Whitepaper:

Artificial Intelligence in the Underwater Battlespace

Commercial in Confidence

15 April 2026

Dr Scott Reed

1 Artificial Intelligence in the Underwater Battlespace

Uncrewed systems are now being deployed for a wide range of underwater battlespace missions including Mine Counter Measures (MCM) operations, Critical Underwater Infrastructure (CUI) protection and Seabed Warfare (SBW). With payload resolution, mission duration and problem complexity all increasing, there is a clear demand signal for AI algorithms that process data at machine speed.

How AI algorithms are effectively developed and applied within the underwater battlespace depends on several key factors that must be considered before designing and deploying this technology.

The first is the specificity of the problem domain that the AI is being asked to solve. For example, using AI to automatically search for specific mine threats within sonar data is a well-defined scope. The constrained problem, coupled with the limited variability in the data being assessed, allows a AI solution to be developed which can perform this task with a high degree of certainty.

The second factor is how data drives operational performance. The first challenge is the availability of data used to train the AI model can be a major issue for Navy applications where data can be scarce and difficult to share with AI vendors. The second challenge is the variability between the data used the train the AI versus the data that the AI will process in the operational environments. As this variability increases, performance of the AI will typically drop.

The third factor is the proximity of the AI to human involvement. Running AI “at the edge” on underwater platforms limits human-machine interaction due to communication challenges. Any AI running at the edge must perform well in isolation given there are limited options for human feedback. Scenarios where a human can collaboratively interact with the AI provides the user more intelligent decision making and data filtering options, reducing the requirement for AI to perform well as a standalone tool. In these scenarios, it is the combined human-machine team that must perform well, with the focus being on how the human interacts with the AI system.

This paper discusses three frontiers of AI within the underwater battlespace.

The first is the use of Automatic Target Recognition (ATR) for MCM where the specificity of the problem is high, the data has limited variability, and the potential for human involvement is limited.

The second looks at how machine learning technology could be applied generally across multiple and varied applications such as protection of CUI and SBW. The reduction in the specificity of the application, coupled with an increase in data variability places a larger demand on effective human-machine teaming to ensure mission success.

The paper concludes by bringing these concepts together, providing a future frontier with an Agentic-AI vision for Navies, where the specific AI approach leveraged will be dependent on the specific application and the three factors discussed above.

2 ATR from data to actionable information

Automatic Target Recognition (ATR) algorithms are used in Mine Counter Measure (MCM) operations to process images from Sidescan Sonar or Synthetic Aperture Sonar (SAS) to detect, classify and localise mine-threats.

Modern ATR algorithms build on Deep Learning technology where large quantities of example data are required to teach the algorithm how to recognise a mine threat.

The success of an ATR algorithm is heavily dependent on the data used to train the model, with three key considerations:

- An ATR algorithm will only reliably detect and classify mine shapes it has already seen within the training data provided. This means an ATR's main utility is the detection and classification of *known* mine threats, for which example data already exists.
- The performance of the ATR will be heavily reliant on the quantity, quality and variability of the training data provided. Deep Learning algorithms typically require 1000s of examples of each object they are required to detect. This requires Navies to prioritise the collection and curation of MCM dataset using training mines.
- Data sets containing mine threats are often classified. This can make it challenging to provide the data to commercial vendors such as SeeByte to allow the ATR algorithm to be trained.

Most notable ATR success stories have come from long-term MCM Programs of Record where the nation involved is committed to conducting regular data collection exercises and providing on-going financial resources to vendors to allow ATR models to be routinely updated and improved.

2.1 Using ATR

The long-term objective for MCM is to use uncrewed systems to autonomously search and clear an area of mine-threats, with minimal human intervention. A key enabler for this vision is to have high performance ATR algorithms and mission level autonomy systems able to run onboard the uncrewed system, detecting and classifying mine-threats at machine speed.

Figure 1 compares the profile for a UUV-based MCM mission with, and without, ATR processing the data on the vehicle. An onboard ATR algorithm removes the need for a human to manually analyse the data, which also removes the need to recover and re-launch the system.

Data analysis can account for 50% of the total MCM mission time, meaning onboard ATR has the potential to significantly reduce mission time. Moving to multi-vehicle operations vehicles (using SeeByte's Neptune Autonomy product) offers a further reduction in mission time, breaking the sequential process detailed in

Figure 1**,** and allowing the different phases of the MCM mission to be run in parallel.

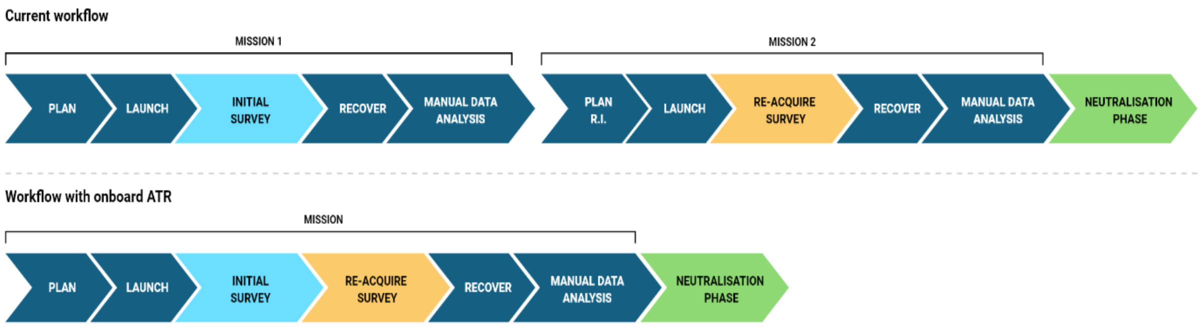

Figure 1: The typical workflow for a UUV-based MCM mission using a single vehicle. Without on-board ATR, the vehicle is recovered after the initial survey so that a human can inspect the data for potential mine threats. The UUV is then re-planned and re-launched to inspect the potential mine threats. With on-board ATR, analysis of the data occurs in real-time by the ATR, removing the need for the human to recover the vehicle and manually process the data.

This potential mission time reduction can only occur if the ATR is "highly performant". A high Probability of Detection (PD) is required to ensure no real mine-threats are missed; a low Probability of False Alarm (PFA) is needed to ensure the ATR does not add time to the mission duration and undoing the benefits of having the ATR in the first place.

The lack of communication bandwidth underwater means a human cannot easily use their expertise while the UUV is operating underwater. This has driven ATR model training to focus on successfully detecting a small set of mine threat types to keep data collection requirements manageable. Current state-of-the-art ATR algorithms can detect and classify four different mine threats in simple environments, producing a false alarm rate similar to a highly trained human team. This performance drops when the seafloor becomes more complicated, and data from these regions must be assessed by a human.

ATR is also used as a Post Mission Analysis (PMA) tool to help the human search through sonar data. An example of this is shown in **Figure 2** where results from the ATR are superimposed on top of the sensor imagery being analysed by the human.

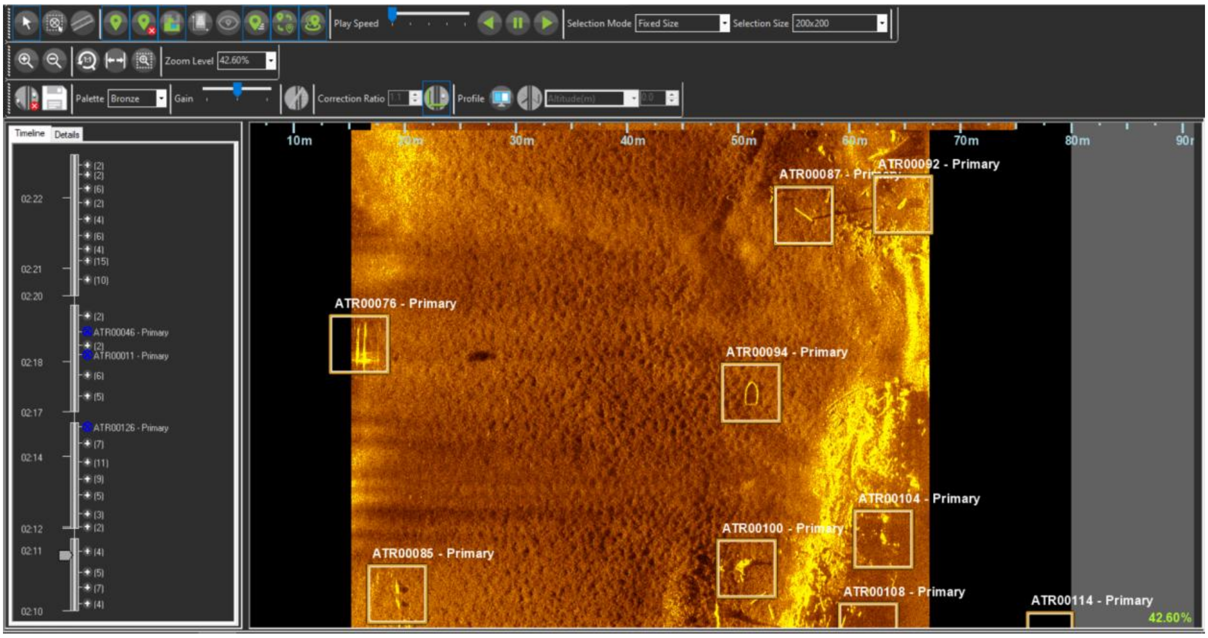

Figure 2: Humans search through sonar data collected by the UUV using a waterfall display. Potential mine threats found by the ATR are superimposed on the sensor imagery for consideration by the human. Image: Sonar data courtesy of Kraken Robotics.

It is more difficult for the ATR to reduce mission time when used in this manner as the human is still required to manually inspect every square meter of the data collected. In PMA mode, ATR is pre-dominantly used as a decision support tool, confirming contacts selected by the human and drawing the human's attention to threats they may have otherwise missed.

Current PMA workflows involving ATR still focus on presenting the results in a static manner, with the ATR output overlaid on the sensor imagery. This approach has been used to build human trust in the ATR performance, allowing the ATR output to be understood without disrupting operations. With ATR performance now at operationally viable levels, the limitations and risks of this workflow must be addressed to allow more effective human-machine teaming where the human and algorithm can interact, collaborate and learn together.

3 Man Made Object (MMO) – arrival of a digital assistant

AI is now critical in applications including Route Survey, Critical Underwater Infrastructure (CUI) Protection and Seabed Warfare (SBW) where Navies are interested in detecting and cataloguing a wide range of objects. Route Survey has a requirement to map all man-made objects within the region; CUI Protection must survey infrastructure with the intent of observing changes that may equate to threats; SBW may be interested in detecting adversarial deploying listening devices and UUVs infrastructure to the seabed.

This broad set of applications reduces the specificity of the problem domain and increases the variability of the data that the AI is required to assess. It becomes more challenging to provide the quantity and quality of data needed to produce an AI model that can successfully deal with each of the scenarios detailed above. Mission success requires the AI to interact with humans, creating an effective human-machine teaming to deal with the increased uncertainty and variability of the data across the applications.

SeeByte's Man-Made Object Detector (MMO) is designed to detect, classify and localise a broad set of objects of interest so that it can be applied across all underwater battlespace applications. Similar to ATR, the MMO produces decision support at machine speed with the following key differences:

- For MMO to be useful across a range of military applications, it must provide a solution capable of detecting a large variety of irregular shaped object types. While an MCM ATR prioritises high precision to reject false alarms by matching a few know threat signatures, the MMO prioritises high recall to ensure no anomaly is missed, inherently accepting more false alarms for mission completeness.
- Machine Learning models require large amounts of example data to perform well. Navies address this for ATR by laying practise mine fields containing the most common mine threats and conducting data collection programs over many years. The diversity, complexity and potential security classification of the objects that a MMO will need to detect means replicating this approach is impractical.

An analogy highlighting the differences between ATR and MMO can be seen in **Figure 3.** In this example, the ATR is trained to detect only cars whereas the MMO is trained to detect anything man-made. It is instantly obvious that the MMO must be capable of detecting a more diverse range of object types, many of which are very different in size and appearance.

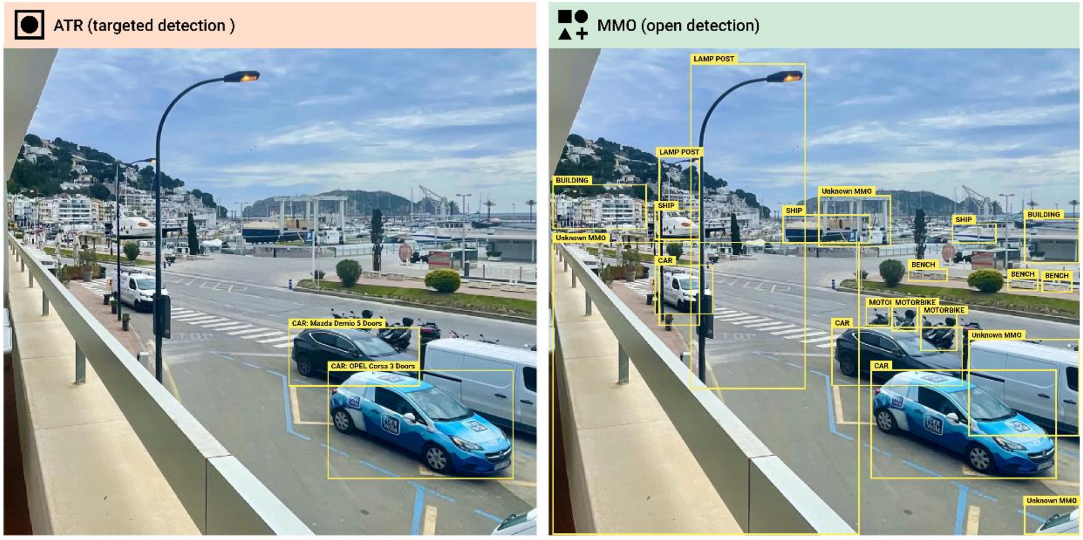

Figure 3: The same scene is processed by a targeted ATR trained to detect only cars, versus a generalised MMO that is required to detect anything that is man-made. Example car data can train an ATR to produce a strong PD/PFA result while ignoring everything else. In comparison, the MMO solution is required to detect a huge range of objects. This is more difficult to train, and the generalisation makes it harder to control false alarm rates.

MMO will be most successfully deployed across military applications through proximity to the human end-user, who will be able to intelligently query and filter the MMO output to achieve the mission result.

A Seabed Warfare officer may wish to quickly interrogate the MMO output to determine if a listening device has been placed within a survey area; on another day they may need to search through the same output to identify if any new IED targets have been placed near a Q-route.

The proposed SeeByte approach builds on the ability for the MMO to process vast amounts of data at machine speed, while leveraging the human's domain expertise (through smart queries) and ability to deal with uncertainty. A simple way of describing this is the human must be able to input their expert knowledge to tailor the MMO output so that the desired mission end-result can be achieved.

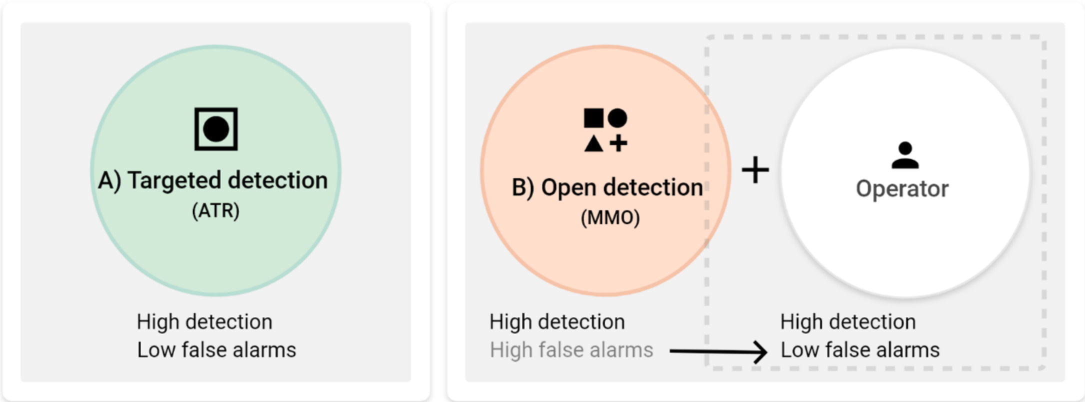

Figure 4: The MMO can rapidly ingest and process large amounts of data looking for a broad array of man-made objects. The human's domain knowledge can be used to intelligently filter the MMO output to achieve the desired result. This can be compared to a traditional ATR model, that is finely tuned to detect a specific set of objects only and has been developed with the intent that it will be able to work in isolation with limited human input.

The second important consideration is how suitable quantities of training data can be collected for the range of applications and object types the MMO may be applied to. SeeByte's MMO uses real data where available but augments the training data using simulation and generative AI.

This has several benefits – firstly, large quantities of training data can be produced for any object, as long as a 3D model of the object can be estimated; this ensures the MMO is extendable and can be updated as new intelligence is obtained; secondly it reduces the need for expensive data collection events, which may also be impractical for classified applications. Some examples of data containing simulated targets, which can be generated rapidly and at volume, are shown below.

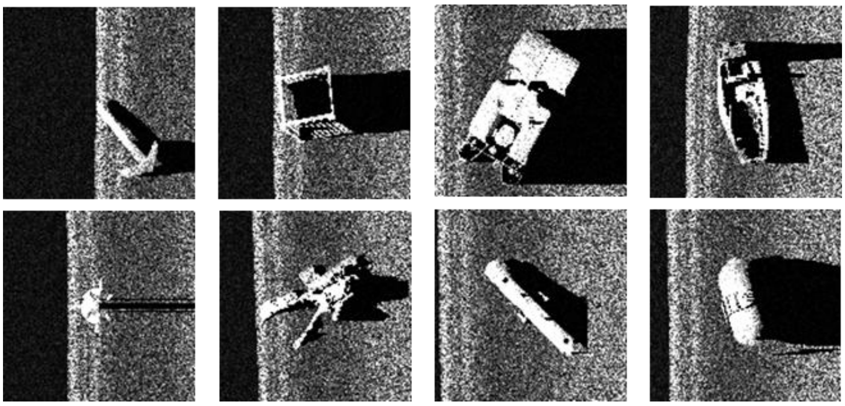

Figure 5: Simulators can be used to generate large amounts of MMO training data. Objects can be inserted into sonar data and realistically visualised. The resolution, appearance and noise models of the sonars are respected during this process. This approach allows the SeeByte MMO to be rapidly adapted to new sensors and new objects.

3.1 Using MMO

The SeeByte MMO acts as a smart digital assistant that the human can interact with to intelligently filter the MMO output based on the specific mission objective. Modern smartphones similarly allow us to rapidly query the vast repository of the internet to extract the information we need.

SeeByte's MMO allows similar operations, where the user conducts military relevant filtering operations based on their expertise and domain knowledge. These concepts can be seen in Figure 6 and 7 below.

In **Figure 6**, the output from the MMO is rapidly filtered based on shape, size and Change Detection concepts, effectively combining the rapid data processing of the MMO with the expertise and domain intelligence of the user.

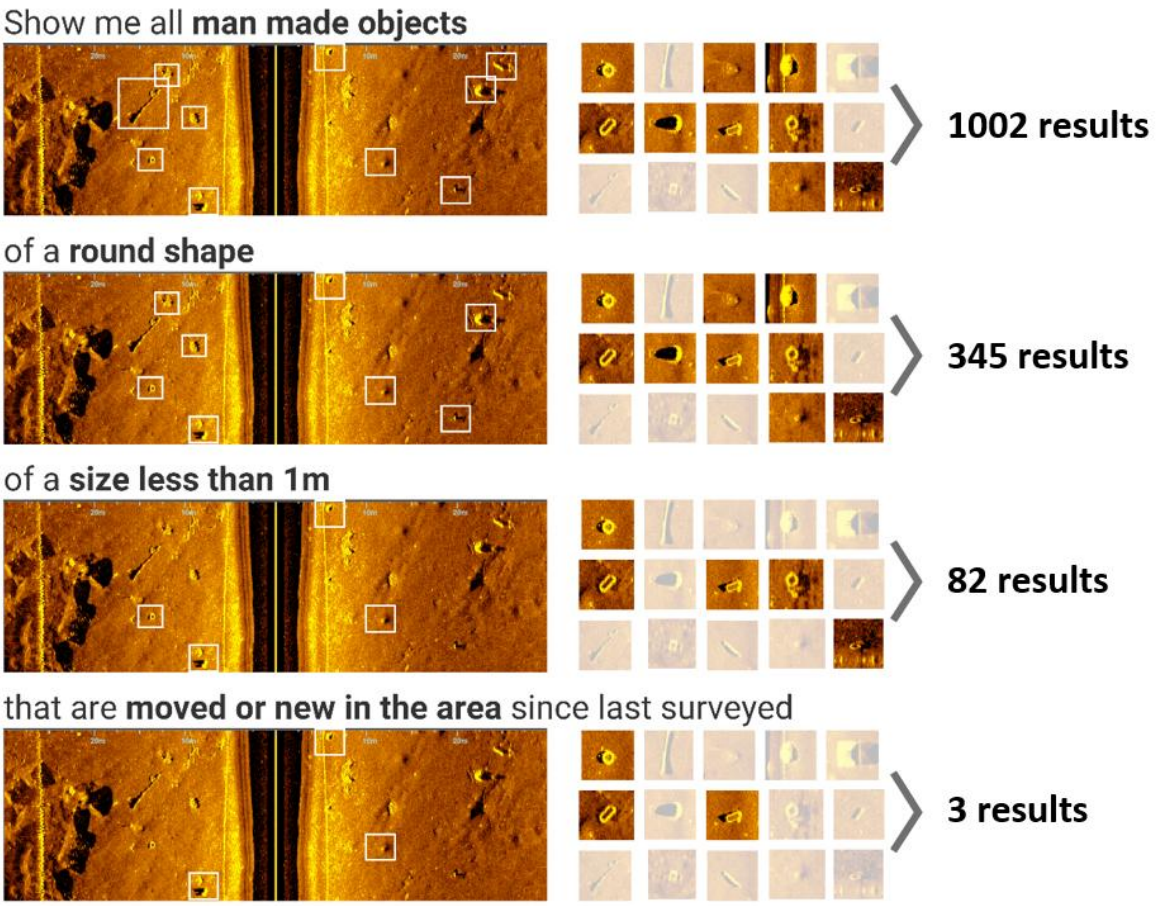

Figure 6: The MMO will act as a digital assistant allowing the human to rapidly query a large amount of data. The human can parse the MMO output using smart filters. The effective human-machine teaming leverages the ability of the machine learning based algorithm to rapidly process data, while incorporating the human's expertise to refine the search output.

Another smart query example available to the MMO user to intelligently search through the MMO output is shown in **Figure 7**. Complex concepts as "object similarity" allow the MMO tool to be adapted to a changing scenario, providing a mechanism for the user to incorporate uncertainty and variability into models. This, coupled with the ability to add new objects or mission intelligence into the MMO system using simulation and Generative AI, makes the MMO a powerful and flexible tool that can be applied across multiple applications within the underwater battlespace.

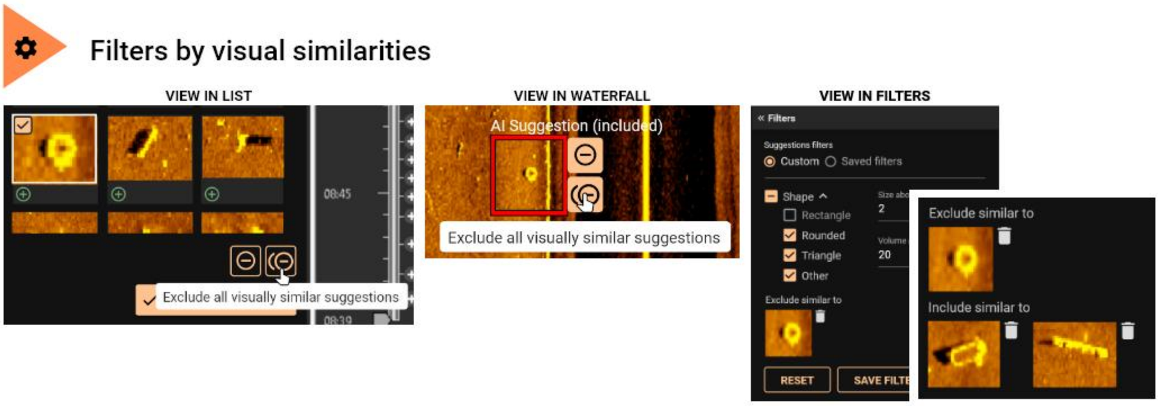

Figure 7: Smart queries allow the human to filter through the MMO output using concepts such as similarity. Querying the MMO to include, or disregard contacts that appear similar to other objects provide the human with a flexible tool that can be adapted to changes in environment or application.

4 What next? An agentic AI vision of agents that think and adapt

AI within the underwater battlespace will gradually move to an agentic AI approach where each agent's skillset and delegation of authority will depend on some of the factors detailed at the start of this paper – domain specificity, data availability, and access to the human.

AI agents that initially require substantial user input due to the complexity of the mission, will learn and improve over time as expertise is transferred from the human, eventually reaching a performance that allows an Agent to run in isolation when communication cannot be assured.

Agentic AI is not a query tool, it is autonomous and can independently assess a situation, plan multi-step actions, and execute them to achieve a goal.

As the Hybrid Navy move to this Agentic AI vision, the AI Data Assurance and defining clear lines of responsibility and delegation between the AI agents and the human end users will become increasingly important.

As the pace of technological adaptation accelerates, the ability to rapidly respond to evolving adversary tactics becomes a defining factor in avoiding overmatch. Modern conflicts have shown adaptation cycles shrinking to as little as 28 days and still tightening as technologies advance.

In this context, to outpace an adversary a solution like SeeByte's MMO are not simply advantageous; they are essential. Navies require systems that can evolve and adapt as quickly as the threat, enabling humans to mine vast datasets, rapidly understand the situation, and deploy updated capabilities at the speed of relevance.

5 Conclusion

This paper has presented SeeByte's journey. For 25 years, SeeByte has pioneered the transition from laboratory AI to operational naval capability. Our journey spans three critical frontiers: field-proven **ATR AI systems** that precisely localise known threats; **MMO AI Assistants** that empower operators to understand vast sensor volumes at machine speed; and our vision for **Agentic AI**, autonomous partners capable of independent situational assessment, multi-step planning, and mission execution.

In a battlespace where technological margins are razor-thin, SeeByte remains the partner of choice for turning complex data into decisive action.