Towards a Predictive Model for T45 energy usage: Initial Insights from analysis of Type 45 Platform Data-sets

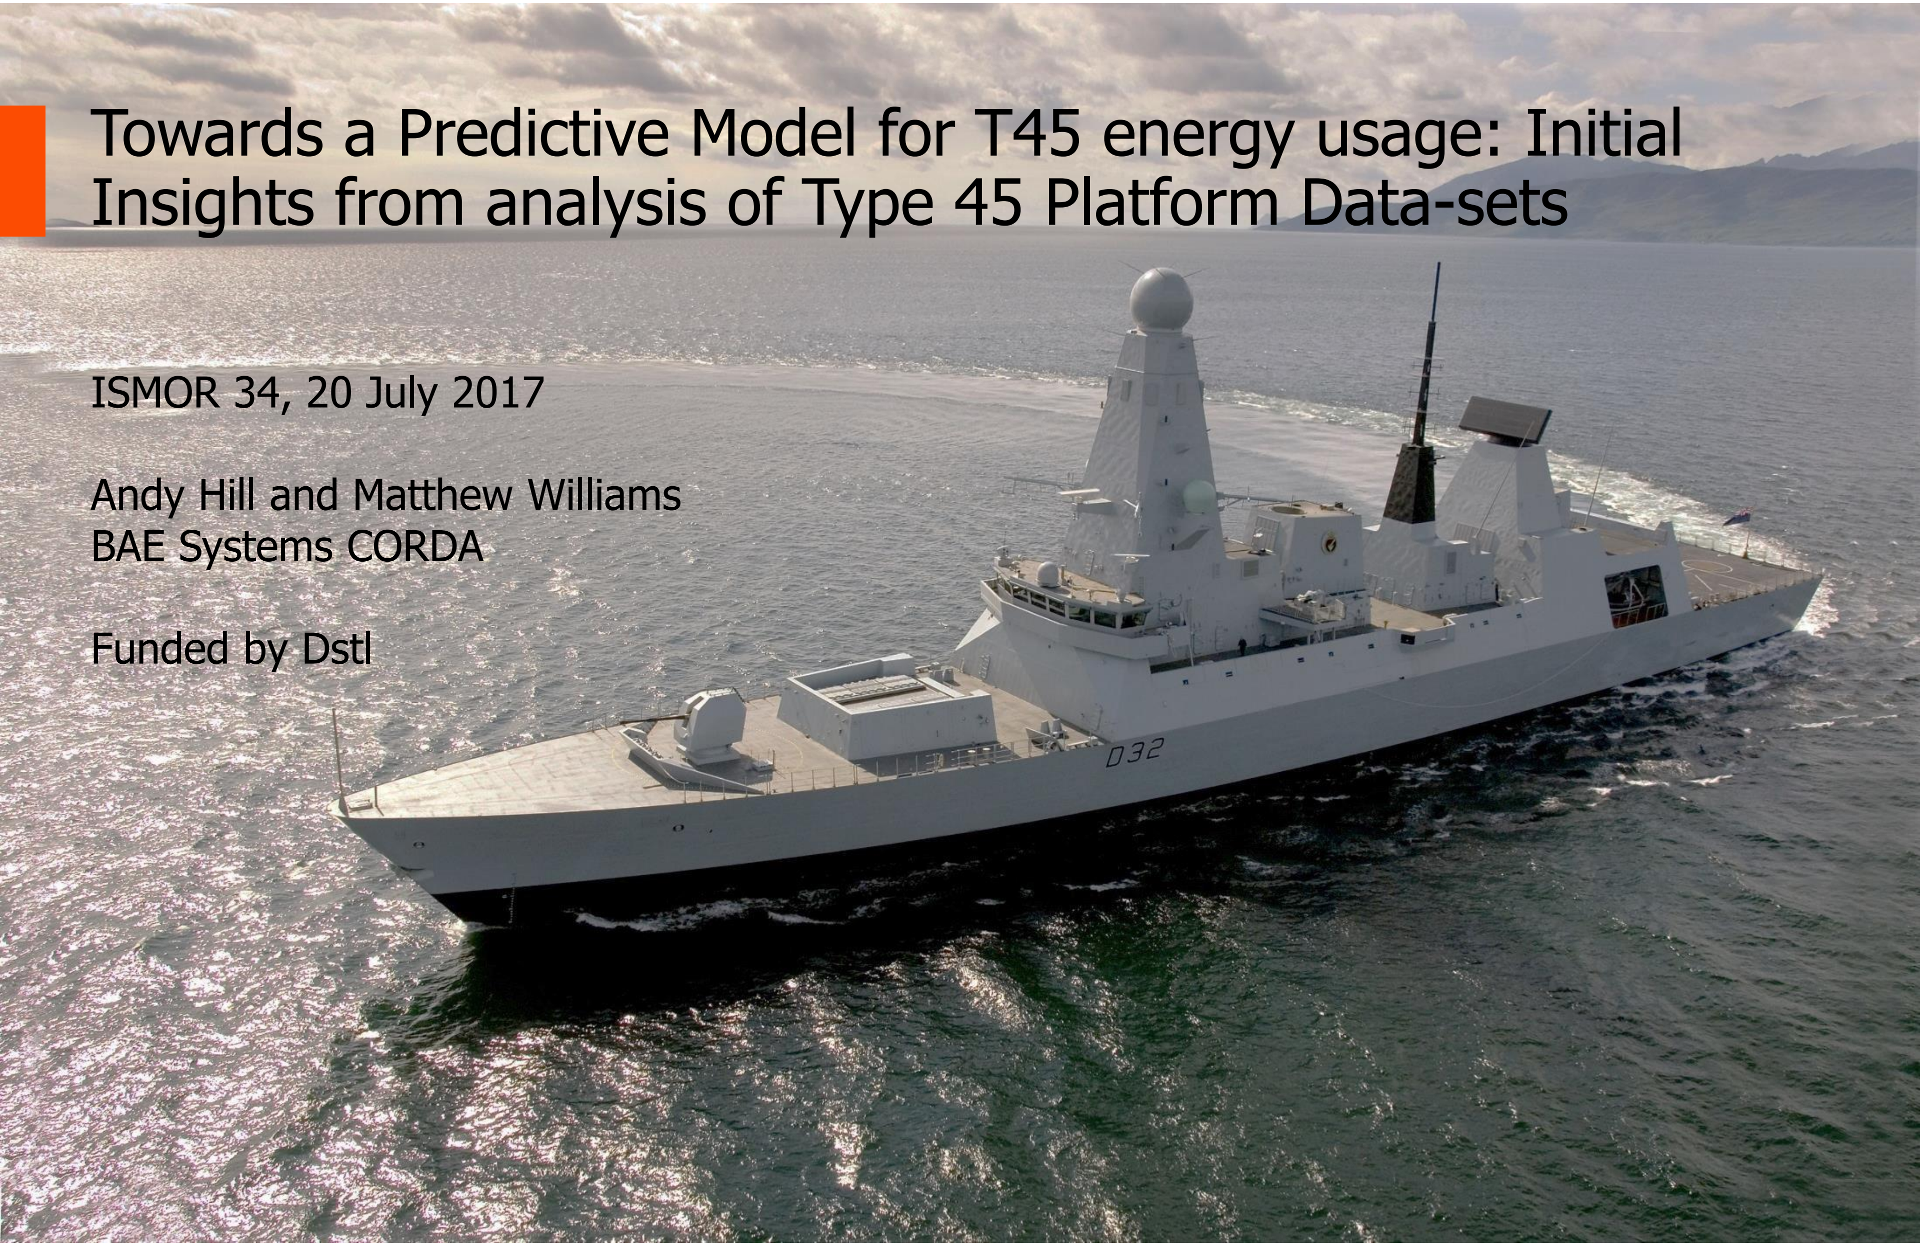

ISMOR 34, 20 July 2017

Andy Hill and Matthew Williams
BAE Systems CORDA

Funded by Dstl

[dstl] BAE SYSTEMS INSPIRED WORK

Background

- With the revolution in “big data” and predictive analytics over the past 5-10 years, it is now theoretically possible to generate novel insights from large data sources
- A warship produces a lot of data to support the operation of the ship – e.g. power consumption, water speed
- These data potentially provide additional value if we can collect and analyse them in an appropriate fashion
  - Either for “real time” intervention – e.g. to inform maintenance and to prevent catastrophic failure
  - Or for trend analysis – e.g. to identify what are the main drivers of fuel consumption

Background

- Dstl is therefore interested in the value and practical application of sensor information on warships
- One major problem with analysis of this kind is separating a signal from a very noisy environmental and operational background
- A BAE Systems-led consortium is developing a system called SEA-CORES to develop a solution to assist maritime operators to understand the trade-off space for optimising fuel and energy consumption across their fleet
- SEA-CORES includes the use of various sensors to assist in capturing the environmental and operational conditions, as well as aiming to capture useful metadata such as weather and ship log data
- By combining data sources it should be possible to:
    - Subtract the noise from the signal
    - Assess the contributions of operational and system factors to energy usage

Our task

- **Dstl asked us to:**
- Explore whether the **impacts of operator behaviours and system defects** on platform energy use could be identified in energy usage data;
- **Assess the robustness of the data** to understand how confident we can be in using the outputs as evidence for decision-making;
- **Provide recommendations for further data and research requirements**.
- Case study for today: breaking down speed-power curve by sea state

Data and assumptions

- Platform Management System data
  - Water speed
  - Propulsive Load
  - Rudder position
- SEA-CORES and ship's log metadata
  - Inclinometers (trim and heel)
  - Sea state, wind speed and lat/long data
  - Wave height and sea state data
  - Event log
  - Fuel load
- Data were collected over a five-month period from a T45 (October 2016 – March 2017)
- A complete dataset was available from 16 October 2016 – 18 November 2016

We used machine learning methods to understand drivers of energy usage

Train

- We merged various data sources
- We split the dataset into training and evaluation subsets
- We trained both a Bayesian Network and Boosted Regression Tree on the propulsive load

Evaluate

- The performance (using mean squared error, MSE) of the trained algorithm was assessed using the evaluation dataset
- The algorithm was refined so that the MSE was minimised

Predict

- The algorithm was then used to predict energy usage under a number of what-if scenarios
- Case study of energy usage as a function of sea state

BAE SYSTEMS
INSPIRED WORK

We chose two Machine Learning methods – BN and BRT - based on a set of relevant criteria

Criterion	Bayesian Network	Boosted Regression Tree	Classical time series analysis	Empirical Mode Decomposition
Recognisable predictor variables	✓✓✓	✓✓✓	✓	✗
Transparency	✓✓✓	✓	✓✓	✓
Flexibility	✓✓✓	✓✓✓	✓	✓✓
Adaptability	✓	✓✓✓	✓✓	✓✓
Accuracy	✓	✓✓✓	✓✓	✓✓
Accounts for interactions	✓	✓✓	✗	✗
Ease of implementation into optimisation software	✓✓✓*	✓✓*	✓*	✓*

* Preliminary opinion

BAE SYSTEMS
INSPIRED WORK

The Bayesian Network model

- We can automatically fit Bayesian Networks to the data
- This network assumes linear relationships and Normal conditional distributions
- Nodes depend on all nodes that have arrows pointing to them
- Arrows show a correlation, not necessarily causation! For example wave height can not affect fuel load, but it is a good predictor for it in the current data

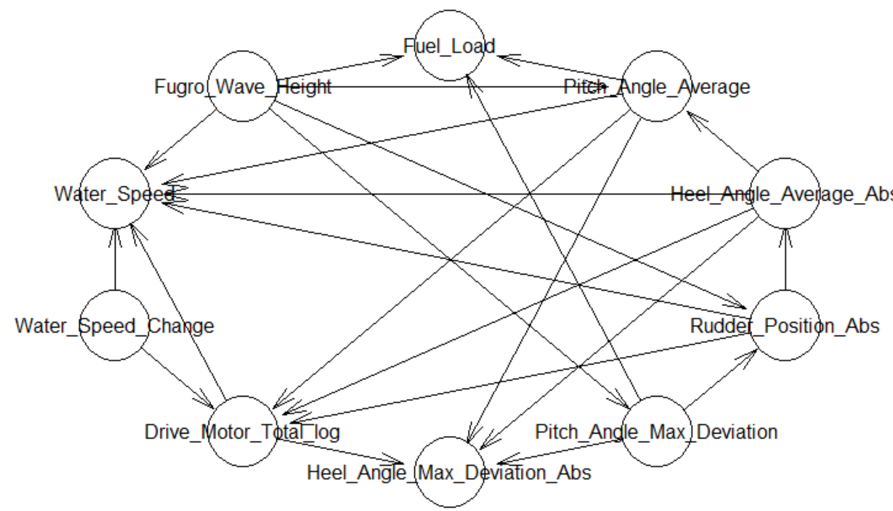

BAE SYSTEMS
INSPIRED WORK

The Boosted Regression Tree (BRT) model

- Boosting iteratively generates an ensemble of decision trees
- New trees are encouraged to become specialists in areas of the data where the ensemble currently performs weakly
- When predicting from the ensemble each tree gets a weighted vote on predicting the value
- This increases the predictive power of the model, albeit at cost of transparency and simplicity

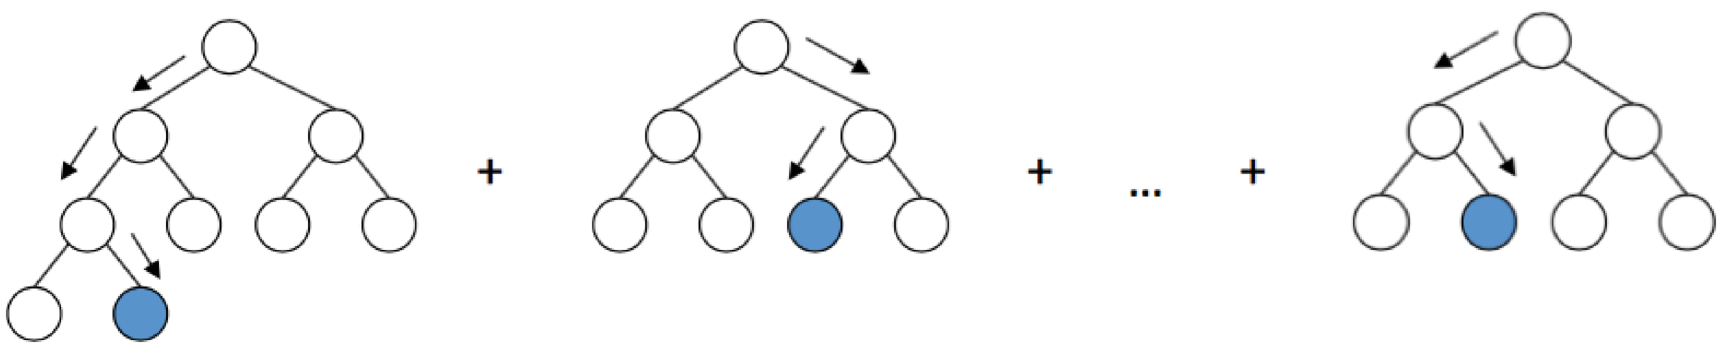

The boosted decision tree produces a better predictive model

- The cross-validated Mean Squared Error was 0.05 - a pretty good fit to the evaluation data in the range where we had plenty of data
- The cross-validated Mean Squared Error for the Bayesian network was 0.091

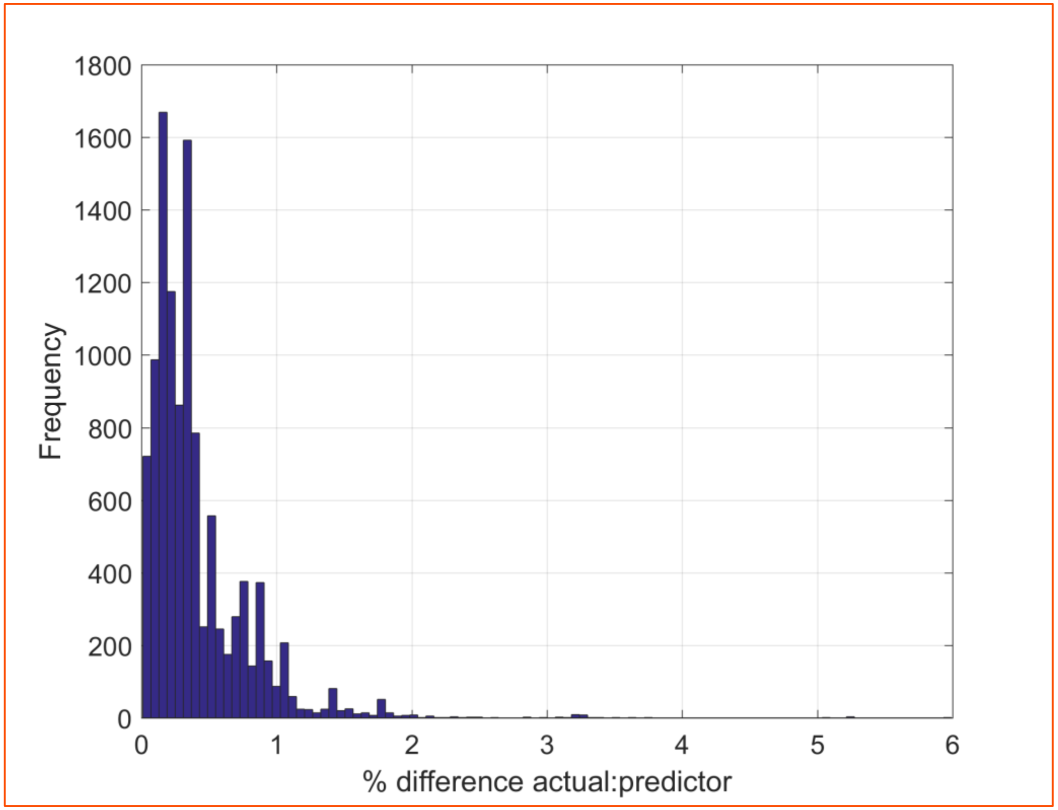

So what?
We can now break the energy signal into constituent parts, taking into account interactions…

Breaking down the noise in the power-speed curve: Results of baseline predictive model (blue), and predictive model when noise from acceleration and rudder removed (red)

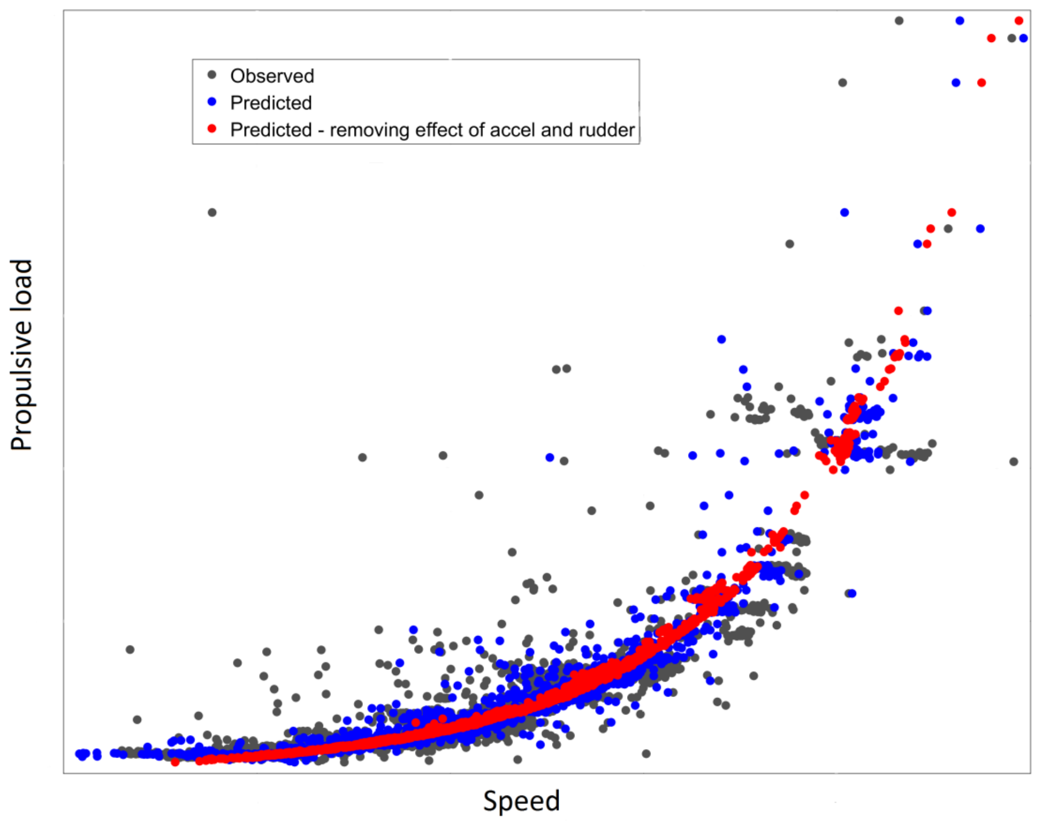

BAE SYSTEMS
INSPIRED WORK

Breaking down the noise in the power-speed curve: Going further, we removed noise due to heel and trim to produce clean power-speed curves by sea state

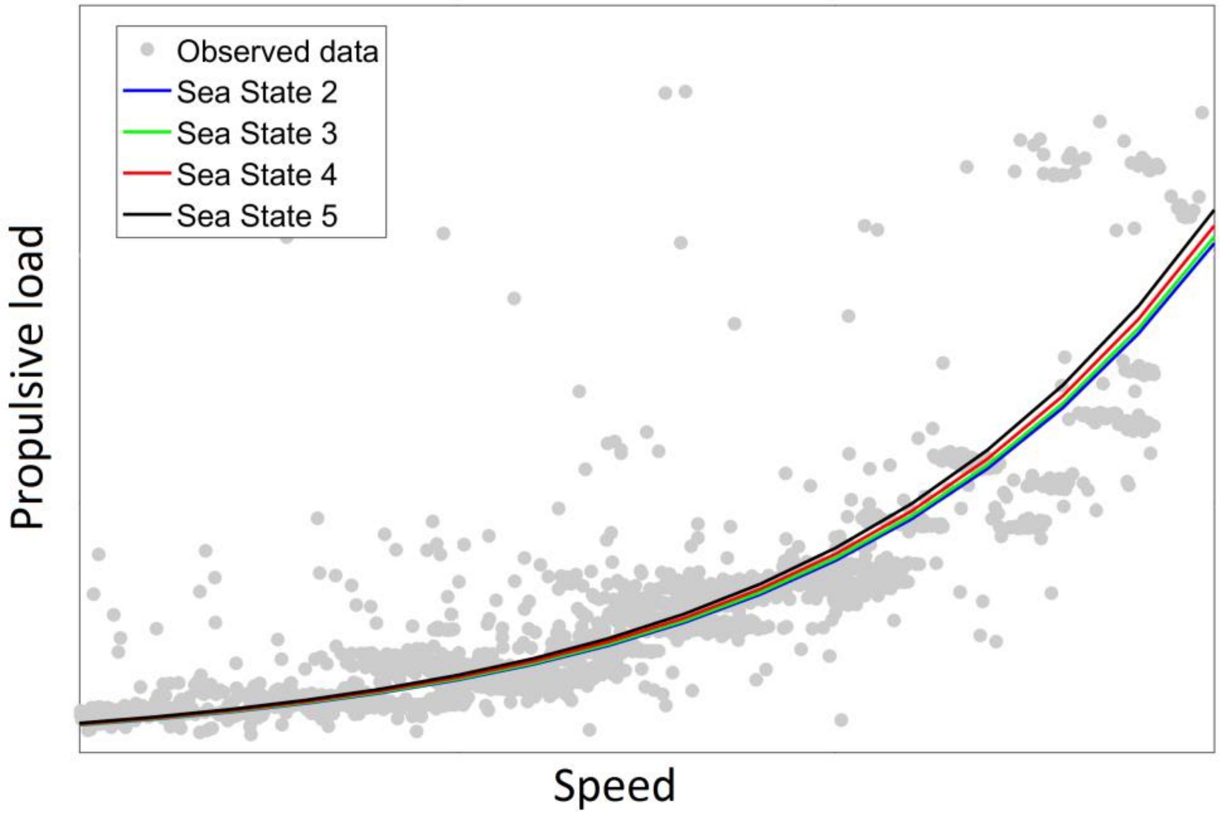

Summary of trend analysis

- Water speed is by far the most dominant variable – this made it quite difficult to extract meaningful contributions from the other variables
- Numerous anomalies in the data appeared like trends from simple visual analysis: these were ultimately determined to be interaction effects
- The boosted regression tree model predicts propulsive load well, and can be used to assess the influence of each variable on propulsive load
- We can now break down the energy signal into its constituent parts and isolate the variables we are interested in – crucial if we are to estimate the impact of operator behaviours and system defects

How to use it: example of trigger trial

- Hull and propeller fouling affects performance
- Trigger trials are monthly tests designed to evaluate extent of fouling – ship does a straight line back and forth at a constant speed. If ship is using more than a set power, a high-speed run is carried out to remove fouling
- However, there are many variables that cannot be perfectly isolated – e.g. system defects, operator behaviour. Also, due to operational requirements trigger trials may not be conducted regularly
- The predictive model could be used to more accurately assess the impact of fouling, and this could be done in real time rather than waiting for an opportunity to conduct a trigger trial – potentially increasing overall energy performance

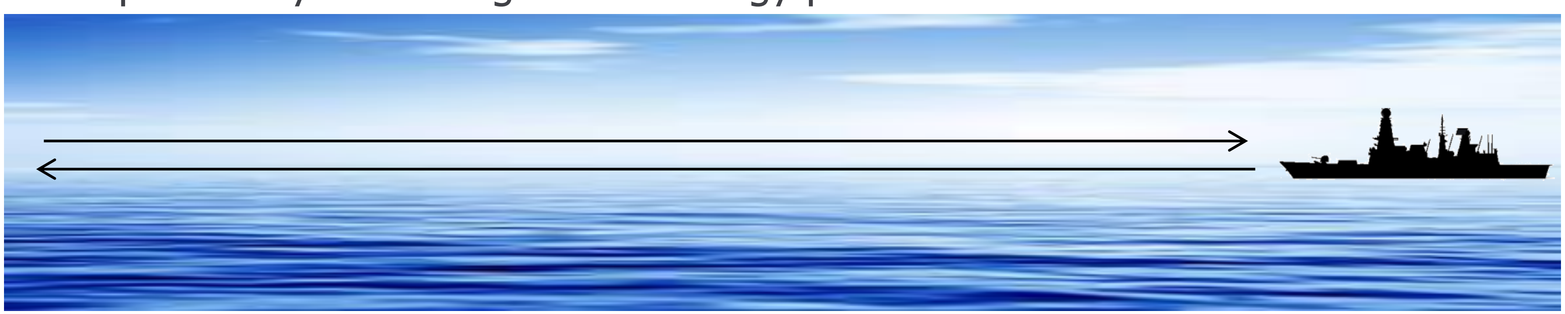

BAE SYSTEMS
INSPIRED WORK

Conclusions

- With only 1-2 months' of complete data, we were able to:
  - Profile in unprecedented detail the energy consumption as a function of ship activity
  - Deconstruct a speed-power curve for the ship
  - Build a reliable predictive model
  - Provisionally quantify the effect of trim and sea state on energy consumption

- More data would allow us to:
  - Quantify extremes of operational activity
  - Include more variables and describe more of the unexplained variance (e.g. windage)
  - Build a predictive model that can be translated into ship optimisation software

Acknowledgements

- This project was funded by Dstl
- We thank BAE Systems Warship Support for the provision of SEA-CORES data