SPE 112212

Spotting a False Alarm—Integrating Experience and Real-Time Analysis With Artificial Intelligence

Roar Nybø, Knut S. Bjørkevoll, and Rolv Rommetveit, SINTEF Petroleum Research

This paper was prepared for presentation at the 2008 SPE Intelligent Energy Conference and Exhibition held in Amsterdam, The Netherlands, 25–27 February 2008.

Abstract

Time series data from past drilling operations are an under-utilized resource in oil-companies. Time series record processes in the well, effects of machinery on the rig, and presents an opportunity to improve upon the models underlying drilling simulations as well as alarm systems.

We present a proof-of-concept for reducing false kick alarms during drilling by combining a physical model with techniques from artificial intelligence. We show that artificial intelligence can be used to learn from the experience implicit in time series data. It then corrects for limitations in the physical model, which results in a more accurate prediction of the mud flow out of the well, and subsequently fewer false kick alarms.

Introduction

Early detection of gas influx into a well traditionally depends on measurements of the flow out and total active volume. An alarm system might take the volume of the mud pits as input and sound an alarm if the change in volume or rate of change exceeds a given threshold. Direct measurements of the flow rate in and out of the well may also be used. This may have the advantage of an earlier detection of gas influx, but relies on accurate flow meters. Typical active volume threshold values are 5-10 bbl and 3 bbl for HPHT wells.

Changes in flow may have a number of benign causes, some of which third generation alarm systems try to take into account [1]. Still, inaccurate prediction of these effects as well as measurement noise, means that false alarms happen frequently on the rig. Such false alarms may lead to non-productive down-time and be a security risk in themselves, as they may divert the attention of the drilling team and erode trust in the alarm system.

The sensitivity of an early-warning system, i.e. how early it can detect a kick, depends on the threshold value. A lower value gives both greater sensitivity and more false alarms, which in practice need to be balanced against each other. False alarms are therefore a major hindrance for early detection of kicks.

More advanced models of the well-rig system is one answer to this problem. Accurate calculations of thermal expansion of the mud and mud column compressibility add improvements. However, more advanced models tend to require more information about the rig and well to be entered into them. Even higher levels of precision may require knowledge of peculiarities of the machinery on the rig, or of other effects that are poorly understood at the moment.

On the other hand, many of these effects will make their mark on the time series recorded during drilling. A promising strategy may therefore be to use the experience inherent in time series from previous wells to increase the precision of our models. The field of Artificial Intelligence (AI) offers a set of tools for “learning” dependencies between variables and making predictions based on past recordings. In this paper we demonstrate a simple combination of AI and traditional modeling which yields fewer false alarms than each method alone.

Alarm system

For the purposes of our simulations, we define a typical alarm system as follows: If the pit volume has changed by more than the threshold value over the last 20 minutes, the alarm state is “on”. Volume change is Flow in – Flow out integrated over time. Note that the 20 minute window does not rule out detecting volume changes over a few seconds, this is mainly dependent on the threshold value.

For a stretch of the time series known to contain no kicks, the percentage of samples where the alarm state is on is equal to what we call the false alarm rate. This serves as a measure of the time the drilling crew must either take unnecessary precautions or ignore the alarm system. On the other hand, it does not mirror the number of "alarm events" on the rig, as this would be dependent on when the driller chooses to reset the alarm system.

Physical model

As our physical model, we employ an advanced dynamic flow model. This model includes physical effects like changes in frictional pressure loss on operational changes, with transition from laminar to turbulent flow, changes in temperature, and changes in fluid volume due to compression and thermal expansion of the fluid.

Artificial Intelligence

Artificial neural networks (ANN) are a common component in many AI systems. These are networks of simple circuits inspired by neurons in the brain. It is a general opinion that ANNs are especially suited to noisy data and ANNs have been repeatedly applied to tasks in the petroleum industry over the last 20 years. Examples include [2-5].

An ANN is trained by presenting it with a set of example inputs and the correct output or answer to each example. This set of examples is called the training set. When trained, its ability to give a correct answer is tested on a separate set of examples where the output has not been revealed to the ANN. (The testing set). A relevant example can be found in [6] where an ANN predict hydraulic pressure based on known properties of the well.

It's important to distinguish between examples that are independent of each other and examples in a time series. In the latter case, each point in time constitutes one example but will also depend on earlier examples. For instance, the bottom hole temperature in a well does not only depend on formation temperature and present pump rate but in fact on the pump rates over the last few hours. Therefore, the ANN needs some kind of memory of recent examples. Recurrent neural networks (RNN), which are ANNs that are allowed to loop back on themselves, have been one way to address this. RNNs have been difficult to train and have had issues of stability, but a recent architecture called Echo State Networks (ESN) partly addresses this. [7, 8]

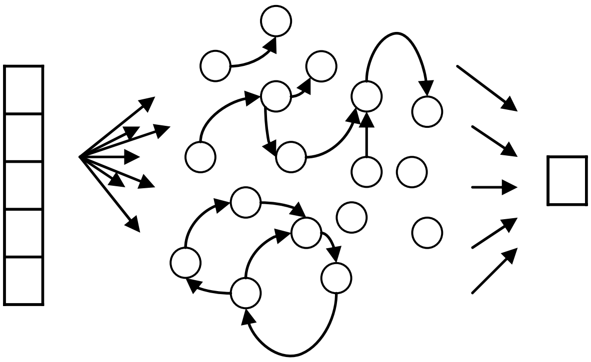

Figure 1: Schematic view of an Echo State Network.

Figure 1 shows a schematic view of an ESN. The input time series examples (left) are fed chronologically into a network of neurons (center). The neurons are connected randomly to each other before training starts and these connections remain fixed while the state of each neuron is updated at each time-step. The propagation of signals between neurons, often in loops, is what gives the ESN a memory of recent examples. Training consists of finding a linear combination of neuron states that approximate the value to be predicted (right). This type of training is much easier than trying to adjust the connections between neurons, which explains the recent popularity of ESN. A drawback is the memory requirements during training. The randomly connected neurons are basically a scattergun approach to finding good RNNs, which requires a large number of neurons. ESN belongs to a wider class of ANNs called reservoir networks; for an overview see [9].

Experiment

The idea of the following example is to illustrate how a combination of AI and an advanced model performs better than each alone. The AI part will account for model inaccuracies and features that are not covered by the model. The improved accuracy of the prediction will make it possible to reduce significantly the frequently erroneous alarms on observed volume changes, and it will be possible to trigger alarms on some real events that would otherwise not be detected because they were masked by other effects.

In this experiment, we have used an early version of the flow model that while being fully functional, had a limited precision in calculating volume flow rates. Our intention is not to have the AI amend something that should have been fixed in the software, but to use this deliberate limitation as a simple example of limitations that will always be present in any physical model.

Our time series represents 22 hours of drilling of one of StatoilHydro's wells. The first two thirds of the time series comprises the training set, which is used to adjust the physical model and train the AI, and the last third the testing set, used for measuring the alarm rate.

The ESN is first trained to predict flow out based purely on pump rates. 10 ESNs, each with 160.000 neurons, are trained on parts of the training set, and the best ESN is selected based on its performance on the rest of the training set before its prediction is measured on the testing set. This prediction is shown in Figure 2. The ESN is able to predict most, but not all, changes in flow out. As seen in Figure 3, the ESN does not solve this by equating flow out with flow in, but learns the time-delay and damping of flow rate surges through the drill string and annulus, as well as expansion of heated mud. Through the time series, the ESN thus learns the effects of formation temperature and mud compressibility without knowing them explicitly. However, it is vulnerable to changes in variables such as bit position, which is mostly independent of pump rates.

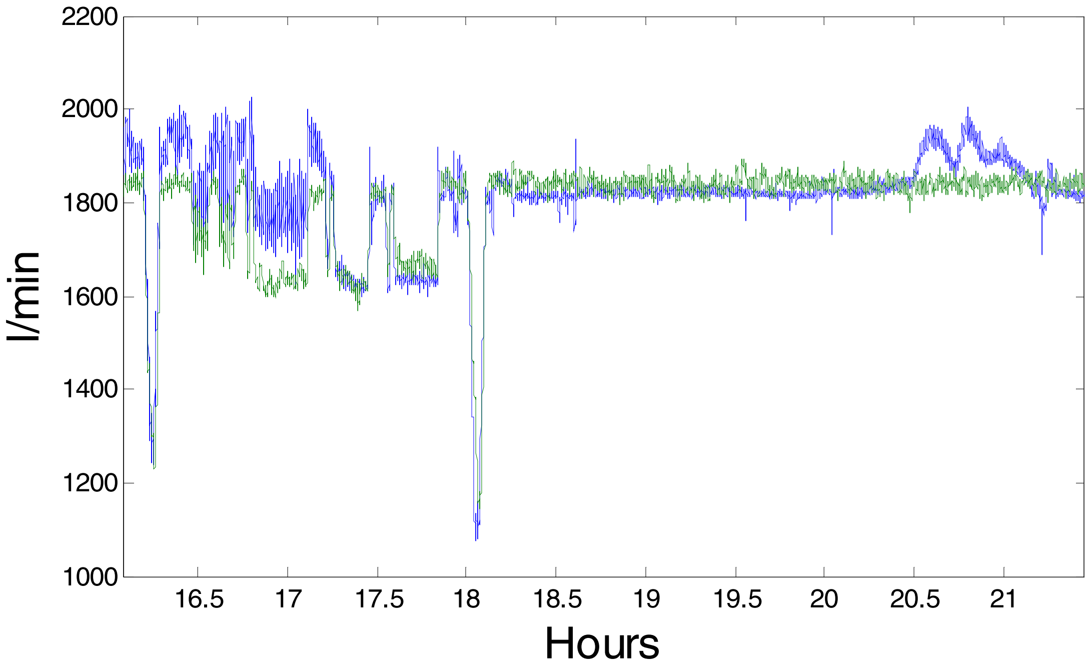

Figure 2: Testing set of measured flow out over 5 hours and flow predicted by the AI (green) using pump rates as input

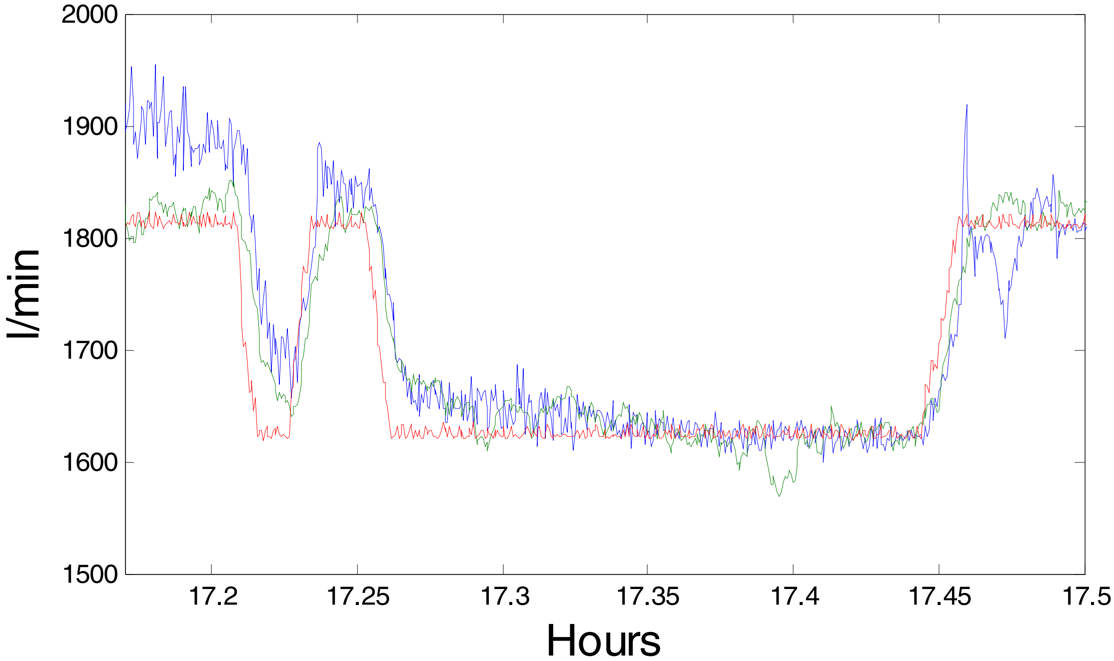

Figure 3: Section of Figure 2. The AI's prediction (green) mostly follows the measured flow out (blue). The pump rate is shown in red.

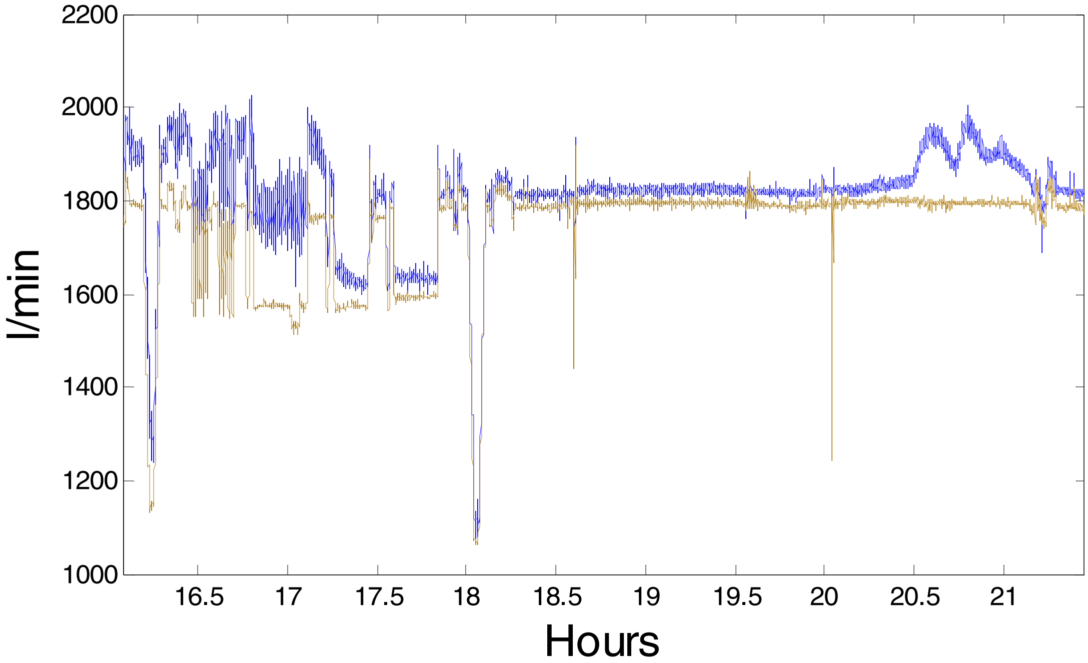

Figure 4: Measured flow out (blue) and flow out predicted by flow model (brown)

In Figure 4, we see the performance of the flow model on the same test set. It is interesting to note that the physical model and the ESN have problems with the same stretches of the time series. If measurements of the pump rate are accurate, this seems to indicate that factors other than the pump rate are responsible for these effects. This in turn means that the AI will need some additional input to improve its prediction.

However, this would present its own problems. We found that we needed a long time series to train the ESN properly. How long depended on how many examples of pump rate changes there were in the time series. If we add more variables, the complexity of the task increases, and according to theory the performance of the AI will in general degrade, unless the number of examples increases to compensate. This is known as the Hughes phenomenon [10], and can also be derived from Vapnik-Chervonenkis theory [11 p. 247-248] [12].

Because of this, we would soon run out of examples to train our AI. We must instead lower the complexity of the task and our strategy here is to remove from the time series what we can already explain, leaving the AI to predict the remaining effects. That is:

Residual ≡ Measured flow – Model prediction of flow

Combined flow prediction = ESN prediction of residual + Model prediction of flow

When we train the AI to predict the residual, the task of predicting damping and delay in flow rates has been delegated to the flow model, and the AI might be able to offer an incremental improvement learning other effects. A similar combination of methods is sometimes called “semi-physical modeling” or “Grey box modeling”. See [13] for a discussion.

A new ESN is now trained with the residual as the desired answer, and both pump rate and a second variable as input. Both an earlier analysis and its performance with the ESN pointed to mud density measured at the outflow as a good choice for the second variable. The mud flow and mud density is measured by a Coriolis flow meter every 2 seconds. Its precision and high sampling rate allows the AI to pick up correlations that would have been unobtainable using traditional measurement equipment. The combined prediction is shown in Figure 5. The improvement is seen more clearly in Figure 6, where we compare the residual with the residual predicted by the ESN. We see that the ESN has indeed picked up on the flow model’s deviation from the measurements, and manage to partly correct for it. In Table 1, we list the alarm rates resulting from the different prediction strategies.

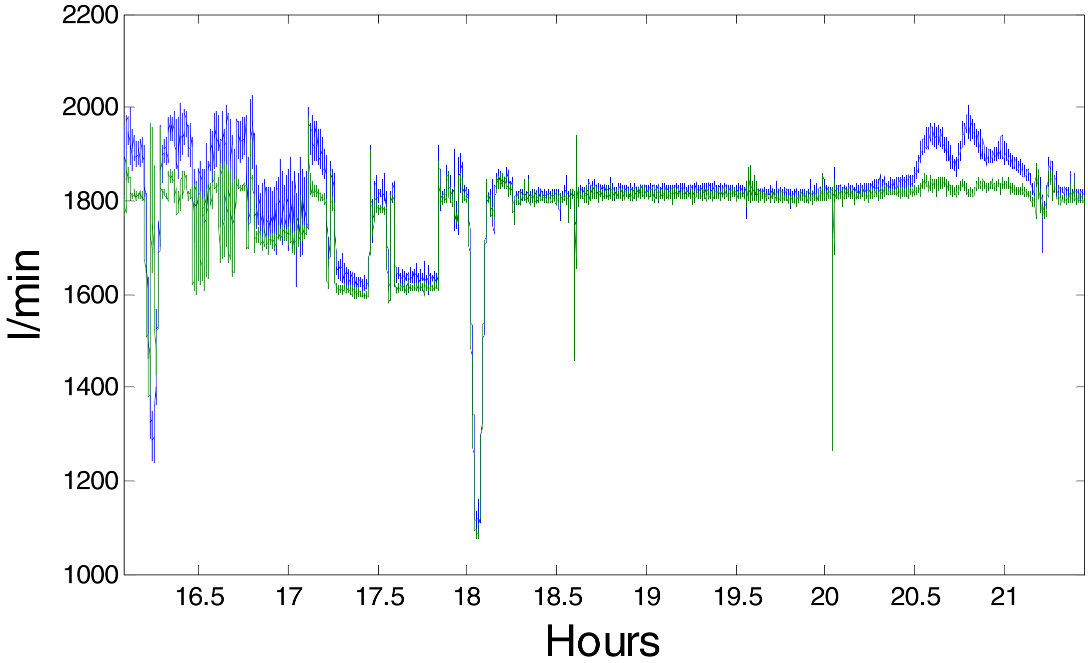

Figure 5: Measured flow out (blue) and flow model + residual predicted by ESN (green)

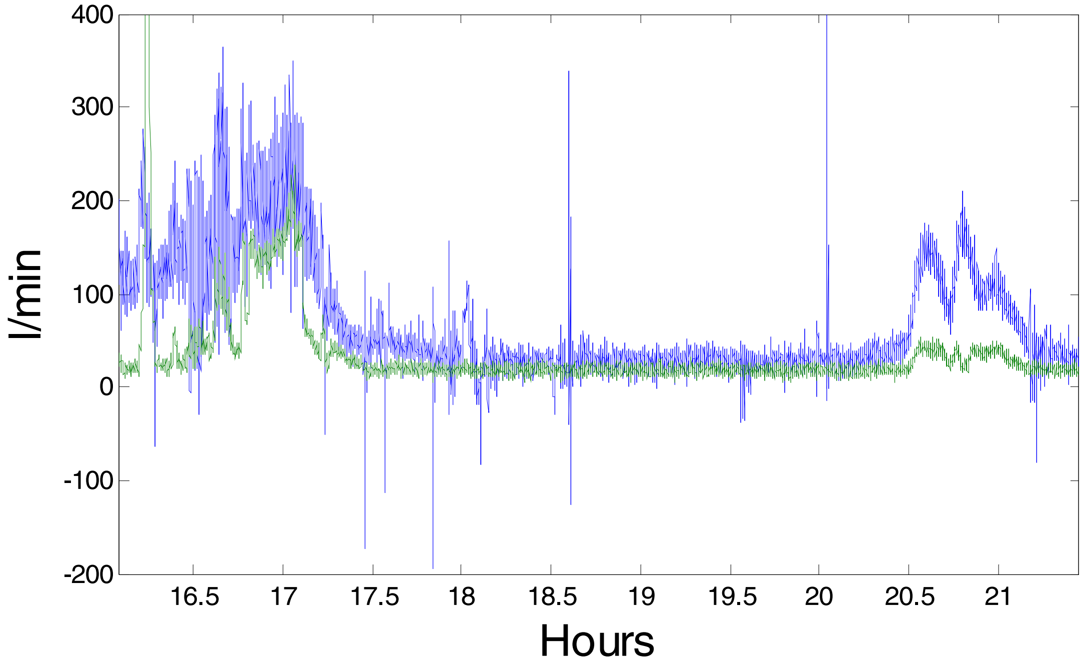

Figure 6: Residual between measurement and flow model (blue) and ESN's prediction of the residual (green)

Table 1

Method	Alarm rate at a sensitivity of 3bbl
No prediction	34 %
Physical model	38 %
AI	57 %
Physical model + AI	28%

Discussion and conclusions

We found that predicting the flow out of the well using a combination of the flow model and an AI yielded better results than applying each method alone. The alarm rate, while still high for high sensitivity, was reduced by a quarter compared to the stand-alone flow model. This serves as a proof of concept for combining physical models with artificial intelligence. It was later revealed that the early version of the flow model we used had a weakness in how it calculated volume flow when mud densities differed along the mud column. This explains why outgoing mud density got picked up by the AI when it was correcting the flow model.

This case also demonstrates things to consider when selecting inputs for the AI. Though the ESN itself is not limited in what variables it can be fed, known correlations and causalities need to be taken into account. This is particularly important with measured outgoing mud density, because the Coriolis flow meter's density measurement would be affected by gas bubbles reaching the surface. How does this affect the alarm system in the case of a real kick? We know that the signature of a kick is an increase in mud flow as the rising gas bubble expands, so that an alarm will have been triggered long *before* gas reaches the surface and influences our AI. We also find that in the case of our particular ESN, a decrease in mud density that could be expected from gas bubbles would itself trigger the alarm system.

We must also consider that the flow rate and mud density are not independent measurements. A Coriolis flow meter is designed to measure flow in units of mass per time. To perform this under changing fluid densities requires a separate measurement of the density [14] [15]. The real-time density measurement also allows the flow meter to give the flow in units of volume per time. Imprecision in the density measurement will therefore influence the flow readings, in such a way that an underestimation of the density results in an overestimation of both mass and volume flow.

Should the flow meter be imprecisely calibrated, it would be possible for the neural network to pick up on this and predict effects in both the flow model and the Coriolis meter. A neural networks ability to correct a Coriolis flow meter has been demonstrated for two-phase flow in a laboratory setting [16]. Here, the AI gave a considerable decrease in the flow rate error, though only on a small range of operational parameters. We conclude that an AI may serve both the role of correcting a physical model and doing quality control on the measurements fed to the model, as long as the proper training sets are used in each case, to separate model limitations from rig-specific effects.

Acknowledgements

We wish to thank ConocoPhillips and the Research Council of Norway for supporting this work as part of a research project, ConocoPhillips for supplying historical time based data from a number of wells, and Hesham Fandy, Halliburton, for answering related questions. We also thank StatoilHydro for permission to use the time series data used in the examples presented in this paper.

References

[1]D. Hargreaves, S. Jardine, and B. Jeffryes, "Early Kick Detection for Deepwater Drilling: New Probabilistic Methods Applied in the Field," in SPE Annual Technical Conference and Exhibition, 30 September-3 October 2001, New Orleans, Louisiana, 2001.

[2]R. A. Arehart, "Drill-Bit diagnosis with Neural Networks," 1990.

[3]V. M. Johnson, and L. L. Rogers, "Applying soft computing methods to improve the computational tractability of a subsurface simulation-optimization problem," *Journal of Petroleum Science and Engineering,* vol. 29, no. 3-4, pp. 153-175, May, 2001.

[4]E. M. Ozbayoglu, S. Z. Miska, T. Reed *et al.*, "SPE 78939, Analysis of Bed Height in Horizontal and Highly-inclined Wellbores by Using Artificial Neural Networks," 2002.

[5]G. Thonhauser, and W. Mathis, "SPE 103211, Automated Reporting Using Rig Sensor Data Enables Superior Drilling Project Management," 2006.

[6]R. K. Fruhwirth, G. Thonhauser, and W. Mathis, "SPE 103217, Hybrid simulation using Neural Networks to Predict Drilling Hydraulics in Real Time," 2006.

[7]H. Jaeger, "The "echo state" approach to analyzing and training recurrent neural networks," German National Research Institute for Computer Science, 2001.

[8]H. Jaeger, "A tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the "echo state network" approach," 2002.

[9]M. Lukosevicius, and H. Jaeger, "Overview of Reservoir Recipes," School of Engineering and Science, Jacobs University, 2007.

[10] G. Hughes, "On the mean accuracy of statistical pattern recognizers," *Information Theory, IEEE Transactions on,* vol. 14, no. 1, pp. 55-63, 1968.

[11] S. Theodoridis, and K. Koutroumbas, *Pattern recognition*, 3rd ed., p.^pp. 248, Amsterdam ; Boston: Elsevier/Academic Press, 2006.

[12] V. N. Vapnik, and A. Y. Chervonenkis, "On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities," *Theory of Probability and its Applications,* vol. 16, no. 2, pp. 264-280, 1971.

[13] U. Forssell, and P. Lindskog, "Combining semi-physical and neural network modeling: an example of its usefulness," in 11th IFAC Symposium on System Identification (SYSID'97), 1997.

[14] K. Plache, “Coriolis-Gyroscopic Flow Meter,” *Mechanical Engineering,* vol. 101, no. 3, pp. 36-41, 1979.
[15] efunda. "Introduction to Coriolis flowmeters," Dec., 2007.
[16] R. P. Liu, M. J. Fuent, M. P. Henry *et al.*, “A neural network to correct mass flow errors caused by two-phase flow in a digital coriolis mass flowmeter,” *Flow Measurement and Instrumentation,* vol. 12, no. 1, pp. 53-63, Mar, 2001.