Upstream Oil and Gas Technology 11 (2023) 100096

Contents lists available at ScienceDirect

Upstream Oil and Gas Technology

journal homepage: www.sciencedirect.com/journal/upstream-oil-and-gas-technology

Stuck pipe prediction from rare events in oil drilling operations

Salvatore D'Amicis [a,1], Marta Pagani [a,1], Matteo Matteucci [a], Luigi Piroddi [a,*], Andrea Spelta [b], Fabrizio Zausa [b]

[a] Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, Milano, Italy
[b] ENI, Drilling & Completion Risk Analysis Department, San Donato Milanese, Italy

ARTICLE INFO

Keywords:
Oil&gas
Drilling
Stuck-pipe
Detection
Prediction
Rare events

ABSTRACT

Stuck-pipe phenomena are relatively rare in drilling operations in the oil & gas industry, but can have disastrous economic consequences, causing costly time delays and sometimes even the loss of expensive machinery. In this work, we develop an event-based prediction model that relates the occurrence of precursor events to the stuck-pipe phenomena. To this aim, the detectors of various types of precursor events that typically anticipate stuck-pipe occurrences are first designed based on the available mudlog data. A Hidden Markov Model (HMM) is then developed to relate these precursor events to actual drilling problems, producing different levels of alarm, with the ultimate goal of predicting stuck pipes. The model has been tested on a dataset of wells with different characteristics, showing positive results.

1. Introduction

In the oil & gas industry, a lot of effort is devoted to avoiding non productive time (NPT) due to faulty operations while drilling. This is motivated both by safety reasons, since some events may harm the team working on the well and/or damage the environment leading to the loss of the rig, and by economical reasons, since costs (in terms of equipment damage and time spent to resume the canonical process) can rise rather sharply during faulty operations. About 25% of NPT can be ascribed to the inability to move, rotate or remove the drillstring from the wellbore [10], a phenomenon referred to as stuck pipe (or sticking). NPT due to stuck pipes does not pose a serious threat to the safety of workers, but the attempts to free the drillstring can lead to environmental damage, and in the worst case the drillstring has to be cut at the casing depth, thus losing the whole Bottom Hole Assembly (BHA), and the well must be sidetracked to continue the drilling, which has a huge impact on costs and causes tremendous delays.

Therefore, it is not surprising that the timely prediction of stuck-pipe events is considered of paramount importance, so that appropriate countermeasures can be put in effect by the drilling team to avoid or at least mitigate the impending drilling issues. This problem is very challenging due to the hybrid nature of the drilling process (which consists of different activities, besides the actual drilling), the variability of geological conditions, and the combined occurrence of adverse events. Consider also that the most reliable source of information comes in the form of measurements associated to the functioning of the surface equipment, which are only indirectly related to the occurrence of stickings. Furthermore, the available data refer to a controlled process with a human in the loop, and as such they reflect the unknown policy of the drilling team. Indeed, depending on its experience and ability, the drilling team may adopt more or less conservative actions, in the presence of the same pattern of observed measurements. As a consequence, the occurrence of a given pattern of precursor events may sometimes lead to a stuck-pipe event (because the team underestimated the issue or took inappropriate or untimely or insufficient actions), and sometimes not.

Numerous works have appeared in the recent scientific literature concerning the application of machine learning methods to data analysis tasks related to the drilling process. In particular, research efforts directed at accurately predicting stuck-pipe events before actually reaching the region of concern have considerably increased after 2006 (see *e.g.* the review article [3]), resulting in the formulation of several models, based on different types of signals.

In [12] physical models of the wellbore and drilling process are combined with machine learning techniques to detect and anticipate anomalies in real-time such as kicks and stickings. In [8] a model is used to predict poor hole cleaning conditions that may lead to a sticking, using the Bottom Hole Pressure (BHP) and the TQA (TorQue - Average)

* Corresponding author.
E-mail address: luigi.piroddi@polimi.it (L. Piroddi).
[1] former master students at the Politecnico di Milano.

https://doi.org/10.1016/j.upstre.2023.100096
Received 25 November 2021; Received in revised form 28 November 2022; Accepted 2 August 2023
Available online 11 August 2023
2666-2604/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

as inputs. In [16] an artificial neural network has been used for stuck-pipe prediction on data from an Iranian oil field. The method uses two different models for static and dynamic stuck-pipe instances, and employs hole geometry information and mud properties besides mudlog data. In [11] a neural network was employed to predict stickings, employing information such as pressure difference, well narrowness, and mud properties.

In [13] two different models (using a SVM and a neural network, respectively) have been developed to predict stuck-pipe events. The models get the configuration of a drilling operation and predict the sticking events. A selection of the input features is carried out, resulting in a model including just five features (nozzle size, the minimum weight on bit, temperature filtrate loss, initial gel, salt). A real-time method is presented in [9] to predict impending stuck pipes with sufficient anticipation to be able to prevent them. This method assumes that there exist two types of data behavior which may indicate an impending stuck pipe: a significant deviation from a model predicted value and a rapid rate of change. Deviations from the model are calculated comparing the actual value measured by a rig sensor and the expected value for that point on the wellbore generated from a hydraulics or torque and drag model. The second detection method, based on the rate of change, detects potential problems by identifying rapid changes in key parameters without reference to their expected value. Rate of change alerts provide valuable insight into deteriorating wellbore conditions that develop rapidly. Both time and depth analysis tools are employed. In [20] a method is developed to detect the precursors of drilling events based on drilling data such as surface data, wellbore geometry data, lithology and downhole measurements from various downhole tools. The drilling events are first extracted from massive drilling data using defined thresholds and/or criteria. Then, the time series of drilling parameters are represented by symbolic aggregate approximation (SAX). The patterns of these SAX strings are clustered by unsupervised learning and then used for pattern recognition with dynamic time warping. Finally, the patterns are classified based on the similarity matrix.

In [5] three data-based indicators are developed, that detect three different physical phenomena associated with the insurgence of a sticking. More specifically, the first two indicators detect linear and rotational mechanical resistance to the motion of the BHA, while the third one targets unexpected pressure surges. A statistical model is then constructed using classical machine learning methods that relates these indicators to the level of risk of incurring into a sticking along the borehole.

Employing precursor events to predict rare faults, as suggested in [5], has been attempted in different contexts, from service problems in computer networks to failure event prediction, using various machine learning techniques (see, *e.g.*, [14,17,18]). For instance, the method proposed in [17] is based on the recognition of similar precursor patterns over short time windows to predict the target events. Unfortunately, this approach is hardly applicable to our case, since very different precursor patterns are observed for each sticking. To deal more effectively with the lack of structure and repeatability of the problem at hand, a genetic algorithm is used in [18] to prediction patterns. In [14], an HMM is used instead to identify and recognize patterns of fault events that lead to failures. Here, the current state is estimated on the basis of the recent fault events, which allows the probabilistic prediction of failure events.

Expanding on the ideas of [5] and adopting the general framework of [14] (albeit with a novel HMM structure), we propose an HMM-based prediction model for stuck-pipe faults, based on the observation of a number of precursor events. The HMM models different alarm states which are associated to precursor patterns observed in a suitable past time-window.

More in detail, a first block of the system is devoted to the detection of precursor events, associated to different physical phenomena, that are likely to occur prior to a stuck pipe event during drilling activities. Notice that detecting these precursor events represents a valuable contribution on its own, in that this information can be of great aid to the drilling team, enhancing its ability to predict and prevent impending problems and also allowing a prompt and appropriate reaction if the stuck-pipe event actually occurs. Indeed, the probability of freeing a stuck pipe rapidly decreases with time, and the enacted countermeasures must match the actual cause of the phenomenon, to avoid worsening the problem [16].

The second block of the system estimates three different levels of alarms by using a Hidden Markov Model (HMM) that takes as input the detected precursor events. The proposed method takes into account mudlog data (the mudlog collects a series of signals, generally measured at rig level, related *e.g.* to the hook load, the block position, the rotation speed of the bit, the torque given by the motors to the drillstring, the mud pressure on the stand pipe, the pumps flow rate), as well as trajectory (*e.g.*, the dogleg severity) and lithology data. The proposed system has been tested on a large dataset of wells, showing good results in simulation, that make it a promising application for a real case scenario.

Overall, the proposed method is capable of dealing with a largely unstructured fault prediction problem, characterized by multiple –and sometimes concurring– causes of the fault, associated to quite different precursor patterns. The method is quite amenable to drilling operators, as some of the precursor events exploited here are known to be related to drilling issues and are often already monitored in the usual practice.

The rest of the paper is structured as follows. Section 2 provides the necessary background on the drilling operations by describing the rig, the drilling process, the stuck pipe issue, and discussing the various data sources (timelog, mudlog, lithology and trajectory) available. Section 3 presents an analysis of the signals before the occurrence of stuck-pipe events and discusses possible precursors (and their detectors). Section 4 describes the proposed HMM-based prediction model by showing the entire pipeline from the data to the outgoing alarm levels. Finally, Section 5 reports the experiments carried out for evaluation purposes and the experimental results of the proposed model, followed by some brief concluding remarks.

2. Preliminaries

2.1. *The drilling process*

Drilling consists in boring a hole (that can be several thousands of meters long) in order to create a well for hydrocarbon production. A composable drillstring is used for this purpose, consisting of a drill pipe that ends in the bottom hole assembly (BHA), which in turn contains the drill bit, *i.e.*, the tool used to crush or cut the rock. The gravity acting on the drill pipe provides the downward force necessary to the bit to break the rock. Rotation is also applied to the drillstring to facilitate the drilling process. The drillstring has to be periodically extended as the wellbore goes deeper, assembling further pipe segments (stands) from the top. Periodically, the drilling is interrupted to fit the wall of the well with a protective casing, for mechanical support and isolation. Overall, the drilling process is a complex sequence of different operations, including drilling, reaming (passing repeatedly along already drilled portions of the hole to clean them), stand changing, casing, and tripping (extracting or reinserting the whole pipe, *e.g.* to replace the BHA). A drilling phase includes all operations between two consecutive casing operations.

An important role in the drilling process is provided by the drilling mud, which is pumped down the pipe and returns to the rig passing through the annulus, *i.e.*, the space between the pipe and the wall of the borehole. Besides removing the drilled cuttings, the mud has other functions, *i.e.* to cool down and lubricate the bit and –most importantly– to hold the bore hole from falling down, by equilibrating the formation fluid pressure with its own hydrostatic pressure.

2.2. The stuck pipe phenomenon

The drilling process may be slowed down or jeopardized by various adverse events, due to well problems (such as circulation losses, stickings, fluid influx, etc.), rig failures, downhole or surface equipment failures, etc. In particular, the occurrence of sticking (or stuck-pipe) phenomena may ultimately result in catastrophic outcomes including pipe breakage, the loss of expensive downhole equipment, and a considerable delay in the drilling operations. The drill pipe is stuck when the drillstring is no more free to move up or down, or rotate, *i.e.*, when the static force necessary to make it move exceeds the capabilities of the rig or the tensile strength of the drill pipe. Various actions can be exerted by the operator in an attempt to free the BHA, at the cost of a delay in the drilling. In the most unfortunate cases, the BHA cannot be freed and must be abandoned, and the drilling is later resumed along a different trajectory.

Pipe sticking can occur due to various physical mechanisms [4]. A differential pressure sticking occurs when the hydrostatic pressure in the wellbore is higher than the formation pressure, which causes a net force pushing the collar towards the borehole wall. In high-angle and horizontal wells, the gravitational force also plays a role in extending the contact between the drillstring and the formation. Differential stickings occur more frequently during pipe connection operations [19]. Mechanical stickings, on the other hand, can have a variety of causes, such as inadequate hole cleaning, chemically active formations, wellbore instability, overpressured formations, the presence of high dip sloughing, unconsolidated formations, mobile formations, under gauge hole or key seating, wellbore geometry anomalies.

2.3. Data sources

Early signals of a poor hole-cleaning conditions can be found in an erratic torque (the string is repeatedly getting stuck in the cuttings, wound up and spun free), an unexplained increase in the bottom hole pressure (which may be associated to a tight spot with packings causing flow restrictions further up the annulus), or an unexpected hook load (if the drillstring rests on a tight packing the hook load is lower than anticipated) [12]. Preventing stuck pipe requires the detection of early warning signs, such as increases in torque and drag, excessive cuttings loading, tight spots while tripping, loss of circulation while drilling. This can be accomplished by monitoring the various sources of data which are available, namely timelog reports, mudlog data, lithology data and trajectory data.

The drilling crew manually annotates the main activities and events happened during the drilling process in a timelog report. This report provides the necessary information to construct a ground truth signal for the detection of stickings, since it contains the operators' assessments regarding all the issues encountered during drilling. Unfortunately, the consistency of these data is sometimes questionable, due to the subjectivity of the operators' evaluations. For example, minor sticking problems may not be reported as such, and other well problems may be reported mistakenly as sticking events. Furthermore, since annotations are taken manually at irregular intervals, they are only loosely aligned with the data, which further hampers the data processing task.

The mudlog records several variables related to the drilling process, either measured or calculated. The measured variables are sampled at a frequency of 50 Hz, and each tuple in the mudlog is reported with a frequency of 0.2 Hz (one sample each 5 s), generally taking the average of the original measurements over consecutive windows of 250 samples. It is important to notice that operator inputs are not recorded: this poses great limits to the machine learning task because it makes it less immediate to distinguish between variations induced by regular operator activities and anomalies. Mudlog measurements are typically collected from the surface equipment, and thus provide only indirect information regarding the conditions at the bottom of the borehole. Further measurements can be taken at the bottom of the well using additional sensors at the BHA, but these data are seldom available and are not accessible in real time (they cannot be used for online processing), and are therefore overlooked in this study.

Table 1
Main mudlog variables.

Variable	Units	Description
DBTM (DBTV)	m	Bit depth along the drilling (vert.) dir.
DMEA (DVER)	m	Hole depth along the drilling (vert.) dir.
BPOS	m	Position (height) of the traveling block
ROPA	m/h	Rate of penetration (avg)
HKLA	Ton	Hookload (avg)
WOBA	Ton	Weight on bit (avg)
RPMA	rpm	Rotary speed (avg)
TRPM	rpm	Bit revolution (downhole)
TQA	kgf·m	Rotary torque applied to the drillstring
MDIA (MDOA)	kg/m^3	Input (output) mud density (avg)
MFIA (MFOA)	l/min	Input (output) mud flow (avg)
SPMT	nr	Total pump stroke rate
STKC	nr	Pump stroke count (cum)
SPPA	kg/cm^2	Stand pipe pressure

The mudlog contains various types of variables, associated with time, depth, the mechanical and fluid dynamics properties of the drilling process (see Table 1). DBTM (DBTV) and DMEA (DVER) are variables measuring the bit and hole depth. BPOS measures the position (height) of the traveling block, which ranges from the bottom to the top of the rig. HKLA measures the total force pulling down on the hook, which equals the drillstring weight, reduced by any force that tends to reduce that weight, such as the friction along the wellbore wall or the buoyancy caused by the immersion of the drillstring in the drilling fluid.

The weight on bit (WOBA) is the amount of downward force exerted on the drill bit and results from all the components of the drillstring (drill collars). To accurately control the amount of force applied to the bit, the driller carefully monitors the surface weight measured while the bit is just off the bottom of the wellbore. When the drillstring is lowered and the bit touches the bottom, more weight is applied to the bit (and less is measured as hanging at the surface). Accordingly, the WOBA is set by the operator to match the characteristics of the formation that must be drilled. Notice that the WOBA is *calculated* from HKLA, RPMA, the drillstring weight, and mud flow variables (hydrostatic force). As such, it is considered reliable only during drilling.

RPMA is the (average) rotational speed of the drillstring, that results from the torque TQA applied by the surface equipment (top drive or rotary table) to spin the drillstring. TRPM is the (average) rotational speed of the bit, which includes the additional rotation applied at the BHA. TQA is a control parameter that the drilling crew employs to maintain the desired rate of penetration (ROPA).

SPPA is the (average) standpipe pressure, which measures the pressure of the mud on the standpipe. The crew sets a target value for the standpipe pressure in view of the lithological information about the wellbore, and the rig automatically controls the pumps in order to maintain the SPPA set point. SPMT denotes the total number of strokes per minute of all the pumps, STKC being equal to its integral over time.

Lithology data concern the geological composition of the formations traversed by the wellbore. These data are not available in real-time but after some delay (the lithology type corresponding to each depth section is determined in practice by examining the mud and cuttings produced during drilling activities). Preliminary studies indicate that lithology data report strata that are heavily fragmented. Trajectory data contain information about the trajectory of the wellbore, such as the inclination (deviation from vertical) and the azimuth angle. Again, trajectory data are not always available in real-time, depending on the tool used for their acquisition.

3. Stuck-pipe precursors

Several studies (see, *e.g.*, [2,6,15]) point out the presence of

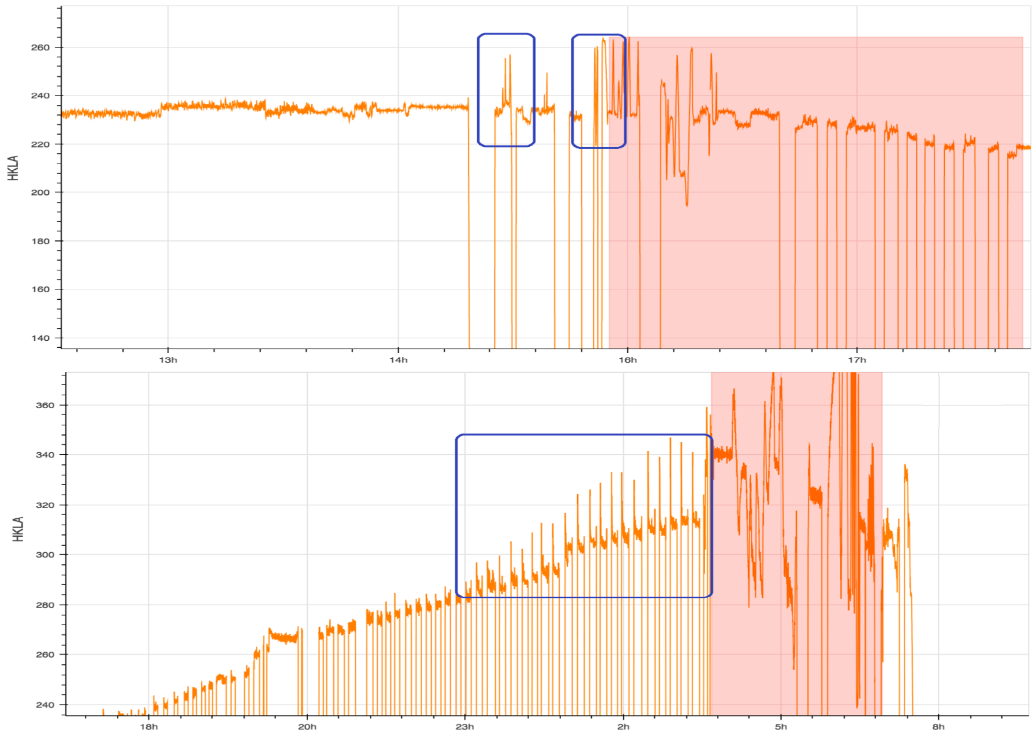

Fig. 1. Two different cases of HKLA overpulls before a sticking event (emphasized in pink). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

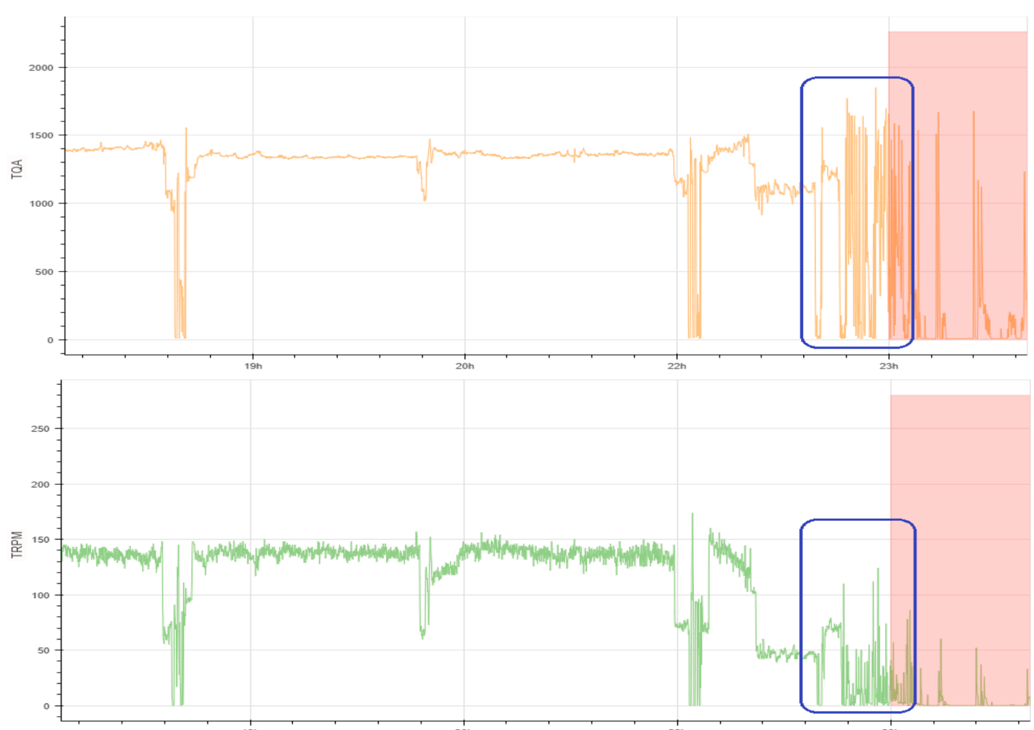

Fig. 2. Torque (TQA, top) peaks and corresponding rotation drops (TRPM, bottom) before a sticking event (emphasized in pink). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

precursors of pipe sticking events. In particular, [15] lists various warning signals for specific drilling issues, directly or undirectly connected to stickings. For example, a differential sticking can be anticipated by an increasing overpull at connections and an increase in torque, while circulation is possible. This event is also favored by conditions such as a stationary pipe, or the BHA being adjacent to a thick sand. Some possible warning signals for inadequate hole cleaning are pump pressure increases and spikes, a reduced overall pull when pumping, excessive overpull on connections and trips, a loss of circulation.

In general, erratic or increasing overpulls, torque increases, and

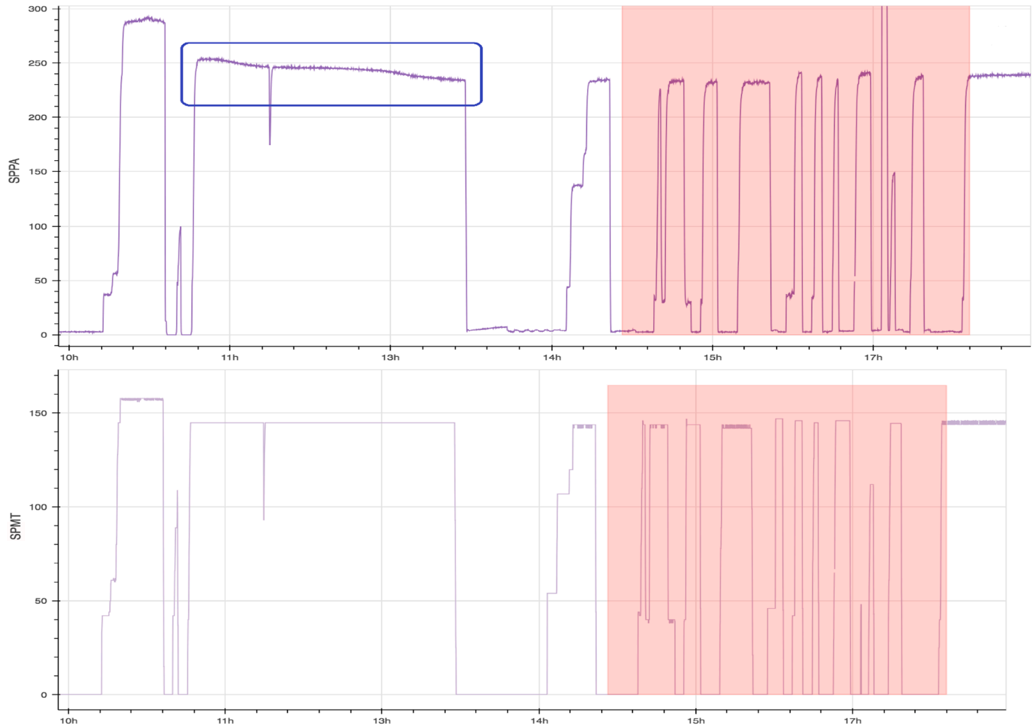

Fig. 3. Slow decrease of the pressure (SPPA, top) in the absence of variations of the mud flow rate (SPMT, bottom) before a differential sticking event (emphasized in pink). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

pressure increases (see Figs. 1–3) are associated to various drilling issues, and are often observed prior to stickings. For example, in the available dataset of 13 wells with 20 sticking events (see Section 5), anomalies related to the HKLA prior to stickings were observed 50% of the times, while torque/rotation and pressure problems appeared 30% and 15% of the times. Besides these warning signals, there are also other elements facilitating a sticking, such as large doglegs, highly deviated wells, and certain types of formations.

A data-based detection of the mentioned warning signs is not always trivial. For example, overpull events cannot always be detected using the signals available through the mudlog. An overpull is the amount of extra force that must be exerted on the pipe to pull it upward, above and beyond its own weight, due to drag and other forces. In principle, it can be computed with the following formula:

$$Overpull = HKL - BF{\cdot}WS, \tag{1}$$

where *HKL* is the measured hook load and *BF*·*WS* is the theoretical hook load, obtained as the product of the buoyancy factor *BF* (used to compensate the loss of weight due to immersion in the drilling fluid) and the weight of the string *WS*. Unfortunately, the information required to calculate *BF* and *WS* (*i.e.*, the mud weight and the drillstring composition and weights, respectively) is not included in the available dataset. In principle, the occurrence of an overpull can be extrapolated by the timelog and some researchers have designed a method for solving this task using natural language processing [1]. This approach has not been attempted here considering the already discussed scarce reliability of the timelog. Conversely, one can try to detect overpulls using the mudlog data, *e.g.*, by computing the difference between neighboring values of HKLA, and looking for large differences before and after a stand connection, or spikes in general.

Another example of precursor event that is difficult to detect from data is the pack off. A pack off occurs when the wellbore around a drillstring becomes obstructed. This can happen for a variety of reasons, the most common being that either the drilling fluid is not properly transporting cuttings out of the annulus or portions of the wellbore wall have collapsed around the drillstring. When the well packs off, there is a sudden reduction or loss of the ability to circulate, and high pump pressures follow. The reaction must be swift, but if the remedial action is not successful, a stuck pipe episode can result. Usually, pack off events are detected by analysing the mud flow rate and the shape of the cuttings. Unfortunately, the latter information is not included in the mudlog, whereas the measurement of the mud flow at the output is considered unreliable for the database available in this study.

The reliability of the precursor events is fundamental in determining the overall accuracy of the stuck-pipe predictor developed in this work, as they constitute its inputs. Some of them are based on thresholds, which are obviously critical design parameters of the process. Their correct detection also depends on a careful identification of stand change operations, which is not trivial to achieve from the available data (although it would be easy to mark them automatically in the data stream during drilling). For example, recognizing correctly these operations is crucial for the detection of overpulls at connections, since these typically take place immediately after a stand change. Stand change detection is not discussed here for reasons of space.

On the positive side, most of these events (if not all) are relatively straightforward to detect using very simple formulas, which is by itself very appealing. It is also important that these events are of direct significance to the process expert, and represent mild anomalies to be carefully monitored.

In the next subsections we define detectors for the main precursor events, that are mainly derived from a subset of mudlog signals, which are consistently available and reliable (at least in our dataset), namely HKLA, BPOS, TQA, RPMA, SPPA and SPMT.

3.1. HKLA-related precursors

The overpull at connections is one of the most frequent precursors of a stuck pipe event. This phenomenon is a warning sign that may indicate inadequate hole cleaning, mechanical stability problems, difficulties due to overpressure, mobile or unconsolidated formation [15]. To spot an overpull, we compare the HKLA values immediately before and after a

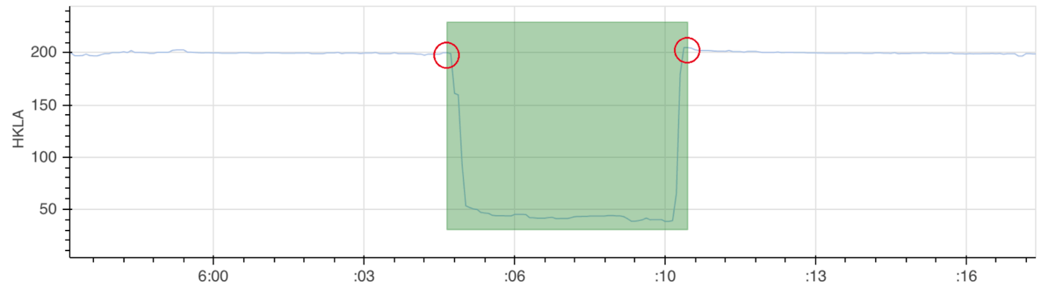

Fig. 4. The two values compared to establish the presence of an overpull.

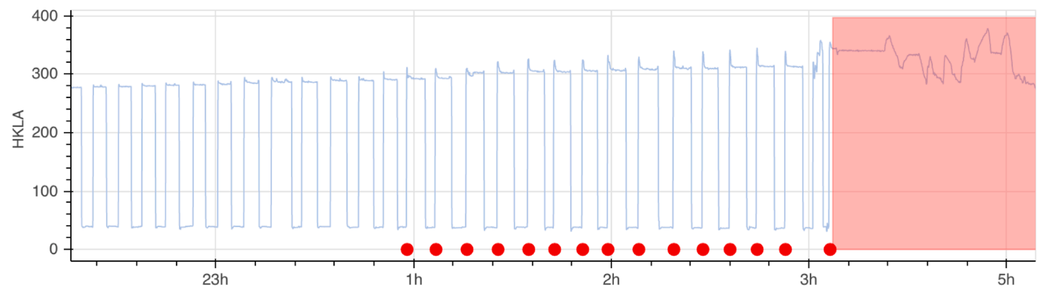

Fig. 5. Detection of overpull events prior to a sticking event.

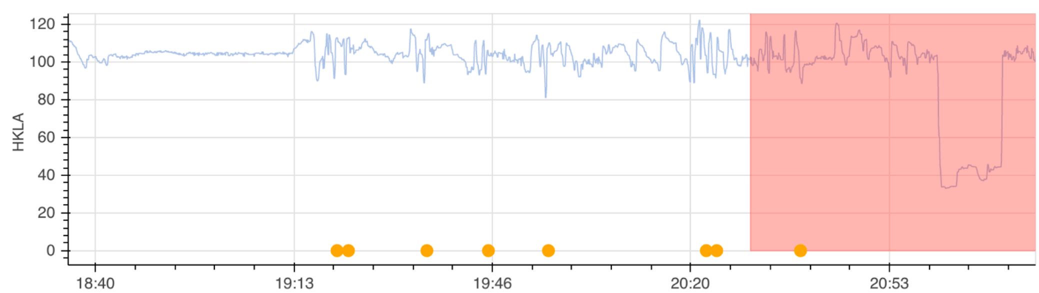

Fig. 6. Identification of variance peaks of HKLA.

connection (see Fig. 4), and issue an *Overpull* warning if the difference between these two values exceeds an heuristic threshold of 15 tons. An example of detection of overpulls before a sticking episode is shown in Fig. 5.

An erratic behavior in the HKLA can also indicate the presence of an anomalous drag between the drill string and the wellbore. For example, spikes of this signal may correspond to an overpull event. For these reasons we also implemented a detector (denoted *HKLA anomalies*) that looks for peaks in the variance of HKLA (the variance is calculated with a rolling window approach), see Fig. 6. During the stand change activity, the detector is idle to avoid misidentifications.

A third index exploits the fact that the HKLA and the BPOS signals become highly correlated when the pipe gets stuck, as acting on BPOS reflects directly on the load measured on the hook, whereas they are essentially uncorrelated in normal conditions (if the drillstring is not impeded, moving it up and down does not affect the hookload). The *HKLA-BPOS correlation index* has been computed as follows: the high frequency content of each signal is first extracted by subtracting the baseline (calculated using a rolling median with a window size of 11 samples, corresponding to 55 s) from the raw signal, and the correlation between the two residuals is computed with a rolling window of 25 samples (125 s). Finally, an event is present if the absolute value of the HKLA-BPOS index is greater than 0.8 (outside stand changes).

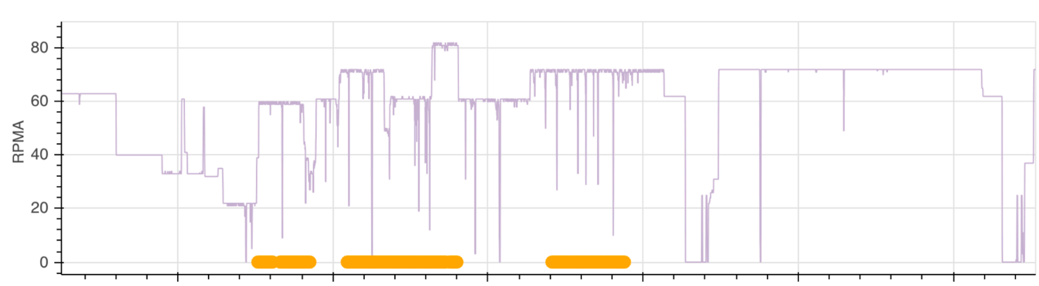

Fig. 7. Identification of stick & slip events on RPMA prior to a sticking.

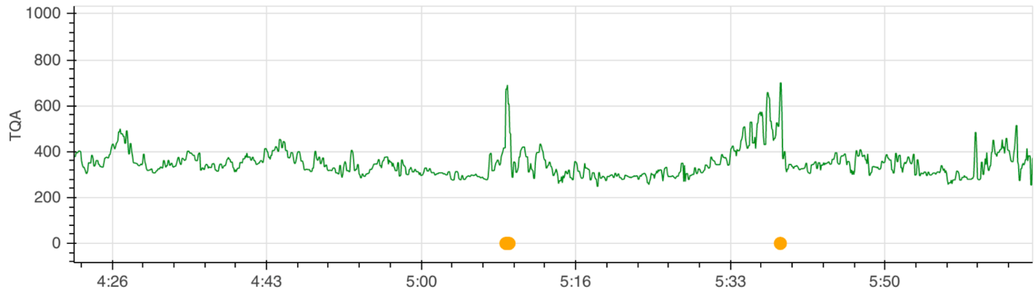

Fig. 8. Two peaks of TQA are signalled by the detector.

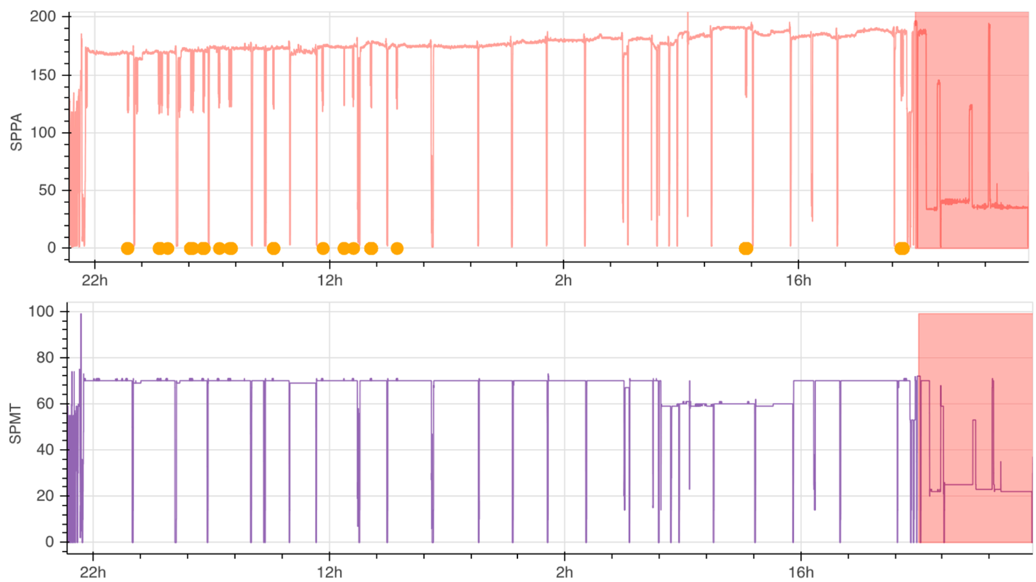

Fig. 9. SPPA-SPMT correlation index evaluated before a sticking event: SPPA and index (top) and SPMT (bottom).

3.2. Rotation-related precursors

While drilling the drillstring is subject to rotation; when resistance to rotation is encountered, the RPMA signal typically drops while TQA takes larger values due to the mechanical torsional stress. This type of phenomenon often takes the form of sudden rotation speed drops during drilling (*i.e.*, when $ROPA > 1$), also known as stick & slip. To target these events, we consider the residual obtained by subtracting a rolling median of the RPMA (with a window of 11 samples) from the raw signal, to remove the signal baseline and emphasize the high frequency behavior.

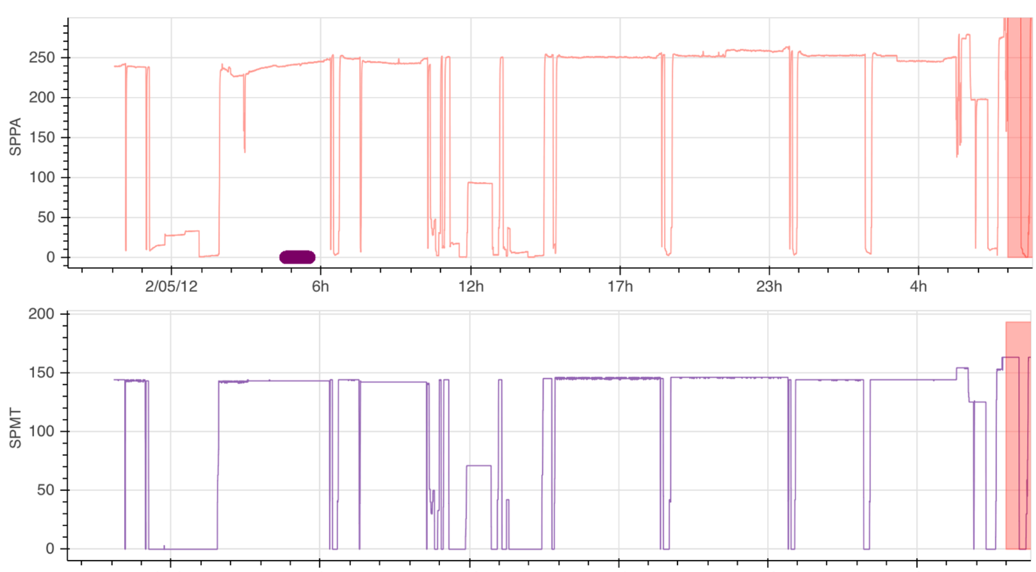

Fig. 10. SPPA slow variation index evaluated before a sticking event: SPPA and index (top) and SPMT (bottom).

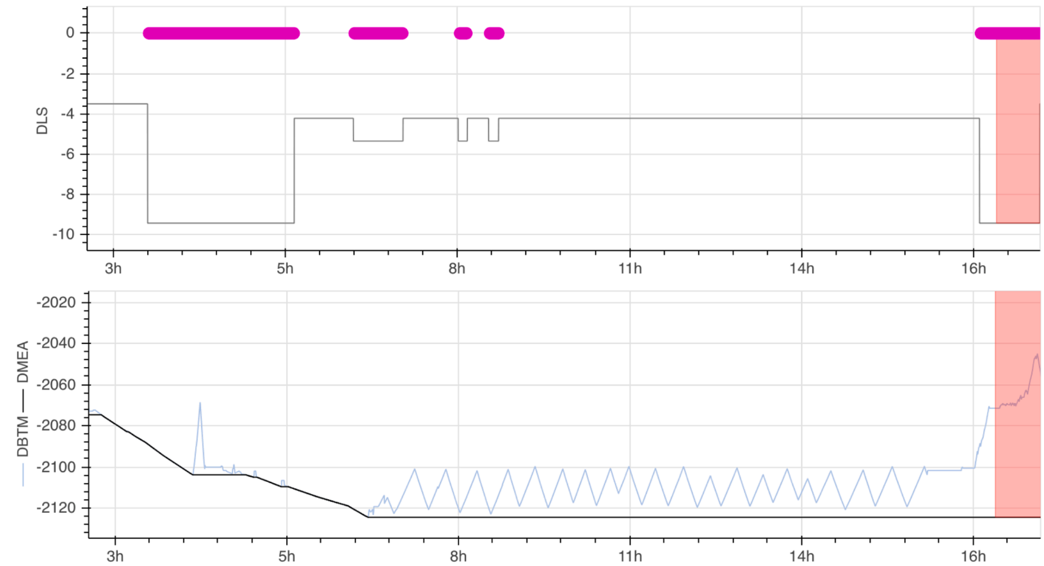

Fig. 11. The *Trajectory variation* detector highlights the presence of abrupt deviations in the trajectory: DLS and detector (top), DBTM and DMEA (bottom).

A *Stick & Slip* event is then generated if there are more than 5 significant values of the residual over a sliding window of 72 samples (360 s), where a value is considered significant if greater than a given threshold (*e.g.*, 15). The performance of this detector is shown in Fig. 7.

The other variable that may reveal rotation issues is TQA, which represents the necessary torque to maintain RPM at the planned value. When the bit is drilling a mixture of different rocks, an erratic behavior in this signal is expected and usually it does not imply any trouble, as long as the data remain in the neighborhood of their baseline. To identify anomalies we compare two low pass filtered versions of TQA, one obtained with a 5-sample moving average filter and the other with a rolling median on a sliding window of 60 samples, both normalized with respect to the maximum value of the signal. A *TQA peaks* event is issued if the difference is greater or equal to an heuristic threshold (0.1) and it is not due to a stand change (see Fig. 8).

Finally, we also consider a combined detector, denoted *TQA-RPMA index*, based on the ratio $\frac{TQA}{RPMA+1}$ (the 1 at the denominator avoids a possible division by 0), normalized wrt. to its largest value. An event is issued when this ratio exceeds a given threshold (*e.g.*, 0.55).

3.3. Pressure-related precursors

In nominal conditions the SPPA and SPMT signals show a similar trend. Accordingly, a lack of correlation between these two signals could indicate a pressure problem. Two indices have been designed to highlight this type of problem.

The first index (*SPPA-SPMT correlation index*) is set to 1 if the standard deviation of SPPA exceeds a given threshold (a value of 20 has been used in this study), and at the same time, SPMT is essentially constant. This last condition is characterized by checking that the standard deviation of SPMT is sufficiently small (*e.g.*, less than 1), while the mean is not negligible (*e.g.*, greater than 10). All the mentioned variables are calculated using a rolling window approach, with windows of 11 samples (55 s) for SPPA and 101 samples (505 s) for SPMT. An example of detection of anomalies using this indicator is shown in Fig. 9.

A second index (denoted *Slow variation*) aims at spotting significant slow variations of SPPA (*e.g.*, drifts), while SPMT is constant. Only data with low standard deviation in both SPPA (less than 1) and SPMT (less than 3) are considered. Both standard deviations are calculated with a rolling window approach, using windows of 400 and 700 samples, respectively. For these data the slope of the SPPA signal is estimated by taking the 400 sample difference of a rolling median of SPPA calculated using a window of 200 samples). An event is issued if the slope exceeds a given threshold for at least 30 out of 70 samples (350 s). The behavior of this indicator is illustrated in Fig. 10.

3.4. Trajectory-related precursors

Wells are often deviated from the vertical orientation to increase the exposure of the producing zones, to intersect a larger number of fractures or to follow a complex geological structure. However, severe doglegs must be avoided to prevent stickings due to key seating. Moreover, the presence of unplanned changes of direction of the well path can highlight difficulties during the drilling activity. The DogLeg Severity (DLS) is a measure of the rate of change of direction of the well path. More specifically, DLS is a normalized estimate of the overall wellbore curvature between two consecutive directional surveys, expressed in degrees per 30 meters. It can be computed as follows:

$$DLS = [arccos(\cos(i_1)\cos(i_2) + \sin(i_1)\sin(i_2)\cos(a_2 - a_1))]\frac{100}{DMD}, \quad (2)$$

where DMD is the derivative of the measured depth between the survey points, i_1 and i_2 are the inclination angles at the upper and lower surveys, and a_1 and a_2 are the azimuth directions at the upper and lower surveys. A *Trajectory variation* warning event is issued whenever the absolute value of DLS is greater than 5. The presence of numerous consecutive events can highlight the difficulty of the drillstring to cross a high curvature in the well (see Fig. 11).

4. HMM-based prediction model

In [14], an HMM is used to identify and recognize patterns of fault events that lead to failures. The fundamental assumption in this approach is that the occurrence of failures can be predicted by identifying special patterns of faults the system is experiencing. The prediction of failures involves two steps: first the current state is estimated on the basis of the recent fault events. Starting from the current state, the future behavior is extrapolated and the failure probability is computed. Faults are modeled as unobservable hidden states producing symbols corresponding to errors, whereas failures are special hidden states that do not produce error messages.

Following the general idea of [14], we also employ an HMM for the prediction of stuck-pipe events, the hidden states of the model being the manifestations of some specific events detectable from the mudlog data.

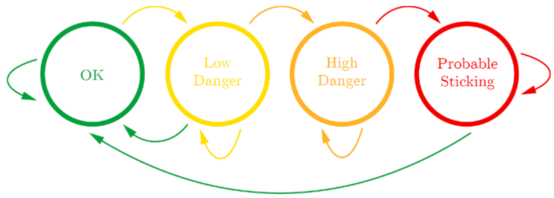

Fig. 12. HMM model structure.

The hidden states are the outputs of the probabilistic system learned by the HMM and correspond to four different levels of sticking risk:

- *Probable Sticking (ST)*: This is the most dangerous state, corresponding to all the data in the time frames associated to sticking events reported in the timelog.
- *High Danger (HD)*: This label is applied to data within the three hours period prior to a sticking event.[2]
- *Low Danger (LD)*: This label is assigned to all the data that are within an hour in advance of a data section marked High Danger, or which are approximately located at the same depth (*i.e.*, in a range [−10 m, 50 m]) of a sticking event, and take place prior to its occurrence. The same label is assigned to data belonging to windows of 120 samples with more than 15 detected events.
- *Normal state (OK)*: This label identifies the safe condition. This state contains all the remaining data.

Having labeled accordingly all the data available for training we have built the HMM depicted in Fig. 12. In particular, if the well is in the OK state, it can remain in the safe condition or upgrade to the Low Danger state. Starting in the Low Danger state, the level of risk can either remain the same, increase to High Danger, or decrease back to OK. The Low Danger state indicates an alert for the drilling crew: if no prompt operations are performed the danger state may degenerate to a High Danger condition, from which no improvement can take place and a sticking is inevitable. When the model is in the Probable Sticking state, it can only linger in that state or return to the safe condition, if the problem is eventually solved by human intervention.

The observations in the HMM are the precursor events discussed in Section 3:

- *HKLA problems (HKL)*: overpulls at connections, HKLA variance peaks, and high correlation events between HKLA and BPOS.
- *Rotation problems (ROT)*: stick & slip events, TQA peaks and TQA-RPMA ratio.
- *Pressure problems (PRE)*: low correlation between SPPA and SPMT, slow variations in SPPA.
- *Trajectory problems (TRA)*: presence of a high curvature in trajectory.

Starting from these precursor events, a 4-element binary vector is defined as a function of time, which has one value for each group of precursors. The bit is set to 1 when at least one event of the corresponding category has occurred in the previous 3 min time frame.[3] These vectors constitute the inputs to the HMM training algorithm.

The transition matrix T contains all the transition probabilities, *i.e.*, it defines the probability of moving from one state to another:

$$T = \begin{bmatrix} p_{OK,OK} & p_{OK,LD} & 0 & 0 \\ p_{LD,OK} & p_{LD,LD} & p_{LD,HD} & 0 \\ 0 & 0 & p_{HD,HD} & p_{HD,ST} \\ p_{ST,OK} & 0 & 0 & p_{ST,ST} \end{bmatrix}, \quad (3)$$

where the probability indices refer to the hidden states with obvious notation (p_{X_1,X_2} denotes the probability of reaching state X_2 from state X_1). Notice that some transitions have been set to zero according to the structure of the model (as depicted in Fig. 12). Furthermore, by construction, each row sums up to 1. Each probability has been estimated as an empirical probability, *i.e.*, as the relative frequency of each change of state.

The emission matrix E defines the probability of an observation being generated from a state. Observations are assumed to be distributed as a Bernoulli distribution, p_{X_j,O_k} defining the probability of emitting O_k when the system is in state X_j. The emission probability p_{X_j,O_k} is also computed with the empirical probability. The relative frequency is calculated for each event O_k that occurs in the hidden state X_j. Notice that in this case rows do not sum up to 1. The emission matrix has the following structure:

$$E = \begin{bmatrix} p_{OK,HKL} & p_{LD,HKL} & p_{HD,HKL} & p_{ST,HKL} \\ p_{OK,ROT} & p_{LD,ROT} & p_{HD,ROT} & p_{ST,ROT} \\ p_{OK,PRE} & p_{LD,PRE} & p_{HD,PRE} & p_{ST,PRE} \\ p_{OK,TRA} & p_{LD,TRA} & p_{HD,TRA} & p_{ST,TRA} \end{bmatrix}, \quad (4)$$

with straightforward interpretation of the indices.

Given T and E and the sequence of observations $O = [o_1, o_2, \ldots, o_t]$, one needs to determine which sequence of hidden states $X = [x_0, x_1, x_2, \ldots, x_t]$ is the (most probable) underlying source of the sequence of observations (also known as decoding task). By predicting the hidden state, one can estimate the risk level of the well with respect to the sticking.

The most common decoding algorithms for HMMs is a dynamic programming method known as the Viterbi algorithm [7]. The idea is to process the observation sequence from left to right, filling out a so called "trellis", each cell of which, $v_t(j)$, represents the probability that the HMM is in the jth state after seeing the first t observations and passing through the most probable state sequence $x_0, x_1, \ldots, x_{t-1}$, given the automaton Λ associated to the HMM:

$$v_t(j) = \max_{x_0,x_1,\ldots,x_{t-1}} P\left(x_0, x_1, \ldots, x_{t-1}, o_1, o_2, \ldots, o_t, x_t = X_j \middle| \Lambda\right)$$

The value of each cell $v_t(j)$ is computed by recursively taking the most probable path that could lead to this cell, *i.e.*, by taking the maximum over all possible previous state sequences. The algorithm works recursively: once the probability of being in every state at time $t-1$ is computed, one can compute the Viterbi probability by taking the most probable path extension leading to the current cell at time t. For a given state X_j at time t, the value $v_t(j)$ is computed as

$$v_t(j) = \max_{i=1,\ldots,N} v_{t-1}(i) T_{ij} E_{jk},$$

where $v_{t-1}(i)$ is the Viterbi path probability at the previous time step, T_{ij} is the transition probability from the previous state $x_{t-1} = X_i$ to the current state $x_t = X_j$, and E_{jk} is the likelihood of observing O_k in the state X_j.

The model has been tested using an *ad hoc* cross-validation approach. Namely, to test the model on a specific well, the data of that well were removed from the dataset for the HMM training phase. In other words, the two probability matrices and the set of observations are extracted from all the available wells, except the tested one. This method was termed *leave-one-well-out*. Furthermore, the data of wells in which a sticking occurred without showing any earlier warning event were also excluded from the training set, but they have been used for validation.

[2] According to drilling experts, a three hours interval is adequate to allow an effective human intervention.

[3] Various values have been tried for the time window size. With values of 20–30 min, the algorithm is not able to predict the sticking events. Conversely, if it is too short (*e.g.*, below 2 min), many false alarms are issued. A 3 min time frame resulted in a satisfactory trade-off.

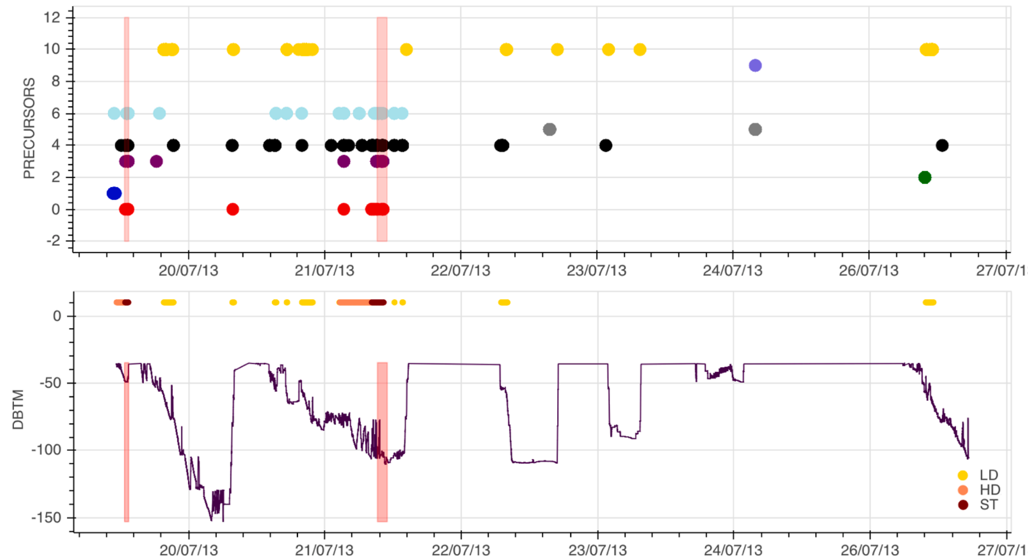

Fig. 13. Well n.7: detected precursors (top), DBTM and predicted level of risk (bottom). Sticking events are emphasized in pink. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

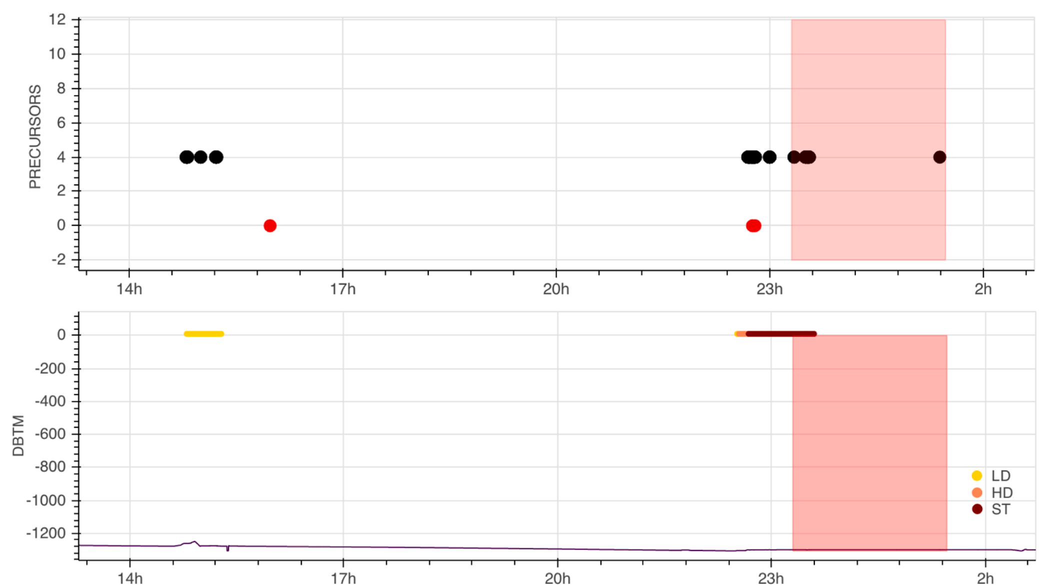

Fig. 14. Well n.2: detected precursors (top), DBTM and predicted level of risk (bottom). Sticking events are emphasized in pink. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

5. Experimental results

The proposed HMM and the associated indices have been tested on a dataset of 13 wells, with a total of 20 sticking events documented in the timelogs. Our model was able to correctly predict 12 sticking events, while the missed stickings are essentially due to the lack of sufficient precursors. For this reason, 5 of these episodes were not included in the training set. All the same, in 6 cases out of 8, at least an LD warning sign was issued. Recall that each well is tested on a model that has been trained on a dataset that does not include any data from that well.

5.1. Prediction examples

The following pictures (Figs. 13–14) show some successful examples. The detected precursors are displayed in the top subfigure, in the following order (from bottom to top):

0. HKLA peaks and high variance portions (red)
1. Stick & slips (blue)
2. SPPA-SPMT low correlation events (green)
3. TQA peaks (purple)
4. HKLA-BPOS high correlation events (black)
5. Slow SPPA variations (grey)
6. Large TQA-RPMA ratio values (light blue)
7. Overpulls at connections (blue-violet)
8. Abrupt variations of the trajectory (magenta)
9. Pack-offs (blue-grey)
10. Alternating layers of sand and shale (yellow)

Both stickings in the first well (see Fig. 13) are anticipated by similar precursor events, mainly related to rotation problems (TQA-RPMA index, TQA peaks) or other issues associated with the HKLA signal (see the HKLA-BPOS index, and the HKLA peaks indicator). The prediction of the sticking of the second well is triggered by the frequent presence of

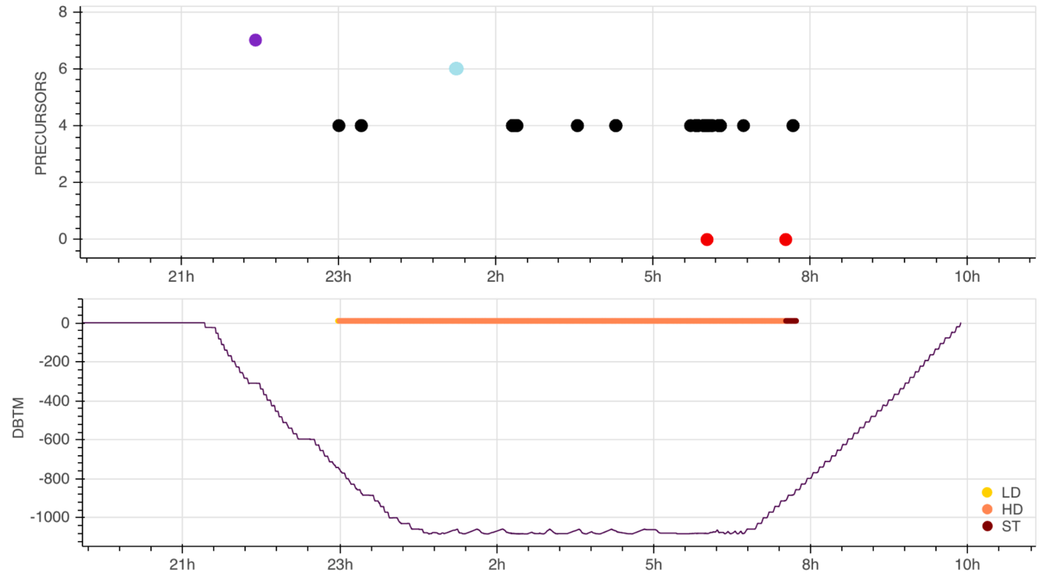

Fig. 15. Well n.4: detected precursors (top), DBTM and predicted level of risk (bottom).

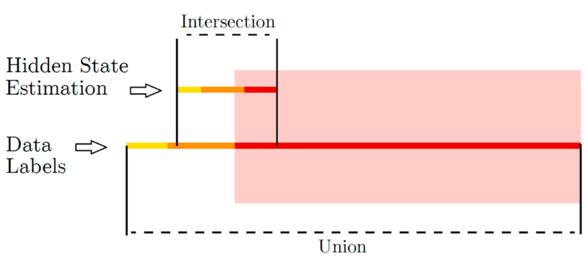

Fig. 16. The Jaccard index is low but the stuck pipe is well predicted.

troubles related to the HKLA signal, especially by its high correlation with the BPOS signal. In all three stuck-pipe cases, the system enters the High Danger state well in advance of the sticking, which would have effectively given the drilling crew enough notice to prevent the event.

Besides the few mentioned failed sticking predictions, a few "false" alarms are also issued. These alarms are often justified in the light of the number of precursor warnings detected. Indeed, it is very well possible that the drilling crew, seeing these signals, and reacting promptly has actually avoided more severe troubles. An example is shown in Fig. 15, where a High Danger state does not anticipate any actual sticking. Various warnings are issued during the observed period, associated to a high correlation between HKLA and BPOS. These difficulties have been so significant that drilling was interrupted for some time (the DBTM is almost constant for 6 h), in reaction to this dangerously evolving condition. Notice also that if the drill pipe is not moved for a long time, the filter cake tends to build up around it and then adds up to the differential sticking force that is holding the drill collars. Presumably, in this case, the operator applied the appropriate manoeuvres to get out of the danger condition.

5.2. Overall evaluation

Common machine learning evaluation techniques cannot be adopted in our case due to the large imbalance of the dataset and the noisy ground truth. Indeed, the target events (the stickings) are highly infrequent. Furthermore, the ground truth is noisy because not all the sticking events have been actually reported in the timelog. For the above reasons, an evaluation metric computed in terms of the standard true/false positive/negative indices is unsuitable. Also the Jaccard index[4] between the model output and the data labels is not really effective in our case. Indeed, the above cited index is computed including the entire time frame in which the sticking occurs, which implies that successful predictions could be evaluated as poor. An example of this is reported in Fig. 16.

The following Table 2 summarizes the results of our experiments. The first three columns (well number, phase number, sticking number)

Table 2
Sticking prediction performance.

Well	Phase	Sticking	Risk level	Outcome
1	4	1	LD	Problems found
1	4	2	HD	Predicted
1	4	3	HD	Predicted
1	5	4	HD	Predicted
2	1	1	ST	Predicted
3	1	1	ST	Predicted
3	2	2	OK	Not predicted
3	2	3	HD	Predicted
4	3	1	LD	Problems found
5	1	1	ST	Predicted
5	2	2	ST	Predicted
6	1	1	ST	Predicted
7	1	1	HD	Predicted
7	2	2	ST	Predicted
8	1	1	ST	Predicted
9	1	1	LD	Problems found
10	1	1	LD	Problems found
11	1	1	LD	Problems found
12	1	1	OK	Not predicted
13	1	1	LD	Problems found

[4] The *Jaccard index*, also known as *Intersection over Union* and the *Jaccard similarity coefficient*, is a statistic used for comparing the similarity and diversity of sample sets. The Jaccard coefficient measures the similarity between finite sample sets, and is defined as the size of the intersection divided by the size of the union of the sample sets.

identify the sticking event. The fourth column shows the corresponding estimated risk level, and the last one reports the prediction outcome. A sticking is considered to be predicted when it is *preceded* by an HD or an ST state. It is partially predicted if anticipated by an LD alarm, and not predicted at all if it is encountered when in the OK state. The last 5 wells in the table report stickings that are preceded by a scarce number of precursors, and have been excluded from the training set. All the same, in all cases but one, at least an LD warning sign was issued.

6. Conclusion

In this paper a new methodology was proposed to address the problem of predicting stuck pipes events in oil & gas drilling operations. A model is developed that monitors the drilling process associating three increasing levels of alert to the current condition. The danger states are triggered by the presence of particular trouble events that are often observed before stickings and are detectable from reliable mudlog signals, and few other data. The precursor warnings are related to the alarm levels by means of an HMM.

The performance of the model is quite promising, with 60% of the stickings predicted with sufficient advance to perform some corrective action. Most of the other stickings are characterized by an almost absence of precursor warnings and are therefore difficult to predict with an event-based approach. Nonetheless, a large percentage of these cases are still associated with a low alarm level.

Given the diversity of the physical phenomena associated to the stickings, one method alone is not sufficient to optimally solve the prediction task, but an ensemble method is necessary. For example, events related to the mud density and flow measures were not considered in this study, due to the lack of reliable signals, although these are considered important precursors especially for differential stickings. Furthermore, the proposed model is only time-based, while often useful information to predict a sticking in a tripping phase can be gathered from an analysis of what happened at that depth well before the sticking actually occurred.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work is part of an ongoing project in ENI concerning big data analytics for NPT (non-productive time) prediction and operational efficiency.

References

[1] M. Antoniak, J. Dalgliesh, M. Verkruyse, J. Lo, Natural language processing techniques on oil and gas drilling data. SPE Intelligent Energy International Conference and Exhibition, Society of Petroleum Engineers, Aberdeen (UK), 2016, pp. 1–6.
[2] L. Bailey, T. Jomnes, J. Belaskie, J. Orban, M. Sheppard, O. Houwen, S. Jardine, D. McCann, Stuck pipe: causes, detection and prevention, Oilfield Rev. 3 (4) (1991) 13–26.
[3] O. Bello, J. Holzmann, T. Yaqoob, C. Teodoriu, Application of artificial intelligence methods in drilling system design and operations: a review of the state of the art, J. Artif. Intell. Soft Comput. Res. 5 (2) (2015) 121–139.
[4] W. Bradley, D. Jarman, R. Plott, R. Wood, T. Schofield, R. Auflick, D. Cocking, Task force approach to reducing stuck pipe costs. SPE/IADC Drilling Conference, Society of Petroleum Engineers, 1991.Amsterdam, Netherlands
[5] A. Brankovic, M. Matteucci, M. Restelli, L. Ferrarini, L. Piroddi, A. Spelta, F. Zausa, Data-driven indicators for the detection and prediction of stuck-pipe events in oil&gas drilling operations, Upstream Oil Gas Technol. (2021).(in print)
[6] T. Bratton, S. Edwards, J. Fuller, J. Murphy, S. Goraya, T. Harrold, H. Nicholson, W. Standfird, B. Wright, Avoiding drilling problems, Oilfield Rev. 13 (2) (2001) 32–51.
[7] G.D. Forney, The viterbi algorithm, Proc. IEEE 61 (3) (1973) 268–278.
[8] T.O. Gulsrud, R. Nybø, K.S. Bjørkevoll, Statistical method for detection of poor hole cleaning and stuck pipe. SPE Offshore Europe Oil & Gas Conference & Exhibition, Society of Petroleum Engineers, 2009.Aberdeen, U.K.
[9] K. Salminen, C. Cheatham, M. Smith, K. Valiulin, Stuck pipe prediction using automated real-time modeling and data analysis. IADC/SPE Drilling Conference and Exhibition, Society of Petroleum Engineers, 2016.
[10] M.A. Muqeem, A.E. Weekse, A.A. Al-Hajji, Stuck pipe best practices - a challenging approach to reducing stuck pipe costs. SPE Saudi Arabia Section Technical Symposium and Exhibition, Society of Petroleum Engineers, 2012.
[11] M.M. Nezhad, S. Ashoori, P. Hooshmand, M. Mirzaee, Stuck drill pipe prediction with networks neural in maroon field, J. Basic Appl. Sci. Res. 2 (6) (2012) 5570–5575.
[12] R. Nybø. Efficient drilling problem detection, NTNU, Trondheim, Norway, 2009. Ph.D. thesis.
[13] H. Rostami, A.K. Manshad, A new support vector machine and artificial neural networks for prediction of stuck pipe in drilling of oil fields, J. Energy Resour. Technol. 136 (2) (2014) 1–4.
[14] F. Salfner, Predicting failures with hidden Markov models. 5th European Dependable Computing Conference, 2005, pp. 41–46.
[15] Schlumberger, Stuck-pipe prevention: self learning course, Schlumberger, Sugar-Land Learning Center, 1997.
[16] S.R. Shadizadeh, F. Karimi, M. Zoveidavianpoor, Drilling stuck pipe prediction in iranian oil fields: an artificial neural network approach, Iranian J. Chem. Eng. 7 (4) (2010) 29–41.
[17] R. Vilalta, S. Ma, Predicting rare events in temporal domains. IEEE International Conference on Data Mining, 2002, pp. 474–481.Maebashi City, Japan
[18] G.M. Weiss, H. Hirsh, Learning to predict rare events in event sequences. 4th Int. Conference on Knowledge Discovery and Data Mining, 1998, pp. 359–363.New York (NY), USA
[19] J. Zhang. The impact of shale properties on wellbore stability, The University of Texas at Austin, 2005. Ph.D. thesis.
[20] J. Zhao, Y. Shen, W. Chen, Z. Zhang, S. Johnston, Machine learning-based trigger detection of drilling events based on drilling data. SPE Eastern Regional Meeting, Society of Petroleum Engineers, Lexington (Kentucky), USA, 2017, pp. 1–11.