the latter processing stages. Furthermore, the detector threshold can be adjusted to control the amount of information reduction.

To facilitate system integration, the classification system has been developed to reside between the clustering and tracking processing stages, *Figure 3.* The features are extracted from the clusters and relate to their signal properties and geometry. The clustering algorithm is essentially identical for different transmission types except for the number of grouping dimensions (i.e. an additional dimension of frequency for narrowband). The features extracted are then passed to the classifier, which has been trained offline, which decides whether or not to reject the cluster; rejected clusters are deemed to be attributed to either false targets or false alarms.

A range of classifier architectures (Fisher Discriminant[3], Non-Linear Disriminant, MLP[4], etc.) has been evaluated using trials data and have all been demonstrated to give similar classification performance characteristics; as a result the current system uses a Fisher Classifier.

It has been recognised that classification features extracted from subsurface target echoes are different to those from surface targets and therefore, to ensure robustness, separate feature subsets and classifier architectures can be used for each, and this can be determined by the launch platform and programmed into the weapon prior to launch.

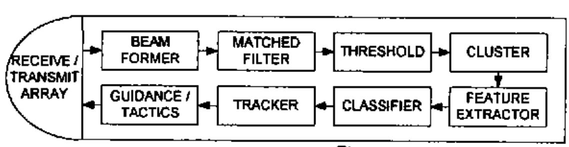

Figure 3 – Location of Classification System

3.3 Classifier Training

The classifier design process uses real active sonar data acquired on sea trials using a representative torpedo sonar array. With knowledge of where the target is, a ground truth is generated to extract the clusters that can be attributed to the target and the remaining clusters are deemed to be from either statistical false alarms or false targets.

This segregated data is then split into two datasets: one for training, and one for testing. The training dataset is presented as input to a classifier design tool to determine the classifier's discriminant function, or rule. This function is then applied to the test dataset and results in a classifier operating characteristic curve, from which an operating point can be chosen for the system. The operating point is chosen so that a significant number of false target clusters can be rejected whilst still retaining enough target clusters in order for the tracking system to form good quality tracks that facilitate target homing.

3.3.1 Classifier Training Challenges

It has been recognised that the process of designing a classification system for torpedo operations in a wide range of scenarios (i.e. different targets and acoustic environments) presents significant challenges; the main challenge being robustness. These challenges can be compounded by changes in transmission and signal processing characteristics during a typical torpedo engagement. For example, as the weapon gets closer to the target its transmission pulse length is reduced.

It is therefore important to understand how classification features vary with environment, target, transmission and processing characteristics. It is equally important that there is enough measurement samples, especially in terms of variety, used in classifier design so that under-sampling effects are minimised[5]. Under-sampling can lead to less accurate / realistic approximations of classifier feature distributions which can ultimately reduce the robustness of the classifier's discriminability (i.e. how the classifier functions when presented with new / unseen samples).

4 SYSTEM EVALUATION THROUGH SEA TRIALS

4.1 Overview

Initial research into the applicability of classification techniques to torpedoes demonstrated the capability of the system in a simulation-based research environment, but, in order to advance understanding and confidence in the techniques, a series of in-water experiments were designed and executed.

These experiments exploited a torpedo-like demonstrator named DasHH-Marlin, *Figure 4*. DasHH (Digital advanced signal processing Homing Head) is a modified torpedo homing head that has full digitisation of all receive array elements together with COTS computers that can perform the necessary signal processing, tactics and guidance. DasHH is fitted to the front of the Marlin autonomous underwater vehicle. Marlin is a similar size to a heavyweight torpedo, but slower and less agile, and has a modular design enabling it to host a variety of payloads, such as DasHH.

DasHH uses an open architecture approach to facilitate the rapid technology insertion and in water evaluation of new techniques, and involves using the same software tools throughout the technology insertion lifecycle (i.e. research->tactical development->simulation & concept studies-> technology demonstrators->beyond); Nereus[6] is used for the full signal processing chain and ODIN[7] is used for guidance and tactics. This reduces risk, time and ultimately technology insertion costs.

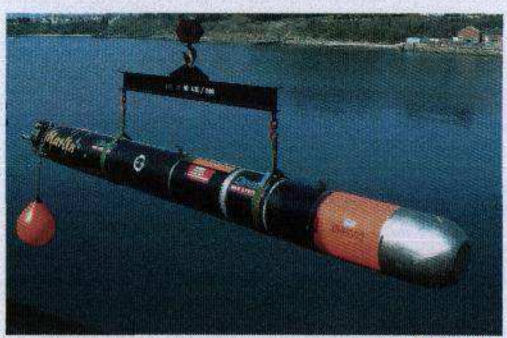

Figure 4 – DasHH Marlin UUV

4.2 Data Acquisition and Classifier Design

DasHH-Marlin was used to acquire element level sonar data against a synthetic target to allow a classifier to be trained. It was not logistically feasible to use a real target (i.e. ship or submarine) and therefore a stationary echo repeater was used. This was configured to retransmit DasHH's active transmissions with a delay and adjustment in echo-level based on a pre-programmed target echo strength (TES).

Broadband data was acquired using a range of target echo strengths in order to achieve a better understanding of the effects of signal to noise on classifier performance. It also ensured that the classifier was not tuned to a particular TES. The classifier operating point, i.e. Fisher Threshold, chosen had a PCC of approximately 0.5 and a PFA of 0.1.

4.3 Classifier Evaluation Experiment

In order to reduce risk, the experiment was developed through the use of off-line simulations using Nereus and ODIN. During these simulations the signal processing software, together with guidance and tactics, were developed from existing torpedo representations to ensure that the experiment mirrored a typical torpedo operation. Once satisfied with the simulations, this software was copied into DasHH's computers.

DasHH-Marlin was programmed with a fixed elliptical TUA and launched inside this TUA, towards its centre in an active search mode using broadband transmissions, since the target is stationary. The target was located in the TUA, but outside of DasHH's sensor coverage to ensure that the TUA was searched and to increase the likelihood of DasHH being seduced by false targets, thus making the experiment more challenging. The acoustic environment where the experiments were performed was challenging: it was extremely shallow (10-12m), and enclosed by harbour walls and breakwaters.

The experiment would be judged to have been successful if DasHH managed to locate the target in its search phase, confirm the target and then home to target proximity.

4.4 Experiment Results

Summary
DasHH Marlin proved an efficient and cost effective platform for the test and evaluation of classification systems for torpedoes. The integration of research software into the vehicle using ODIN and Nereus was seamless and problem free.

A total of five approaches on the target were performed and the majority were successful; a successful approach is shown in *Figure 5*. The dots (small and red) represent clusters that have been rejected by the classifier as being deemed to be from either statistical false alarms, or false targets. The squares (large and green) are clusters that have been accepted by the classification system and deemed to originate from the target.

From the figure it can clearly be seen that performance of the classifier is very good: DasHH detects and correctly classifies the target in its active search phase, confirms the target and homes. It is not seduced by false clusters at all because only a few clusters are accepted by the classifier and these are not consistent enough over a number of pings for a good quality track to be formed.

Classifier Performance

Averaged over all approaches, the classifier is estimated to have a performance of: PCC=0.66 and PFA=0.003 and this compares favourably with the original design performance (PCC=0.5 and PFA=0.1). Performance, in terms of the confidence of the classification decision, as calculated from the confusion matrix in *Table 1*, is also impressive at over 95% (i.e. there is a >95% probability of any cluster being classified correctly).

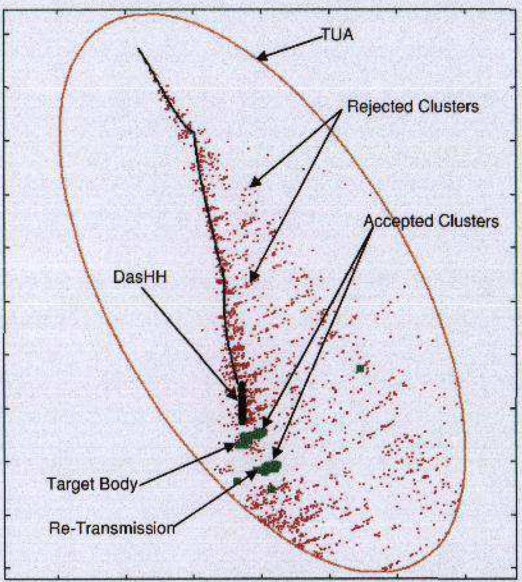

Figure 5 – Experiment 1: DasHH homes on target

		True	False
Target Clusters	True	946	494
False Clusters	False	32	12316

Table 1 – Confusion Matrix summarizing the six target approaches

Effect of SNR on Classifier Performance

To design a robust classification system It is important to understand how it performs at a variety of signal to noise ratios (SNR) that can be caused by a combination of target type and acoustic environment (ambient noise levels, propagation etc.). *Figure 6* shows a scatter plot of how the

Fisher Discriminant, the output of the classifier, varies with SNR for target clusters (upper plot) and false clusters (lower plot). The data for different SNR was achieved by using different pre-programmed TES for the target, and different ranges, but in the same acoustic environment. From this it can clearly be seen that the separation between target and false target classes increases with SNR and therefore classification performance improves as the SNR increases.

The poorer class separation at lower SNR can be attributed to a number of factors including detection performance. For example, at low SNR the detection performance decreases and therefore this reduces the amount of information from which to extract features from and this in turn effects feature estimation. Using a synthetic target may also reduce the amount of structure in the received signal and also impact feature extraction and estimation. It would be expected that a similar trend to that shown in *Figure 6* would be observed with real targets, but would possibly be less marked.

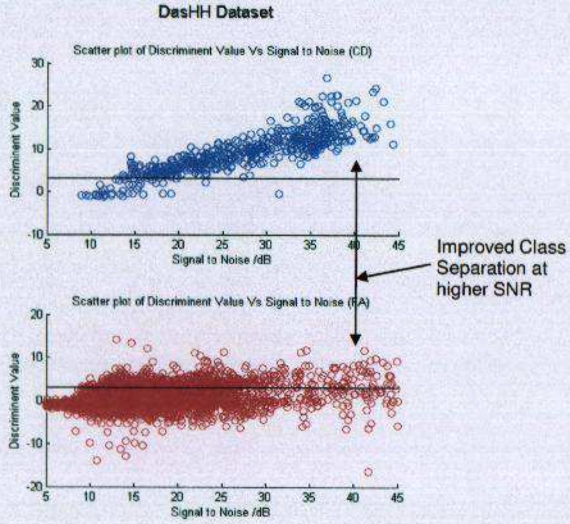

Figure 6 - Effect of SNR on discriminant value

5 CONCLUSIONS

A single-ping classification system has been successfully developed using a whole-system approach, and its capability evaluated for torpedo operations in littoral environments. The evaluation of the system has culminated in a number of successful closed loop homing experiments, which have been based on real torpedo system architectures, in a shallow water environment with a stationary synthetic target. In general, the experiments have been successful with DasHH Marlin homing on the target without being seduced by false targets. The in-water performance has agreed well with that obtained through the classifier design process and the confidence of the classification decisions is impressive and in excess of 95%.

Although the performance of the classification system is very good, it has been recognised that it is dependent on the SNR: performance is better at higher SNR. It is believed that such an artefact can largely be attributed to a lack of detections which reduce the amount of information at the feature extraction stage, and also the fact that the target signature is synthetic. Further analysis of data from threat targets is required to make a more complete assessment of the effect of SNR on the performance of the classification system.

6 FUTURE RESEARCH

With proof that the classification system can work well in a shallow water environment against a synthetic target, future work should address in more detail the operational robustness of the system in a number of independent (unseen) environments, and against representative threat targets. Such research and development will help to identify whether or not single-ping classification should be included in next generation underwater weapons.

At present further research is being conducted to develop single-ping classification into a mutli-ping classification system, where classification information is integrated along target tracks, and the discrimination decision is made once the tracks have existed for a certain number of pings. It is hoped that multi-ping integration will improve the quality of the classification decisions and also allow further classification features to be extracted.

7 REFERENCES

1. *"Spearfish – a Torpedo for the Future" – G Holt, Underwater Systems Group, Marconi Electronic Systems Ltd*
2. *"Formulation of an Integrated Robust Design and Tactics Optimization Process for Undersea Weapon Systems" – A P Frits, School of Aerospace Engineering, Georgia Institute of Technology, 2004;*
3. *"Statistical Pattern Recognition" – Andrew Webb,1999;*
4. *"Interpretation & knowledge discovery from the MLP network: opening the black-box" – Computing & Information System Management Group, Cranfield University;*
5. *"Nightmare at Test Time: Robust Learning by Feature Deletion" – Globerson, MIT, Cambridge, MA, USA;*
6. *"Nereus v2.2 User Guide" – S C Thorpe, QINETIQ/D&TS/SEA/SUM0601475, 2006;*
7. *"ODIN User Guide v4.0" – L F Dart, QINETIQ/S&E/OSA/TR011285, 2004*