IADC/SPE-212481-MS

Automated Detection of Rig Events from Real-Time Surface Data Using Spectral Analysis and Machine Learning

Timothy S. Robinson and Olav Revheim, Exebenus

Abstract

The authors present a method for automated, high-fidelity detection of rig events characterized by complex temporal signals, such as downlinking, or wave-induced heave affecting floating rigs. These can adversely impact other systems utilizing relevant data streams, for example downlinking via mud pulse telemetry can interfere with detection of pressure changes that might indicate hole cleaning problems. Identifying these events using classification techniques applied to time-domain data is difficult, hence spectral (frequency domain) techniques, combined with Machine Learning (ML), were applied to solving this problem. Surface measurements from a variety of wells, fields, regions, service companies and operators were used to develop and validate the detection methods. Data was preprocessed using time-frequency analysis, and then input to discriminative classifiers to identify rig events of interest.

For downlinking state detection, high recall and precision scores (both >93%) were achieved on independent holdout well data, and thus false positive rates were low. Successful detection was demonstrated on wells separate from the training data, hence the method is expected to generalize to new well operations. The detection method enhances situational awareness, and can actively support other software in improved automated decision-making by providing operational context in real-time, such as suppression of false warnings from monitoring pressure or modelled ECD for detecting signs of poor hole cleaning. These techniques are not limited to downlinking or heave detection, and can be applied more generally to scenarios with complex periodic signals.

Introduction

Well operations, such as drilling, in the upstream oil and gas industry are expensive endeavors. Costs may range from several tens to several hundreds of thousands of dollars per day, and a well's productivity may be hampered by delayed or failed operations. These operations can suffer from high percentages of non-productive time, often in the range of 10 to 20% of the total operations time. In addition to costliness, incidents leading to non-productive time may also pose risks to health and safety or result in environmental damage. To optimize operations, several steps are usually implemented, including carefully tracking the activities (rig states) executed on the rigs to benchmark rigs, crews, processes and methodologies, and applying analytical, physical or statistical models to the operations to either optimize operations or identify risks. Correctly identifying rig states is important for situational awareness, risk analysis, performance benchmarking and modelling tasks. For benchmarking, it is essential to accurately identify the start and end times of particular rig states or activities. For operational modelling, whether in real-time or through simulated environments, the operational context provided by the rig state is in many cases key to producing accurate estimates, and consequently to high-quality decision making.

Adoption of ML techniques to support Oil & Gas industry applications is increasing. The authors direct the reader to a selection of review papers for a broader perspective on this topic (Hanga 2019; Olukoga 2021; Sircar 2021). Example topic areas where ML has been applied include stuck pipe risk detection (Meor Hashim 2021a; Meor Hashim 2021b; Meor Hashim 2021c; Nakagawa 2021; Robinson 2022a; Othman 2022), optimizing rate-of-penetration (Batruny 2019; Singh 2019; Singh 2021; Ambrus 2022; Berrocal 2022; Cao 2022; Robinson 2022b; Singh 2022), and generation of synthetic drilling data for enhanced well planning (Batruny 2022).

The authors present a method combining spectral analysis and Machine Learning (ML) techniques for automated, high-fidelity detection of rig events characterized by complex temporal signals, such as downlinking via mud pulse telemetry, or wave-induced heave effects on offshore floating rig sensors, occurring during poor weather conditions. A provisional patent based on this work has been filed with the U.S. Patent Office (Robinson 2022c). These rig events can be problematic for other systems consuming the relevant data streams, for example downlinking can interfere with detection of pressure changes that could indicate hole cleaning problems. Downlinking is used for communicating instructions between the surface systems and downhole tools, typically using mud pulse telemetry. This is very commonly used when drilling with Rotary Steerable Systems (RSS). The communication is done by generating a series of pulses measurable by a predefined sensor, usually pressure, but in some cases torque, where varying the drillstring rotary speeds is used to convey information. This work exclusively focuses on downlinking via pressure pulses transmitted through the drilling mud.

When uncompensated, wave-induced heave perturbs the hookload sensor readings; this is typically the case near both the beginning and ends of stands where disabling active heave compensation is standard practice. This is a source of uncertainty for systems monitoring the hookload shortly after picking up the drillstring, when there is no compensation. In addition, when monitoring for large hookload changes in response to small block movements, such as overpulls, it is essential to understand if those hookload changes are driven by waves, or indicative of restrictions on drillstring movement.

This work builds on the authors' previous publications on downhole equivalent circulation density (ECD) estimation techniques, and their use for detecting hole cleaning risk symptoms based on unexpected ECD increases or fluctuations (Robinson 2022a). The ECD estimation model used standpipe pressure as an input, hence pressure fluctuations from downlinking events were a source of false warnings, and motivated the development of automated downlinking detection capabilities for handling these and enhancing the hole cleaning risk detection system.

Related work in this topic area includes that of Fei (2022) on discriminating between "phantom" and authentic downlinks at the downhole sensors, while D. Cao et al. (2020) previously demonstrated positive results from a method for detecting downlinking events using Machine Learning using the time-domain standpipe pressure data. However, their approach was computationally intensive, requiring first-differencing the pressure data, conversion of the level values and first-differenced pressure data into two 100 by 3600 pixel pseudo-images, then use of two deep convolutional neural networks (CNN) based on ResNet-34 (He 2016) to estimate preliminary classes. Both CNNs needed to predict a positive class for the final prediction to be positive. For real-time detection of downlinking during live operations, processing time is an crucial factor to consider to avoid lag. Hence, in this work the methodology was designed to achieve similarly high detection performance metrics via a more lightweight approach, suited to real-time operations.

Methodology

Certain rig states are difficult to reliably identify algorithmically from their impacts on signals in the time domain, however can be more readily detected from their frequency domain representations. Downlinking and wave-induced heave are good examples of such rig states; these exhibit distinctive signatures when examined in the frequency domain, as can be seen in Fig. 1, which shows standpipe pressure measurements containing several downlinking events, and a corresponding spectrogram showing how the frequency content of the signal varies over time. The spectrogram was calculated using the Short-Time Fourier Transform method. Classifying in the frequency-domain is also advantageous, due to relative insensitivity to phase differences between particular signal components, compared to time-domain analysis.

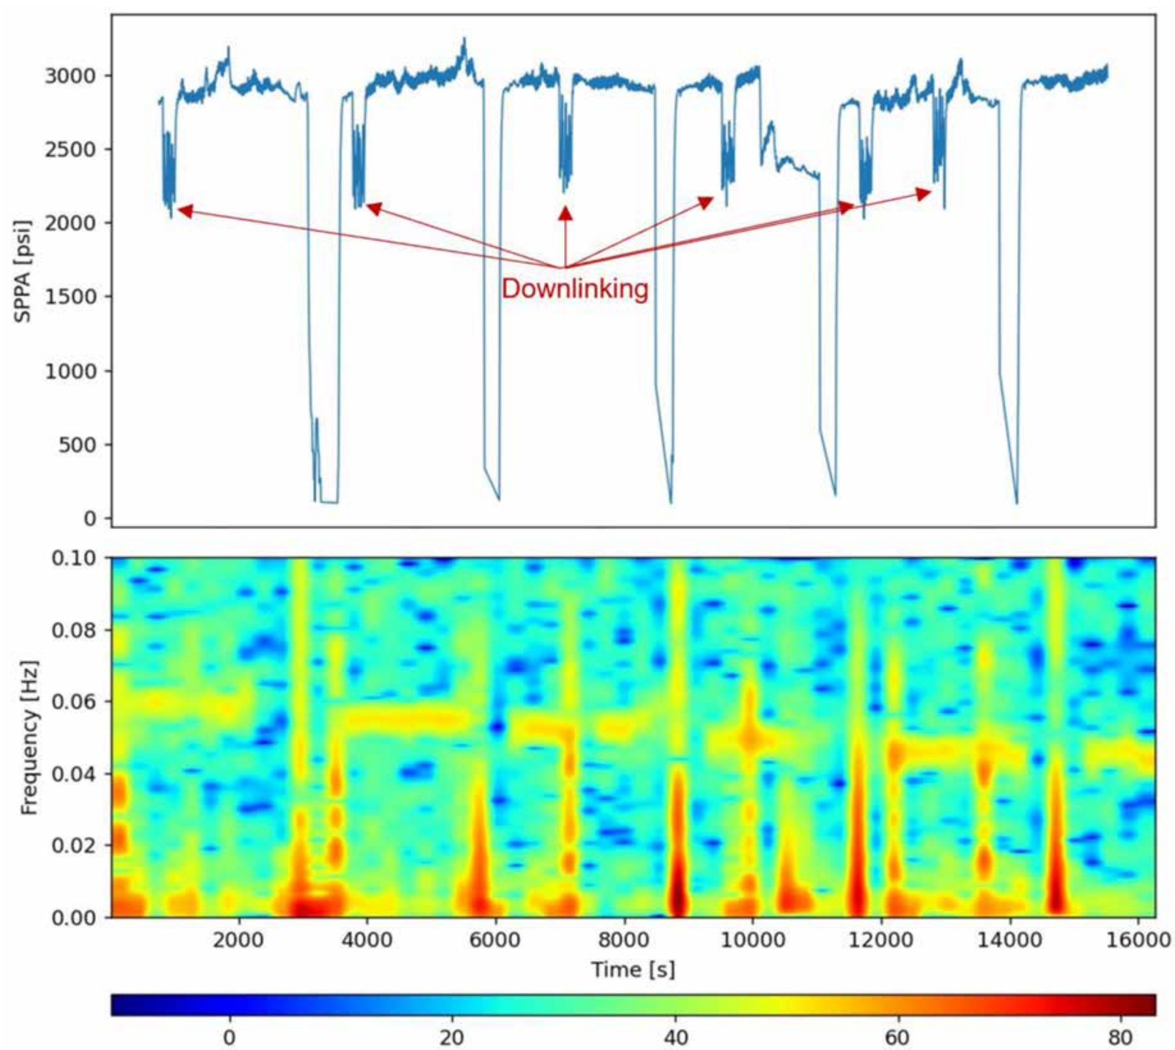

Figure 1—Standpipe pressure (SPP) measurements from an interval with multiple pressure disturbances corresponding to downlinking events (annotated), displayed as a time-series signal, as well as a spectrogram representing how power in the signal varies as a function of frequency, over its time span. Spectral power densities (in arbitrary units) are indicated by the color scale.

Visual inspection reveals distinctive signatures in the distributions of spectral power density at different points in the signal. These signatures contain certain spectral components with high power densities, visible as red-orange zones in the colormap, which are present for downlinking events, but not in other time intervals. This observation was utilized for discriminative modelling purposes, with a working hypothesis that an ML classifier receiving the distribution of spectral power from a windowed standpipe pressure signal would be able to effectively distinguish between windows containing downlinking states, and pressure measurements not containing any downlinking states. Note that the methodology is not limited to the use-cases of heave and downlinking, and can be applied to detection of complex periodic signals in other scenarios, if the signatures of the rig states are recorded within the relevant sensor data. Data sampling rates must be sufficiently high to avoid aliasing and resolve frequency components of interest; according to the Nyquist theorem, sampling rates should be at least twice the frequency of the highest frequency component under consideration. For example, to properly resolve frequency components of 0.07 Hz or below (the frequency range containing the downlinking signatures), a data sampling frequency of at least 0.14 Hz is required, corresponding to a data point measured every ~7 seconds. As service companies are typically able to provide data at rates of 0.2-1 Hz (5 seconds or 1 second between measurements), the sampling rate requirements should not cause issues in most drilling operations.

Standard surface measurements were used in this work; standpipe pressure data were used for detecting downlinking, while hookload data were used for identifying wave-induced heave effects. Temporal signals were pre-processed using time-frequency analysis techniques, into a format that was easily compatible with discriminative classifiers, using either ML or explicit rules-based logic, depending on the particular application. Signals were sliced into (overlapping) time windows, prior to calculating spectra for each window via Fast Fourier Transforms. This replicated a sliding window, with a set of spectral features generated at each time step. These features were used as inputs for a binary classifier, which discriminated between signals containing events of interest (the positive class), and those that did not. By sliding the signal windows forward by configurable time increments, for example a 5-20 seconds, multiple estimates could be generated for each downlinking event present in the data, while providing control on how often a new spectrum is calculated.

In the case of downlinking, a window containing an event was defined as one where a configurable proportion of the window's data (e.g. 20-50%) was measured during a downlinking event; an example of a 5 minute-long window containing downlinking is shown in Fig. 2 (a). This was calculated using standpipe pressure data with labelled downlinking intervals. A spectrum calculated from the windowed data using is shown in Fig. 2 (b). These were further processed by aggregating the spectral power densities into a configurable number of discrete, equally-sized, bins, exemplified by Fig. 2 (c). An ensemble of gradient-boosted trees (XGBoost) (Chen 2016) was used to detect downlinking based on the preprocessed features, namely spectral power densities for 14 discrete frequency bands. A larger number of frequency bands provides higher spectral resolution, at the expense of a higher dimensional space to be modelled by the classifier.

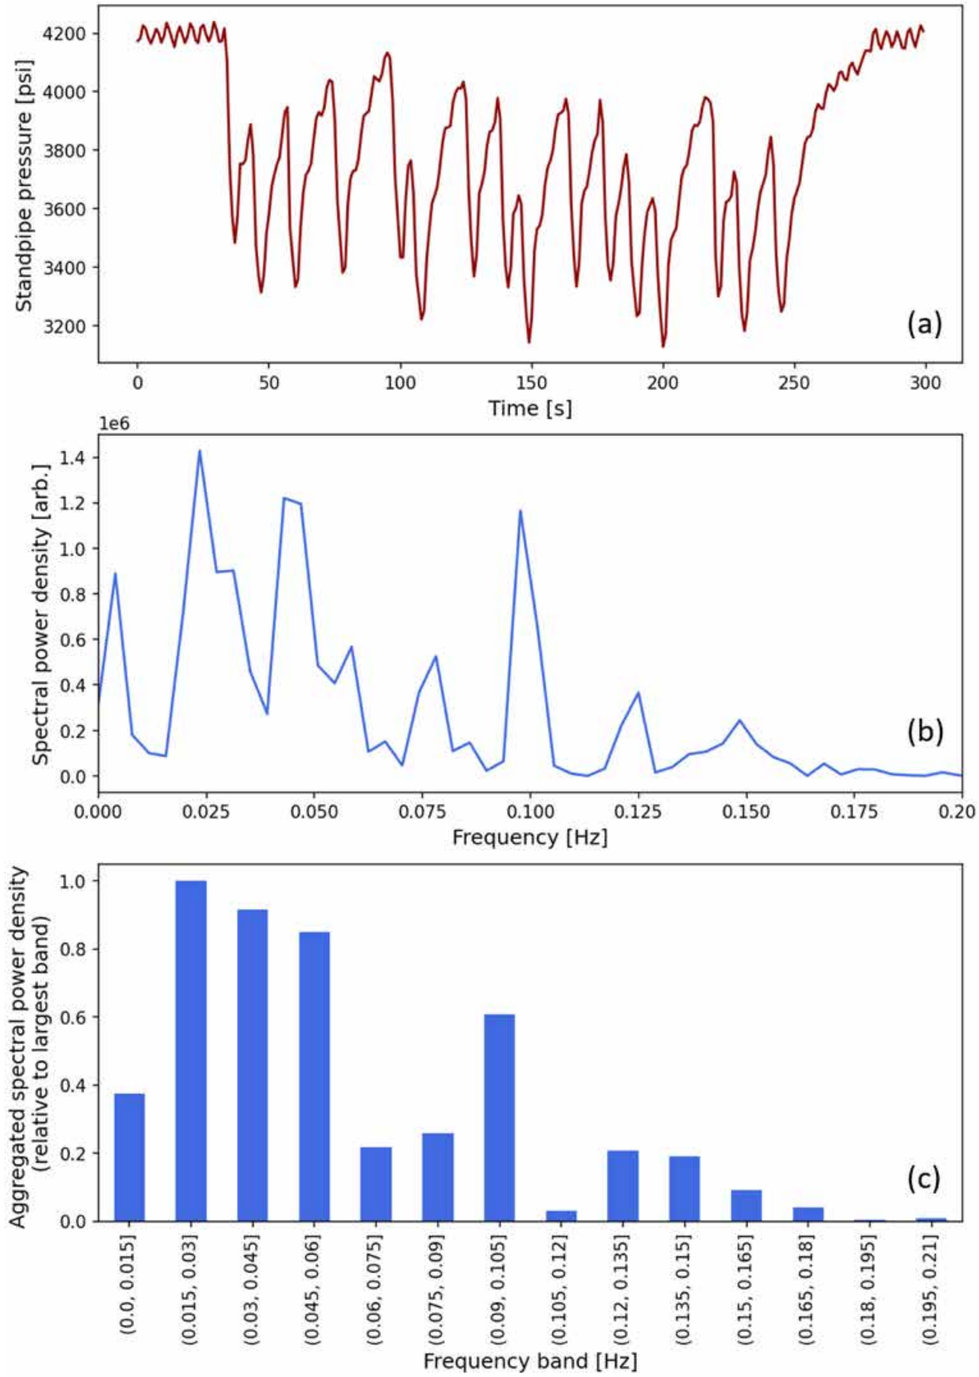

Figure 2—(a) Example slice of standpipe pressure data containing a downlinking event. (b) Spectrum corresponding to the pressure signal. (c) Spectral power density aggregated into discrete bins as a preprocessing step for modelling. The aggregated power density values were normalized relative to the largest value.

This method is computationally-lightweight compared to that of D. Cao et al. (2020), as it uses smaller amounts of input data to models (14 floating point numbers), efficient preprocessing steps commonly used in Digital Signal Processing, and a single tree-based ensemble model rather than a pair of deep CNNs based on ResNet-34 (He 2016), each consuming 360,000 numerical values in the form of 100 by 3600 pixel images. According to He's original paper, the ResNet-34 model performs 3.6 billion floating point operations (multiply-add) on a forward-pass. Assuming use of a processor that can perform 30 billion operations per second, in the theoretical best-case scenario, a single forward pass of one ResNet-34 model would take 120 ms, almost certainly an underestimate of the true time needed. While this workload could be distributed across multiple processors, based on this runtime estimate, there is a clear incentive to develop more lightweight methodologies suited to 1 Hz real-time data streams.

Heave detection was achieved using a rules-based approach, inspired by Moskovitz's (1964) work on the power spectra of ocean waves, where spectral power of ocean waves tends to be distributed within a specific frequency band spanning 0.06-0.12 Hz. An example of real hookload data from a tripping stand, taken from a floating rig located in the North Sea, and perturbed by wave-induced heave, is shown in Fig. 3, with the spectrum calculated from the plotted time-series shown below on a logarithmic scale (base 10). A large spectral component is clearly visible at ~0.075 Hz, which falls within the aforementioned frequency band. The largest spectral component of a hookload time-series falling within this band indicates an effect of wave-induced heave on the floating rig. Hence, a simple rules-based system examining the distribution of spectral power within a hookload signal can effectively discriminate between signals with, and without, heave effects.

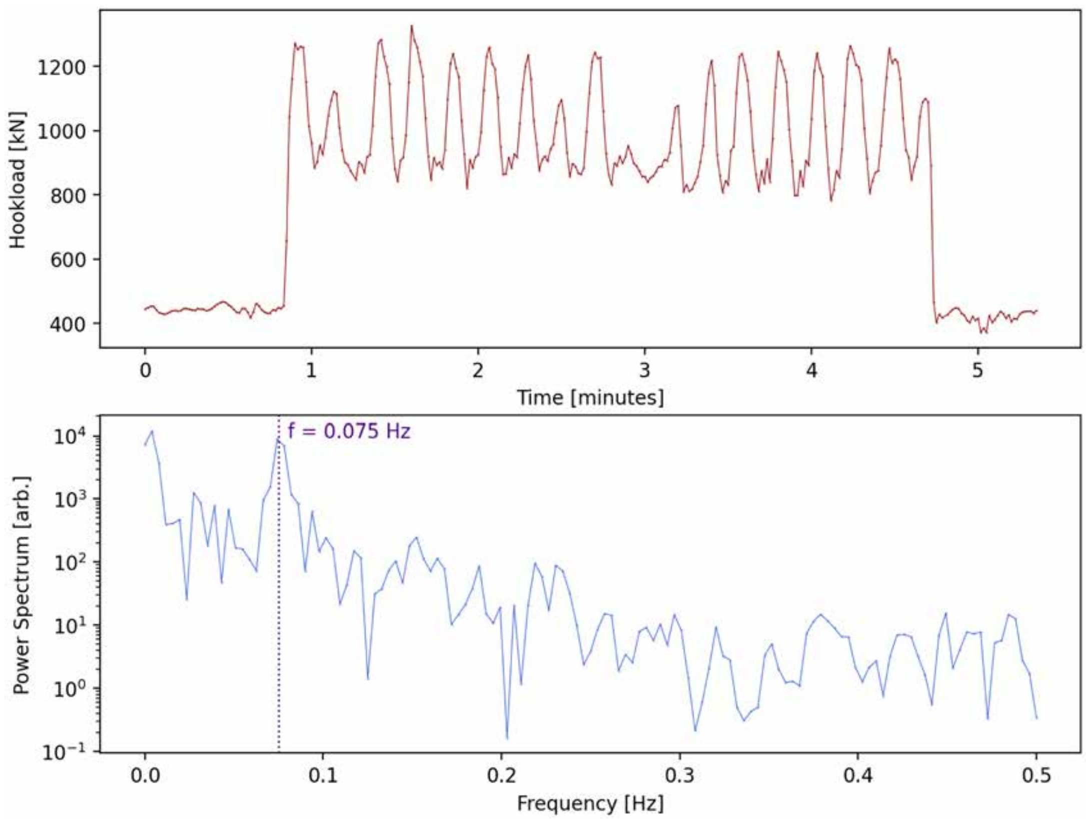

Figure 3—Wave-induced heave effects on hookload measurements taken from a floating rig located in the North Sea, and the spectrum associated with this time-series. The largest spectral component not associated with connection/ disconnection of the drillstring matches the frequency band of ocean waves described in Moskovitz's (1964) work.

Active heave compensation is typically not enabled at the beginning and end of a stand, thus the heave signatures are only notable within hookload measurements during these intervals. The last few minutes of data were selected from each stand, prior to disconnecting the drillstring, and used to infer the presence of heave effects based on the power spectrum of the sliced signal. This assists monitoring staff in distinguishing between sharp changes in hookload such as overpulls, and wave-induced changes which are very unlikely to indicate stuck pipe risks.

A global dataset containing information from a variety of wells, fields, regions, service companies and operators was curated to support development and validation of the detection methods. Surface measurements with sampling rates between 0.2-1 Hz were used; downhole measurements were not required. To demonstrate generalization and suitability for out-of-the-box usage, validation was done on independent wells from those contained in the training datasets; this approach minimizes the risk of information leakage between the training and validation datasets, a common pitfall of practitioners using random-sampled train-test split approaches without due consideration for the problem context. A set of 28 wells with labelled downlinking intervals was used to develop the classification model; 17 for training, 5 for validation and early-stopping training, and 6 additional wells containing a total of 510 downlinking events were reserved for a holdout test set. Performance on the holdout wells was only evaluated after completing hyperparameter tuning and selecting the best performing model on the validation set.

Due to data availability constraints stemming from the seasonal nature of ocean waves, the detection method was most extensively tested on detecting downlinking events. These are commonly observed during drilling operations, particularly when drilling with an RSS, hence performance statistics could be readily calculated and model generalization tested on multiple independent wells.

Results and Discussion

A downlinking detector implemented in Python was tested with 1 Hz data-streams. The average execution time for updating time window data, recalculating a spectrum and generating an updated downlinking status was 941 ± 84 µs, measured using IPython's %timeit magic command run on a standard laptop computer. This runtime is sufficiently fast to identify downlinking events from 1 Hz data streams without experiencing lag in real-time operations. Similarly, positive results were achieved on wells with 0.2 Hz data sampling rates, though computation time requirements are less stringent in this scenario. At lower sampling rates, e.g. 0.1 Hz (1 sample per 10 seconds), the reduced information content of the signals degrades the classification effectiveness of the method, and increases the risk of aliasing problems. Hence, the lowest pressure data sampling frequency within the scope of this work was 0.2 Hz.

An example illustrating the downlinking detection model in action on a drilling interval (taken from a well in the holdout test dataset) is provided in Fig. 4; here, a time-series of (raw) standpipe pressure data from the drilling stand is shown, with the corresponding downlinking statuses estimated by the binary classification model presented below. A status of 1.0 means the model estimates a downlinking event is present within the windowed data, while conversely a status of 0.0 means no downlinking event was detected. The downlinking statuses were timestamped according to the *end-time* of the pressure data window used to generate spectra and model inputs, hence positive statuses generally appear towards the middle or end of time intervals spanned by downlinking events.

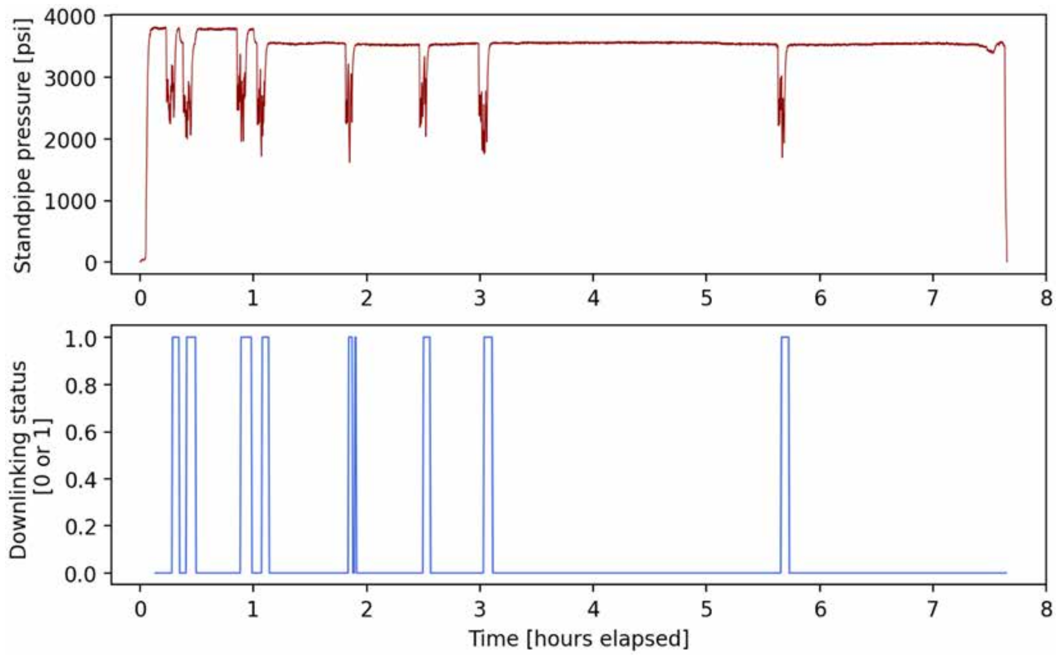

Figure 4—Standpipe pressure data series from an example drilling stand containing 8 downlinking events, and the corresponding downlinking status estimates from the binary classification model (1.0 = downlinking, 0.0 = not downlinking). The data was taken from one of the holdout test set wells, which were not used in model fitting. The timestamps of the status estimates correspond to the end-time of the sliding windows used as raw data.

Tests on independent holdout test wells were successful, indicating that the method can generalize and reliably be applied to new wells. This allows it to support other analytics challenges, such as pressure monitoring and ECD modelling for detecting signs of poor hole cleaning. For downlinking, high recall and precision scores (both above 93%) were achieved, clearly demonstrating the tool's capability to correctly identify downlinking, while maintaining low false positive rates. When considering downlinking detection performance with respect to its usage in suppressing false warnings from a hole cleaning risk detection system, precision and recall scores are effectively inverted; a high recall for downlinking means that few false risk warnings will be triggered by downlinks (high precision), while high precision means a low likelihood of suppressing genuine hole cleaning risk warnings, which leads to higher recall from the hole cleaning risk detection system.

Summary statistics evaluated on the holdout test dataset are provided in Table 1. In order to imitate operational usage of a downlinking detector most closely, recall values were calculated based on whether each downlinking event was detected or not, hence if multiple sliding time windows containing the same downlinking event each resulted in positive class predictions, this would still only count as a single correctly identified event. Precision was calculated by considering the proportion of positive class predictions that were generated from time windows which actually contained downlinking events.

Table 1—Summary statistics from evaluating the downlinking detection model performance on data from the 6 holdout test wells.

Metric	Value (3 s.f.)
Total no. of events in dataset	510
No. of events identified	479
Recall	0.939
Precision	0.935
F1 score	0.937
Accuracy	0.973

From Table 1, it follows that 479 of the 510 (~94%) false warnings likely induced by downlinking events could be suppressed by incorporating downlinking detection into the hole cleaning risk detection tool described previously (Robinson 2022a), without compromising on its general sensitivity to risk symptoms. In classification problems, a common consideration is the tradeoff between model sensitivity and specificity; usually, a consequence of reducing false positive rates is reduced sensitivity. However, by exploiting domain knowledge regarding underlying causes of false warnings, and using an additional model explicitly targeting these to augment the risk detection process, this paradigm can be broken.

With ~94% precision, only a very small proportion of flags from the downlinking detector do not correspond to genuine downlinking events, hence the likelihood of wrongly suppressing a "true" warning from the hole cleaning risk detection tool is also low. This is an important consideration for any risk detector, as missed risks may be costly if incidents occur, particularly in the context of pack-offs leading to a stuck pipe and extended non-productive time spent on remedial actions.

The downlinking detection performance statistics obtained were similar those reported by Cao (2020), but without the computational complexity of their method that could reduce suitability for real-time usage with rigs streaming data at 1 Hz sampling rates. However, it should be noted that the evaluations were not done on identical datasets, so only an approximate comparison should be made between the summary statistics for the two approaches.

Conclusions

The rig event detection methodology presented in this work enhances situational awareness, and can actively support improved automated decision-making by providing information on operational context to other software applications in real-time. Automated inference of rig states can contribute to the foundations required for reaching the goal of autonomous rigs. The methodology used only surface measurements, readily available from the majority of drilling operations at the required sampling rates. High-performance detection and generalizability was demonstrated for downlinking events observable in standpipe pressure signals, as well as the capability to recognize wave-induced heave influencing hookload measurements, using sufficiently computationally lightweight methods to support usage in live operations. The techniques presented in this work are not limited to the use cases of downlinking or heave detection, and can be applied more generally to scenarios with complex periodic signals, which are challenging to recognize solely through analysis in the time-domain.

References

Ambrus, A., Alyaev, S. Jahani, N. et al 2022. "Rate of Penetration Prediction Using Quantile Regression Deep Neural Networks", Proceedings of ASME 2022 41st International Conference on Ocean, Offshore and Arctic Engineering, Hamburg, Germany, June 2022. OMAE2022-79046. doi: https://doi.org/10.1115/OMAE2022-79046

Batruny, P., Yahya, H. Omar, A. et al 2019. "Drilling in the Digital Age: An Aproach to Optimizing ROP Using Machine Learning", Abu Dhabi International Petroleum Exhibition & Conference, SPE-197157-MS. doi: https://doi.org/10.2118/197157-MS

Batruny, P. and Robinson, T., 2022. "Unsupervised Machine Learning: A Well Planning Tool For The Future", Proceedings of the ASME 2022 41st International Conference on Ocean, Offshore & Arctic Engineering, June 5 – June 10, 2022, Hamburg, Germany, OMAE2022-78423. doi: https://doi.org/10.1115/OMAE2022-78423

Fernandéz Berrocal, M., Cao, J. and Sui, D. 2022, "Evaluation and Interpretation on Data Driven ROP Models From Engineering Perspectives", Proceedings of ASME 2022 41st International Conference on Ocean, Offshore and Arctic Engineering, Hamburg, Germany, June 2022. OMAE2022-78752. doi: https://doi.org/10.1115/OMAE2022-78752

Cao, D., Hender, D. Ariabod, S. et al 2020. "The Development and Application of Real-Time Deep Learning Models to Drive Directional Drilling Efficiency", IADC/SPE International Drilling Conference and Exhibition, Galveston, Texas, USA, March 2020. SPE-199584-MS. doi: https://doi.org/10.2118/199584-MS

Cao, J., Ren, H. and Sui, D. 2022. "Global Optimization Workflow for Offshore Drilling Rate of Penetration With Dynamic Drilling Log Data", Proceedings of ASME 2022 41st International Conference on Ocean, Offshore and Arctic Engineering, Hamburg, Germany, June 2022. OMAE2022-79747. doi: https://doi.org/10.1115/OMAE2022-79747

Chen, T. and Guestrin, C. 2016. "XGBoost: A Scalable Tree Boosting System", Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, California, USA, 2016, pp. 785–794. doi: https://doi.org/10.1145/2939672.2939785

Erge, O., Özbayoğlu, M., Özbayoğlu, E. 2022. "A Deep Learning Model for Rate of Penetration Prediction and Drilling Performance Optimization Using Genetic Algorithm", Proceedings of ASME 2022 41st International Conference on Ocean, Offshore and Arctic Engineering, Hamburg, Germany, June 2022. OMAE2022-79623. doi: https://doi.org/10.1115/OMAE2022-79623

Hanga, K. M., Kovalchuk, Y. 2019. "Machine learning and multi-agent systems in oil and gas industry applications: A survey", *Computer Science Review*, Vol. **34**, 100191. doi: https://doi.org/10.1016/j.cosrev.2019.08.002

He, K., Zhang, X., Ren, S. and Sun, J. 2016. "Deep Residual Learning for Image Recognition", IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778, doi: https://doi.org/10.1109/CVPR.2016.90

Li, F. 2021. "Investigation into phantom downlinks and development of a prevention algorithm in rotary steerable system", *Petroleum* Vol. **7**, Iss. 3, pp. 349–355, Sept 2021. doi: https://doi.org/10.1016/j.petlm.2020.09.001

Meor Hashim, M. M. H. et al 2021 (a). "Performance Improvement of Wells Augmented Stuck Pipe Indicator via Model Evaluations". Presented at the IPTC International Petroleum Technology Conference 2021. doi: https://doi.org/10.2523/IPTC-21455-MS

Meor Hashim, M. M. H. et al, 2021 (b). "Case Studies for the Successful Deployment of Wells Augmented Stuck Pipe Indicator in Wells Real Time Centre". Presented at the IPTC International Petroleum Technology Conference 2021. doi: https://doi.org/10.2523/IPTC-21199-MS

Meor Hashim, M. M. H. et al 2021 (c). "Utilizing Artificial Neural Network for Real-Time Prediction of Differential Sticking Symptoms". Presented at the IPTC International Petroleum Technology Conference 2021. doi: https://doi.org/10.2523/IPTC-21221-MS

Moskowitz, L. 1964. "Estimates of the power spectrums for fully developed seas for wind speeds of 20 to 40 knots", *Journal of Geophysical Research*, Vol. **69** (24), p.5161–5179, 1964. doi: https://doi.org/10.1029/JZ069i024p05161

Nakagawa, Y., Inoue, T. Bilen, H. et al 2021. "Approach for Real-Time Prediction of Pipe Stuck Risk Using a Long Short-Term Memory Autoencoder Architecture". Paper presented at the Abu Dhabi International Petroleum Exhibition & Conference, Abu Dhabi, UAE, November 2021. SPE-207805-MS. doi: https://doi.org/10.2118/207805-MS

Olukoga, T. A., Feng, Y. 2021. "Practical Machine-Learning Applications in Well-Drilling Operations", *SPE Drill & Compl* Vol. **36**, Iss. 04, pp.849–867. SPE-205480-PA. doi: https://doi.org/10.2118/205480-PA

Othman, E. B., Gomes, D. Tengku Bidin, T. E. B. et al 2022. "Application of Machine Learning to Augment Wellbore Geometry-Related Stuck Pipe Risk Identification in Real Time". Paper presented at the Offshore Technology Conference Asia, Virtual and Kuala Lumpur, Malaysia, March 2022. doi: https://doi.org/10.4043/31695-MS

Robinson, T.S., Gomes, D. Meor Hashim, M. M. H. et al 2022 (a). "Real-Time Estimation of Downhole Equivalent Circulating Density ECD Using Machine Learning and Applications", IADC/SPE International Drilling Conference and Exhibition, Galveston, Texas, USA, March 2022. SPE-208675-MS. doi: https://doi.org/10.2118/208675-MS

Robinson, T.S., Batruny, P. Gomes, D. et al 2022 (b). "Successful Development and Deployment of a Global ROP Optimization Machine Learning Model", Offshore Technology Conference Asia, Kuala Lumpur, Malaysia, March 2022. OTC-31680-MS. doi: https://doi.org/10.4043/31680-MS

Robinson, T., Revheim, O., 2022 (c). "Frequency Based Rig Analysis", *U.S. Provisional Patent filed September 2022*.

Singh, K., Yalamarty, S. S. Kamyab, M. et al 2019. "Cloud-Based ROP Prediction and Optimization in Real Time Using Supervised Machine Learning", SPE/AAPG/SEG Unconventional Resources Technology Conference, Denver, Colorado, USA, July 2019. URTEC-2019-343-MS. doi: https://doi.org/10.15530/urtec-2019-343

Singh, K., Yalamarty, S., Cheatham, C. Cheatham, C., et al 2021. "From Science to Practice: Improving ROP by Utilizing a Cloud-Based Machine-Learning Solution in Real-Time Drilling Operations". SPE/IADC International Drilling Conference and Exhibition, Virtual, March 2021. SPE-204043-MS. doi: https://doi.org/10.2118/204043-MS

Singh, K., Siddiqui, F. Braga, D. et al 2022. "ROP Optimization using a Hybrid Machine Learning and Physics-Based Multivariate Objective Function with Real-Time Vibration and Stick-Slip Filters". IADC/SPE International Drilling Conference and Exhibition, Galveston, Texas, USA, March 2022. SPE-208751-MS. doi: https://doi.org/10.2118/208751-MS

Sircar, A., Yadav, K., Rayavarapu, K. 2021. "Application of machine learning and artificial intelligence in oil and gas industry", *Petroleum Research* Vol. **6**, Iss. 4., pp. 379–391. doi: https://doi.org/10.1016/j.ptlrs.2021.05.009