

MATERIAL CLASSIFICATION BY DRILLING

Diana LaBelle, John Bares, Illah Nourbakhsh

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Underground coal mining is one of the most dangerous occupations. Years of effort have been dedicated to researching methods of characterizing mine roof and floor for improving the mining environment. This research investigates using a neural network to classify rock strata based on the physical parameters of a roof bolting drill. This paper presents our methodology, as well as early results based on drilling experiments conducted in the laboratory using a custom poured concrete test block. We have classified, with a trained network, the five layers of the test block with less than 5% error.

Keywords: mining automation, rock classification, neural networks, drilling, coal interface detection.

1 INTRODUCTION

Underground coal mining is one of the most dangerous occupations. The failure of structural supports accounts for approximately 400 injuries and 10 deaths each year. One method of increasing mine stability is to drill and bolt the weak mine roof after a portion of the coal seam has been removed. Mine ground control has been thoroughly researched for the last several decades. Yet, mine stability problems, such as roof falls and rock bursts, continue to kill or injure people every year. Over half of the most recent fatalities have occurred under supported roof [8].

Mine workers have limited information about the lithology of the rock surrounding the coal seam. If a local and detailed lithology of the surrounding rock could be determined, mines could better assess the effectiveness of the roof bolts, alert miners to local hazards, augment ground control plans, and thus, greatly improve mine safety. Years of effort have been dedicated to researching methods of characterizing the mine strata in order to improve safety.

This research investigates using the physical parameters of a roof bolt drill, such as torque, thrust, rotary speed and penetration rate, to train a neural network to classify rock. Section 2 discusses relevant coal mining background, defines the problem and reviews research that uses a similar approach to ours. Section 3 presents a detailed description of the method in which the drill information is used to classify materials while section 4 presents the experiments and results. Then we present our latest experimental tests and results. Finally, future research plans are outlined in section 5.

2 BACKGROUND

The most common means of attaining lithological information before mining begins is with pre-mine exploratory drilling. However, this is expensive and

therefore, sparsely spaced. Core logs give limited information about the coal bed and surrounding strata and with limited accuracy. Drill cores miss local geologic anomalies that pose a hazard to miners. Furthermore, the process of mining continuously changes a mine's structural conditions. The overall safety of a mine could be greatly enhanced with real-time monitoring of local structural conditions as the mining progresses. Figure 1 shows the cross-section of a typical coal mine corridor.

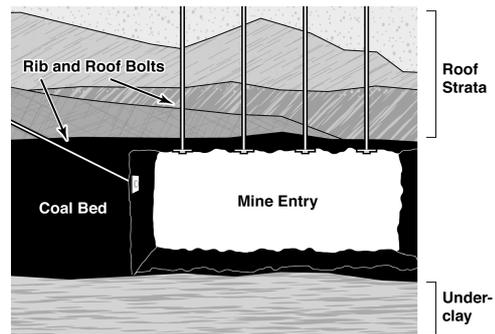


Figure 1. Underground Coal Mine Cross-Section

One hazard is severely fractured or delaminated roof layers that cannot support their own weight. Often, weak strata cannot withstand the stresses produced by large overburden, and as a result, the roof and floor heave and the ribs bulge, constricting the mine opening.

The most common hazard detection techniques are reactive in nature, such as extensometers, and sense changes that signal a hazard is imminent. Often, the experienced mine worker is the best hazard detector, because of his familiarity with the feel and sound of the machines. But the mine worker is also the most vulnerable of all detectors. There has been little success in an effective pro-active measurement device of local ground conditions *in situ* such as ground penetrating radar, ultrasonic sensors, and instrumented roof bolts.

In contrast, there has been extensive research into all types of precision drilling. This research has provided

valuable insights into drilling analysis. Ramirez and Thornhill [5] monitored printed circuit board drill wear using drill force spectrum analysis and a sensor fusion technique to combine this data and reported that in some cases, cutting forces are related to chip segmentation frequency, which depends on the physical properties of a material. Kaburlasos et al. [2] used learning techniques to associate drilling thrust and torque recorded during surgery with the thickness of an inner ear bone being drilled, in order to predict drill breakthrough and prevent damage to other ear tissue.

There have been only a handful of groups that have researched intelligent drilling for mining and construction applications using rotary drills. Leighton et al [4] used an instrumented rotary blasthole drill to correlate drill performance and strata with blasting variables to plan open pit blasting. They were able to discover an empirical relationship between drill thrust, penetration rate and a material's resistance to breakage. Scoble et al [6] used drill monitoring to verify coal-rock boundaries in a surface coal mine using a rotary blasthole drill. They calculated the specific energy of drilling (SED)—an estimate of the rotational and linear energy used to drill—and compared it to known strata strength and delineation.

King and Signer [3] have interpreted data from a coal mine drill using learning techniques. They used unsupervised learning to classify underground coal mine roof drill data into clusters and then used the classifications to train a neural network on other drill data. Similarly, Utt [11] used a neural network to classify rock based on SED. His work assigned a soft, medium, or hard to the quantified drill information so that a miner could be warned of weak strata when bolting up the mine roof. Both King and Signer and Utt chose an automated approach to characterizing coal mine strata, but they made assumptions that differentiate their work from this research. Both categorized strata by relative strength estimated from the drill parameters. Each group had positive results, but primarily for the cases in which two layers of strata had widely differing compressive strengths.

Many researchers have used SED as a way of characterizing strata. This is acceptable if one wants to estimate the relative strength between the layers and other geological features. However, SED depends on how finely the rock is ground at the bit. Using SED to estimate the relative strength between materials could be misleading, particularly if the manner of drilling is not constant. Furthermore, SED alone cannot be relied upon to distinguish between two materials because there are cases where different strength materials are similar in strength (coal and shale have overlapping ranges of compressive strength) or seem similar in strength (such as a strong

material being fractured or drilling at different rates for each material). By differentiating between materials in addition to providing an estimation of some physical characteristics *in situ*, one can classify a material with a higher degree of confidence.

3. APPROACH

In this research, we intend to use data from an instrumented mine drill to classify a small set of materials that are typically found around a coal seam. This must be done as mining takes place, without requiring a mine worker to perform classification, and regardless of the drill, operator or local conditions. Our approach is motivated by the fact that drill response is known to be correlated with the properties of the material being drilled. For example, researchers have verified that material properties such as abrasiveness, hardness, and strength directly affect the drilling process [10].

The physical properties of a material include mechanical properties (such as strength, hardness, abrasiveness, and porosity), electrical properties, and molecular structure to name a few. Physical properties can have subtle or pronounced effects on the bit-rock interaction. For example, two materials can have similar compressive strengths and require a similar amount of drilling energy, but have very different grain sizes and wear the drill bit at different rates. A material's properties are also affected by *in-situ* conditions such as confining pressure, temperature, moisture content, the presence of gases and the process of mining itself.

The process of drilling is complicated to physically model. There are a large number of variables that influence the drilling process. The factors that affect drilling originate from the drill, the material, and the environment in which the drilling takes place. Variables such as drill string stiffness, drill bit geometry and wear, method of flushing, and machine condition can significantly affect the performance of the drill and the bit-rock interaction [9]. These variables also differ between drills and operators.

Considering all of the geologic, environmental, and mechanical variables, drilling quickly becomes a very large, real-valued, multivariate data set. Often, the complex relationships between these dynamic variables are not well-understood or even known. For these reasons, this drilling application is a good candidate for machine learning. An early survey of these methods suggested that a neural network is an appropriate learning algorithm to use.

A neural network is composed of layers of interconnecting nodes as shown in Figure 2 [7]. The drill sensor values are the input layer and the material identity (shown in binary form) is the output layer.

Every node is connected by a constant real-valued number, or weight, to all the nodes in the succeeding layer. The number of hidden layers and nodes in the layer determines the complexity of the network and the functions it can represent. The goal of training the network is to minimize the error of classifying the material. Therefore, training is an iterative process: propagate inputs through the network, calculate the error between network output and actual output (or target), and adjust weights by backpropagating a function of the network error. The cycle continues until a minimum error is attained or a user-defined number of training iterations is reached.

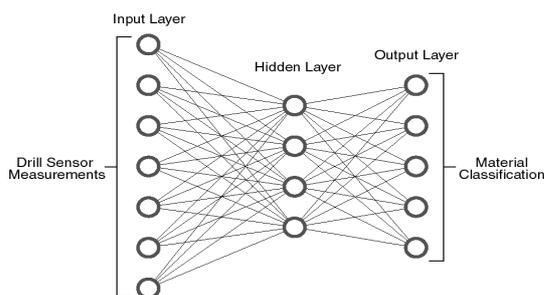


Figure 2. Neural Network Configuration

To help the neural network make use of the complex relationships between all of these variables, we have augmented the drill sensors with additional “virtual” sensors. These sensors are not physical sensors, but functions of the drill’s sensors. They can represent complex relationships between drill behavior and material properties. The information from the virtual sensor is another drill parameter and another variable for a neural network to use.

The following sections present an iterative approach of gathering laboratory and coal mine drill data, establishing a set of real and virtual sensors, training a neural network to classify the drilled materials, and evaluating the classification results.

4. EXPERIMENTS

4.1 Experimental Apparatus.

Our drilling apparatus consists of a portable, hydraulically-powered, manually-operated, water-cooled coal mine drill instrumented with sensors, data acquisition hardware, and a laptop computer (see Figure 3 below). The electronic hardware is isolated from the drill so that it can operate in a real mine environment. The data acquisition system is in a waterproof box, with one cable running to the sensors and another cable connecting to the laptop which can be taken several feet away from the actual drilling site.

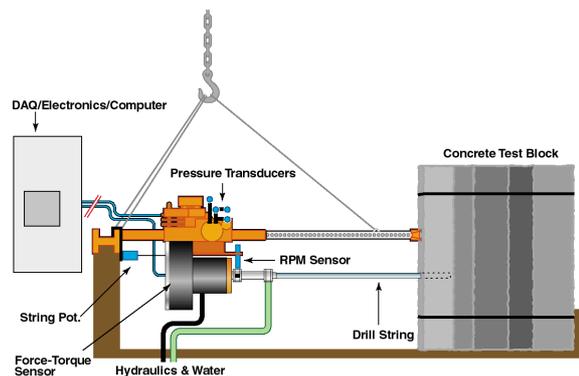


Figure 3. Laboratory Drill Apparatus and Setup

The drill parameters that are recorded are torque, thrust, rotary speed, hydraulic pressures and drill position. A highly accurate six-axis, decoupled force-torque sensor is connected in-line between the drill motor and the drill carriage. Hydraulic pressures of the thrust and rotation motors inlet and outlet are recorded. The reason for this is to assess the feasibility of measuring thrust and torque values from less expensive sensors for a real-world system. The rotary speed is measured using a magnetic sensor and a collar with four embedded magnets attached to the spinning drill chuck. Using LabVIEW the sensor readings are captured at a constant rate, and drill penetration rate is calculated off line, using the drill bit position readings at known, fixed time steps.

Our laboratory drilling set-up includes an adjustable frame to support the drill as it drills horizontally through layered concrete test blocks. The drill is supported vertically with cables. When a hole is drilled, the drill mast is expanded between the steel frame and the concrete test block. The linear and rotary movement of drilling is controlled manually while the computer controls the data acquisition. The thrust motor valve is held fully opened, while a hydraulic restrictor valve is used to keep the flow rate at a maximum value if with the goal is to keep the penetration rate constant. The rotation motor valve is held fully opened, but the flow is not controlled. To keep the drill hole as clean as possible from drill fines, the water is turned on full-flow each time a hole is drilled to minimize unintended regrinding.

4.2 Data Collection and Processing.

We have gathered data on about 40 holes drilled into a three foot thick concrete test block. The 3’x3’x5’, 8,000 lb test block has five layers of concrete of different strengths and materials. Each concrete mix was tested for compressive strength. The physical characteristics of each layer are given in Table 1. We drilled holes into the concrete test block in a rough grid pattern. We also drilled and recorded about 30 holes at the Bruceton Coal Mine in Pittsburgh.

Layer	1	2	3	4	5
Concrete Mix	Grout	Lime stone	Fly-ash	H.E.S.	H.E.S.
Comp. Str. (psi)	1,900	5,600	1,300	4,500	4,300
Thickness (in)	11	5	9	4	7

Table 1. Concrete Test Block Characteristics (H.E.S. is High Early Strength concrete)

It took, on average, 90 seconds to drill a hole into the concrete test block. A typical data file has between 60,000 and 100,000 data points, each with seven real-valued sensor readings: force and torque from the force-torque sensor, hydraulic pressures at the thrust and rotation motor inlets and outlets, drill position and rotary speed. Figure 4 is an example of sensor data recorded while drilling a hole through five layers of concrete (abscissa is drill position, ordinate is sensor readings)

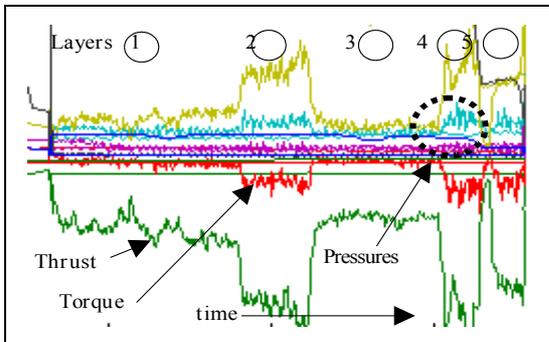


Figure 4. Concrete Drill Hole Sensor Recordings for One Drill Hole.

The drill sensor data used in these experiments is post-processed. The output of the string potentiometer is filtered with a capacitor, while the force-torque sensor uses a low-pass filter. The conversion of magnet pulses to rotary speed uses software filtering. The size of each drill data file is reduced for training the neural network and to facilitate in the analysis of real and virtual sensors. After calculating the drill penetration rate, each drill data file is sub-sampled by 1% (a plot of the sub-sampled file is very acceptable when visually compared to its parent file). These files are further processed by choosing data points from clean segments of each material and leaving out areas of material transition or drill start or stop points. The cleaned files have between 500 and 1500 data points comprising a segment of sensor values for each material that was drilled in that particular hole. Each is labeled by hand and normalized over the range of sensor values. Finally, each segment of data is collapsed into a single data point in N-dimensional space, where N is the number of fields. The fields of a data point are, for example, average thrust over the drill segment, and N is different for each experiment.

4.3 Neural Network Training and Testing

We are using Netlab's [1] back propagation algorithm with a two-layer feed forward neural network. The class labels are converted to binary numbers for neural network training and testing, and then converted back to an integer classification using a best-of-N voting scheme.

The following experiments have been designed to determine if simplified drill data from concrete blocks can be classified into the 5 hand-labeled materials, and to do a sensitivity analysis of drill parameters. A neural network with no hidden units was trained and the poor results, averaging 80% classification error, indicate that there are non-linear relationships in the drill sensor data. Subsequent networks used 4 hidden units and were trained over a range of iterations.

We used several sets of attributes as inputs to the neural network (see Table 2) including measurements from the physical sensors (labeled p) as well as virtual sensors (labeled v). To prevent the network from learning material class solely based on position of the drill (since every hole was drilled in the same concrete test block), the drill bit position was not used as an input to the neural network.

Attributes	Experiment												
	1	2	3	4	5	6	7	8	9	10	11	12	
Mean													
Thrust (p)	■	■	■	■	■	■				■			
Torque (p)	■	■	■	■	■	■					■		
RPM (p)	■	■	■	■	■	■					■		■
Penetration (p)	■	■	■	■	■	■							■
Rotary In (p)	■								■	■			
Rotary Out (p)	■								■	■			
Thrust In (p)	■								■	■			
Thrust Out (p)	■								■	■			
Thrust Diff. (v)											■		
Rotary Diff. (v)												■	
Std. Dev.													
Thrust (v)	■												
Torque (v)	■												
Penetration (v)	■												■
Thrust Diff. (v)													■
Rotary Diff. (v)													■

Table 2. Attribute Sets Used in the Experiments.

There are 14 drill hole data files used for the experiments. Each data set is generated by randomly choosing 11 of the 14 files for training and the remaining 3 for testing. Each experiment begins with training and testing 100 unique data sets on a neural network. An average test error rate is computed from the test error rates of the 100 data sets. This is

repeated with a range of values for the ‘number of iterations’ parameter. The iteration value with the lowest average error rate is reported.

4.4 Experiment Results

The average error rates for each experiment, and a breakdown of the error rates by material are shown in Table 3.

Experiment	Error Rates by Material (%)					Avg
	1	2	3	4	5	
1	0	0	1.5	15.1	5.8	4.5
2	4.2	4.1	13.9	33.3	27.5	15.2
3	48.9	21.3	60.5	83.6	84.4	59.8
4	12.7	6.8	32.2	11.1	48.6	22.3
5	0.2	0.9	9.9	25.8	37.1	14.8
6	0.1	0.8	6.3	29.8	32.3	13.8
7	44.8	31.7	63.2	61.7	68.3	54.0
8	0.0	14.5	5.8	34.7	66.7	24.3
9	14.0	8.6	8.6	87.2	91.8	42.0
10	57.6	16.2	43.4	93.9	90.5	60.3
11	76.8	62.6	63.8	86.0	97.2	77.3
12	54.8	47.7	77.2	92.3	97.1	73.8

Table 3. Error Rates by Material for Each Experiment

Experiment 1 used all of the physical sensor values—thrust, torque, rotary speed, penetration rate and motor inlet and outlet pressures—as well as a number of the virtual sensors in its attribute set. This ‘base’ set of attributes had the lowest average classification error rate of the 12 experiments, with 4.5% error. This result verifies that a neural network can classify the five materials and it also serves as a value with which to compare results of the other experiments.

Figure 4 shows the classification error rates for a range of iterations. Increasing the number of iterations improves the classification accuracy until about 90 iterations, where it levels off. Material 5 consistently has the highest error rates. This is true for all of the experiments which are described later.

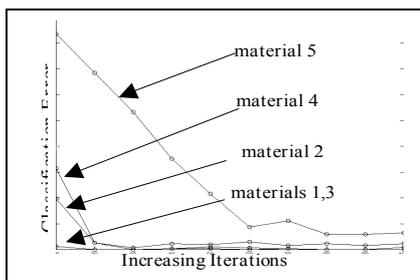


Figure 4. Network Classification Error Rates vs. Number of Iterations Using Experiment Group 1. Each line represents a material.

Figure 5 below is an image representation of the network’s confusion matrix. The confusion matrix is a measure of how well the network performed on each material class and shows where the misclassifications occurred. The ordinate axis is the network’s classification of the data sets in experiment 1. The columns show the percentages of the material classes they actually are. For example, materials 1 and 2 were correctly classified 100% of the time, and material 3 was misclassified as material 4 about 10% of the time.

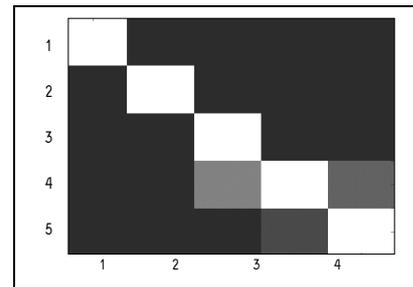


Figure 5. Experiment 1 Confusion Matrix (lighter colors represent higher numbers).

The purpose of experiments 2 through 6 was to evaluate the the relative discriminatory power of the classic drill parameters, thrust, torque, rotation rate and penetration rate. The network classification error rates for experiments 2 through 6 are shown in Table 3. Experiment 2 used all four drill parameters, while the next four experiments removed thrust, torque, rpm, and penetration rate, one at a time. Removing the rotary speed or penetration rate caused the classification error to increase mildly. Removing drill thrust significantly impairs the network’s ability to classify the materials. This may be due to the fact that rotation rate and penetration rate were relatively constant compared to thrust and torque, so that these two parameters became even more critical in classifying the materials

Experiments 7 and 8 replaced the sensor values from the highly-accurate force-torque sensor with the sensor values from the less expensive pressure transducers installed at the inlet and outlet ports of both hydraulic motors. Experiment 7 achieved a poor classification rate with 54% error. Experiment 8, however, trained the network using 5 additional virtual sensors and the average error rate dropped to 24%, less than half that of experiment 7. This suggests that the addition of virtual sensors to the sensor suite may make it possible to use the inexpensive and robust pressure sensors in lieu of the expensive and delicate strain gauge sensor to classify mine roof rock.

A single drill parameter (force, torque, rotary speed and penetration rate) was used to train a network in Experiments 9 through 12. Training the network with only drill thrust measurements produced the lowest error rate. This result agrees with the results of Experiment 3 which showed that of the four same parameters, removing drill thrust from the training caused the most increase in error. Closer examination of the results of Experiments 9 through 12 reveals that although drill thrust is the most important parameter in classifying materials 1, 2 and 3, all four drill parameters—thrust, torque, rotary speed and penetration rate—appear to be equally poor at classifying materials 4 and 5. This is additional evidence that there is probably some other drill parameter that can better distinguish these two materials.

Experiments 1, 7 and 8 indicate that the standard deviations of the drill parameters are useful in increasing the degree of accuracy in the classification of the concrete test materials. One explanation may be that the concrete mixes range in the size and concentration of aggregate, causing the drill vibrations to be different for each material. This effect is better represented in the larger, unsampled data files, and may prove to be even more useful where changes in the vibrations over time can be tracked. Experiments 7 and 8 also highly suggest that virtual sensors are useful and that they improve the classification accuracy of similar materials.

5. FUTURE WORK

In the experiments presented in this paper, the data files used to train and test the neural network were reduced to a few points. This gave the network very clean data to train with. The next step is to test the network on the larger data sets. It is likely that for the larger datasets we will need to discover and use additional virtual sensors. These experiments also used a single network architecture and did not vary the error function or search algorithm used to tune the weights. Subsequent experiments should apply model selection, so that the best network architecture and training algorithms can be determined prior to exploring virtual sensors. Future experiments will focus on classification with noisy drill data and understanding relationships between drill and material. Future research also includes classifying materials with data gathered at an underground coal mine.

Acknowledgements

We would like to thank Dr. Andrew Moore and Daniel Nikovski for their help with neural networks, NIOSH and Robotics Institute people involved, and

S.T.M. for graphics. This research is supported under NASA contract NCC5-223.

REFERENCES

- [1] C. Bishop, I. Nabney, *Netlab Software Web Page*. <http://www.ncrg.aston.ac.uk/netlab/>
- [2] V. Kaburlasos, V. Petridis, P. Brett and D. Baker, "Learning a Linear Association of Drilling Profiles in Stapedotomy Surgery." *Proc IEEE Int Conf on Robotics and Automation*, May, 1998.
- [3] R. King and S. Signer, "Using Artificial Neural Networks for Feature Detection in Coal Mine Roofs." *Proc 8th Int Conf on Comp Methods and Adv in Geomech*, Vol. 3, May, 1994.
- [4] J. Leighton, C. Brawner and D. Stewart, "Development of a Correlation Between Rotary Drill Performance and Controlled Blasting Powder Factors." *CIM Bulletin*, Vol. 75, No. 844, August, 1982.
- [5] C. Ramirez and R. Thornhill, "Drill Wear Monitoring in Circuit Board Manufacturing Using Drilling Forces and Their Spectra." *ASME J. of Elec Packaging*, Vol. 114, September, 1992.
- [6] M. Scoble, J. Peck and D. Kennedy, "Drill Monitoring Investigations in a Western Canadian Surface Coal Mine." *Proc 2nd Large Open Pit Mining Conference*, April, 1989.
- [7] T. Mitchell, "Machine Learning." *McGraw-Hill*, 1997.
- [8] NIOSH, *Mining Safety and Health Research Web Pages*, <http://www.cdc.gov/niosh/mining/>, 1998.
- [9] R. Teale, "The Concept of Specific Energy in Rock Drilling." *Int J. of Rock Mech and Mining Sci*, Vol. 2, 1965.
- [10] J. Paone and W. Bruce, "Drillability Studies; Diamond Drilling." *U.S. Dept. of Int., BuMines RI 6324*, 1963.
- [11] W. Utt, "Neural Network Technology for Strata Strength Characterization." *Proc IEEE Int. Joint Conf. on Neural Networks*, Washington DC, July, 1999.