Using a sprint approach to analyse offshore maintenance records

Chris Bell, Consultant, DNV GL, Cromarty House, 67–72 Regent Quay, Aberdeen, AB11 5AR

Matthew Celnik, Principal Consultant, DNV GL, 7 West Nile Street, Glasgow, G1 2PR

DNV GL presented at the Hazards 29 conference, a computational method for analysing large numbers of offshore Safety and Environment Critical Elements (SECE) maintenance records. That demonstrated a computationally efficient way to analyse thousands of records using machine learning techniques in a few seconds accurately. The verifier or maintenance and reliability engineers could then apply a focussed approach to verification, reducing manual effort and potential for human error or bias in sampling. This increased project efficiency and allowed DNV GL to make more useful findings and recommendations to the operator.

Following on from that work, DNV GL worked with a larger dataset related to operation of the fire and gas (F&G) detection, fire water and related systems. The objective was to facilitate a "systems understanding" by considering the constituent SECEs and how they interact. Those SECEs include:

- F&G panel / logic;
- Emergency generator;
- Diesel generators for fire pumps;
- Emergency power distribution board;
- Essential service distribution board; and
- Firewater pumps.

We used a cross-industry process for data mining (CRISP-DM) sprint methodology to analyse available data. CRISP-DM is an iterative approach whereby the team works closely with the data owner in short "sprints" to explore the data sources, determine potential outcomes and develop models for data analytics, including the machine learning model presented previously.

This paper details the CRISP-DM process carried by DNV GL in conjunction with the asset operator. We combined analysis of the manual inspection/maintenance records and the automated system logs. This demonstrates how analysing records at a system level can give valuable insight to both the verification teams and the operator of offshore assets, by utilising all available data sources.

Introduction

Today's increased focus on digitalization is linked to the oil and gas industry's drive toward greater cost efficiency and understanding of systems. Operators with successful programs must think in terms of total life cycle costs economics and should build teams that will benefit from the application of such technology.

The increase in digitalisation across the oil and gas sector offers offshore operators the chance to automate high-cost, error-prone tasks in which the cumulative effects of inconsistency and analytical error can adversely impact safety. For instance, as part of any assets' assurance process, it can be instructive to review maintenance records for insights, particularly trending issues and identifying potential improvements. The goal is to ensure the asset is performing safely and effectively, with high reliability while adopting the most cost-effective strategies for all maintenance work.

Following on from the preliminary data review (Celnik and Bell 2019) DNV GL was asked to analyse a larger dataset related to operation of the fire and gas (F&G) detection, fire water and related systems on the same offshore installation. The operator wanted to improve their "systems understanding" by considering the constituent SECEs and how they interact. This requirement is necessarily vague; the client was not sure what data they had, nor really what they should be looking for. This makes planning a traditional project, with defined objectives and deliverables, very difficult.

Instead, we proposed to use a CRISP-DM (cross-industry standard process for data mining) sprint methodology. The sprint approach requires close working and regular feedback with the client, to ensure they derive maximum benefit from the process. This paper describes the sprint methodology in detail, as well as presenting some of the finding of our specific sprint.

Methodology – the CRISP-DM approach

Prior to the project kick-off, there was no defined objective end point for this work, as neither the asset operator nor DNV GL had prior view of the available data sources or preliminary client queries. As the objective was loosely defined, we chose a CRISP-DM sprint methodology to work with the client to analyse the available data. CRISP-DM is an iterative approach whereby the team works closely with the data owner in short "sprints" to explore the data sources, determine potential outcomes and develop models for data analytics. The key cycle steps as shown in Figure 1 are:

1. **Business understanding**: develop and understand the business/user requirements.
2. **Data understanding**: review and understand the data; what anomalies are present in the data, and can it be used to provide insights?
3. **Data preparation**: prepare the data to be analysed; remove anomalies, reshape as required, merge with other sources etc.

4. **Model preparation**: develop algorithms and models to read the prepared data, and provide statistics, plots or insights to be evaluated.

5. **Evaluation**: evaluate the model's effectiveness and ability to provide the answers required? Does it raise more questions that should be answered before proceeding further?

Figure 1: CRISP-DM key cycle steps

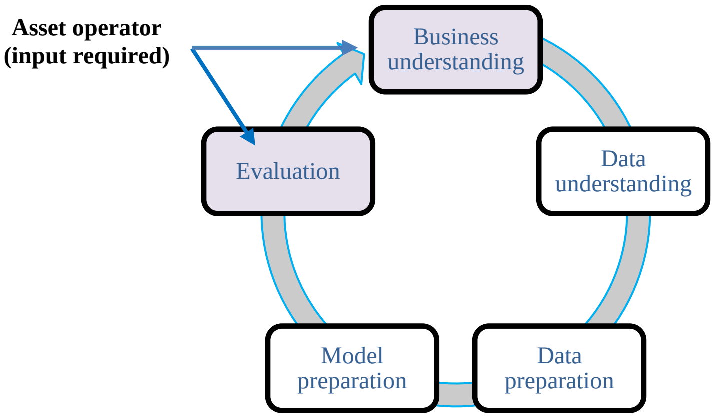

After each evaluation step, the team reviews the business requirements with the client. Hence, the requirements are subject to change during the project life cycle. This is a challenge for a traditional project setup, which may suffer from hierarchical structure and an inability to control scope creep. Also, most engineering organizations are not structured to facilitate such approaches; they are not *agile*.

The software industry has applied agile development techniques for a long time with great success. Here we demonstrate application of a similar approach for data mining to handle the change in business requirements as the project evolves. We believe this works well, and is an efficient and effective method to answer critical business questions. Essentially the project team can shortcut management debate cycles and reduce requirements gathering and solution design into a single week.

Application

Here we describe a recent project undertaken by DNV GL which applied the sprint approach.

The project schedule included a 4-day iterative sprint with all team members in the same place. The table below gives the team members.

Table 1: Team roles and requirements.

	Person	Role	Availability
1	**Sprint lead**	To facilitate and co-ordinate the sprint activities, with ability to help with the data preparation, model preparation and evaluation	4 days (full attendance)
2	**Operator technical authority (TA)**	To develop the understanding and requirements for the operator's business	4 days (partial attendance) - Kick-off meeting - Daily evaluation meetings
4	**Operator data manager**	To provide a single point of contact between the project team and the client data stores. This person had the ability to request access to data held by the business as the project evolved.	4 days (full attendance)
5	**Verification manager (DNV GL)**	To develop the understanding and requirements for the ICP.	2 days: - Kick-off meeting - Summary meeting
6	**2 × DNV GL data analysts / domain experts**	To carry out the data preparation, model preparation and evaluation in conjunction with the operator. Ideally the data analysts and domain experts are the same people, but as these are specialist skills usually they are not.	4 days (full attendance)

The CRISP-DM timeline was as follows:

- Half-day kick-off workshop with asset operator to develop the business understanding and agree sprint objectives. This stage identified the project success criteria, that is "what does 'good' look like."
- Three-day iterative sprint to explore the data; prepare and test suitable models and evaluate.

- Final evaluation meeting with operator to discuss results and agree next steps.
- Preparation of summary report.

Potential outcomes identified prior to kick off included comparison of perceived versus actual systems performance (as determined using data analytics model), a review of system availability and recommendations to enhance verification activities.

The following subsections provide details of each sprint stage during the project, noting that each was repeated multiple times during the week.

Business understanding – defining the scope

For the sprint to be successful, it must deliver something of value to the client. Hence, we need to understand the client's needs. The business understanding steps enables us to do this by asking the following questions:

1. What output do we want from the analysis or tool?
2. What value can be gained by the business in getting that output?

These two high level questions are essential in driving the scope of work and the project life cycle. During the kick off meeting, we identified initial sprint outcomes with key stake holders. We generated three "user stories" to guide the sprint data analysis steps. We reviewed the user stories daily with the client.

User stories are key elements of an agile approach to software development but are well applied to data analyses too. The user stories help the entire team understand what data to target for analysis, for whom, why, and when. They also help non-technical team members understand the purpose of the analysis and encourage their participation. They define the desired outcomes in terms of the main stakeholder requirements. Our sprint used Microsoft Azure DevOps (Microsoft 2019) to create and track tasks assigned to each user story, which tool is more commonly used to develop software products. Many standard practices and tools available for software development transfer easily to a "traditional" engineering context.

Using this method, the sprint team can think assign priorities to each task which reflect the expected value for the user, complexity, dependencies and other business requirements. There is no need to mark items as being "out of scope", instead, the team can assign a lower priority to those ideas that are difficult to achieve within the sprint schedule or less related the outcome required, while moving up the user stories that are important. This enables future work and knowledge gaps to be captured, which can facilitate future management of change.

The table below shows the user stories developed for this sprint.

Table 2: CRISP-DM user stories

	Stakeholder / user story	**Details**	**Potential outputs**
1	**Operator technical authority (TA)** wants a better understanding of how the firewater system is operating.	Specifically of interest is the component availability and reliability, and the system-level availability accounting for all components. In addition, a better understanding of the information in the maintenance management system (the work order history) is desirable.	Availability figures for each component, and time-series plots showing overlapping periods of component outage. (**High priority**)
2	The **operator TA** wants a better understanding of how failures, inhibits and isolations are reported (or not).	In particular, the TA wants to identify knowledge gaps to focus offshore workforce record keeping. Potential issues include: • Master key left in the inhibit panel • Paper records not kept or not fully kept • Failures logged as passes, e.g. fail-fixes, potentially due to difficulties using the management system (it might suit the workflow, e.g. mark as a "pass" then raise a corrective order). There may also be system interaction issues, e.g. operator cannot check detectors situated near generators when running; power generation is prioritized over detector assurance.	Identified areas where there are unrecorded faults. For example, long inhibits or process system faults not corresponding with maintenance records. (**Medium priority**)
3	The **independent verifier** wants to focus effort on deficiencies in the assurance process.	Potentially data analytics can: • Take away some manual effort (no need to manually review data) • Focus attention on potential issues	List of items to review, based on the outputs above. (**Low priority**)

Using the above stories, we split work amongst the team into three main areas:

1. Apply the machine learning (ML) classification algorithm from the preliminary work (Celnik and Bell 2019) to other systems;
2. Using the process system logger data, calculate statistics related to inhibits, alarms and faults logging by the system; and
3. Using process data and work order history, plot equipment availability history relative to the performance standard criteria.

While we did identify other, lower priority, tasks, this sprint did not consider due to time constraints. They were nevertheless tracked in Azure DevOps for future work.

Data Understanding, data preparation and model preparation

We made extensive use of Python (version 3.6) (Python Software Foundation 2016) coding in Jupyter Notebooks (Project Jupyter 2019) for the data analysis. We managed all code using a Git (GIT n.d.) repository, another tool repurposed from software development which enables a distributed team to work on the same code. This paper does not discuss in detail the code implemented or the management of the code development.

The operator provided the following datasets to be used in the sprint either immediately before the sprint or during:

Table 3: Data provided for analysis

No.	Dataset	Used in CRISP-DM?
1.	Block diagram of fire water system parts	Yes
2.	Work order history, all items, 2012-2019	Yes
3.	Process system event logger (ESD) September 2018 to August 2019	Yes
4.	Process system event logger (F&G) September 2018 to August 2019	Yes
5.	Equipment tag list, all items	Yes
6.	Tags for key equipment in fire water system	Yes
7.	Work order description list	Yes
8.	CRO handover logs, July 2019	No
9.	Weather history January 2017 – August 2019	No

The equipment tags tie together all the datasets, though the operator only provided the tag list on sprint day three. Nevertheless, we were able to reconcile the datasets but identified two issues with the tags:

1. They are not grouped by system, so required further processing to assign to system categories; and
2. The datasets do not use a single consistent tag format.

Table 4: Principal tag formats of fire and gas system components

Equipment	Tag format (example)
Gas detector	Asset-GD-XXX
Flame detector	Asset-FD-XXX
Heat detector	Asset-HD-XXX
Smoke detector	Asset-SD-XXX
Methanol detector	Asset-MD-XXX
Emergency generator	Asset-EG-XXX
Polar generator	Asset-PG-XXX
RGT generator	Asset-RGTG-XXX
Solar generator	Asset-SG-XXX
Emergency SWB	Asset-PESWG-XXX
Mains SWB	*Unknown*

Equipment	Tag format (example)
FW pump	Asset-FWP-XXX
FW ring main	*Unknown*
Jockey pump	Asset-JP-XXX
Switchgear control & protection (UPS)	Asset-PU-XXX, Asset-UPS-XXX
Telecommunications (radios & SOLAS) (UPS)	Asset-RSU-XXX
Telecommunications (PAGA, PABX & helideck CAP437 lighting) (UPS)	Asset-PAG-XXX
ESD/SCADA/F&G (UPS)	Asset-ESD-XXX

Using this information, with additional text interpretation code, the tag list was categorized by system.

This allowed a quick comparison of related items. Only the items above were used; items such as sensors and valves were not categorised. These items were instead considered when reviewing the inhibits, alarms and faults in the process system logger (ESD) logs.

The process system logger data is more complex, as the records do not contain any tag information. Instead, the tags were inferred from the records where possible. For example, the table below gives typical log records for each equipment group:

Table 5: Example process event log for F&G system

Equipment	Example process log (F&G) record			Resultant tag
	Event date time	**Plant area**	**Tag description**	
Gas detector	01/07/2019 01:20:14	ZONE **103**	GD-**C** HIGH ALARM	Asset-GD-XXX
Flame detector / IR LOS detector *	01/07/2019 23:29:52	ZONE **1**	IR-**B** FAULT P.C.B.	Asset-FD-XXX
	07/05/2019 10:11:05	ZONE **4**	IR-**A** ALARM	Asset-FD-XXX
Heat detector	01/07/2019 08:28:57	ZONE **10**	HEAT-**A** FAULT P.C.A. CLEARED	Asset-HD-XXX
Smoke detector	07/07/2019 10:17:00	ZONE **206**	SMOKE-**A** ALARM	Asset-SD-XXX
Emergency generator	06/07/2019 07:22:36		**EM. GEN.** RUNNING	Asset-EG-XXX
FW pump	02/07/2019 18:05:05		F/P-**A** AVAILABLE	Asset-FWP-XXX
Jockey pump	01/07/2019 22:37:32		JOCKEY PUMP **A** AVAILABLE	Asset-JP-XXX

** There are several tags associated with flame detectors and it is not clear from process system logger which are correct. For example, in zone 1 potential tags include; "<asset_name>-FD-001A|B", "<asset_name>-FD-002A|B", "<asset_name>-FD-003A|B". The process system logs do not enable precise identification of the IR/flame detector tags.*

Only the rows containing the equipment groups listed above were processed during the sprint. The equipment tags used in this dataset often did not match the format used in the other data sources. For example, a pressure indicator "asset1-PI-001A" in one data set may appear as "asset1-system1-PIH-001-A" in the other dataset. This required additional data sources to be reviewed and assumptions to be made when merging data across different datasets.

Work order review

The figures below show a heat map of work order counts (left) and the weighted counts by number of items/tags (right). While there are relatively more smoke detector work orders, there are also many detectors in service, hence the weight counts are lower. The mains power and firewater ring main show a relatively higher number of work orders per tag than other systems.

Figure 2: Work order counts per system

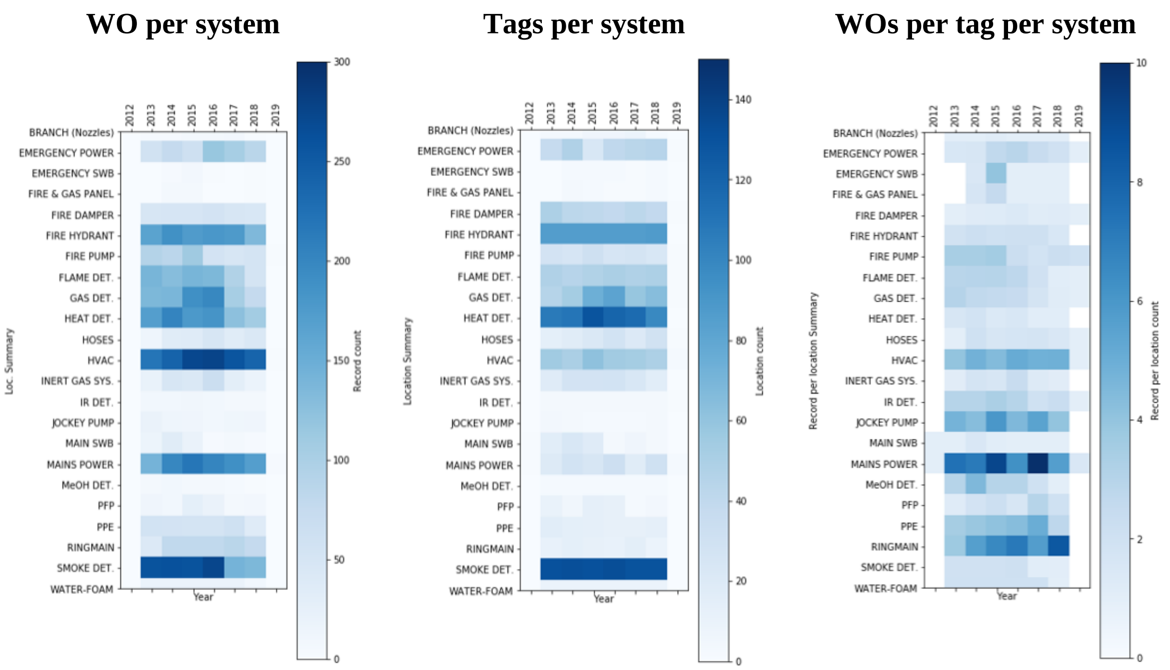

Deferrals

Considering deferrals (deviation of target finish date from actual finish date), we see that heat detectors and mains power have the largest number of deferred work orders. These values are not weighted by tag. There are relatively many heat detectors, so perhaps the deferral rate for these items is less critical. However, the deferrals rates for mains power and HVAC are substantially higher than average, especially given the relatively few equipment tags in those systems.

Figure 3: Work order deferrals over 6 months

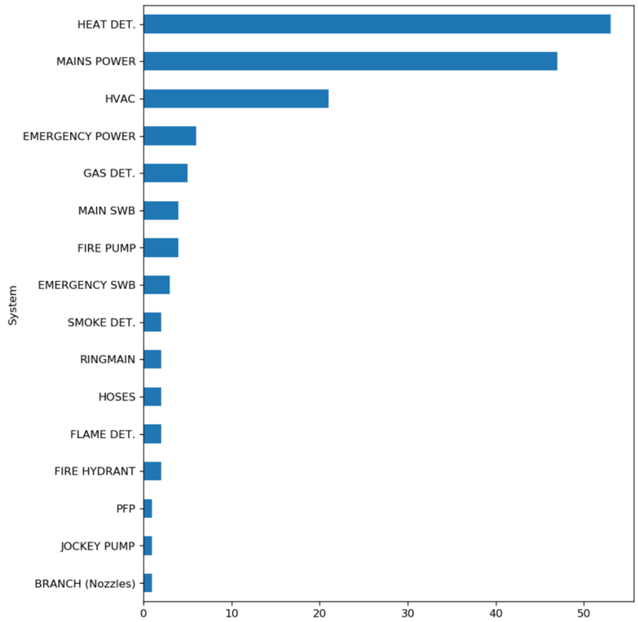

Pass/Fail classification

We repeated the work order classification work from the preliminary study (Celnik and Bell 2019), this time applied to all equipment items, not just detectors.

The figure below shows the initial distribution of PASS/FAIL records. As before, the data is highly skewed towards PASS. This requires additional steps in the classification algorithm to prevent bias.

A significant number of records are unclassed (NULL below), meaning they cannot be used directly to estimate availability or reliability. The ML algorithm predicts PASS/FAIL for unclassed work orders.

Figure 4: Initial (as reported) PASS/FAIL counts

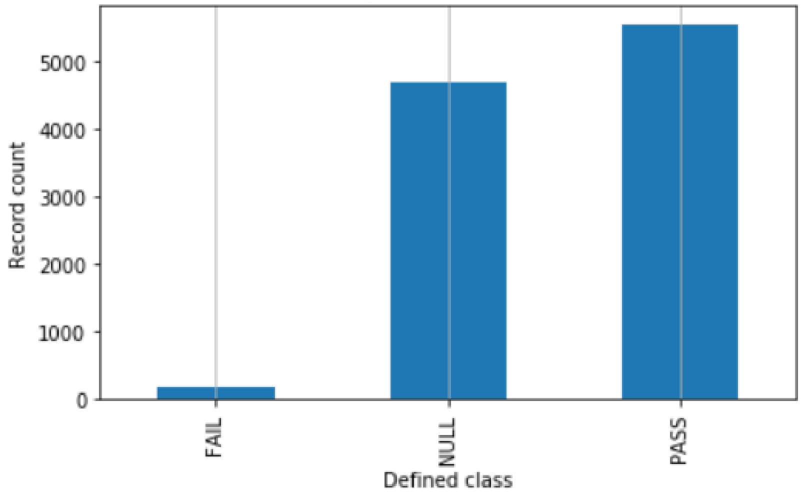

We again tried several models and cleaning workflow to achieve an acceptable prediction accuracy of the classification model. Figure 5 below presents results for the most suitable model demonstrating a 96% accuracy when comparing against the entire WO dataset to the as-reported values, including the test data. We did not manually verify the as-reported values, so the model could be mirroring inaccuracies in the original data.

Figure 5: Prediction accuracy of best PASS/FAIL model

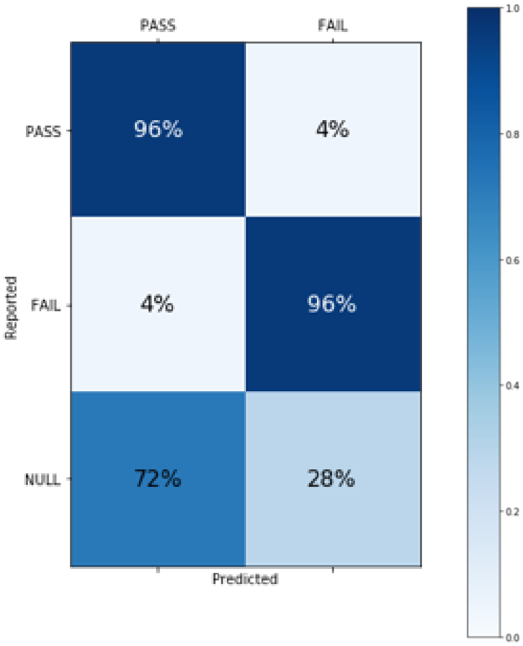

Using these predictions, we could estimate availability of equipment over time as done for the preliminary study. However, there are some caveats with that approach:

- Other than assurance routines, it is not guaranteed that a PASS or FAIL record means the system is left in a functioning state. It may simply mean the work order was completed as per the description.
- The work order history does not relate directly with the performance standard criteria, hence gives a measure only of component availability, not system availability.
- The work order history as provided did not include creation dates. In principle, we could have considered an equipment item as unavailable between a corrective order creation date and its completion date, should that information be available.

Because of these caveats, this work instead attempted to determine availability using the process system logs (see later section).

Work order timeline

To aid understanding of the work order history, we attempted to create timeline figures to illustrate live work orders, and periods of unavailability (after failed assurance routines). The below figure shows the plot generated for the fire pumps.

These timelines enable rapid identification of key features of the work order history, which can aid the operator and the independent verifier. For example:

- Periods of high activity are identifiable.
- Failed assurance routines are clearly marked, along with periods of impairment/unavailability.
- Consecutive periods of high activity or unavailability across related equipment items are clear, which can then be compared to the performance standard criteria.

Figure 6: Work order time line

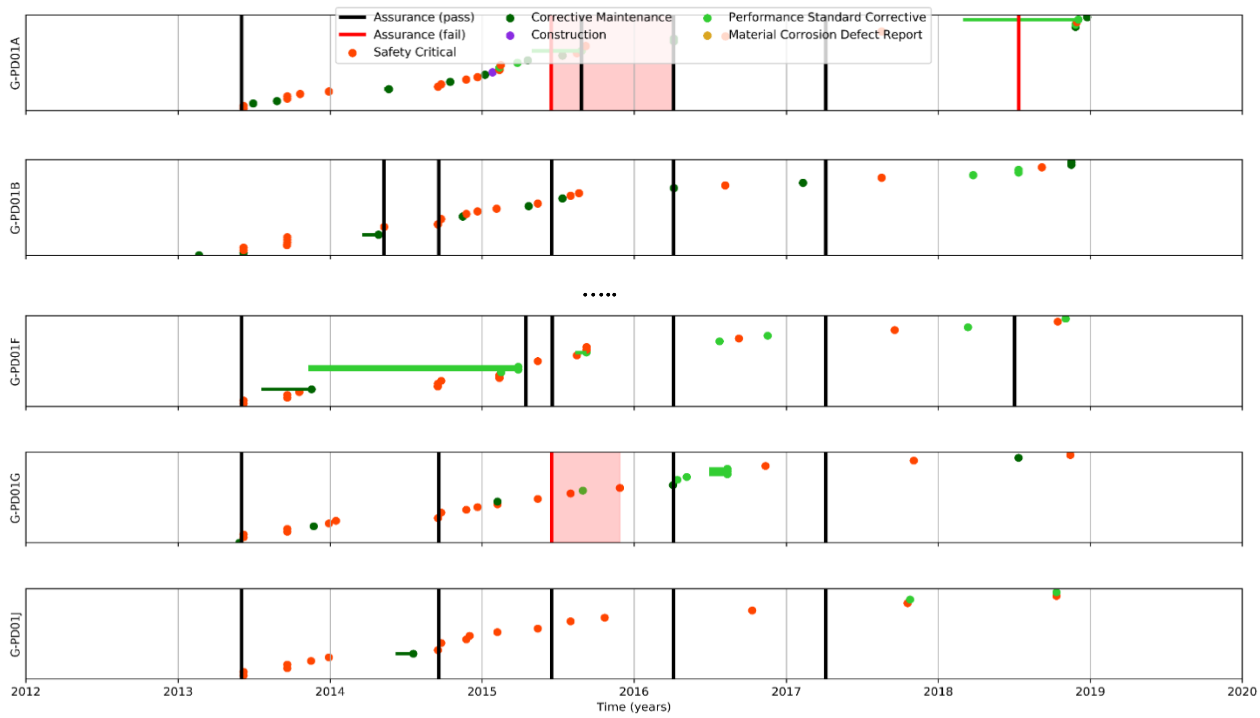

Each dot represents one work order, coloured by type. They are separated vertically only to make them easier to see, otherwise the y-axis scale has no purpose. Lines preceding a work order illustrate the difference between the actual start and finish dates (if the dataset included creation dates, this could instead illustrate the difference between creation and finish dates to show unavailability). Vertical lines are the assurance routines, coloured black for PASS and red for FAIL. The plots are shaded red between reported FAIL records and the next reported PASS; as there are many records with a PASS/FAIL classification, this shows only a few some periods.

The figure below shows the same data after running the ML classification algorithm. This presents a more complex picture of availability, potentially with non-assurance work orders resulting in a FAIL, though we must still consider the caveats noted above when interpreting this data.

Figure 7: Work order time line (after ML pass/fail classification)

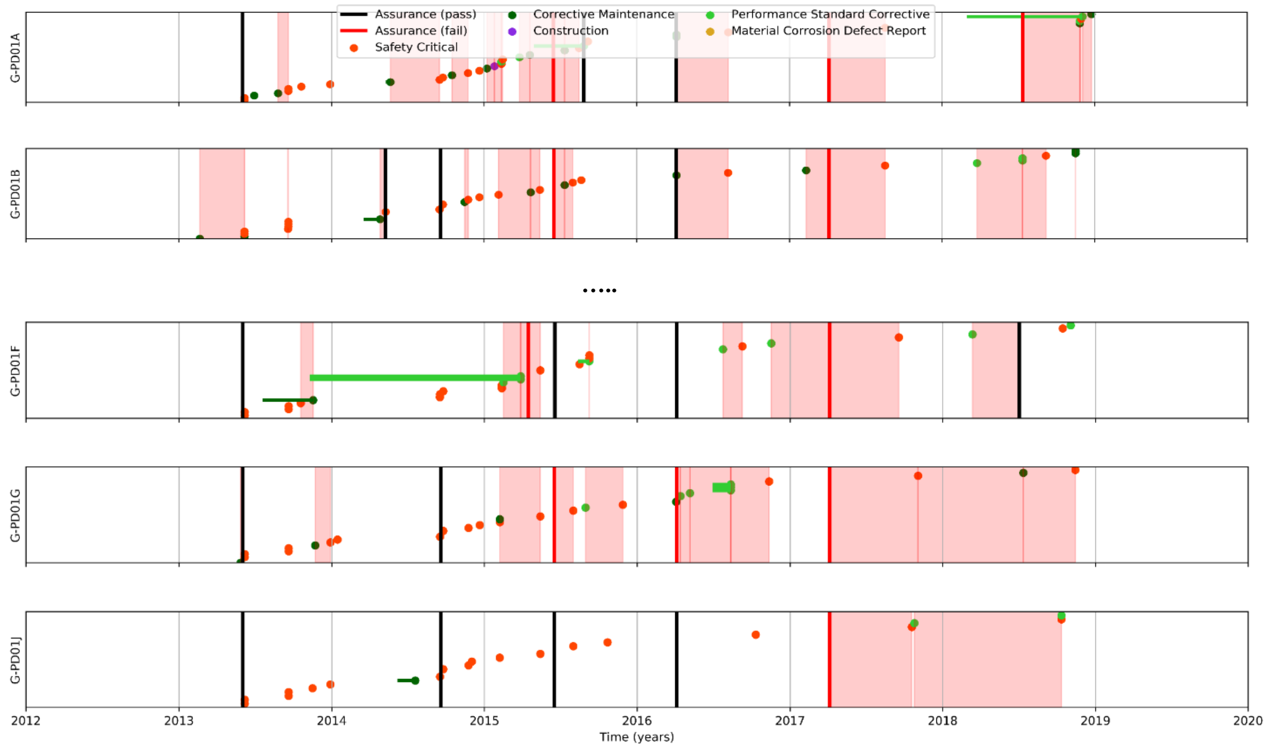

The above figure also illustrates a mis-prediction by the ML algorithm. The 2017 assurance test shows failure on all nine pumps. However, the relevant work order text suggested they passed (though did not explicitly state as such). The text for all pumps was the same:

> "5/4/17 SC AB
> Fire pumps tested with 100% & 150% orifice plates to obtain a flow rate of minimum 280 m3/ph at min 12.5 Bar"

We also note, the 2017 assurance routines were initially unclassed PASS/FAIL in the work order history. Either way, this information is useful to a verifier to check the records, highlighting useful application of the above figure.

Inhibits review

To identify long inhibit durations, we processed the process system (ESD) logs. The figures below show the number of inhibits at different durations. As expected, the majority of the inhibits are short duration, under one day, though some last several months.

Figure 8: Inhibit summary

(a) Inhibit duration summary **(b) Inhibit frequency summary**

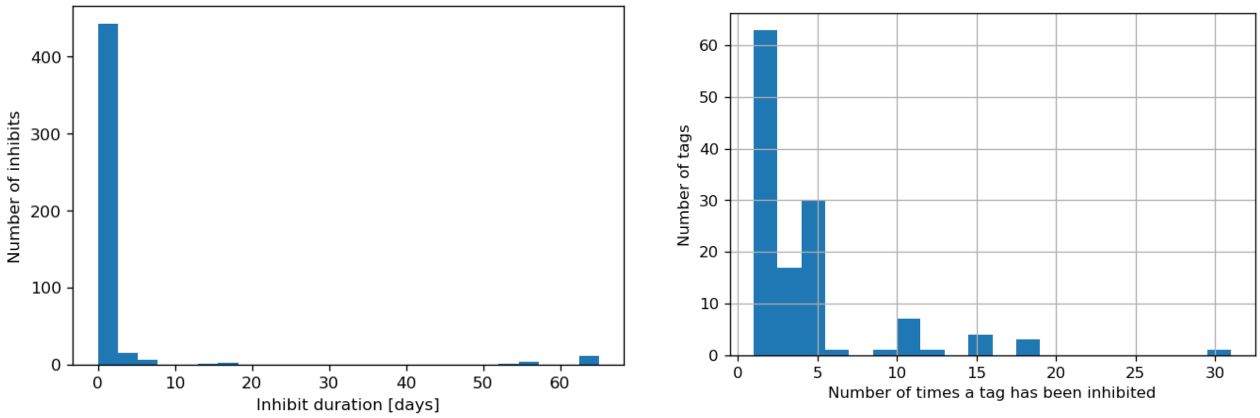

From the datamining carried out the authors were able to list the longest inhibited equipment (~66 days). The tags appear related, so this is most likely related to a long-duration maintenance job. Most items are infrequently inhibited, 1 to 5 times in the year. However, some items are inhibited much more frequently. The most frequently inhibited item (31 times in the year) is the pressure switch low on one of the asset risers.

We also plotted inhibits by group. Most inhibits are in the ESD class 2 group, followed by the compressors.

Figure 9 below shows inhibits by time of day. There are clear peaks around the start and end of the day shift. This suggests inhibits are left on for the shift duration rather than for as short a time as possible. There is also a peak at 04.00 which is unexplained.

Figure 9: Inhibits by time of day

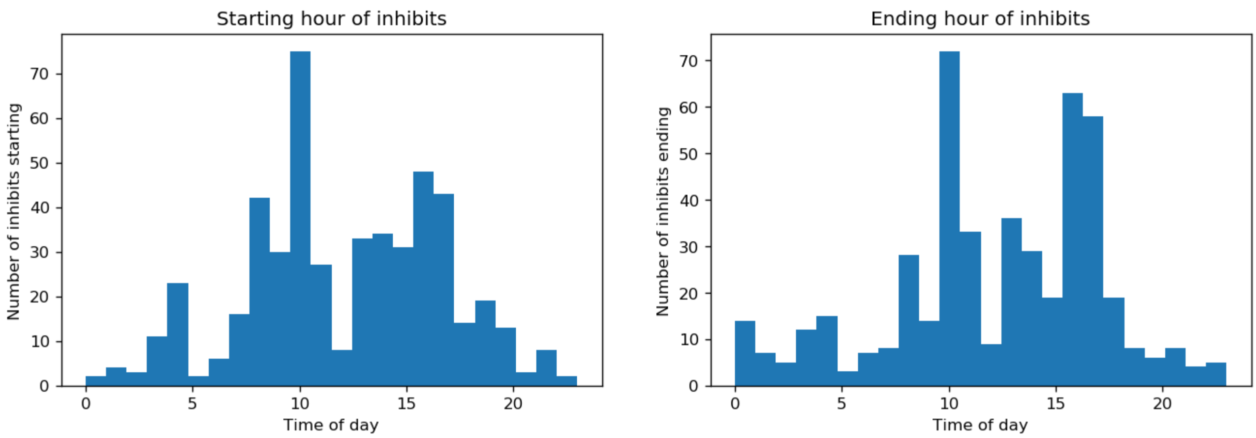

Evaluation

F&G detector faults

To investigate unreported faults, we started by analysing the process system (F&G) logs for the detectors. The detectors have two channels each; we assume both channels must fail to make the detector unavailable – this complicates log interpretation as the logs include failure messages for both A and B channels.

Figure 10(a) below shows the F&G detector failure duration distribution. Most failures are short, a few seconds or minutes, but some are very long with the longest (a smoke detector) lasting almost a month.

Figure 10(b) below shows the failure count distribution, to identify detectors with the highest fault frequency. Most detectors demonstrate a low fault count, though many have over 50 faults (approximately one per week), and some far more. The most frequently in-fault detector is a gas detector in zone 1 with 352 faults; this is almost one per day.

Figure 10: F&G detector fault plots

(a) F&G detector fault frequencies by duration **(b) F&G detector fault frequencies by count**

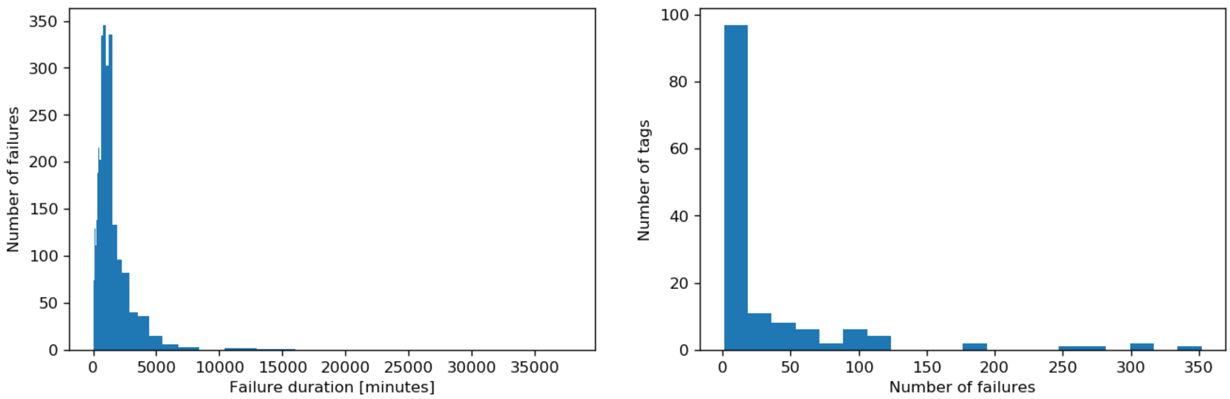

Figure 11(a) below shows total detector failures per week, which has an increasing trend throughout the year. Plotting the same data against time of day (Figure 11(b)) shows an unexpected result: there is a definite increase in detector failures at the start and end of the day shift. This is unexplained. There are also fewer failures overnight than during the day. This suggests there is a human factor in the failure rates.

Figure 11: F&G detector fault trending

(a) F&G detector faults per week **(b) F&G detector faults by time of day**

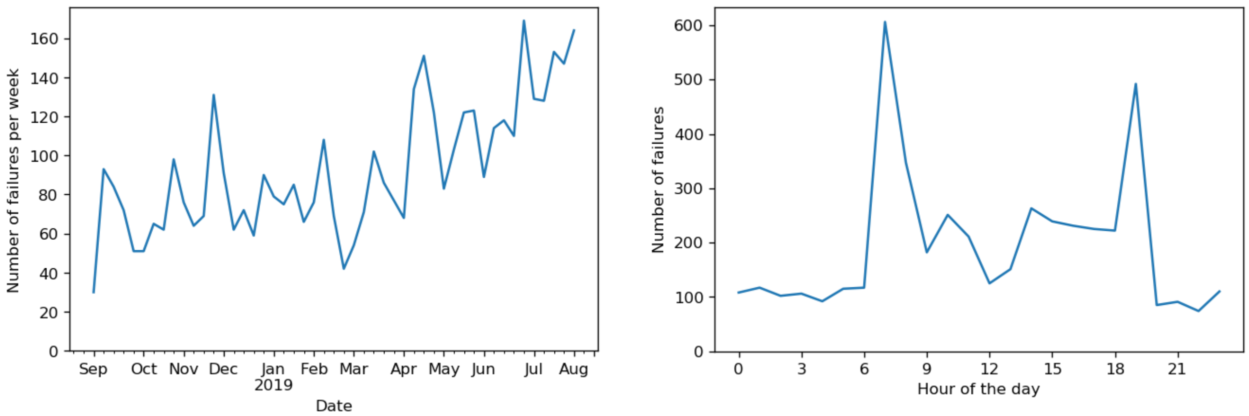

Availability

The process logs enabled the authors to plot component uptime over time. This leads to a calculation of system availability, if we have knowledge of the performance standard requirement. To discount spurious faults, outages under 15 minutes duration were neglected from the analysis below.

Zone 1 gas detectors

Using Process system logger data for faults, the number of gas detectors available in each zone over time was assessed. The number of available detectors never falls below 10. The performance standard does not explicitly state a lower limit for gas detector coverage, hence we have not applied an availability criterion here.

Fire pumps

We performed a similar analysis for the fire pumps. The asset operator advised that the performance standard requires five out of nine pumps to be made available. This analysis determined there were several periods during which fewer than five pumps were available totalling approximate 1.5 days over the year. For approximately 90 minutes, there were no fire pumps available. This gives a system availability (5oo9) over the year of 99.6%. The figure below shows the results for the number of fire pumps available at any given time.

Figure 12: Fire pumps (number available)

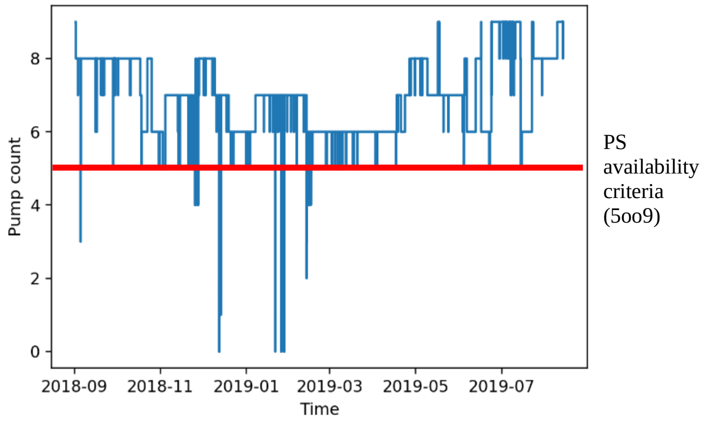

Conclusions

This study applied a sprint methodology often used in data science applications when the precise outcomes are undefined. This differs from a "traditional" project approach, which is more suited when the required outcomes are known. There is a risk that the sprint does not produce meaningful or useful results, though that is mitigated as the rapid sprint cycles enable effort to be re-focussed at regular intervals or stopped if it reaches a natural end-point.

In this study, the project requirements were loosely defined, hence the first step was to set appropriate goals (as stakeholder user stories) The table below summarizes the goals and the study outputs.

Table 6: CRISP-DM conclusions

	Stakeholder / story	Study outputs
1	**Operator technical authority (TA)** wants a better understanding of how the firewater system is operating.	We worked through the system tags to group together similar components into systems, as per the flow chart in appendix A. This was a non-trivial task as there was no pre-existing database relating system tags, just the tag list. Also, tag names are not used consistently across all datasets. We reapplied the previous study machine learning method to the entire work order history, to identify records as pass/fail and identify potential mis-classifications and other problems with the recorded data. We produced timeline figures showing the work order history, including periods of PASS/FAIL. It is simple to stack timelines for related equipment (e.g. all fire pumps.) This enables quick identification of periods of high activity, lengths of deferrals and overlapping periods of assurance fails. Using the process logs, we plotted periods of equipment availability, aggregated by system type. We showed illustrative plots for the zone 1 gas detectors and the fire pumps. This enables review of periods when the performance standards are not met, as well as calculation of the system availability.
2	The **operator TA** wants a better understanding of how failures, inhibits and isolations are reported (or not).	We analysed the process system logs to produce statistics related to system faults and inhibits. This identified some outlier data points, which could be indicative of underlying problems. In particular, regularly occurring faults and long duration faults. We noted a general trend of increasing F&G detector faults over time and, curiously, a marked increase in faults at the start and end of the day shift (07.00 and 19.00).

	Stakeholder / story	Study outputs
		Reviewing the system inhibit logs, shows a few long duration inhibits, presumably related to ongoing maintenance. There is a noticeable spike in inhibits being started/ended at the beginning and end of the day shift; this suggests items are being inhibited for entire shifts rather than for as short a time as possible.
3	The **independent verifier** wants to focus effort on deficiencies in the assurance process.	The graphs and figures produced as part of this study can be easily reproduced to aid verifiers' development of the verification scope of work. Establishment of a focussed sample for verification review will allow for a more detailed review of root cause assurance process failures within a similar timeframe.

This paper shows that applying a CRISP-DM framework to the process safety data held by asset operators, is a great way to get a high level overview of the data and unlock its potential quickly and effectively.

This work used the equipment tags to group together similar items to aid a system-level understanding. However, tag use is not consistent between datasets. In particular, the process system logs use a different naming convention that is sometimes difficult to map directly to the equipment tags. We recommend all systems use the same naming convention with a common standardised format where possible, or a separate table mapping between systems is maintained. This is a common problem seen in the process safety industry where many systems are set up and configured independently of each other. Indeed, data safety guidance (The Data Safety Initiative Working Group (DSIWG) 2018) on this issue has already been published but this approach is more focussed on the implementation of new systems and does not address the problem faced by the majority of operators today where the systems are already in place and are not configured appropriately.

As shown in this paper CRISP-DM sprints can be used to quickly identify data that is available and potential use cases, trend systems, gain a high level understanding of operations and uncover system problems and inconsistencies like tag formatting between different systems with operators effectively. The CRISP-DM provides a way to recognise the challenges faced by operators early on so solutions can be built to gain deeper insights to how process safety systems are being managed allowing improvements to be made.

References

Celnik, Matthew, and Chris Bell. 2019. "Automated review of offshore maintenance records." *IChemE Hazards 2019.*

GIT. n.d. https://git-scm.com/about.

Microsoft. 2019. Accessed 12 30, 2019. https://azure.microsoft.com/en-gb/services/devops/.

Project Jupyter. 2019. https://jupyter.org/.

Python Software Foundation. 2016. https://www.python.org/downloads/release/python-360/.

The Data Safety Initiative Working Group (DSIWG). 2018. "Data Safety Guidance." Version 3.0.

Software

This study used publicly available Python packages for data-processing, machine learning and plotting. The table below lists all packages used.

Table 7. Python software packages used

Software package	Version	Software package	Version
Python	3.6.6	Numpy	1.13.3
Matplotlib	2.0.2	Pandas	0.23.4
NLTK (natural language toolkit)	3.3.0	Scikit-learn	0.19.2
Jupyter Notebook		Azure DevOps	

Acknowledgements

This study uses real offshore maintenance data provided by a DNV GL client. While we have anonymised the data for presentation, we would like to express thanks to the data provider for facilitating this work.