BETTER SHIPS, BLUE OCEANS

Operational speed power analysis

Using Bayesian modelling

Report no.	:	70138-7-RD
Date	:	March 2025
Version	:	1.0 Final

Operational speed power analysis

Using Bayesian modelling

Ordered by : MARIN

Reported by : J. van den Berg

Version	Date	Version description	Checked by
1.0	March 2025	Final	F. Sterlini

Contents

Figures

1 Introduction

The speed power relation is one of the key performance indicators for ship fuel efficiency. To determine the effectiveness of fuel saving measures, detection of changes in the speed power relation is often the most direct means. This can be done by dedicated trails as described in Report 70138-8-PAS. However, dedicated trails are not always feasible in view of time and cost.

This report describes an approach to determine the speed power relation from operational data. This provides a cost effective method to detect changes in fuel efficiency, either due to fuel saving measures or due to e.g. maintenance issues.

The report is split into four main parts:
- Section 2 provides an overview of the type of signals that can be used for the speed power analysis
- Section 3 describes the preprocessing that is applied to the operational data before it can be used in the actual analysis
- Section 4 describes the general setup of the analysis
- Section 5 illustrates the analysis approach using two datasets from general cargo vessels

Finally, Section 6 summarises the main conclusions.

2 Relevant signals

The speed - power relation is a key indicator of a ship's efficiency. It relates the relative forward velocity with respect to the water, to the propulsion power required for that stationary speed. Procedures to test and measure the speed - power at model scale, in still water, are well established and standardized. Standard procedures to determine the speed - power at full scale using dedicated trails are also available.

To monitor the long term efficiency of a ship in the actual operational conditions, no standard approach is available yet. There are a number of limiting factors:

- Reliability of data: Keeping the required measurement equipment available and calibrated over a long period of time can be a challenge
- Environmental conditions: In the operational use, the ideal testing environment (negligible wind and current) is rarely available
- Ship loading conditions: The operational use of the ship usually dictates changes in e.g. loading conditions that affect the required power

All these factors can obscure small, but significant, changes in the speed power relation. In Section 3 and Section 4 we will describe a procedure to deal with these confounders. In this section we intend to outline the possible measurement inputs that can be used for a speed - power analysis on operational data.

2.1 Speed

The two most commonly available measures of the vessel speed are the speed over ground and the speed through water.

2.1.1 Speed over ground

The most straightforward measurement of the vessel speed is the speed over ground. All modern ships are equipped with a GNSS sensor, offering a reliable measurement of speed over ground.

2.1.2 Speed through water

Instead of the speed over ground, speed through water is commonly used. This refers to the vessel speed relative to the ambient current. While this measure is less relevant to the vessels mission (getting from A to B in a certain timespan), it does lead to a cleaner speed - power relation because it eliminates the disturbance from current. Two methods to obtain the speed through water are commonly used:

- Measurement by means of a speed log (Doppler or electromagnetic type) or similar instrument. The quality of such a measurement will depend on the placement of the sensor, hull shape and sensor calibration (it will need to look through a flow that is disturbed by the vessels presence). Report 23200-12-TM describes an extensive analysis of speed log accuracy in operational conditions
- Reconstruction based on the speed over ground vector and the current vector. The current vector is rarely measured and is usually taken from hind-cast data

2.2 Power
The power consumption of a ship at stationary speed is usually dominated by the propulsion power. To determine the speed/power relation, we are usually interested in the isolated propulsion power. Although this is not the best measure of the final fuel bill, it is the least disturbed measure of the hydrodynamic performance.

2.2.1 Shaft torque and rpm
The most direct method to obtain the propulsion power is to measure the propeller shaft torque (e.g. by strain gauge) and shaft speed (rpm). This will include some bearing and seal losses, but these will be small at typical cruising speeds (order of 1%).

2.2.2 Propeller trust
Another method to determine the propulsion power is the resistance multiplied by the speed through water. As measure for resistance, one might be able to measure propeller thrust through the shaft. However, propeller thrust is rarely measured at full scale and if it measured, it has not yet proven reliable. Next to this, to get to resistance, a thrust deduction factor t is needed as well.

2.2.3 Engine power
In some cases, the engine management system provides an output of the delivered engine power. This is a reconstruction based on a large amount of engine monitoring equipment. Apart from the uncertainty in the reconstruction, this figure can also include additional bearing and gearbox losses and there may also be a shaft generator absorbing power for auxiliary use that does not arrive at the propeller.

2.2.4 Fuel rate
The last option to get an idea of the propulsion power is the fuel rate. While this is the most direct measure of the fuel consumption, it has significant uncertainty when used as an indicator of the propulsion power. Not only does it include the gearbox and generator losses and can be plagued by measurement uncertainty of the flow meters; the fuel quality has a large variability as well and the engine efficiency will also vary with maintenance.

2.3 Draft and trim
The draft and trim can have a significant influence on the required propulsion power at a given speed. If the draft and trim vary during the vessel operation, the draft and trim dependencies should be included in the fit to reduce scatter and to learn the actual dependency.

Typically, there are two sources of draft information:

- Fore and aft draft sensors
- Visual inspection of the fore and aft draft marks on the hull, reported in the noon reports

When available, draft sensors provide a continuous time trace. The visual inspection values are only available at much larger intervals and are prone to human interpretation and book keeping errors. The draft sensors can be distorted by the bow wave at forward velocity.

The trim can be determined using the difference between the fore and aft draft.

2.4 Weather

Three weather aspects can have a significant influence on the required propulsion power at a given vessel speed: wind, waves and current. The current is mostly due to ocean circulation and tidal current, which are independent of the wind. The wind forced current can have a dependency on the wind. The waves are commonly (somewhat arbitrarily) split into wind driven (correlated to the present wind at the present location) and swell (long wave components from weather elsewhere that travelled into the area).

2.4.1 Wind

The apparent (relative to the vessel) wind velocity and direction has a direct influence on the vessel resistance through aerodynamic drag loads on the above water structure. In addition, the true (earth-fixed) wind velocity and direction can be used as a marker for the expected wind driven waves. Most vessels are equipped with an anemometer that gives a direct measurement of the relative wind velocity and direction. Using this combined with the speed over ground and heading from the GNSS system and/or ship's compass, this can be converted into absolute wind velocity and direction. Sometimes this is integrated into the wind logging and the distinction between absolute or relative wind is not always clearly indicated. The anemometer mounted on the ship can however be disturbed by the flow interaction with the deck top structures, affecting accuracy of apparent wind measured and true wind derived from it.

An alternative to onboard wind measurements is hind-cast data. Weather models can provide an estimate for the local undisturbed wind. These models operate on relatively coarse grids in space and time and interpolation is needed to provide the appropriate values at the vessel's positions over time. This interpolation can be problematic when there are strong local weather variations, especially when the vessel is sailing near shore.

2.4.2 Waves

The most dominant effect of waves on the required propulsion power is a higher order effect known as wave added resistance or wave drift loads. The change in resistance is mostly relevant in bow waves. The strongest effects are usually at higher frequencies (due to wave reflection) and around the maximum pitch response period.

In operational conditions, wave buoy measurements are not a viable way to obtain the local wave conditions since the buoy will not travel with the ship. An onboard wave radar can provide local measurements, but the measurements will typically be disturbed by the bow wave, spray and the reflected/diffracted waves due to the presence of the hull. Depending on the used radar technique, the accuracy of the measured wave height can be problematic as well.

Alternatively, hind-cast data can be used to get an indication of the local wave conditions. This has the same accuracy concerns as wind conditions from hind-cast data.

To translate the wave conditions to the corresponding shift in required propulsion power two approaches can be used:

- Use a hydrodynamic model to compute the expected change in resistance and correct for it. A range of model fidelities is available with a corresponding range of required input detail
- Fit the wave - resistance relation on the operational data. This requires sufficient cross variability in the data to isolate the wave influence from other influences as e.g. draft variation

2.4.3 Current

The ambient current is needed when the (more reliable) speed over ground needs to be translated to speed through water. If a reliable speed through water measurement is available, information on the current has little added value. In principle, hind-cast data for the current can be used to construct speed through water from speed over ground. However, in most large datasets, the long term current is zero mean and not correlated to the other disturbances. In the absence of a speed through water measurement it is usually a better idea to just use the speed over ground and only use the hind-cast data on the local current to check against bias.

3 Domestication

Compared to a tailored test campaign, operational data is typically much larger in size (long term monitoring) and the data is ‘undomesticated’:

- The quality of the measurement equipment is less well known
- Source locations and output units of the sensors are often poorly documented
- Equipment failures and repairs and their exact timing are not always documented
- Sensor locations may change over time
- Acquisition is often done using multiple vendor specific systems that often lack proper time synchronization
- Acquisition is continuous, but only a part of the dataset will be relevant for the performance analysis of a particular type of operation

Before the data can be used in the actual analysis, it needs to be domesticated. This usually entails the following steps:

- Check for outliers and inconsistencies
- Reduce the data size and derive characteristic parameters
- Select the relevant parts of the data

3.1 Outlier detection

A common approach to remove outliers is to look at the sample distribution. By comparing the extreme values to e.g. quantile values, sample values that are outside the typical distribution can be spotted. For measurement time series, this is not always adequate. An example is given in Figure 1: relative to the signal, the spike around 2.6 seconds is an outlier. However, the sample value at the spike fits well within the overall sample distribution. Traditional outlier removal would fail in this case.

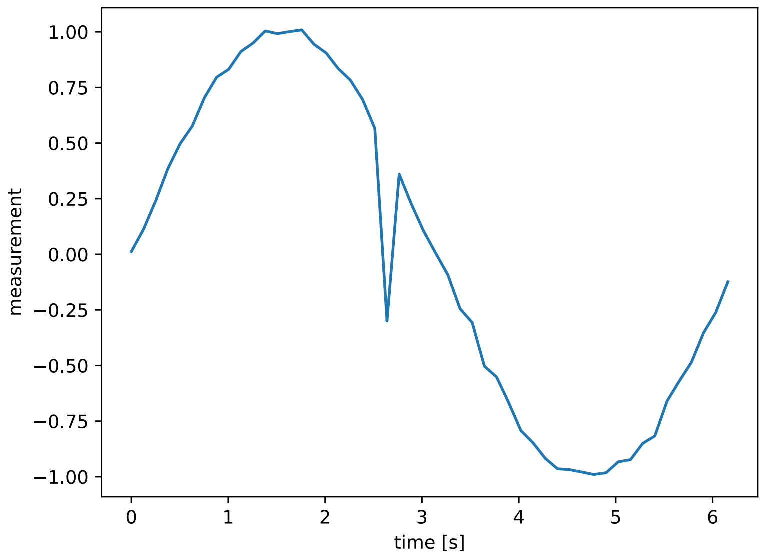

Figure 1: Example harmonic signal with noise and an inward spike. With respect to the signal, the spike is an outlier, but the value is well within the overall sample distribution

To construct a more reliable detector that also catches the ‘inlier’ spikes, we use the principle that the consecutive samples of a measurement time trace should not be uncorrelated. The discrete measurement samples are taken from an underlying smooth continuous process. To test if a sample is an outlier, we use the following procedure:

- We fit a local linear trend through the nearest neighbours of the sample (excluding itself)
- We compute the absolute difference between the sample value and trend value
- We normalize this difference by the standard deviation of the neighbour samples with respect to the fitted trend

As long as the sample rate is sufficient that local neighbourhood can be approximated by a linear trend, the standard deviation of the neighbour samples with respect to the trend are a measure for the local noise level in the signal. The normalized absolute difference is an indication of whether de sample fits within this noise distribution. If it does not, it is considered an outlier and removed.

3.2 Windowing

To reduce the data size and derive characteristic parameters, we use a windowing approach. We impose e.g. 20 minute windows over the time trace and extract the characteristic value for each window. A common approach is to use the median value to represent the data in each window. The median value is robust under variations in sample rate and automatically suppresses the influence of outliers.

However, using just the median to represent the window may result in loss of essential information on internal variations. For example, in a speed power analysis the signal could appear stationary at the 20 minute window level, but is not stationary at the original 10Hz sampling level. One approach to address this, is to isolate the stationary sections at the 10Hz level, before down-sampling afterward. This implies doing non trivial analysis at the full data size.

Here, we use a different approach. In line with the window median approach, we use tiled windows over the time trace. In each window we fit a linear trend on the valid samples. Using that trend we represent each window by the following characteristic values:

- The mean of a fitted trend, which can be used in the same way as the median
- The slope of a fitted trend, which can be used to detect a trend within the window
- The standard deviation with respect to the fitted trend, which gives an indication of the level of short term variation e.g. turbulence or noise
- The minimum and maximum value with respect to the fitted trend, also indicators of the deviation of individual samples from the window trend
- The ratio of Not a Number values over valid samples can be used as a quality indicator

Depending on the question at hand, we can use a subset of these characteristics to select the relevant windows (Section 3.3) and as input for the actual analysis (Section 4).

3.3 Stationary sailing window selection

Operational speed power data has a large amount of scatter. Part of this scatter is due to manoeuvring. For a sea going vessel, a large part of the data can be classified as ‘stationary sailing’:

- Constant significant speed

- Constant rpm
- Constant course

By selecting the windows where these criteria are met, we remove the port calls and instationary data points from the analysis. The general idea is that the open water, stationary sailing, data has a more consistent speed power relation.

To establish the thresholds for the minimum speed, and the maximum variations on speed, rpm and course, we look at the probability distributions of the window mean, slope and standard deviation as described in Section 3.2.

4 Analysis

After domestication and selection of the stationary sailing windows, operational speed power data still has a large amount of scatter. This is due to variations in the sailing conditions, some of these variations are known and some are not. Typically we can expect:

- Draft and trim variations
- Water depth (keel clearance) variations
- Weather variations (wind, current and waves)

Depending on the particular dataset, we may have data of these confounding factors. Ideally, the analysis should be able to account or compensate for them. For example, if we know the draft variations, we can either fit the dependency of the speed power relation on the draft, or we could use a prediction model to compensate for the draft dependent variations beforehand. The latter approach requires a reliable prediction model. Here, we will use the data driven approach.

4.1 Bayesian modelling

A common method to fit a relation to a multi variable problem is regression. This results in a best fit. Here, we use Bayesian modelling instead. This leads to a best fit that is identical to the regression result. However, Bayesian modelling also keeps track of uncertainties and provides a likelihood of the complete range of possible fits.

We start from a parametrized model that can predict the measured data. The vector y denotes the measured data and the vector θ the set of model parameters. For a given set of parameters θ, we can compute the probability of finding y: $P(y|\theta)$. Using Bayes rule this can be rewritten into the likelihood of particular parameter values θ, given the data points y:

$$P(\theta|y) = \frac{P(\theta)P(y|\theta)}{P(y,\theta)} \propto P(\theta)P(y|\theta) \tag{1}$$

Using Equation 1, we can use Monte Carlo simulations to evaluate the likelihood of a range of θ values. If we have no prior knowledge of the parameters θ, we can assume a uniform distribution for $P(\theta)$. Instead of brute forcing all possible θ values, we use Monte Carlo Markov Chain sampling.

4.2 General model setup

For a speed power analysis, we start with a basic speed power law:

$$p = bv^3 \tag{2}$$

The most basic model for our speed power data becomes:

$$p_i = N(\mu = bv_i^3, \sigma = r) \tag{3}$$

For each sample pair v_i, p_i, the p_i can be predicted by a normal distribution with a mean μ and a standard deviation σ. The mean is given by the power law. The coefficient b and the noise level r are unknown model parameters.

Fitting this model to actual speed power data, will result in likelihood probability distributions for b and r. The distribution of b can be interpreted in terms of an error band around the fitted speed power curve. The maximum likelihood value of r gives an indication of the expected spread around these error bands of the individual data points.

4.3 Wind speed dependency

The first extension we make to the basic model (Equation 3), is to include a wind speed dependency. Wind speed is a reliable indicator of weather severity. We have lower confidence in a data point when the wind speed is high. In the model this is reflected by writing:

$$r = r_0 + v_{w,i} r_w \tag{4}$$

We replace the single model parameter r by a combination of a base noise level r_0 and a wind speed dependent noise level r_w. The result will be that we automatically get an indication of how much uncertainty is added by the wind and we get a stronger weighing of the low wind speed samples in the fit of the speed power coefficient b.

4.4 Draft dependency

The next extension we can make is to add a draft dependency. We denote the draft of a sample by d_i. Next we use:

$$b = b_0 + b_d d_i \tag{5}$$

For a linear dependency of draft. In theory it should be a power of d_i, but in practice that usually does not improve the fit for typical draft variations.

If the data contains sufficient independent draft variations (as in, e.g. speed and draft vary independently), we get an indication of the draft dependency.

Trim dependency can be included in a similar fashion.

5 Example data sets

Within the green deal project "Operational data driven advice," two operators were found willing to share a large operational dataset to test the analysis approach presented in Section 4. To protect their commercial interests, the datasets are presented anonymously in this public report.

5.1 Dataset 1

The first dataset contains data for a fleet of general cargo vessels. Six of these vessels are sisters that should have similar speed power relations. For this analysis we will focus on these sister vessels.

5.1.1 Window selection

After deriving the 20 minute characteristics as described in Section 3.2, we first check the probability density curves for the most relevant signals:

- The speed through water
- The shaft rpm
- The course over ground

The corresponding plots are shown in Figure 2.

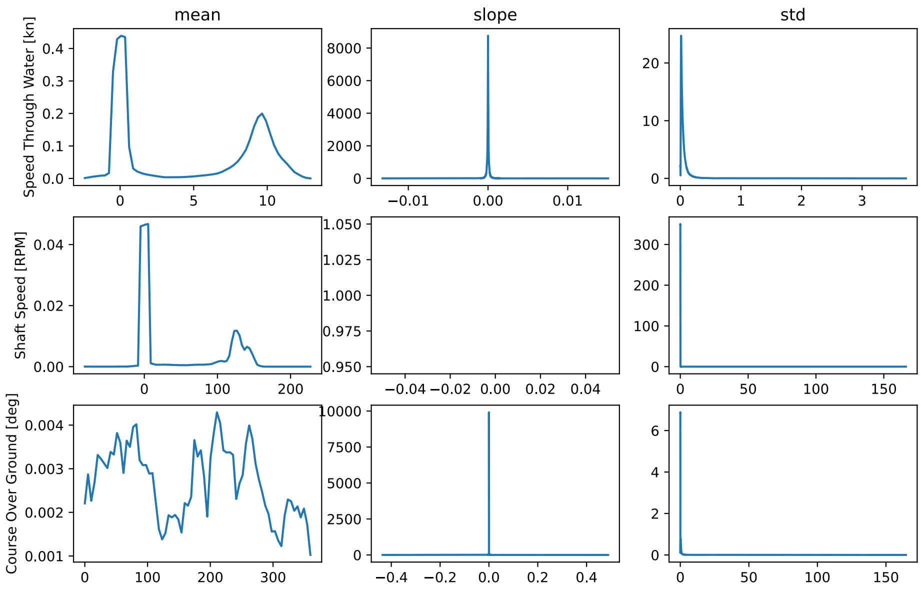

Figure 2: Dataset 1, 20 minute window characteristics distributions for all windows

Based on these results we set the minimum speed at 5 knots and the minimum shaft speed at 80 rpm. In addition we apply maximum standard deviations for the shaft speed (3 rpm), the speed through water (0.2 knots) and the course over ground (2 degrees).

After this selection we can look at the distribution of the remaining 25 thousand samples. The corresponding plots are shown in Figure 3. There are no significant trends (slope) in the remaining samples. And we use this set as the stationary sailing samples for further analysis.

The raw speed power points are shown in Figure 4. Clearly, the amount of scatter is too high for straightforward interpretation.

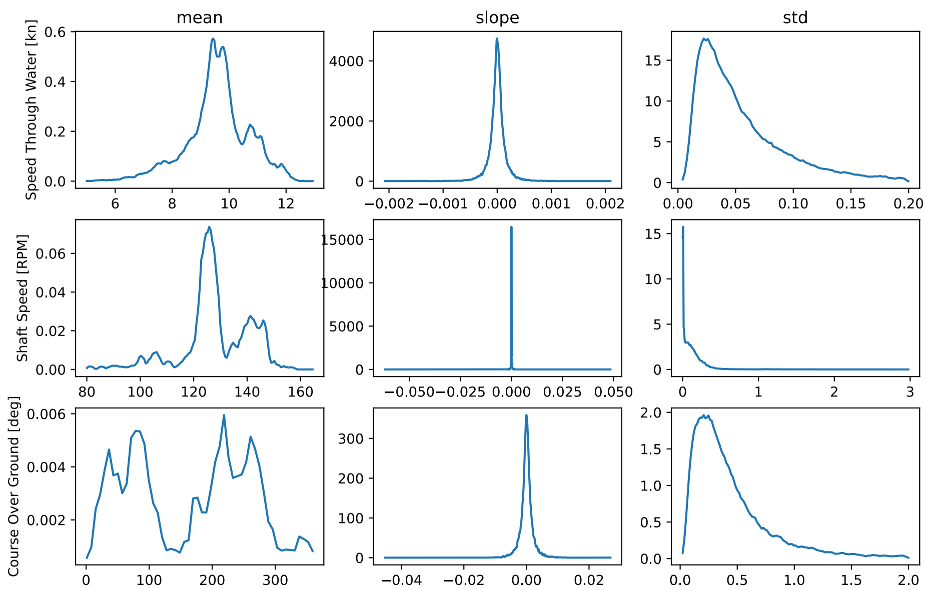

Figure 3: Dataset 1, 20 minute window characteristics distributions for stationary sailing windows

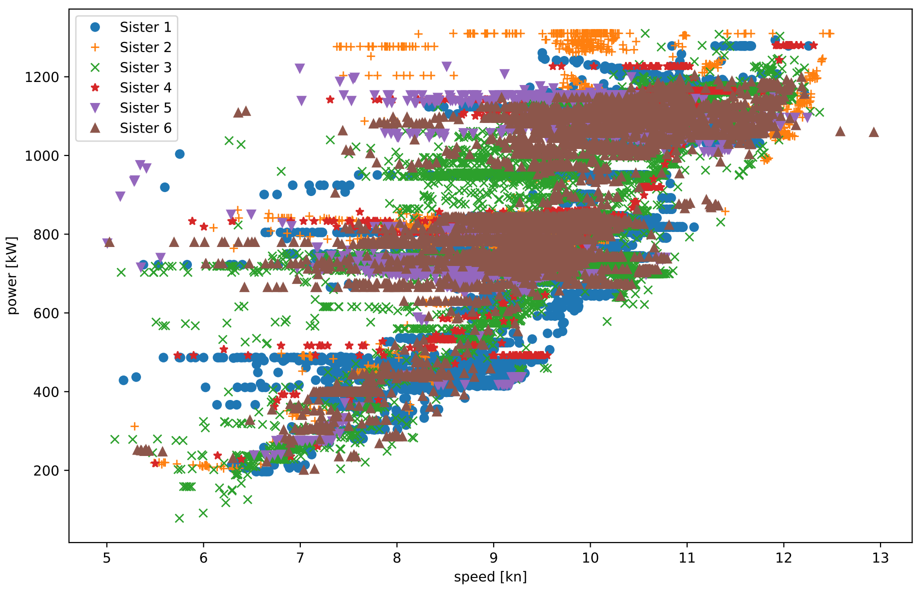

Figure 4: Raw stationary sailing speed power data points per sister vessel

5.1.2 Speed power fit

As a first check, we use the speed power law (Equation 2) to determine b coefficient per sample:

$$b_i = \frac{p}{v^3} \tag{6}$$

If we plot this per vessel over time (Figure 5), we can check for obvious trends. In this case there are none.

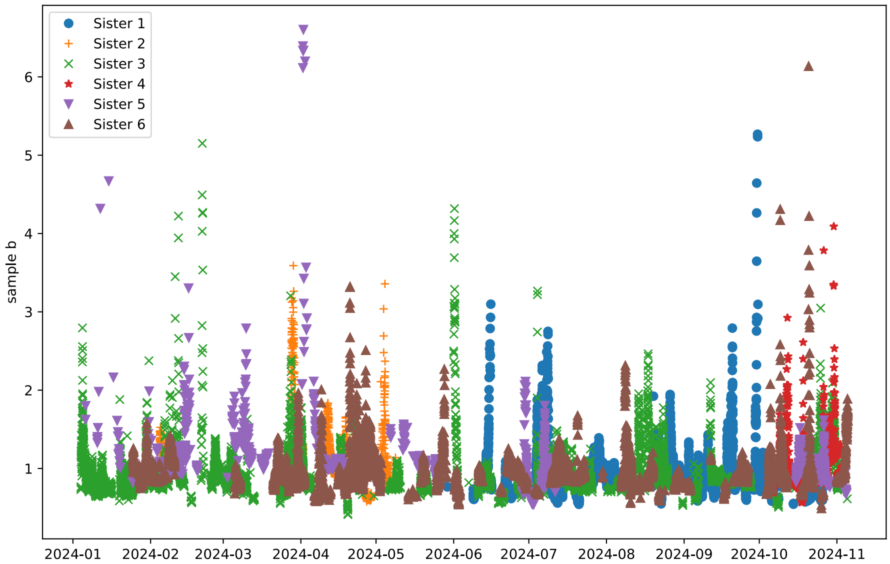

Figure 5: The speed power coefficient over time per sister vessel

If we fit a speed power model with wind speed and draft dependencies (Equation 3 with Equation 4 and Equation 5), we get a joint likelihood for all model parameters. Figure 6 shows cross sections of this joint probability per parameter pair. From this we can conclude:

- b_0 is independent from r_0 and r_1 (a clean circular cross correlation)
- b_d is independent from r_0 and r_1 (a clean circular cross correlation)
- b_0 and b_d are cross correlated, a larger value of one combined with a lower value for the other gives a similar predicted power. Fortunately the maximum likelihood (bright yellow) is restricted to a small centre
- r_0 and r_1 are cross correlated: the overall scatter in the power can be blamed on base uncertainty or wind dependent uncertainty to some extent

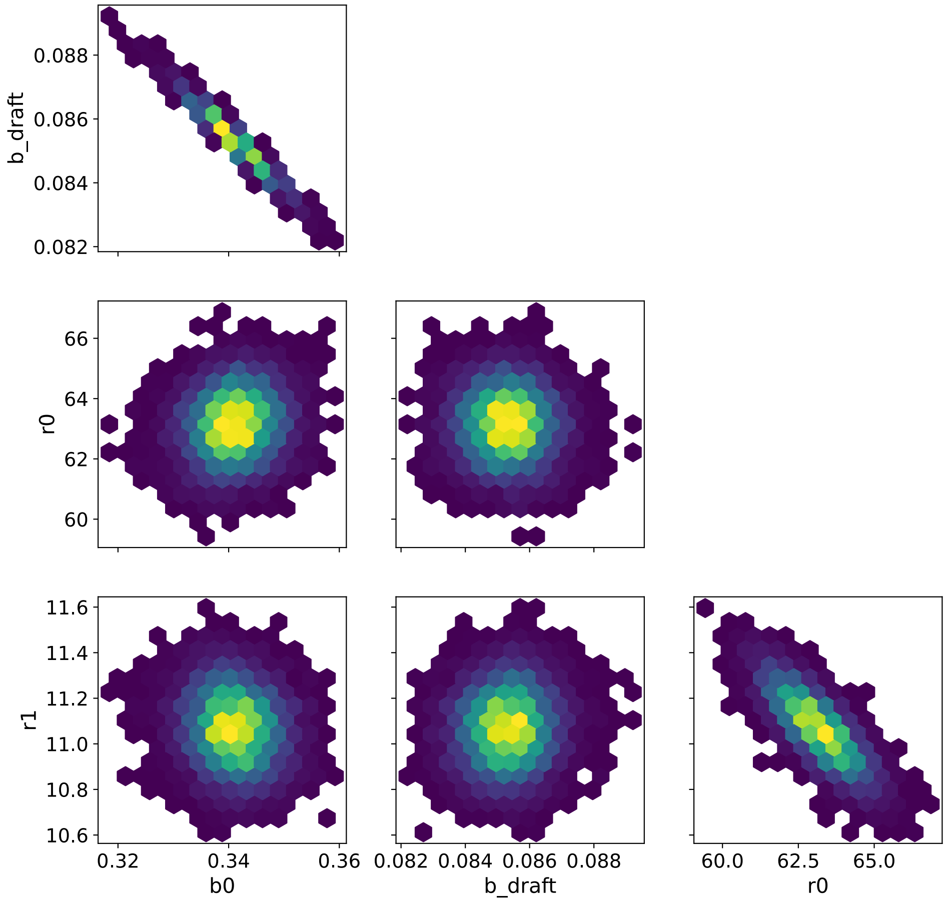

Figure 6: Joint probability distributions of the fitted speed power model for dataset 1

More relevant to the application, is the translation of the likelihoods of b_0 and b_d into the corresponding speed power relations for operational and ballast draft. This is shown in Figure 7, including the uncertainty bands.

Instead of treating all vessels as identical, we can also fit a model where each vessel has its own b_0, while sharing all other parameters with the sisters. The resulting b_0 estimates are shown in Figure 8. Sister 2 is clearly an outlier. This could be because confounding factors are not sufficiently compensated. For example; the vessel might be sailing in consistently worse conditions than the others and this leads to a bias in the results. However, it could also indicate a maintenance issue with either the instrumentation (too low speed or to high power reported) or e.g. fouling.

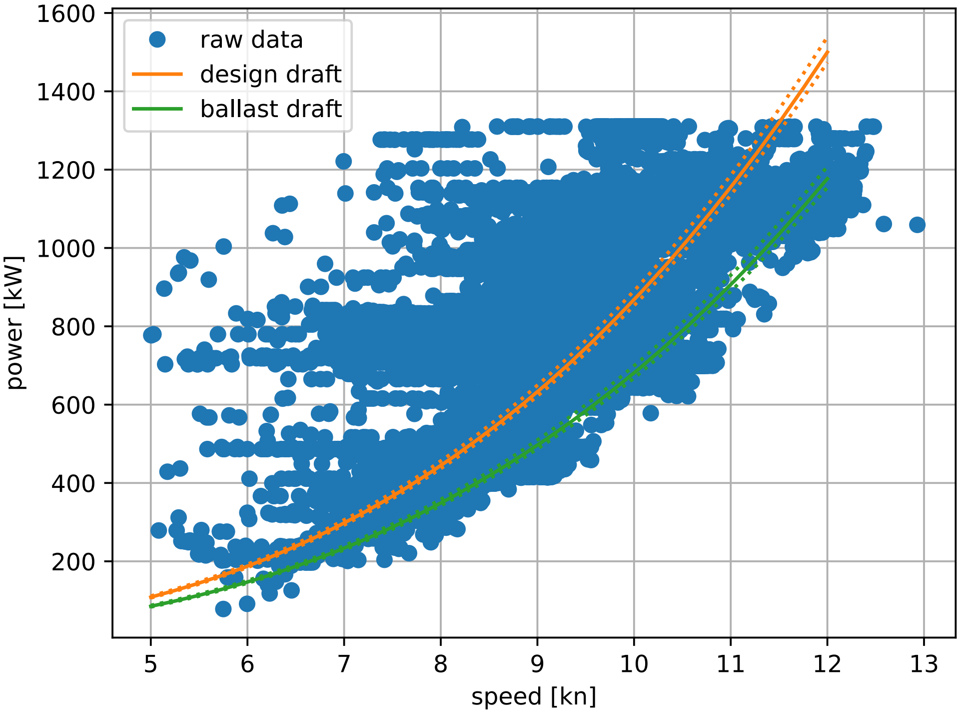

Figure 7: Resulting estimated ballast and operational speed power curves including error bands

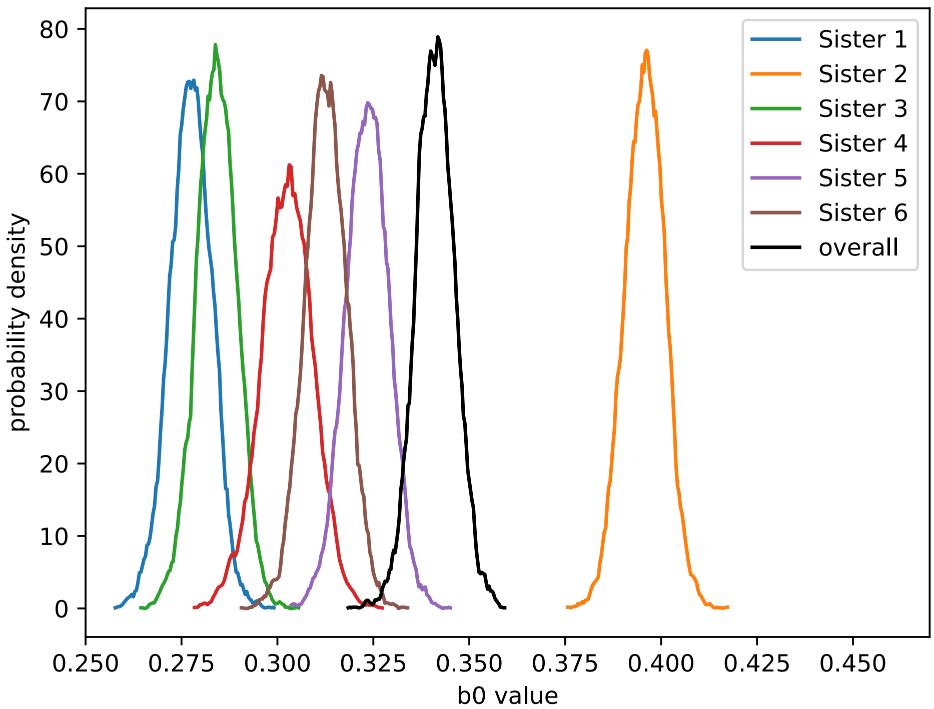

Figure 8: Speed power coefficient likelihood per vessel and overall

5.2 Dataset 2

The second dataset contains data of one general cargo vessel. It is a direct drive vessel without shaft torque measurement. As such we need to rely on the fuel flow measurement. Based on information from the opperator, the recorded fuel flow rate contains the net instantaneous flow. As such, it can be used as a measure for the required power.

5.2.1 Window selection

After deriving the 20 minute characteristics as described in Section 3.2, we first check the probability density curves for the most relevant signals:

- The speed through water
- The shaft rpm
- The course over ground

The corresponding plots are shown in Figure 9.

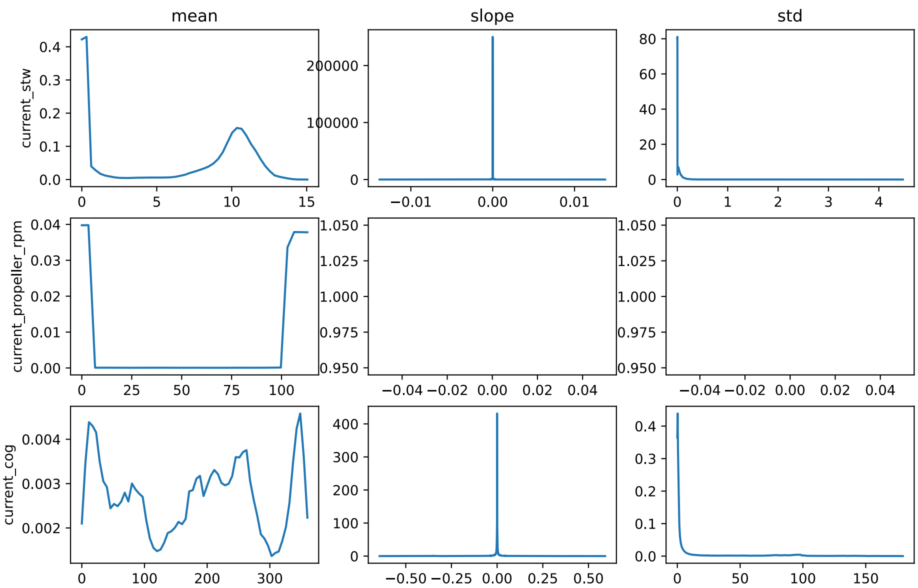

Figure 9: Dataset 2, 20 minute window characteristics distributions for all windows

Based on these results we set the minimum speed at 5 knots and the minimum shaft speed at 80 rpm. In addition we apply maximum standard deviation for the course over ground (2 degrees).

After this selection we can look at the distribution of the remaining 31 thousand samples. The corresponding plots are shown in Figure 10. There are no significant trends (slope) in the remaining samples. And we use this set as the stationary sailing samples for further analysis.

The raw speed power points are shown in Figure 11. Clearly, the amount of scatter is too high for straightforward interpretation.

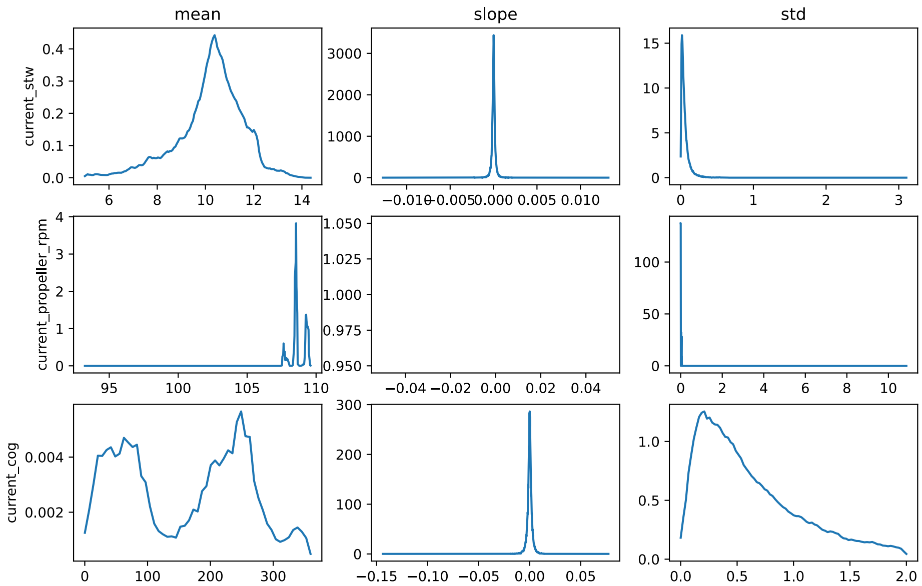

Figure 10: Dataset 2, 20 minute window characteristics distributions for stationary sailing windows

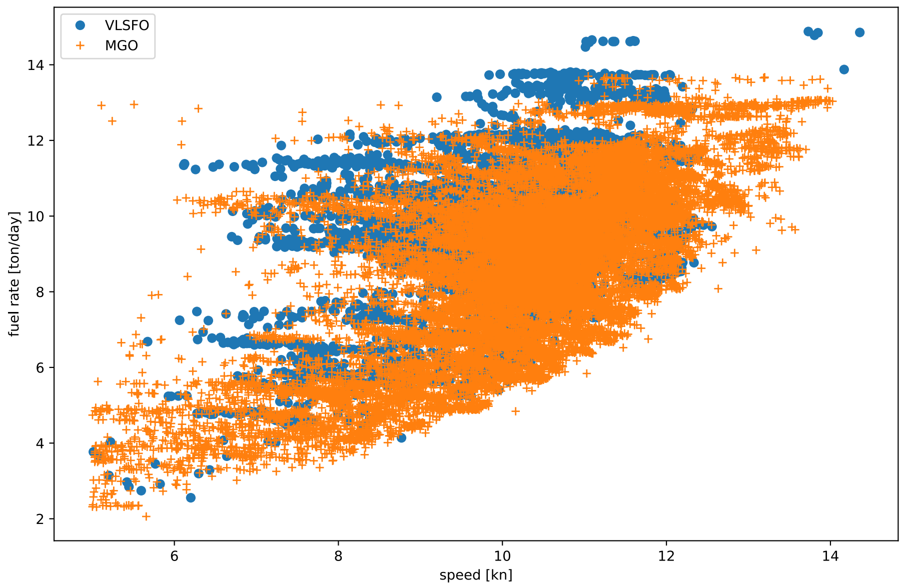

Figure 11: Raw stationary sailing speed power data points per fuel type

5.2.2 Speed power fit

As a first check, we use the speed power law (Equation 2) to determine the per sample b coefficient:

$$b_i = \frac{p}{v^3} \tag{7}$$

We use the fuel flow rate as a substitute for the power. If we plot this per fuel type over time (Figure 12), we can check for obvious trends. In this case there are none.

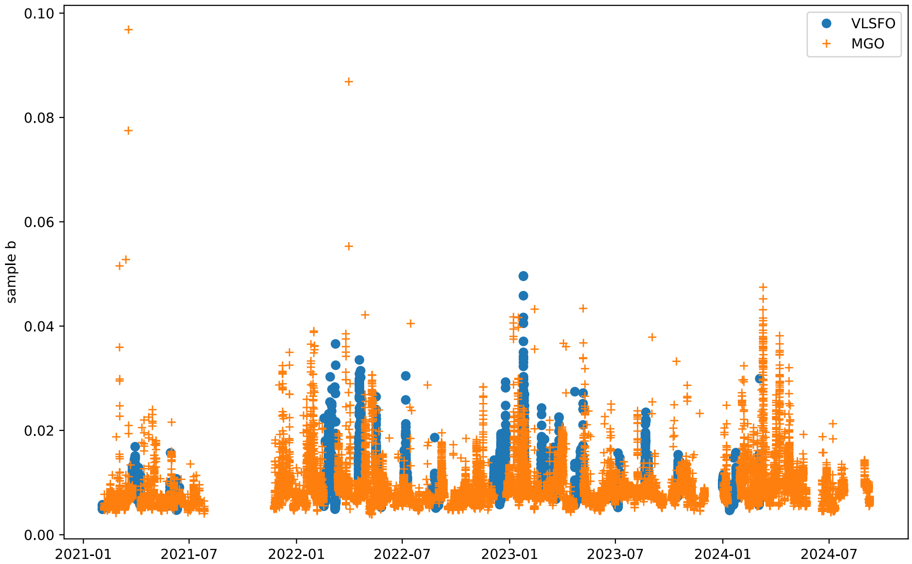

Figure 12: The speed - fuel rate coefficient over time per fuel type

After fitting a speed power model with wind speed and draft dependencies (Equation 3 with Equation 4 and Equation 5), we can compare the maximum likelihood fit to the raw data. This is shown in Figure 13, including the uncertainty bands. Towards the lower speeds, there is a clear misfit. While the fuel rate should be a good indicator for the engine power, the engine power may not be a good indicator for the propulsion power. In this particular dataset, we also have a recording of the power take off by the shaft generator. This part of the engine power is used for the auxiliary systems and not for propulsion.

To take this confounding factor into account, we extend the speed power model (Equation 3) to compensate for the shaft generator power:

$$p_i = N(\mu = as_i + bv_i^3, \sigma = r) \tag{8}$$

The factor a is needed, because the shaft generator power and the fuel flow rate have different units. The shaft generator power and the propulsion power are sufficiently uncorrelated in the data set to get a reliable fit of the constant a.

With the best estimate a value, we can subtract the influence of the shaft generator from the fuel rate. Figure 14 shows the compensated speed power together with the extended model fit curves. Compensating for the shaft generator leads to a much better fit.

Based on the extended model, we can also fit separate b_0 values for each fuel type. The energy density of the fuel will vary and we should at least account for the known differences. The resulting b_0 estimates are shown in Figure 15. The VLSFO fuel is significantly less efficient in terms of flow rate to speed than MGO. (Cost and environmental impact could reveal a different picture.)

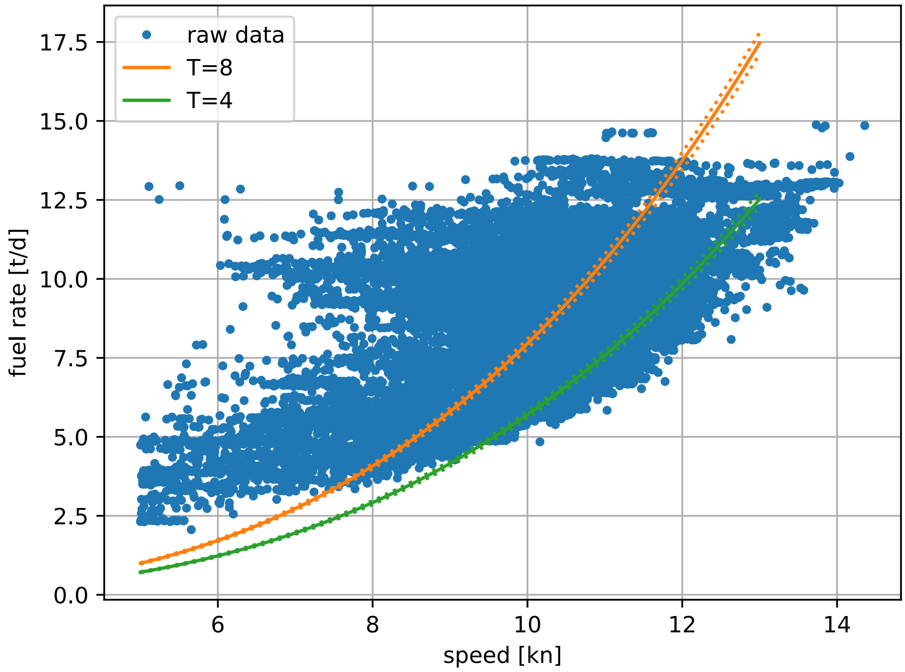

Figure 13: First estimated ballast and operational speed - fuel rate curves including error bands. At the lower speeds, there is a clear misfit

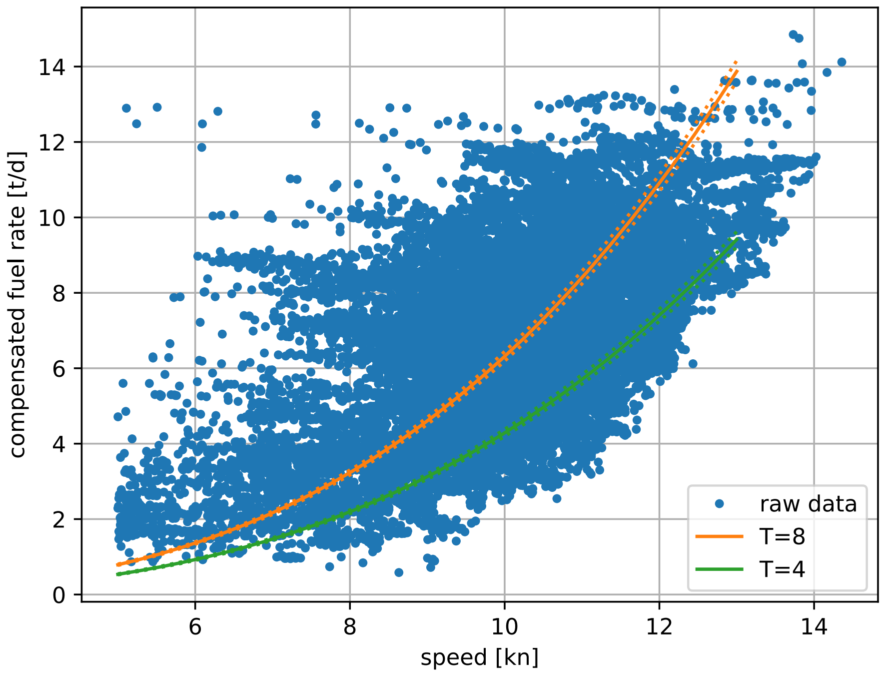

Figure 14: Estimated ballast and operational speed - fuel rate curves including error bands. The fuel rate is corrected for the shaft generator power, leading to a cleaner fit

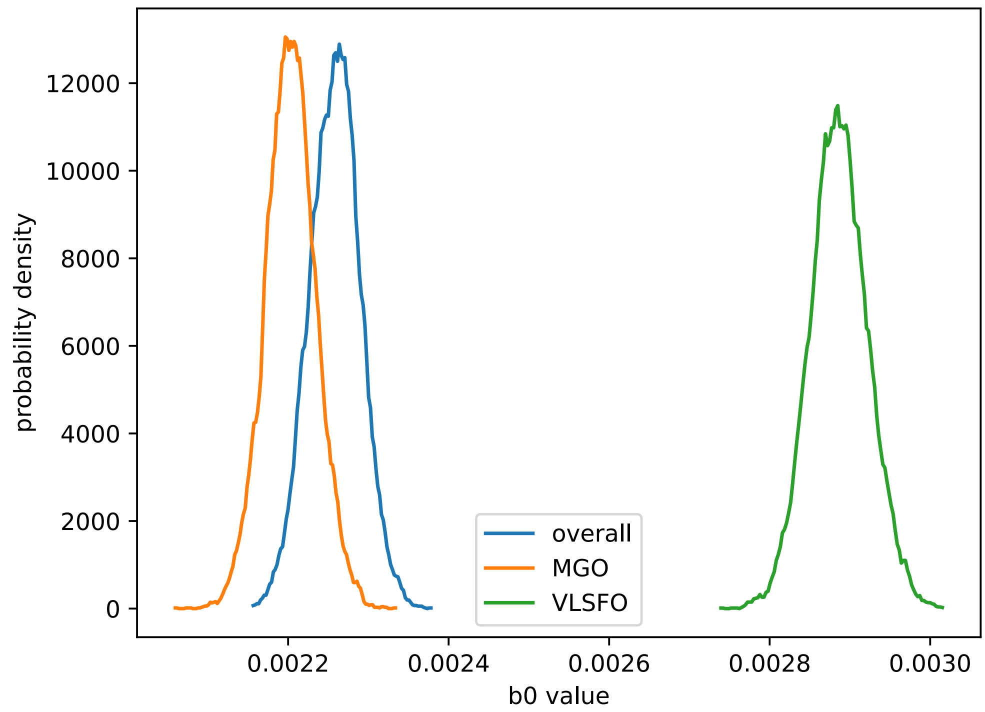

Figure 15: Speed - fuel rate coefficient likelihood per fuel type and overall

6 Conclusions

In this report we presented a Bayesian approach to determine the speed - power relation from operational data. The results show that operational data can be used to determine the speed power relation as well as changes in this relation:

- We can distinguish design and ballast draft
- We can reveal outliers between ships

However, there are practical concerns:

- The scatter in the operational data is large, as such long term measurements are needed to get a reliable fit
- Good communication with the vessel operator is needed to ensure the correct interpretation of measured values
- A good model helps in suppressing the impact of the scatter, but to set up a good model, understanding the dominant underlying physics is crucial
- Part of the scatter can usually be explained by confounding factors. Where possible, the model should include the dominant confounding factors. This requires some input on the time dependent magnitude of each confounding factor. If the confounding factors are sufficiently uncorrelated to the speed - power, we can suppress the related scatter in the fit. In this report we corrected for:
  - Draft variations
  - Weather induced uncertainty, using the wind velocity as a measure of the magnitude
  - Shaft generator power take off, by fitting its influence on the fuel flow rate

Wageningen, March 2025

MARITIME RESEARCH INSTITUTE NETHERLANDS

dr. ir. H.J. Prins,

Manager Research & Development

MARIN
P.O. Box 28

6700 AA Wageningen
The Netherlands

T +31 317 49 39 11
E info@marin.nl

I www.marin.nl