Lithological Classification by Drilling

Thesis Proposal

Diana LaBelle
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
dlabelle@cs.cmu.edu

Abstract

There are many drilling tasks in which drill monitoring is used to improve the quality of a product: detecting tool breakage in manufacturing drilling, exploratory drilling for oil and natural gas reservoirs, collecting soil samples on Mars with a robotic drill. However, in many applications, a human is partially or entirely responsible for controlling and analyzing the interaction between the drill bit and the drilling medium. This research is exploring intelligent drilling that can be applied to multiple applications.

I propose to develop a methodology with which to build a working lithological classifier for different drilling applications. The methodology is a template with which a classifier can be built using the proprioceptive sensors of a drill, and it includes: (1) acquiring drill hole sensor data; (2) using expert knowledge from a mine or drill operator to generate a feature set from drill sensors; (3) automating the selection of a subset of drill features which maximize classification rates and minimize the sensitivity of the classifier to drill and environmental variables; (4) using machine learning to classify rock and detect material boundaries; and (5) extending the classifier to operate with a data stream. The methodology and it's adaptability will be tested with coal mine drilling, but another application, such as well drilling or surgical drilling, will be investigated.

The methodology uses a neural network to classify material lithology where the inputs to the neural network are sensed drill parameters such as thrust, torque, rotary speed and penetration rate, as well as information derived from these sensors over time. In preliminary experiments, five different layers of concrete and three layers of coal mine strata were classified using the sensors attached to a coal mine drill. These results suggest that drill parameters can be used to classify rock strata, and that using additional features derived from the drilling parameters significantly improves classification accuracy.

1 Introduction

Drilling in an underground coal mine, drilling human bone and drilling rock on Mars each involve drills and materials that are very different, operating in very different environments. Machine drills typically operate at high rotary speeds and with high cutting forces, whereas the sensors used on a surgical drill, for drilling tiny bone of the inner ear, measured force to accuracies measured in ounces. Mining and construction drills are characterized by extremely large forces, thousands of pounds, and relatively low rotary speeds, a few hundred rpm. Extra-terrestrial drills are designed to be lightweight, low-powered, and very slow. These very different drilling applications have common needs: they either require or would benefit from the ability to monitor proprioceptive drilling parameters, such as thrust, torque, rotary speed, and penetration rate, to improve the outcome of the drilling process.

In recent years, hip replacement surgery has become less invasive, and with smaller incisions, the surgeon has less visual feedback. Doctors want to develop smarter drills to prevent unnecessary damage, such as drilling into softer tissue, removing too much bone, or breaking through the bone altogether. Surface mining operations can increase the amount of coal recovered while minimizing damage to underlying rock if they have an accurate measurement of the boundary between coal and other rock. Many mines rely on the drill operators to detect the coal-rock boundaries. Mars rock is drilled remotely. Preventing the drill from seizing up is critical to the success of the mission.

In many situations such as remote drilling on Mars or at Chernobyl, the human is not in direct contact with the drill, but he or she must control the drill. This is difficult even with good human-computer interfaces. Almost immediately after drilling begins, an operator is drilling blindly, usually with minimal information about the material being drilled. A skilled operator drilling on-site monitors the drill behavior manually, and can react to a variety of problems. In deep drilling, the operator is not at the drilling site and can only react to drill behavior after some time lag between the sensors and the computers. In order to make timely and safe decisions about drill control, drill operation needs to be partially or fully automated. Natural gas drilling is highly directional, and the drill must be guided, in part, by downhole sensors. Because the drill operates at great depths, local sensing and automated control is desired.

In all of the applications discussed above, gaining information about the material being drilled and about the state of the drill is either required or would be highly beneficial. My research focuses on drilling for the underground coal mining industry, but one aim is to progress intelligent drilling for multiple drilling applications. Underground coal mine drilling is discussed in detail in the remainder of this paper.

The mechanics of drilling is very complicated to model, and there are an enormous number of variables that influence the drilling process. For example, drill bit geometry, drill bit wear, the amount of drill hole flushing, material composition and environmental variables can significantly affect the dynamics of the entire drilling system. Machine learning—a neural network in particular—is an appropriate method of classification for systems with complex relationships between many variables.

There are only a handful of groups that have researched intelligent drilling for the mining, construction or medical fields. Of this, much of the research still required a human to interpret the drill data or optimize the classifier, or the classification abilities of the drill were limited. This research aims to develop a methodology with which to build a real-time material classifier that can be used for multiple drilling applications, and that can adapt to various drill and material types. I will develop and test the methodology on real data obtained by drilling in underground coal mine, and on data from drilling through layered concrete blocks. I will also perform an analysis of the accuracy of the classifier with changing drilling environments.

2 Motivation

Underground coal mining is one of the most dangerous occupations. Approximately 42% of the fatalities in underground coal mines are a result of ground instability. The failure of structural supports still account for approximately 400 injuries and 10 deaths each year. Over half of the most recent fatalities have occurred under structurally supported roof [33]. Before an underground mine is opened, exploratory drilling helps engineers evaluate the conditions expected underground and design the proper ground control needed to support the mine roof and ribs. Exploratory drilling is expensive, and therefore, sparsely spaced. Core logs give limited information about the coal bead and surrounding strata and with limited accuracy. Drill cores miss local geologic anomalies such as faults in the rock that pose a serious hazard to miners [53]. Furthermore, as mining progresses, the structural conditions change.

There are telltale signs of poor structural conditions in a coal mine [30]. One dangerous situation 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 tools are ones that sense changes in the mine that signify a hazard is present. For example extensometers measure the amount of roof sag, and and instrumented roof bolts indicate the stresses in the strata. Often, the experienced mine worker is the best hazard detector, with a keen sense for the 'feel and sound' of the machinery and the mine itself. But the mine worker is also the most vulnerable of all detectors.

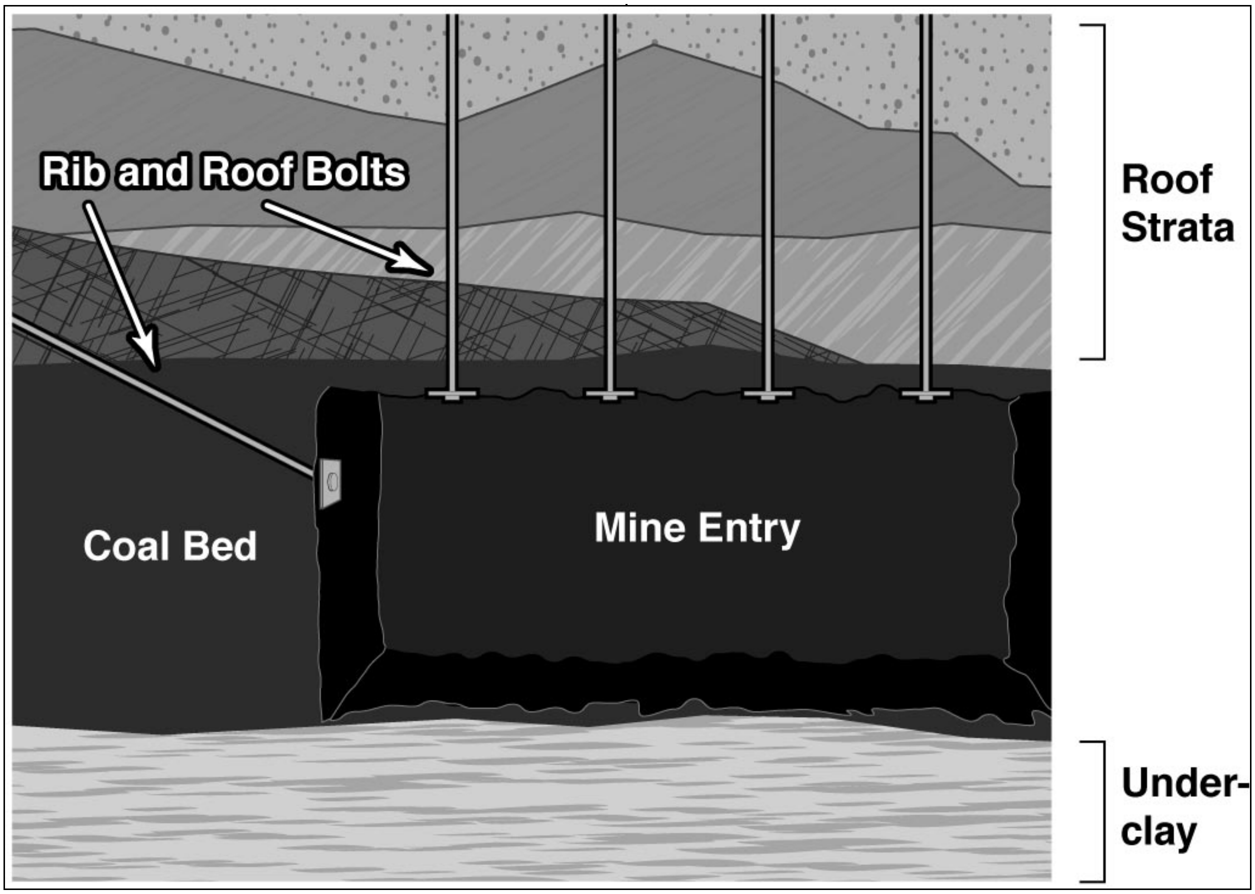

Figure 1. Underground Coal Mine Cross-Section

Figure 1 shows the cross-section of a typical coal mine corridor. One method of increasing mine stability after coal is removed from the coal seam is to drill and bolt up the roof with long steel bars anchored into strong rock. It is critical that the roof bolts be anchored in *strong, solid* rock. Underground mines with extremely problematic under- and overburden may require a remnant layer of coal to be left on the roof and floor to maintain the structural integrity of the rock under or over it. This technique is the most effective when a consistent thickness of coal is left on the roof and/or floor. In terms of productivity, many mining companies leave large quantities of coal behind because they lack the tools to accurately and quickly measure the coal thickness. Miner operators often find it difficult to accurately detect the thickness of the coal and control the mining machine accordingly.

Mine workers have limited information about the lithology of the rock surrounding the coal seam. They need the ability to identify the mine strata—much of which has very similar physical properties. For example, both coal and shale are very commonly found in in close proximity (or adjacent to each other) in the same mine. These two materials have ranges of compressive strength that greatly overlap. Coal and shale also look similar. Drill operators may not be able to tell what they are drilling into, or even if it is strong enough to support the bolted roof. If a local and detailed lithology 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. To be effective, this must be done as mining progresses and with a high degree of confidence. To be attractive to the mining industry, this information gain has to be cost-effective. One way to achieve this is to use existing equipment—a roof bolter drill.

3 Problem Statement

There does not exist a lithological classifier based on proprioceptive drill parameters that can reliably distinguish between similar strength materials in a realistic environment and in real-time.

4 Problem Discussion and Background

One approach to coal-interface detection, or CID, is to analyze the thermal images of the mining machine cutter teeth to detect when they penetrate the softer coal and begin to cut the harder rock above, making the cutter teeth warmer. Another approach is to analyze the vibrations in strata. Mowrey used accelerometers, geophones and piezoelectric films to monitor vibrations in the mine roof and walls and on the mining machine [32]. The methods above require the roof rock to be penetrated in order to detect the coal-rock interface. Other non-invasive methods include natural gamma radiation, ground penetrating radar and ultrasonic sensors, all of which can detect the boundary between coal and other roof rock [33]. Unfortunately, each of these methods has it's own shortcomings: they are too slow to be used in real-time, or they can only be used in mines with specific strata conditions. These sensors are also very expensive—not an easy sell to the mining industry.

My approach to determining rock lithology, and therefore, the coal-rock interface, is to use the information from an instrumented roof bolt drill to classify coal mine strata. This is motivated because drill response is known to be correlated with the properties of the material being drilled. Numerous drilling investigations, some of which are reported in following sections, have provided empirical results verifying the relationships between drill parameters and material properties. Roof bolt drills are in most mines, and roof bolting is done on a scale that is appropriate for obtaining a detailed mapping of the surrounding strata. This knowledge can greatly enhance the safety and productivity of coal mining.

4.1 Drill-Material Interaction

There are many technical challenges in classifying mine strata with a mine drill. An underground coal mine is a harsh environment, with extremely powerful machinery working in an area barely large enough for it to fit. Coal miners operate their machinery in a rough manner. Roof bolt drills are over-torqued. Drill steels are bent in half or mis-aligned while spinning at 200 rpm. Drill motors, drill chucks and drill steels have significant mechanical free play, and therefore sensor readings will be noisy. For the most part, mine strata are heterogeneous, containing thin laminations of various materials. The drill, the rock, and the environment strongly affect the dynamics of rock cutting. The challenges of modeling the drilling process are discussed in detail in this section.

4.1.1. Drill Mechanics

The mechanics of rock drilling are very complicated to model. Compared to metal cutting, relatively little is known about the basic mechanism of rock cutting. Researchers may not agree on the model for bit-rock interaction, but the basic mechanisms are similar for drag bits, roller-cone bits, and coring bits. As a bit tooth penetrates the surface of the rock, a pocket of crushed rock is formed beneath the tooth. When the penetration force exceeds the crushing strength of the rock, a crack forms and propagates back to the rock surface, creating a chip [37].

There are a large number of variables that influence the mechanics of drilling. These variables originate from the drill mechanism, the geology of the mine, the mechanical properties of strata, and the variations between drills and drill operators [15][35][34][48]. For example, the depth of the mine can significantly affect the dynamics and mechanics of the bit-rock interaction. Some researchers have postulated that at very high confining pressures, rock deforms in a quasi-plastic manner rather than elasitically [29]. This type of deformation changes the geometry of the crushed section, the size of the chip, and the chipping frequency. Garnier and van Lingen presented evidence showing that drilling rate decreases with increasing confining pressure, and increases with cutting angle of the bit [15]. Somerton reported strong correlation between drilling variables and penetration rates. Specifically, drilling chips became smaller as drill bit wear progressed [48].

4.1.2. Material Properties

The physical properties of a material include strength, hardness, abrasivity, porosity, grain size, conductivity, microscopic structure and fluid properties, to name a few. After decades of research, it is widely accepted that drill behavior during drilling is related to the mechanical properties of a material. Using constant drilling forces and bit type, Gstalder and Raynal reported useful relationships between rock hardness, Young's modulus and the velocity of sonic waves through the rock [18]. In one of their drillability studies, Paone and Bruce demonstrated that drill behavior changes with various types of concrete [10]. They also reviewed over 30 papers on rock drilling experiments. Several experiments were successful in using the rock properties measured in the laboratory to predict drilling rates. One study correlated drill penetration with material hardness and the number and size of quartz grains present in the rock. Researchers agree that relationships between rock and drill are difficult to theoretically define, and indeed, most have been drawn from empirical data. However, the common conclusion was that rock properties and drill parameters are related.

Making it even more difficult to model the bit-rock interaction is that the physics vary between different types of rock and even on different samples of the same rock. Strata are usually riddled with microcracks, mineral deposits and discontinuities. Two rocks with similar strength may have very different grain sizes and therefore, the drill bit will cut and wear differently. Somerton investigated the effects of heterogeneous rock on drilling energy. He analyzed the drilling chips from rock with two or more mineral constituants and concluded that the weaker constituent is broken into smaller chips. Therefore, as the makeup of a rock mass varies, so will the drilling energy. Furthermore, even the physical and mechanical properties of rock are not invariant: strata are affected *in situ* by conditions such as confining pressure, temperature, moisture content, the presence of gases and with the process of mining in itself.

4.1.3. Drill-Material Relationships

A common characterization used in rock drilling is the specific energy of drilling or SED [51]. SED is a measure of the rotational and linear energy needed to drill a volume of material, and is defined as:

$$\mathbf{E = F/A + 2\pi NT/Au}$$

where F is the drill thrust force in lbs., A is the area of the drill hole in sq. in, N is the rotary speed in rev/s, T is the drill torque in in-lb., and u is the drill feed rate in in/s. SED is often compared to the compressive strength of a material because these two numbers share the same units and a great deal of research has empirically verified a relationship between them [25][28][36]. In an extensive rock drilling study, Teale reported that drilling energy is not constant unless the rock chip fragment size is constant: several experiments showed that the amount of energy required to break a rock depends on the size of the rock fragments, or the amount of re-grinding that takes place [51]. Maurer studied the effects of hole cleaning on drill efficiency and concluded that if the drill hole is not cleaned sufficiently, re-grinding occurs and drilling rates fall below a theoretical rate for ‘perfect cleaning’ [28].

Most researchers have used SED as a method of characterizing strata [25][45][55]. This is acceptable if one *only* wants to learn the relative strength between the layers of strata and other geological features (such as voids, faults and clay veins). SED is at best *only* an estimate of the compressive strength of rock and critically depends on how finely the rock is ground at the bit. Using SED as an estimate of the strength of a single material would be misleading and perhaps dangerous [28]. Furthermore, SED by itself cannot be relied on to distinguish between two materials because there will always be cases in which different strength materials look similar in strength (such as a strong material being fractured or the operator changing the manner in which he or she drills). Therefore, regardless of their similarities or differences and regardless of variations in the drilling environment, additional relationships between rock and drill must be used to distinguish between materials,. By classifying rock strata *and* using measured drilling parameters, one can determine, for example, that the rock is classified as sandstone, but it’s relative strength hasn’t changed much from the last material, and therefore, it may be fractured.

4.2 Neural Networks for Material Classification

With all of the geologic, environmental, and mechanical variables, drilling is a process which produces large, multivariate data sets. Considering the number of variables that affect the drilling process, and the complex relationships between these variables, machine learning is an appropriate method of accessing the information that will make it possible to classify rock and detect material boundaries [30][44].

An early survey of and experiments with several machine learning methods, including decision trees, Bayes’ classifiers, and instance-based learning, suggested that a neural network is an appropriate learning algorithm to use. Decision trees are best suited for problems that are represented by discrete-valued functions. They also perform best with input variables that are not inter-dependent. The drilling in this research has a large number of variables and highly complex relationships between those variables exist. The decision trees performed poorly in experiments. Instance-based learners, such as k-nearest neighbors, most often had the problem of overfitting to the training data, and thus did not perform well with new data. Bayesian Learning is a probabilistic approach to classification, that requires prior probability knowledge and discrete variables. A Bayes’ classifier would probably perform equally as well as a neural network, but the fact that the scale of the sensor readings may change non-linearly would make implementation problematic. Bayes' classifiers also performed poorly in early investigations.

The primary reasons that a neural network is well-suited to this material classificaion problem are:

- the input vector is real-valued data from multiple sensors
- the dynamic, non-linear relationships between mechanical, geologic, and environmental variables are not well-understood or even known.
- drill sensors will contain noise and/or errant information.
- network outputs can be probabilistic and thus information about the confidence of the classification is available.

An overview of the basic principals of neural networks is provided in the appendix.

4.2.1. Framework for Classification

The drill data points are represented in a multi-dimensional drill *feature* space. Features can be drill parameters (force, torque, penetration rate, etc.) and functions of the drill parameters, for example SED. Different materials will form distinct clusters in the n-dimensional drill feature space if the number of dimensions is large enough to uniquely define all of the materials under consideration. Incidentally, an analysis of the spatial and temporal characteristics of the data clusters within the drilling space may yield additional information such as possible rock mass fracturing, variations within a material, or drill bit wear.

To enable the neural network to interpret the complex relationships between the drilling variables, I 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 non-linear relationships between drill behavior and material properties such as specific energy of drilling. Additionally, they may capture complex, hidden relationships between the drill and the material, for example, a time-series analysis of drill forces may quantify a material's force "signature". The information from the virtual sensor becomes another parameter, or feature, in the drilling space and another training variable for the neural network. For example, the mean and standard deviation of a sensor over a finite time step or a function of one or more sensors can all be virtual sensors.

4.2.2. Classification by Drilling

There has been extensive research into precision drilling—such as manufacturing drilling, milling and machining—in which the drill is monitored in order to improve the outcome of the process. This research has given some valuable insights into drilling analysis. A few studies, briefly described below, introduce the relationships between drill parameters (or other drill-related sensors) and the material being drilled.

Ramirez and Thornhill 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 [42]. Rangwala and Dornfeld monitored drill condition with a neural network trained on selected features from the power spectra of force and acoustic emission [43]. They employed the Sequential First Search algorithm to select the most relevant features (frequencies from the power spectra). There are a number of other studies on intelligent drilling for the manufacturing industry [8][19][27][41].

Kaburlasos et al. used learning techniques to associate drill thrust and torque, recorded during surgery, with the thickness of the inner ear bone being drilled [23]. The goal was to learn the drilling profile for the entire drilling time and use that to predict drill breakthrough and act accordingly to prevent unnecessary damage to ear tissue. In the same surgical procedure, Brett at al. measured force and torque to determine the state of the drilling process, drill and the stiffness of the stapes at constant feed rate. They used the resulting force and torque data to detect when breakthrough was about to occur [7].

4.2.3. Feature Selection

One important element of building an accurate classifier is to choose the features that best represent the specific drilling application for classification purposes. This sometimes requires more information than the physical sensors alone. Manually determining appropriate features is time consuming, and may be acceptable for a small number of inputs or for one time use, but there are too many variables to do this manually for large data sets with varying characteristics, especially if the drill data may come from different sources at different times. For these cases, a methodology for network design is called for. Various classification research projects and their associated feature extraction methods are described below.

Wong et al. used a neural network to predict the porosity of oil bearing rock [57]. They improved approach of classifying rock by using extra information from the human interpreted core data as input to a neural network. A neural network was trained and tested on downhole log data using both the traditional input parameters to classify rock and an improved method which included data on the hydraulic properties of the rock, called lithofacies, as input features. A separate neural network was first trained to classify the well log data into lithofacie classifications. Core data was used to verify those classifications. Lithofacie was then used as an input, along with the well log data, to a network to predict porosity. Their research used weight visualization curves as a method of optimizing the network structure for classifying the lithofacies. The network weights measure the contribution of each node to a node in the next layer, and number of nodes is optimized in this way. The additional feature improved the porosity predictions, and this was verified by sonic logs not used in the training phase. The neural network reached 95% correct classification. The authors also proposed a technique for successfully designing a network architecture for their applicaton, however, the generality of the network and the of technique was yet to be investigated.

Rangwala and Dornfeld used intraclass Euclidian distance measures to select features that are most sensitive to tool wear and least sensitive to process parameters [43]. This distance measurement is a representation of the scatter within classes and between classes. (Classes are the clusters in the multi-dimensional feature space). They chose features using the signal-to-noise ratio as search criterion for a sequential first search.

Benaouda et al. demonstrated a method of using a neural network to assign rock lithology to the missing sections of a recovered ocean floor core. The network was trained on borehole data from various resistivity, lithodensity, sonic, natural gamma and geochemical sensor data [4]. Principal Component Analysis (PCA) and Cluster Analysis (CA) identified six significant dimensions to the data and suggested five distinct classes, many of which had significant overlapping regions in sensor data. They compared the network results with various Discriminant Analysis (DA) techniques, concluding that neural networks were better at classifying the missing rock. They reached upwards of 93% classification accuracy with test data and good qualitative correlation with non-depth-matched cores. Most of the analyses were interpreted by a geologist who subsequently fine tuned the network parameters. Benaouda et al. did not present a method for training a network on new borehole data.

Preston et al. correlated four physical properties of rock—seabed grain size, shear strength, bearing strength and porosity—with their recorded acoustic profiles [39]. The goal of their research was to empirically estimate properties of sections of seabed with acoustic data instead of using core samples or other expensive instrumental analysis. They first used multivariate analysis to extract features from the acoustic signals and used unsupervised clustering to classify rock data. Canonical correlation analysis demonstrated strong correlations between sets of acoustic and linear combinations of the rock properties taken from core data. (The analysis also showed very good correlation between the four physical properties of rock.) These linear functions were used to classify a test set of acoustic profiles into categories of rock. This relationship could then be used to estimate the four rock properties from new acoustic data. Although the work of Preston et al. was not directed at actually classifying rock, their work was important in that it verified relationships between mechanical properties of rock and sensor data collected from acoustic echoes form seabed strata, another method of observing characteristics of rock. Equally as important, their research successfully demonstrated the use of machine learning to find these relationships.

Perez et al. used a neural network to classify rock from a copper mine. They selected the network inputs from a database of color images of the rocks. Although they do not use a drill as the source of data to train the neural network, their research is further evidence that different kinds of physical properties of a rock, such as those sensed in a non-destructive manner, can be used in classification. They extracted 130 features from each rock image, such as rock radius, perimeter, diameter, geometrical moments of inertia, geometrical variation statistics, etc. They used a genetic algorithm to find a minimum set of features that would maximize the network classification rate while minimizing the number of network inputs [38]. They were able to reduce the number of inputs by 50% and still obtain classification rates better than 90%. In addition to using neural networks to classify rock, they developed a unique method to automatically search for the best rock features with which to train. The best feature sets were found through an iterative search using the accuracy of the trained networks as the criteria for judging the usefulness of the feature sets. Perez et al. chose to use a genetic algorithm, but there are many other types of search that would be appropriate.

My proposed methodology for building a classifier will include automated feature selection from a set of features generated by an expert driller. In this research, I will employ PCA, a feature selection method that is commonly used with neural networks.

4.2.4. Mining Drilling and Material Interface Detection

Early drill parameter investigations in mining and construction were limited in effectiveness and sophistication—the drilling parameter recorders did little if any analysis and rarely was it done real-time. For the most part, the drill operators were responsible for interpreting the drill data from a strip chart at the end of each day. One of the earliest instruments to be packaged specifically for recording drilling parameters in ground engineering was the Enpasol by Soletanche Enterprises. It estimated selected rock properties from the drilling rate, thrust and mud pressure, and output strata denominations based on human interpretation of the drill data [39].

Leighton et al. used an instrumented rotary blasthole drill to correlate drill performance in different materials with blasting variables in order to better plan open pit blasting [26]. They were able to discover an empirical relationship between the Rock Quality Index—a parameter calculated with drill thrust and penetration rate and representing a material's resistance to drilling forces—and the Blasting Powder Factor. They accomplished this by performing a regression analysis on the drill and blast data from several satisfactory blasts. This work was performed with a particular drill and blasting powder. The authors acknowledged that drill data had biases toward operator performance and data interpretation.

Scoble et al. used drill monitoring to verify coal-rock boundaries in a surface coal mine using a rotary blast hole drill [45]. They calculated SED and compared it to known strata conditions. They reported that specific energy of drilling correlated well with the drilling of different materials, and concluded that drill thrust was the most important parameter when rotary speed and torque were held constant. They acknowledged that for the specific conditions of their test, the coal-rock compressive strength differences were significant which made the strata layers obvious. They also felt that small-scale features in the so-called homogeneous strata significantly affected drill parameters. In addition, the layering of the strata was known beforehand and drill data was interpreted by hand.

King et al. interpreted data from a coal mine drill using learning techniques [24]. They used unsupervised learning to classify underground coal mine roof drill data into clusters of geologic features represented by relative SED. The strata they drilled through consisted of different types of sandstone, with a wide range of tested strengths between 12,000 and 25,000 psi. They attempted to classify several other local features present in the strata, such as voids and clay veins which would also be clustered into a 'strength' category. The authors' conclusion was that unsupervised classification was able to discover patterns in roof bolter drill data. Depending on the distance cutoff values used to classify data points in the feature space, they obtained both nine and 16 classes of SED. Later, King and Signer used the previously formed classes to train a neural network [25]. They used drill torque, thrust, rpm and penetration rate as inputs to a neural network. (Actually, two neural networks were used—one with 9 classes, and one with 16 classes—and they compared the results.) Again, the classes represented major geologic features in the roof strata, derived from SED. King and Signer only performed a visual inspection of the network's classification next to the strata in a core sample. They concluded that the "networks did a good job of matching major features," however no quantitative analysis was performed.

Utt continued the work of King and Signer and used a neural network to classify rock based on SED [55]. He assigned each data point to one of 32 categories of rock strength, which was calculated from drill parameters. This data was used to train a neural network to assign a label of soft, medium, or hard to the drill data. The goal was to warn a miner 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 many things differentiate their work from this research. Their neural network categorized some artifacts of the strata by relative strength estimated from the drill parameters. Their research used strata with a wide range of compressive strengths (although none of the layers was labeled with a true strength measurement). It is not always the case that roof strata will have very different characteristics: coal and shale have overlapping ranges of compressive strength and are often found together in a mine. Both teams used a very small amount of data to train and test their algorithms. The combined training data from five drill holes only amounted to 617 data points. It is important to note that neither study classified the rock, but only sorted geologic features by the relative amount of energy used to drill it. These geologic features were not quantitatively examined using ground truth features from a core log.

4.3 Real-Time Classification of Materials

None of the aforementioned research attempted to classify strata in real-time, or from a data stream. In this research, I plan to classify strata that are typically found above or below a coal seam as sensor data streams into the network. In a real mining application, classification and CID must be achieved as mining takes place to maximize the benefits of real-time structural analysis within a mine. Furthermore, it is desirable to do this without requiring a mine worker to perform any classification, and regardless of the drill type, drill operator or location of the coal mine. My approach to extending the classifier to real-time operation is explained in detail in the Proposed Research section.

5 Preliminary Experiments And Results

5.1 Experimental Apparatus.

The experimental 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 2 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.

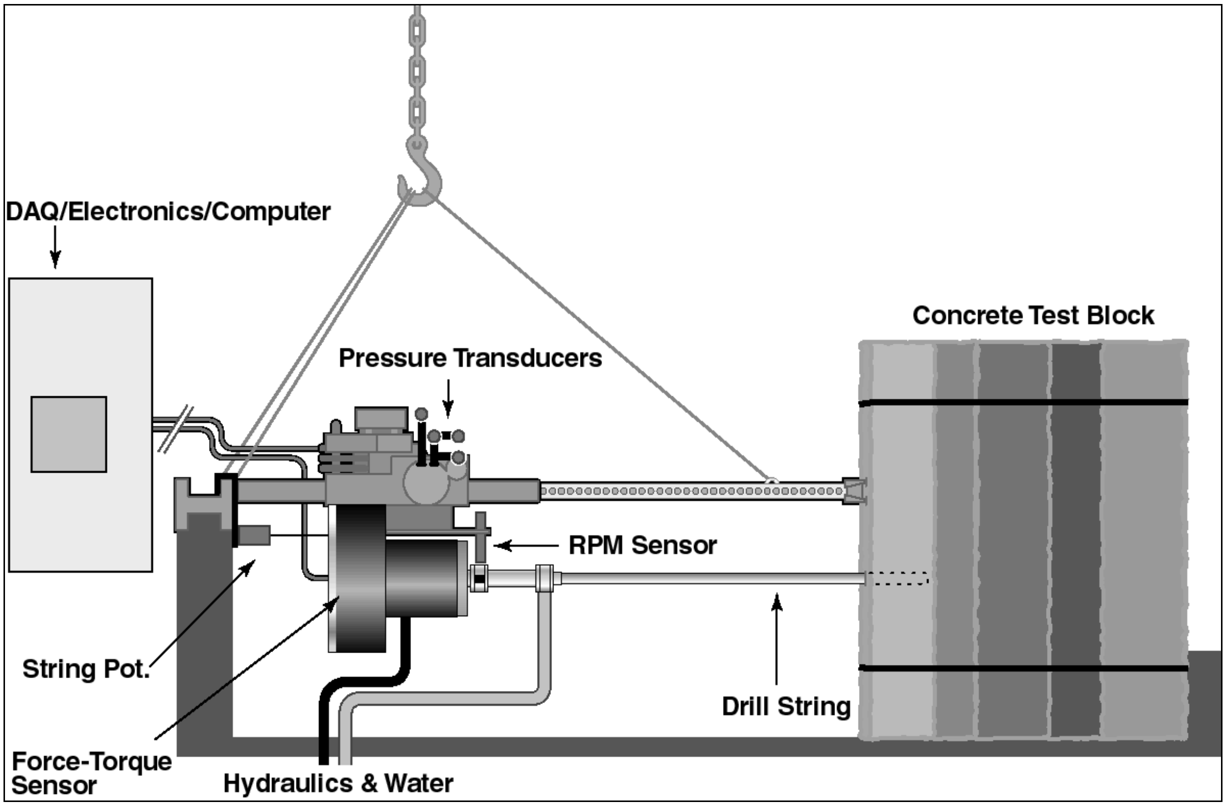

Figure 2. Laboratory Drill Apparatus and Setup

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. This is a redundant sensing scheme since motor pressures are proportional to thrust and torque. It has been implemented for the purpose of assessing the feasibility of measuring thrust and torque values from less expensive sensors that are more appropriate for a real world drilling system. The rotary speed is measured using a magnetic sensor and a collar with 4 embedded magnets attached to the spinning drill chuck. Currently, the drill hole recordings are being captured at 1 kHz, and drill penetration rate is calculated off line, using the drill bit position readings at known time steps.

The laboratory drilling set-up includes an adjustable frame to support the drill as it drills horizontally through layered concrete test blocks. The drill is vertically supported with cables. When a hole is drilled, the drill 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 some maximum value, with the goal of keeping the penetration rate as constant as possible. The rotation motor valve is fully open, but the flow is not controlled so rotary speed varies with the load on the system. To keep the drill hole as clean as possible from drill fines, the flushing water is set to full-flow each time a hole is drilled, the reason being to keep the drilling energy at a constant minimum as much as possible. Our intention was to keep the penetration rate constant, and as a result, the range of the force-torque sensor was at times exceeded by the high thrust needed to maintain that penetration rate.

5.2 Data Collection and Processing

To date, I have gathered data on about 40 holes drilled into a 3 foot thick concrete test block. The 3'x3'x5', 8,000 lb test block has 5 different layers of concrete, each being mixed to achieve different strengths and to simulate different types of rock. 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. I also drilled and recorded 30 holes at the National Institute for Occupational Safety and Hazards (NIOSH) Bruceton Coal Mine in Pittsburgh.

Layer	*Mix*	*Strength (psi)*	*Thickness (in)*
1	Grout	3,310	11
2	Sandstone	9,900	5
3	Limestone	2,300	9
4	High Early Strength	8,050	4
5	High Early Strength	7,580	7

Table 1. Concrete Test Block Characteristics

It takes about 90 seconds to drill a hole into the concrete test block. Figure 3 is an example of sensor data recorded while drilling a hole. We use LabVIEW software to acquire data from the seven sensors at a rate of 1000 Hz. A typical data file has between 60,000 and 100,000 data points, each with seven real-valued sensor readings.

Penetration rate is calculated off-line because the present data acquisition software cannot filter, process and record data at the desired rate. In the future, data acquisition software and hardware can be custom designed for very fast data acquisition and processing.

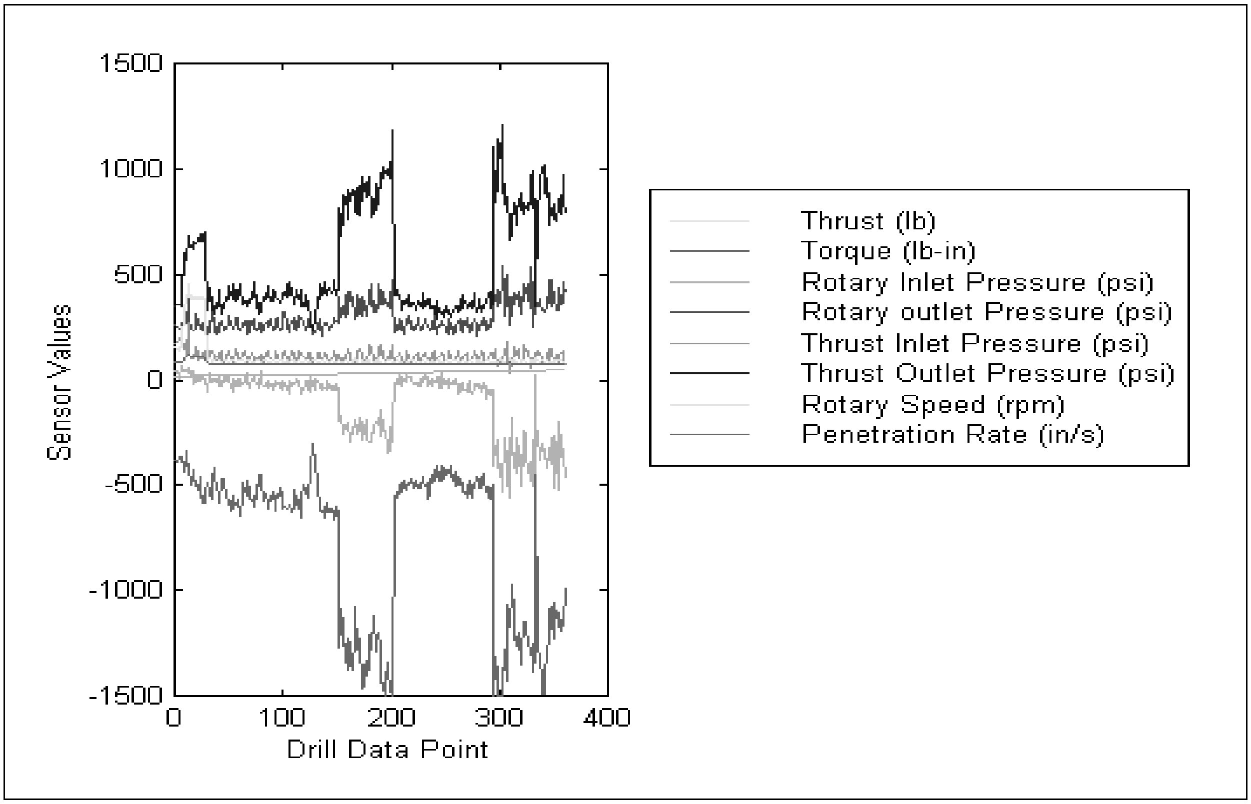

Figure 3. Concrete Drill Hole Sensor Recordings for One Drill Hole.

The drill sensor data used in these experiments is post-processed. The output of the string potentiometer is filtered with a capacitor and the force-torque sensor has a low-pass filter. The conversion of magnet pulses to rotary speed uses software filtering. Drill data is processed to reduce the size of the files, to prepare it as input into a neural network, and to facilitate in the analysis of real and virtual sensors (see Figure 4).

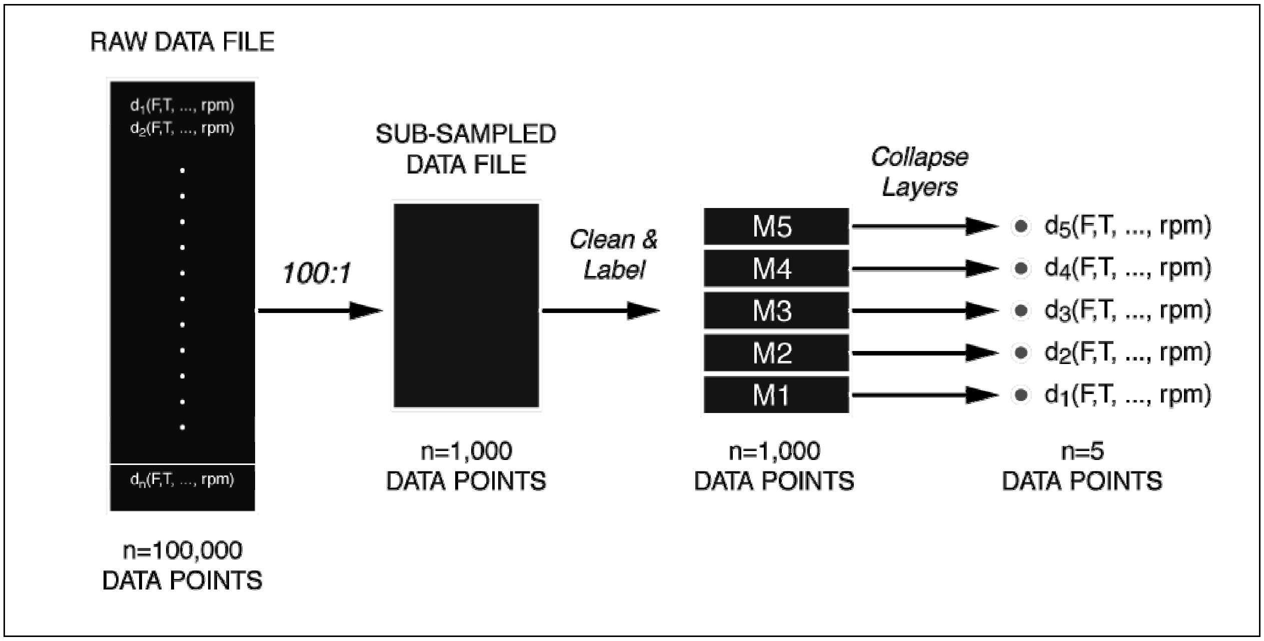

Figure 4. Data Processing Flow Chart

We begin with calculating the drill penetration rate at each data point using drill position and a known time step. 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, 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. Each is labeled by hand and normalized over the range of sensor values in that particular drill file. Finally, each segment of data is collapsed into a single data point. A data point contains many fields which are average sensor values or functions of the sensor values over the drill segment.

5.3 Neural Network Training and Testing

The neural network has been implemented with Netlab [6] and uses scaled conjugate gradient search, cross-entropy error function and softmax outputs (see appendix for more information). 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, in which the class with the highest probability is the winner.

The present experiments have been designed to determine if the drill data from concrete blocks can indeed be classified into 5 separate materials. A neural network with no hidden units was trained and the poor results, averaging 80% classification error, suggest 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 two layer 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 or in the same mine roof), the drill bit position was not used as an input to the neural network.

Drill Sensor	*Type*	*Experiment*														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Mean Value																
Thrust	p	■	■		■	■	■			■				■	■	■
Torque	p	■	■	■		■	■				■			■	■	■
RPM	p	■	■	■	■		■	■	■			■		■	■	■
Penetration Rate	p	■	■	■	■	■		■	■				■	■	■	■
Rotary Inlet Pressure	p	■						■	■					■		■
Rotary Outlet Pressure	p	■						■	■					■		
Thrust Inlet Pressure	p	■						■	■					■		■
Thrust Out let Pressure	p	■						■	■					■		
Thrust Pressure Difference	v								■						■	■
Rotary Pressure Difference	v								■						■	■
Standard Deviation																
Thrust	v	■														
Torque	v	■														
Penetration Rate	v	■							■							
Thrust Pressure Difference	v								■						■	■
Rotary Pressure Difference	v								■						■	■

Table 2. Attribute Sets Used in the Concrete and Coal Mine Experiments.

There are 14 drill hole data files used for the concrete 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 in this paper.

5.4 Experiment Results

The average error rates for each experiment, and a breakdown of the error rates by material are shown in Table 3. The error rate of material one is the percentage of time material one was misclassified as another material and is an average of the error rates of material one for the 100 data sets.

	Error Rates by Material (%)					
Experiment	*1*	*2*	*3*	*4*	*5*	*Avg.*
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 Concrete 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 rate. 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 5 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.

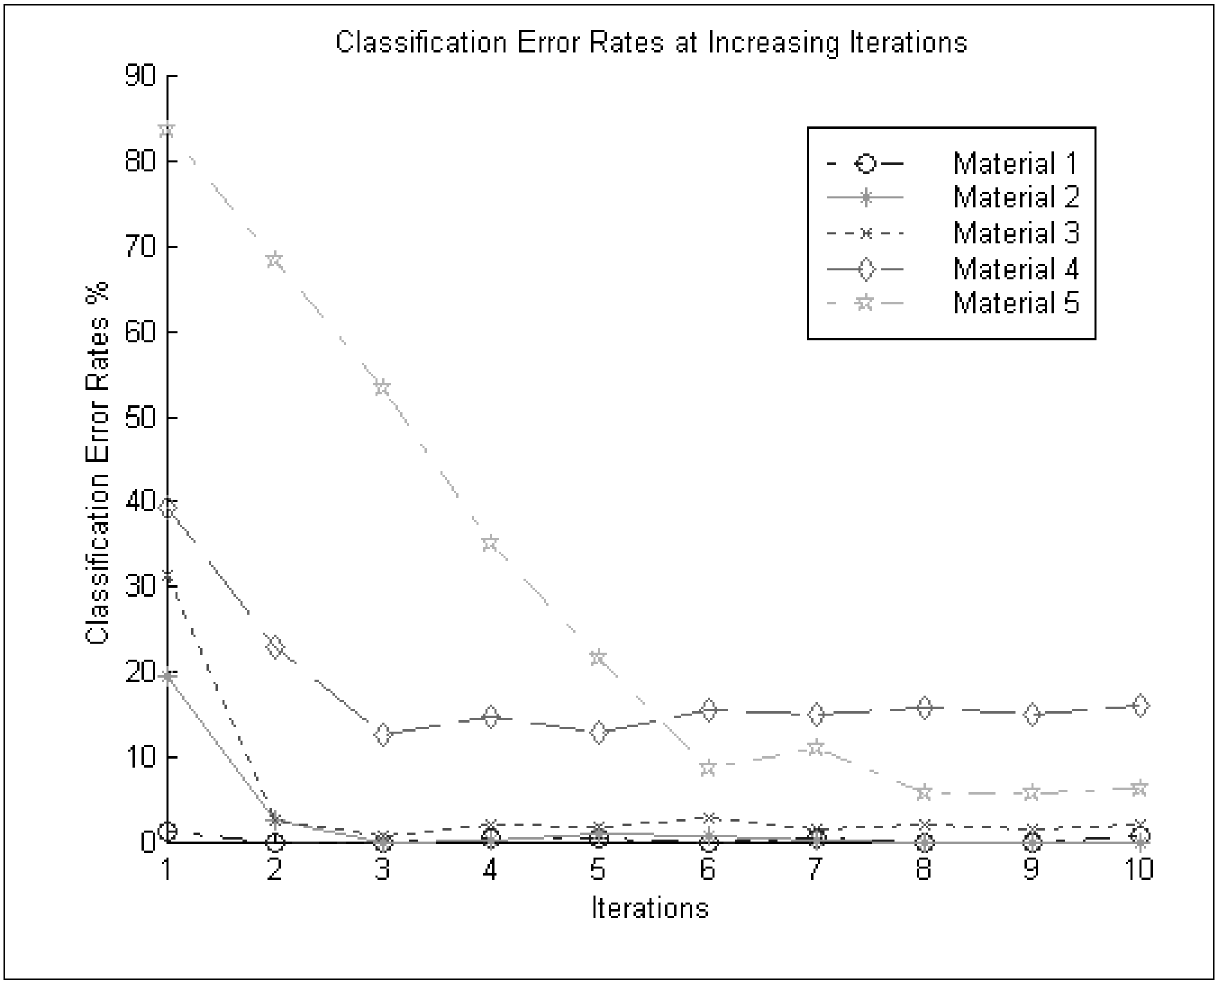

Figure 5. Network Classification Error Rates vs. Number of Iterations Using Experiment Group 1. Each line represents a material.

Figure 6 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 classification of the data sets in experiment 1. The ordinate axis is the true material class. Material 3 was misclassified as material 4 about 1.5% of the time, and material 4 was misclassified as material 3 about 9% of the time. Materials 4 and 5 were confused with each other about 6% of the time. The two concrete mixes of materials 4 and 5 are very similar in strength and makeup. The average error rate over all materials was 4.5%.

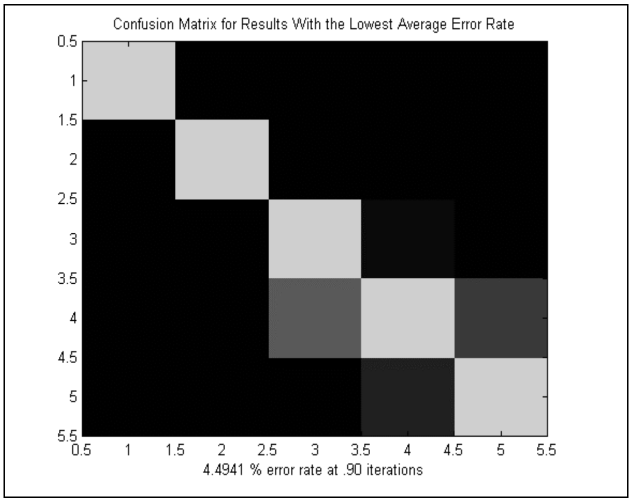

True Class / Network's Classification

	1	2	3	4	5
1	100	0	0	0	0
2	0	100	0	0	0
3	0	0	98.5	1.5	0
4	0	0	9.1	84.9	6.0
5	0	0	0	5.8	94.2

Figure 6. Experiment 1 Confusion Matrix Image (L) with Classification Percentages (R).

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. Since torque divided by penetration per revolution is theoretically constant for a given material, the thrust became more important in identifying a material.

Figure 7 shows a scatter plot of the normalized, labeled materials from one data set with the thrust-torque relationship and the thrust-penetration relationship. The clustering of the 5 concrete materials in these plots indicate that the combination of drill thrust and penetration rate is the are better at discriminating between the five concrete mixes. However, in addition, the thrust-penetration relationship can resolve the ambiguity in the thrust-torque plot between materials 1 and 3 as well as the ambiguity between material 2 and materials 4 and 5.

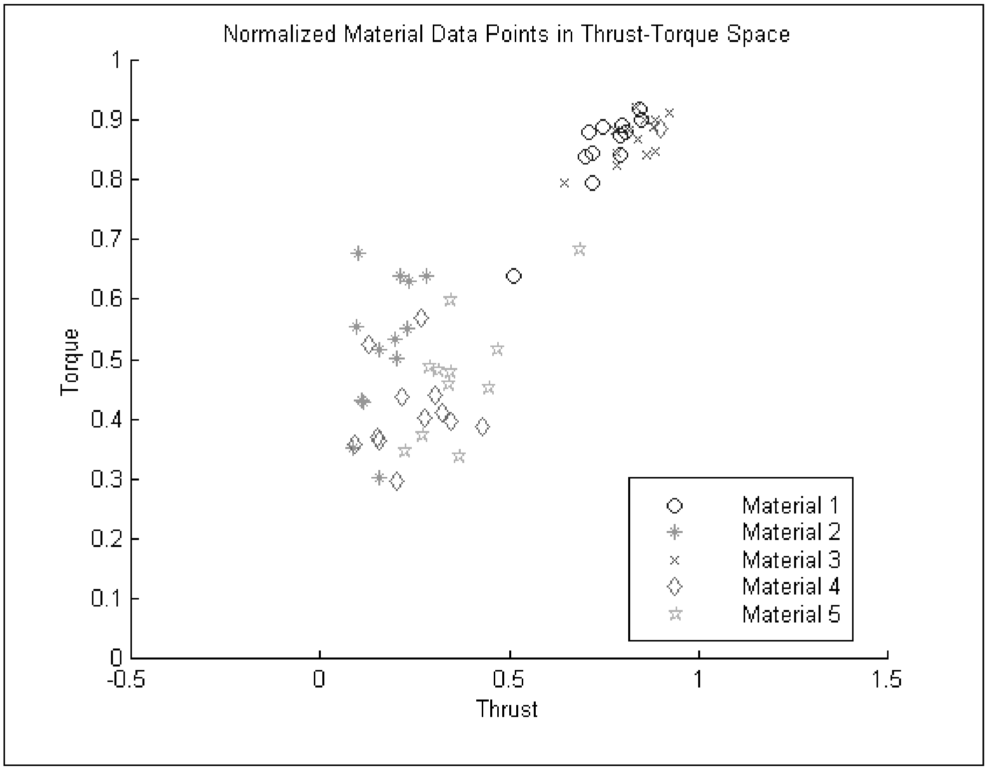

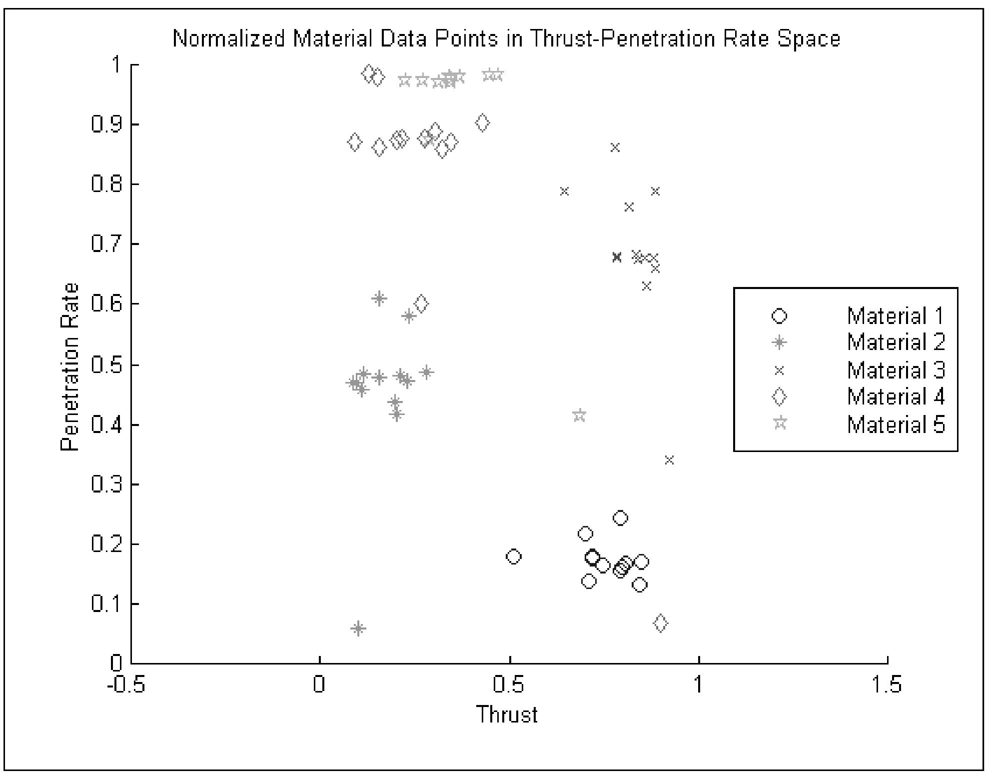

Figure 7. Experiment 7 labeled data sets: Thrust vs. Penetration Rate (Lower) and Thrust vs. Torque (Upper) .

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 was used to train a network in Experiments 9 through 12. Drill thrust appears to be the most important drill parameter in terms of having the lowest error rate, and this agrees with the results of Experiment 3. Closer examination of the results 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 in classifying materials 4 and 5.

Experiments 1, 7 and 8 indicate that the standard deviations of the drill parameters are useful in increasing the accuracy of the classification of the concrete test materials. This may be because the concrete mixes range in their size and concentration of aggregate, causing the drill to chatter at different amplitudes and frequencies for each material. This effect is better represented in the larger data files, and may prove to be even more useful. Experiments 7 and 8 also demonstrate that virtual sensors are indeed useful and may make a significant contribution to the accurate classification of similar materials.

5.5 Coal Mine Data

The experiments using the collapsed data from the concrete test block were performed primarily to test whether or not a neural network could even classify materials using drill parameters in a very simple format. The goal of this research is to correctly and confidently classify coal mine strata in real-time. Specifically, the classifier must detect boundaries between two very similar materials, coal and shale, using real coal mine drill data. The coal mining environment introduces more variables to the problem such as material and machine variations, environmental variables and data uncertainty.

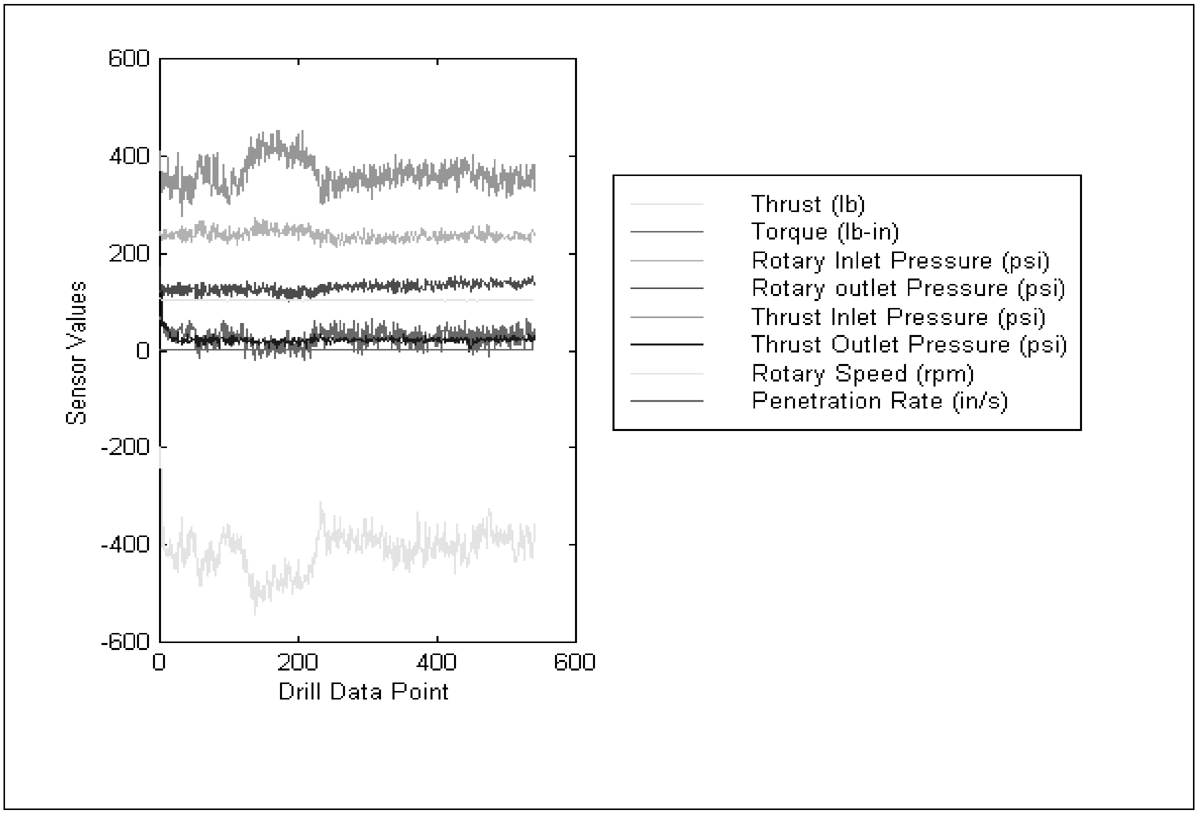

Figure 8. Coal Mine Drill Hole Sensor Recordings for One Drill Hole.

Figure 8 shows recorded sensor values for one hole drilled in the NIOSH experimental coal mine. Data was collected from roof strata consisting of coal, shale and clay. Figure 9 depicts the lithology of the sample drill hole. Table 4 lists the compressive stengths of the three materials.

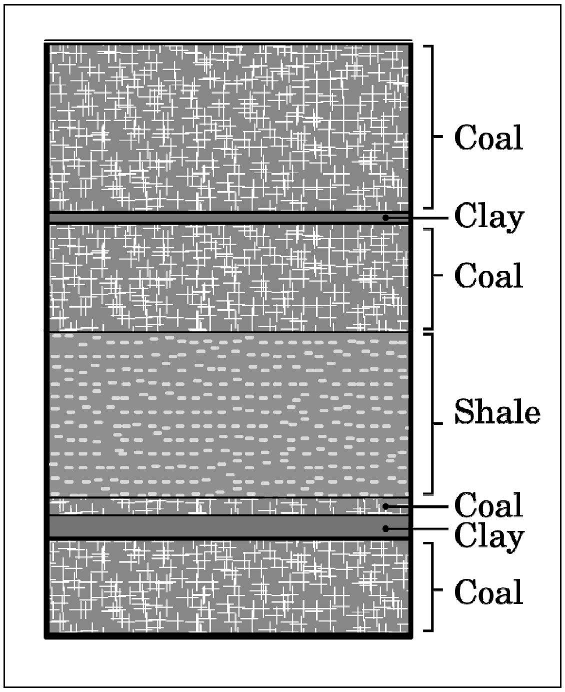

Figure 9. **NIOSH Coal Mine Strata for drill hole in Figure 8. The labeled strata corresponds to the drill data (data point zero is the bottom layer of coal)**

Material	*Strength (psi)*
Coal	900-1,000
Shale	4,000-8,000
Clay	3,000-6,000

Table 4. **Compressive Strength of the NIOSH coal mine strata**

Four out of the 30 coal mine drill files are currently labeled. I performed ten experiments with the hand-labeled data files, using 4 different attribute sets to train a neural network to classify the materials. Table 2 (page 20) shows experiments 13, 14 and 15 attribute sets and Table 5 shows the results of these experiments.

Experiment	*Error Rates By Material (%)*			
	Coal	*Shale*	*Clay*	*Average*
13	38.0	58.0	100	65.3
14	35.0	25.0	100	48.3
15	**0.0**	**0.0**	**68.0**	**22.6**
	Coal	*Non-Coal*		*Average*
15a	**0.0**	**0.0**		**0.0**

Table 5. Error Rates by Material for Each Coal Mine Experiment

In experiment 15, coal and shale were accurately classified. In experiment 13, no virtual sensors were used—only the physical drill sensors. Experiment 14 added several virtual sensors, and experiment 15, which had the most number of training features, added the most virtual sensors. As expected, experiment 15 performed the best, resulting in misclassifications of clay only (coal and shale were not misclassified). Classification errors may have accompanied the data sets with clay because these thin veins of material and could have been mislabeled or too small for the sensors to detect. Experiment 15a classified the same data set as experiment 15 except that there were only two classes—coal and non-coal. With two classes, all of the data points were classified correctly. Since there were only 4 coal mine drill holes were used in training, an accurate assessment of classifying mine data will only be possible when the full number of data sets are trained and tested.

6 Proposed Research Plan

The results of my initial work suggest that a neural network can be used to classify rock using drill parameters, and that using virtual sensors can improve the classification accuracy. The remaining research focuses on developing a methodology for building a working, real-time drill classification system that can be used for multiple drilling applications.

The primary goals of this research must answer the following questions:

- How will rock and drill variables affect classification rates?
- How can classification rates and confidence in the classifications be improved?
- What is the best method for selecting drill parameter features?
- What is the best way to detect material transitions?
- How will the neural network be modified to enable classification of a data stream?

The general research activities involved in answering these questions are:

- Perform sensitivity analyses on the drilling classification system as a whole, to obtain highest possible classification rate and confidence.

There are a great number of variables in the drilling process, even for the small data set I recorded in controlled conditions. The classification system should operate in the rough and often unpredictable conditions of a coal mine, and under the less-than-caring manner in which the drills are operated. For example, if an operator decides to drill faster all of a sudden, will this be detrimental to classification? How will classification rates be affected if there is sensor degradation? One shortcoming of the analyses is that there are only a limited set of variables that can be varied. Therefore, it will be a challenge to simulate as many conditions as possible to test. If some variables do affect classification rates, how can they be improved?

- Develop a methodology for automating the selection of drill parameter features and training and optimizing a neural network to classify materials and detect material transitions.

The aim is to build a system that takes a large number of previously generated features, and finds the best ones with which to train the network. The criteria for selecting the number of features can be based on the impact that the feature has on classification as seen from PCA. The features are evaluated and selected only once in the methodology. In this research, I generate features by hand from a fixed set of physical sensors. One issue about feature selection is that the best feature sets for classification and boundary detection may be different. In that case, the final feature set must contain all of these features in order for the network to train properly.

- Extend the method to classification of a data stream and test on new coal mine data

The approach to real-time classification is to simulate it using the data sets already obtained in real-time. There will undoubtedly be features used to classify that are time dependent, such as averages and standard deviations of sensor values. The approach to this is to use a sliding window technique. One technical challenge is choosing the window size: some materials may have smooth signals and averages can be computed with a smaller window; other materials may require a wider window because the variations in the signals are slower or more complicated. A fixed window size can be used with the trade-off of less classification accuracy. Or confidence can be measured to determine if the classification should be rejected.

6.1 Sensitivity Analysis

Initial experiments made the assumption that the drilling equipment, the manner in which the drill is operated, and the type of material does not affect the neural network's classification accuracy. The results were good, in part, because many variables were kept constant (for example, I drilled all of the holes in a consistent manner, the water flushing was set to the maximum flow rate for each drill hole, and the same 5 materials were drilled through in the same order for each concrete run). The proposed research will include experimental analyses to determine how each of the drilling variables affects classification rates, with and without confidence measures.

6.1.1. Experimental Analysis

Several experiments will be conducted to assess the effects that the drill dynamics have on material classification. These analyses will be performed with concrete data and with real application data sets. This will require the manufacture of additional custom concrete blocks. The proposed research experiments will include:

- Varying the drill parameters such as penetration rate and rotary speed, and changing from highly accurate force and torque sensors to less accurate pressure transducers.
- Examining classification results between the same material from two different holes, during progression of drilling within a single material layer, and between adjacent layers that are of different materials.
- Examining material classification with respect to changing drill dynamics such as changing the drill friction within different layers by drilling the same concrete block from the opposite direction, varying drill string stiffness by using various length drill steels, and varying the energy of drilling by varying the penetration rate.
- Examining material classification results with sensor inconsistencies (accuracy, range, resolution) by corrupting sensor data.

6.1.2. Confidence Measures

In some drilling applications, very high classification rates may not be directly attainable. However, a higher classification rate may be obtained if a data point is accepted as valid only when the classification has a high confidence rating. All of the experiments described earlier in this document have used an all-or-nothing voting scheme in which the class with the highest probability is chosen as the network output, and there is no indication of the confidence of the network's classification. Confidence is a real number beween zero and one that measures the difference between the two highest class probabilities, P_1 (highest) and P_2 (2^{nd} highest) output by the network. It is calculated as:

$$\text{Confidence} = 1 - (P_2/P_1)$$

I performed ten experiments to determine if classification error rates can be reduced by using confidence values to filter out the least confident classifications. The experiments used the trained network from a single test/train run. The classifier was trained and tested in the same manner as with the concrete and coal mine data, but with the intention of obtaining high error rates so as to better demonstrate the results. I increased classification error rates by limiting the number of training iterations to 120, and limiting the number of training data sets to 20 (compared to 100 data sets used in the concrete experiments). Error rates were further increased by using the weakest drill sensors as determined from the previous experiments. In fact, it is possible that the network was trained so poorly that it did not learn a useful model for some of the materials.

I used several different values to filter out the classifications with the least confidence. If the confidence was below a given value, the data point was rejected. The results of the confidence thresholding experiments are shown in Table 6 below.

	Error Rates (% Original Data Classified)				*Top 2 Class Probabilities* ↓
Confidence	Zero	0.5	0.7	0.9	Zero
Material 1	3.8 (100)	3.8 (100)	0 .0 (96.2)	0.0 (92.3)	0.0 (100)
Material 2	78.9 (100)	87.8 (71.9)	93.3 (52.6)	100 (24.6)	35.1 (100)
Material 3	10.5 (100)	7.3 (96.5)	3.8 (93.0)	1.9 (91.2)	7.0 (100)
Material 4	64.2 (100)	75.0 (45.3)	72.2 (34.0)	60.0 (9.4)	39.6 (100)
Material 5	83.9 (100)	90.9 (71.0)	94.1 (54.8)	100 (22.6)	35.5 (100)
Average	48.3 (100)	53.0 (76.9)	52.7 (66.1)	52.4 (48.0)	23.4 (100)

Table 6. Classification Error Rates Using Confidence Threshold

These experiments demonstrate the validity of using confidence thresholds to improve the classification accuracy. Without confidence thresholding, classification error rates averaged 48.3%, worse than randomly classifying the materials. As the confidence rate was increased, the error rate significantly decreased for materials one and three—as much as an entire order of magnitude. However, the error rate for the other materials remained at a significantly high level. This suggests that the classifier learned a correct model for materials one and three but did not have enough training to learn a model for materials two, four and five. It is also possible that there was too much error for the network to learn those materials. Even with the strictest filtering, with confidence at 0.9, more than 90% of the data points from materials one and three were still classified, whereas the amount of data points classified from materials two, four and five sharply decreased to below 10% at times. Table 6 also shows that the error rate is dramatically decreased when the true class is represented by either the first or second highest class probability. In other words, nearly 75% of the time, the correct class is one of the top two class probabilities output by the network. This suggests that with the right drill sensors and confidence requirements, all five materials can be classified with low error rates.

6.2 Classification Methodology

In the preliminary experiments presented, the data files that were used to train and test the neural network were highly simplified (see Figure 10). First, for every drill file, the transition data between two materials was removed, and only the data from each contiguous layer of material was hand-classified. Then, each drill feature for a contiguous layer of material was reduced to one number by averaging it over the layer (giving, for example, average thrust for material A). The purpose of this data simplification was only to verify that a neural network could indeed classify several different materials using drill sensors as network inputs. There are three critical aspects of the drill data that will be addressed in the development of a classification methodology for a real-world lithological classifier.

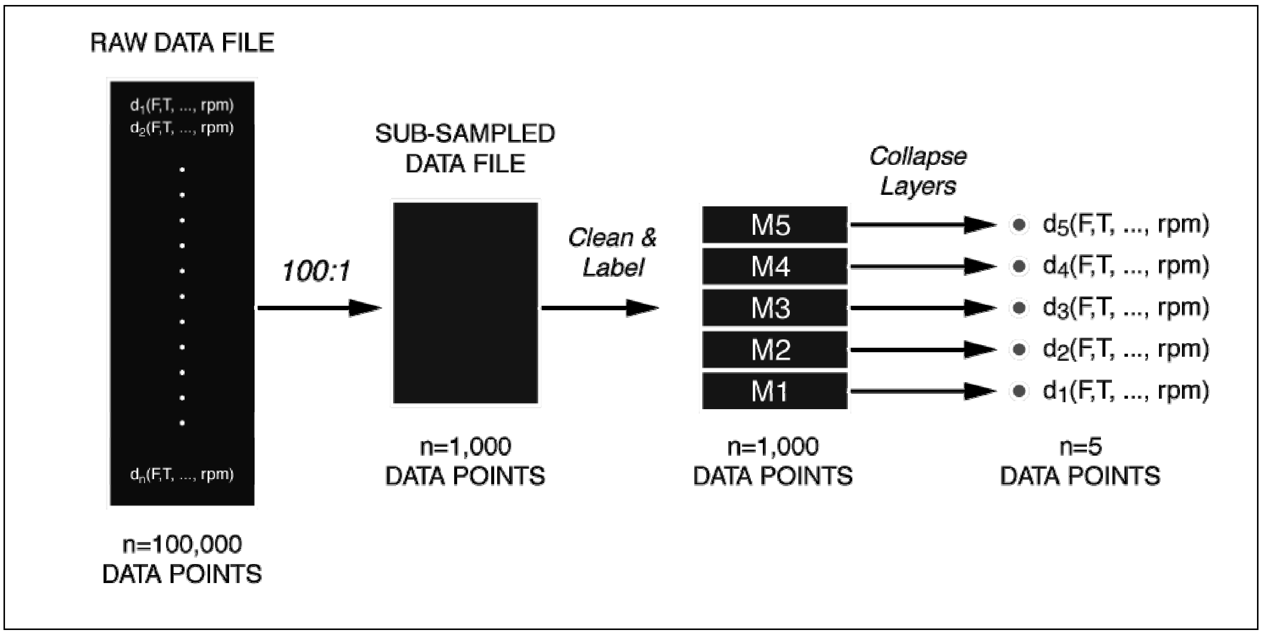

Figure 10. Data Pre-processing Flow Chart

The first task is to expand the data used in training and testing the neural network to include *all* of the sensor data in a layer, not just an average of each sensor over the entire layer. The consequence of this is that one does not know all of the statistical information for a layer of material beforehand, rather the data is fed into the training algorithm incrementally. Therefore, features such as sensor averages and variances must be computed incrementally, using a temporal ‘window’ of data.

The second task is to retain the material transition points during classification. In real-time classification, these data points are present, but the fact that they represent a transition from one material to another is not known beforehand.

The third task is to address the topic of feature selection. The drill features best suited to classify a set of materials will not be the same for all drilling applications or even for different instances of the same application. The optimal feature set will most likely depend on both the materials being drilled and the drill itself. Therefore, one cannot assume there is an optimal feature set for a drilling application or for a certain type of drill. Knowing which features to use in the first place or how to generate them is an additional challenge. One could use brute force to calculate a very large set of features from the available sensor data, but this would be a formidable task. Fortunately, it has already been established that drill parameters are dependent on material properties, and relationships between these variables is a good place to start. I propose to generate features in a scientific manner, using both theoretical and empirical knowledge about drill mechanics, mechanical properties of materials, and drill-rock interaction models.

Figure 11 below shows a diagram for the proposed methodology. The left half of the diagram represents the training phase in which data is pre-processed, features are selected and a learning model (a neural network) is trained. The right half of the diagram represents classification of a data stream. The data is not pre-processed, but the selected features stream into the classifier via a temporal window. The output of the model is the material probability vector from which the material is classified and the material boundary is detected. The material probability vector is calculated from a probabilistic model (to be determined) based on drill features over time. This is discussed in the next section.

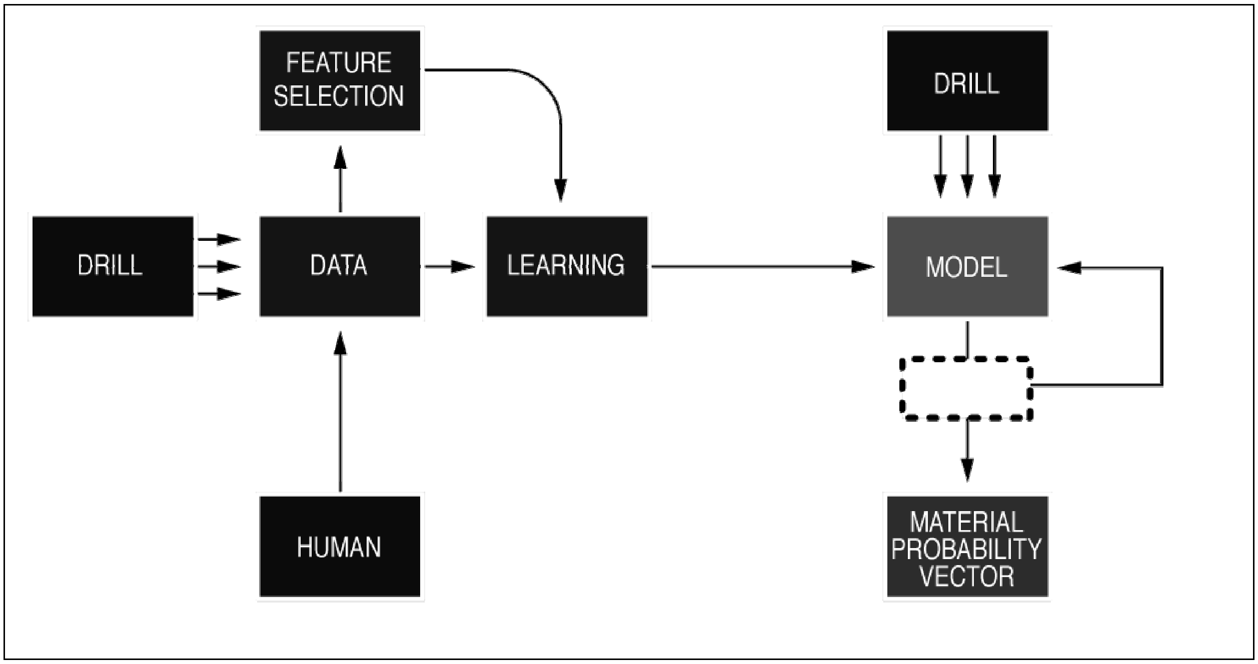

Figure 11. **Diagram of the Classification Methodology**

6.3 Classifying a Data Stream

There are many things that make real-time classification different and more challenging than off-line classification. Fortunately, applying and testing real-time classification does not require live drilling—the drill data sets which were recorded in real-time can be streamed to a classifier, thus 'simulating' real-time classification. The major challenges of classifying a data stream and proposed approach are discussed below.

The experiments performed thus far use data sets consisting of a single data point of condensed information about each layer of material in a drill hole. In a realistic environment, data is streaming in constantly. As a result, some of the features used in the initial experiments, such as averages and standard deviations cannot be computed for a single point. However, the time series nature of streaming data can be used. The network can make use of the history of the data by using a window of data to compute features.

- Real-time classification places constraints on how the drill data can be used in classification. For example, one cannot aggregate all of the material data points to calculate drill features because the sensor data is being received incrementally by the classifier. I propose to continue to train a network with cleaned and condensed data sets but to use a sliding window of sensor data to calculate features for the task of classification. The danger in using chunks of data to train the network is that chunks of data from the same material may form clusters in the feature space, and a neural network may overfit to this more complicated pattern of data rather than representing a more general function.

- As a consequence of data streaming, a smaller amount of information is contained in a chunk of sensor data. As levels of noise in the sensors vary, the size of the data window that is needed to capture critical behaviors is likely to be different for each material and even for each window of the same material. An insufficient amount of data will produce inaccurate features, and further reduce the accuracy of the classifier. I propose to employ a voting scheme to reduce noise in the network output.

- In many drilling applications rapid material boundary detection is desired. This research will investigate some ways in which material transitions can be detected. Material transitions can be addressed in several ways, soe of which are described below:

    - Consider material transitions as separate classes, or
    - Use a separate material transition 'detector' such as a fitting a step function to an input parameter, or
    - Detect the transition via a probabilistic model, using the temporal characteristics of the data.

  Material interface detection may be defined as going from a high-confidence material probability vector in one material, to a lower level of confidence for a while, and then back to a higher confidence in the classificatins as the next material is penetrated. In the temporal setting, probabilities from the network and the confidence of a classification can be very telling and may well aid in detecting material boundaries. Temporal voting will help to filter noisy network outputs, resulting in better overall material interface detection rates.

6.4 Drilling Applications

This research proposes to test the training methodology on concrete drill data as well as real underground coal mine drill data. The goal is to classify a small subset of strata typically found in a coal mine, such as coal, shale, sandstone, mudstone and clay. In addition, the classifier must be able to distinguish between two very similar materials such as coal and shale. The concrete and coal mine data will also be used to port the classifier to classifying a data stream. There is a long list of materials, drill types and drill sensors used in drilling applications today. I will investigate another application, such as well drilling or bone drilling, with the goal of making an initial assessment of the portability of the methodology to different drilling applications.

6.5 Success Criteria

This research will be evaluated using several criteria. These are:

- Classification rates, especially for two similar materials
- Material boundary detection
  - false positive detections
  - false negative detections
  - successful detections
- New knowledge about material classification using a drill
- Scalability to other applications

7 Schedule

7.1 Detailed Tasks and Duration of Work

All subsequent experiments are performed using a data stream.

1. **Complete experiments with initial coal mine data. (8 mos.)**
   - Borescope remaining drill holes in the NIOSH mine.
   - Label remaining 25 coal mine data files (label rock strata and note suspected transition data points).
   - Sample data files to reduce size.
   - Modify training programs to use a sliding window of averages of sensor data to test and train a network.
   - Develop a model for detecting material transitions
   - Train network using all drill sensors, real and virtual, to obtain a baseline network classification performance with which to compare subsequent training runs.
   - Choose the algorithm to search drill data for the best features (i.e., PCA)
   - Write programs to evaluate and select the best features to train with and use this *final feature set* for subsequent training runs.

2. **Collect additional concrete data for sensitivity analysis (2 mos)**

Existing concrete block

- Make any necessary repairs to the drill and data acquisition hardware/software.
- Turn concrete block 180 degrees in order to drill the block from the other direction (or pour another layered concrete block if needed).
- **Experiment (a)**

Drill 4 holes in the existing concrete block. These holes will be compared to holes drilled in the other direction.

New Concrete block

- Pour two new concrete blocks with two different layers, size 3' x 3' x 3', with which to perform sensitivity analyses. The concrete layers should be similar enough to simulate similar materials, but the strength must be different enough to detect with the drill parameters, for example 2 sandstone mixes at a 1000 psi strength difference.
- Test concrete samples to determine compressive strength.
- **Experiment (b)**
  Train networks using all drill sensors, real and virtual, to obtain a baseline network classification performance on the new concrete blocks, with which to compare subsequent training runs.
- **Experiment (c)**
  Drill 6 holes, 3 holes each in both directions, keeping the manner of drilling as constant as possible (i.e., set maximum constant flow rate to drill penetration and rotation motors) to examine how friction from the drill string, which increases with the depth of the hole, affects classification. The holes drilled in the opposite direction can evaluate the dynamic effects with changing material positions as well.
- **Experiments (d,e)**
  Drill 6 holes while varying one drill parameter. Drill 3 holes each with varied penetration rate and rotary speed.
- **Experiments (f,g,h)**
  Drill 6 holes, 2 each with 3 different drill string lengths to capture changes in drill dynamics (i.e., drill string stiffness, wobbling action).
- **Experiments (i,j,k)**
  Drill 6 holes, 2 each with 3 different flushing rates. Drill in a consistent manner so that other drill variables will not affect SED. Examine how flushing rates affect classification rates (and SED).

3. **Perform sensitivity analyses on all concrete data in simulated real-time (2 mos)**

- Train and compare classification rates for varied thrust, penetration rate, rotary speed, friction, flushing rates, and drill string length (**experiments d-k**) with the baseline classification rates (**experiment b**). Examine the effects of parameter variations on classification rates within each experiment (i.e., determine how variations in thrust affect classification rates.)
- Test and train using all of the holes from the new layered concrete block (**experiments b-k**). Compare this with a training run with no deliberate parameter variations (i.e., pick one thrust, one flow rate, etc.)

4. **Final Experiments (2 mo)**

   - Train using real data sets and test on data sets with a disabled sensor to evaluate the effects of missing data on classification rates.
   - Test methodology on two other coal mine data sets of 10 to 20 holes each, from different mines.
   - Evaluate the accuracy of material transition detection.

5. **Write Dissertation (4 mos)**

7.2 Schedule

Task	Oct, 00	Jan, 01	Apr, 01	July, 01	Oct, 01	Jan, 02	Apr, 02
Concrete Data Collection			■				
Concrete Sensitivity Analysis				■	■		
Data Stream, Boundary Detection Programs	■	■					
Coal Experiments (Initial Data) with data stream		■	■	■			
Coal Experiments (New Data)					■	■	
Final Experiments (Coal and drill sensor sensitivity)						■	■
Dissertation							■

Table 7. Condensed Schedule

8 Contributions

- Methodology for building a material classifier with a drill
  - classification of mine strata
  - detection of material boundaries
  - feature selection
  - classifying data streams
- Increased knowledge and greater applicability of material classification in realistic drilling environments.
- Detection of material boundaries, especially between very similar materials such as coal and shale and clay.

- Increase productivity and efficient use of natural resources
- Building a database of coal mine roof drill strata and concrete test block drill data

9 Appendices

Neural Network Overview

The fundamental unit of a neural network is the perceptron. A perceptron and its connections to other perceptrons is modeled after the neurons in the human brain. A neural network is composed of layers of interconnecting perceptrons, or *nodes*. Each node in one layer is conneced to every node in the succeeding layer with a unique real-valued number, or *weight*. The value at a node is the sum of its inputs mulitplied by their respective weights. A neural network configuration is shown in Figure 12 below.

The simplest network has two layers, an input layer and and output layer, with each output node being the sum of the inputs multiplied by their weights. This network structure yields only linear functions at the outputs. A typical network contains an input layer, one or more middle or *hidden* layers, and an output layer. Known as a multi-layered perceptron, this network can yield non-linear functions at the outputs. The number of hidden layers and the number of nodes in each layer determines the complexity of the network and thus, the complexity of the functions it can represent.

The perceptron thresholded unit (the output is a step function) is discontinuous and the gradient of the network error, which is needed in order to search for the minimum error, cannot be computed. A sigmoid threshholded unit, shown in Figure 12 below, takes the output of a perceptron unit and applies a continuous smoothing function so that the gradient of the error can be computed and the network weight space can searched for the minimum error. The outputs of these functions are from zero to one and monotonically increasing. They are sometimes called a squashing function because they map a large input domain to small range of outputs. Some continuous smoothing fuctions are the hyperbolic tangent function and the sigmoid function, a logistic function.

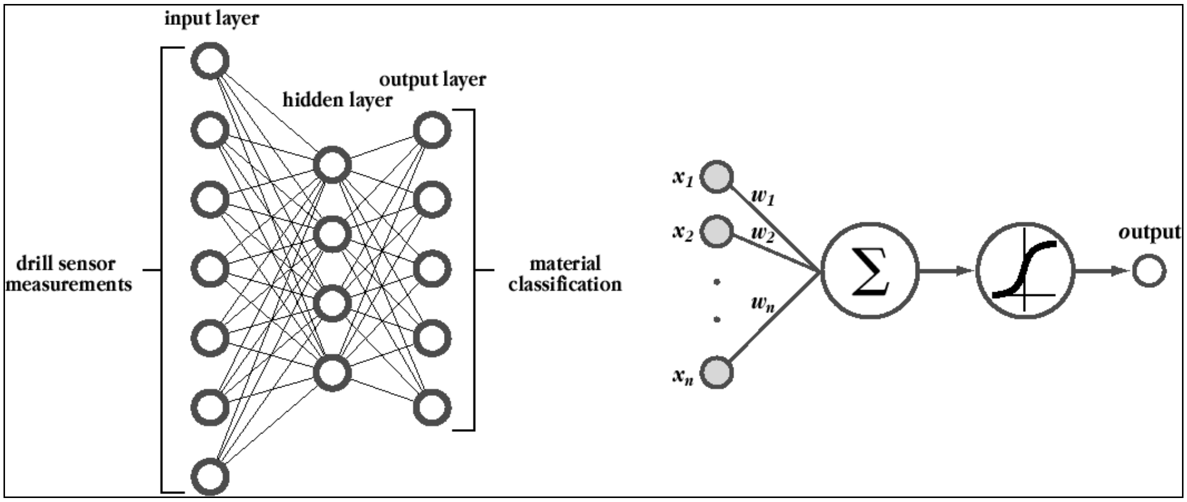

Figure 12. **Neural Network Configuration**

The the purpose of training a network is to minimize it's classification error, or output error, when presented with a vector of input values. Training is an iterative process of tuning the network connections for this purpose. For the backpropagation learning algorithm the procedure is: propagate inputs through the network, calculate the error between network output and actual output (also called the target), step through the weight space in a direction that will reduce error, propagate a function of this error backward through the network and adjust each layer of weights with weight update rule. Although the backpropagation is not guaranteed to find the globally minimum error, in practice it performs very well in spite of this [30].

There are a number algorithms one can use to search the error space—parameterized by all possible combinations of weights—for a minimum error. Gradient descent adjusts the network weights using directon of steepest descent along the error surface (for linear functions). Conjugate gradient is a sequence of line searches: the first step is in a direction negative of the gradient and subsequent updates choose a different direction to make sure the previous minimization is not "undone". There are also several functions one can use to calculate error in the network hidden and output units. For a non-linear system, such as this drilling system, with binary outputs, it has been shown that the outputs are best represented as probabilities [30]. The cross-entropy models the network outputs as the probability, or maximum likelihood that each of the binary outputs is the true class. The benefit of this is that more information can be gleaned from this output representation. For example, two binary digits having close probabilities may represent a change from one class to another (i.e., from one material to another). The neural network I am using minimizes cross-entropy. Cross entropy is obtained by using the softmax error function:

$$-\sum t_d \log o_d + (1-t_d) \log (1-o_d), \ \forall d \in D$$

Where t_d is the target value for training example *d*, and o_d is the network's probability estimate for training example *d*.

A training cycle continues until either a minimum in the error function is attained or a user-specified number of training iterations is reached. The larger the number of training iterations, the more discontinuous, or non-linear the decision surfaces can become. Too many iterations may lead to overfitting, where the network weights adjust specifically to the training data and reduce the network’s ability to generalize the decision surfaces for data representative of the application. A separate validation (or test) data set can be monitored during training to make sure that there is no overfitting (characterized by a decrease in the training set error with an increase in the test set error). Small data sets that may not include all of the representative training examples, usually use a k-fold cross-validation approach, where there are k training runs, and each run has separate training data and validation data.

10 References

[1] B. Allotta, F. Belmonte, L. Bosio and P. Dario, “Study On A Mechatronic Tool For Drilling in the Osteosynthesis of Long Bones: Tool/Bone Interaction, Modeling and Experiments,” *Mechatronics*, Vol. 6, No. 4, pp 447-459, 1996.

[2] Fred Aminzadeh, Simon Katz, and Keiti Aki, “Adaptive Neural Nets for Generation of Artificial Earthquake Precursors,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 32, No. 6, November, 1994.

[3] Andrea Baraldi and Flavio Parmiggiani, “A Neural Network for Unsupervised Categorization of Multivalued Input Patterns: An Application to Satellite Image Clustering,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 33, No. 2, March, 1992.

[4] D. Benaouda, G. Wadge, R.B. Whitmarsh, R.G. Rothwell, and C. MacLeod, “Inferring the Lithology of Borehole Rocks by Applying Neural Network Classifiers to Downhole Logs: An Example From the Ocean Drilling Program,” *Geophysics J. Int.*, 136, pp 477-491, 1999.

[5] Christopher M. Bishop, “Neural Networks For Pattern Recognition,” *Oxford University Press*, 1995.

[6] Christopher Bishop and Ian Nabney, *Netlab Software Web Page*, http://www.ncrg.aston.ac.uk/netlab/, 1999.

[7] P. N. Brett, C. A. Fraser, M. Hennigan, M. V. Griffiths and Y. Kamel, "Automatic Surgical Tools for Penetrating Flexible Tissues," *IEEE Engineering in Medicine and Biology*, p. 264, May/June, 1995.

[8] Ratna Babu Chinnam, "Neural Networks For Performance Reliability Prediction," *Department of Industrial Engineering and Management*, North Dakota State University, Fargo, ND 58105

[9] Valentina Colla and Benedetto Allotta, "Wavelet-Based Control of Penetration in A Mechatronic Drill for Orthopaedic Surgery," Proceedings IEEE International Conference on Robotics and Automation, May, 1998.

[10] R.A.Cunningham, "An Empirical Approach for Relating Drilling Parameters," *Journal of Petroleum Technology*, p. 987, July, 1978.

[11] Jaak J. K. Daemen and Richard A. Schults, eds., *Proceedings 35th U. S. Symposium on Rock Mechanics*, 5-7 June, 1995

[12] Jayantha Ediriwickrema and Siamak Khorram, "Hierarchical Maximum-Likelihood Classification for Improved Accuracie," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 35, No. 4, July, 1997.

[13] Donald S. Frankel, Robert J. Olson, Shiela L. Frankel and Sallie W. Chisholm, "Use of a Neural Net Computer System for Analysis of Flow Cytometric Data of Phytoplankton Populations," *Cytometry* 10, pp. 540-550, 1989.

[14] D. F. Galloway, "Some Experiments on the Influence of Various Factors on Drill Performance," *Transactions of the ASME*, February, 1957.

[15] J. Garnier and N. H. van Lingen, "Phenomena Affecting Drilling Rates at Depth," *Petroleum Transactions, AIME*, Vol. 216, p. 232, 1959.

[16] Paul F. Gnirk, J. B. Cheatham, Jr., "A Theoretical Description of Rotary Drilling for Idealized Down-Hole Bit/Rock Conditions," *Society of Petroleum Engineers Journal*, p 443, December, 1969.

[17] Igor Grabec, "Chaotic Dynamics of the Cutting Process," *International Journal of Machine Tools and Manufacturing*, Vol. 28, No. 1, pp 19-32, 1988.

[18] S. Gstalder and J. Raynal, “Measurement of Some Mechanical Properties of Rocks And Their Relationship to Rock Drillability,” *Journal of Petroleum Technology*, p 991, August, 1966.

[19] Shane Y. Hong, J. Ni and S. M. Wu, “Analysis of Drill Failure Modes by Multi-sensors on a Robotic End Effector,’ *Transactions of the ASME,* Vol. 118, p 672, November, 1996.

[20] Pau-Lo Hsu and Wei-Ru Fann, “Fuzzy Adaptive Control of Machining Processes With A Self-Learning Algorithm,” *Transactions of the ASME, Journal of Manufacturing Science and Engineering*, Vol. 118, November, 1996.

[21] H. M. Hughes, “Some Aspects of Rock Machining,” *International Journal of Rock Mechanics and Mining Science*, Vol. 9, pp. 205-211, 1972.

[22] Gary W. Johnson, “LabVIEW Practical Programming, Practical Applications in Instrumentation and Control,” *McGraw-Hill*, 1994.

[23] V. Kaburlasos, V. Petridis, P. Brett and D. Baker, “Learning a Linear Association of Drilling Profiles in Stapedotomy Surgery,” *Proceedings IEEE International Conference on Robotics and Automation*, May, 1998.

[24] Roger L. King, Michael A. Hicks and Stephen P. Signer, “Using Unsupervised Learning for Feature Detection in a Coal Mine Roof,” *Engineering Applications of Artificial Intelligence*, Vol. 6, No. 6, pp. 565-573, 1993.

[25] R. King and S. Signer, “Using Artificial Neural Networks for Feature Detection in Coal Mine Roofs,” *Proceedings 8th International Conference on Comuter Methods and Advances in Geomechanics*, Vol. 3, May, 1994.

[26] 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.

[27] T. I. Liu and K. S. Anantharaman, “Intelligent Classification and Measurement of Drill Wear,” *Transactions of the ASME, Journal of Engineering for Industry*, Vol. 116, p. 392, August, 1994.

[28] W. C. Maurer, “The “Perfect-Cleaning” Theory of Rotary Drilling,” *Journal of Petroleum Technology*, p. 1270, November, 1962.

[29] William C. Maurer, “Shear Failure of Rock Under Compression,” *In Proceedings of Society of Petroleum Engineers Conference on Drilling and Rock Mechanics*, University of Texas, June, 1965.

[30] T. Mitchell, “Machine Learning,” *McGraw-Hill*, 1997.

[31] Noel N. Moebs, ‘Mine Design: Geologic Studies in Ground Control.’ U. S. Department of the Interior, *Proceedings: Bureau of Mines Technology Transfer Seminars*, Pittsburgh, PA December 6-7, 1983, Denver, CO, December 8-9, 1983.

[32] G. L. Mowrey, “Applying Adaptive Signal Discrimination To Vibrational Coal Interface Detection.” *SME Annual Meeting*, Salt Lake City, UT, February 26-March 1, 1990.

[33] NIOSH, *Mining Safety and Health Research Web Pages*, http://www.cdc.gov/niosh/mining/, 1998.

[34] Y. Nishimatsu, “The Mechanics of Rock Cutting.” *International Journal of Rock Mechanics and Mining Science*, Vol. 9, pp. 261-270, 1972.

[35] H. D. Outmans, “The Effect of Some Drilling Variables On The Instantaneous Rate of Penetration.” *AIME Petroleum Transactions*, Vol. 219, p. 137, 1960.

[36] James Paone and W. E. Bruce, “Drillability Studies: Diamond Drilling.” U.S. Department of the Interior, *Bureau of Mines Report of Investigations 6324*, 1963.

[37] B. Paul and D. L. Sikarskie, “A Preliminary Theory of Static Penetration By a Rigid Wedge Into a Brittle Material.” *Transactions, Society of Mining Engineers*, p. 372, December, 1965.

[38] Claudio Perez, Aldo Casali, Guillermo Gonzalez, Gianna Vallebuona and Ricardo Vargas, “Lithologocal Composition Sensor Based on Digital Image Feature Extraction, Genetic Selection of Features, and Neural Classification.” *Department of Electrical Engineering and Department of Mining Engineering*, University of Chile, 1999.

[39] Paul Pfister, “Recording Drilling Parameters in Ground Engineering,” Geodrilling, October, 1985.

[40] Jon M. Preston, William T. Collins, David C. Mosher, Roland H. Poeckert and Ron H. Kuwahara, “The Strength of Correlations Between Geotechnical Variables and Acoustic Classifications,” *Proceedings of Oceans ’99*, p 1123.

[41] K. Ramamurthi and C. L. Hough, Jr., “A Hybrid Approach For Robust Diagnostics of Cutting Tools.” *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 24, No. 3, March, 1994.

[42] C. Ramirez and R. Thornhill, “Drill Wear Monitoring in Circuit Board Manufacturing Using Drilling Forces and Their Spectra,” *ASME Journal of Eletronic Packaging*,” Vol. 114, September, 1992.

[43] S. Rangwala and D. Dornfeld, “Sensor Integration Using Neural Networks for Intelligent Tool Condition Monitoring,” Journal of Engineering for Industry, Vol. 112, p 219, August, 1990.

[44] Stuart Russell and Peter Norvig, “AI, A Modern Approach,” Prentice-Hall, 1994.

[45] M. Scoble, J. Peck and D. Kennedy, “Drill Monitoring Investigations in a Western Canadian Surface Coal Mine,” *Proceedings 2nd Large Open Pit Mining Conference*, April, 1989.

[46] M.C. Shaw, P.A. Smith and N.H. Cook, ”The Rotary Cutting Tool,” Proceedings of ASME Annual Meeting, Atlantic City, New Jersey, November 25-30, 1951

[47] P. Sklavounos and M. Sakellariou, “Intelligent Classification of Rock Masses,” *Artificial Intelligence in Engineering*, p. 387.

[48] Wilbur H. Somerton, “A Laboratory Study of rock Breakage by Rotary Drilling,” *Petroleum Transactions, AIME*, Vol. 216, p. 92, 1959.

[49] Changyi Sun, Christopher M. U. Neale, Jeffrey J. McDonnell, and Heng-Da Cheng, “Monitoring Land-Surface Snow Conditions from SSM/I Data Using an Artificial Neural Network Classifier,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 35, No. 4, July, 1997.

[50] Russell H. Taylor, Leo Joskowicz, Bill Williamson, Andre Gueziec, Alan Kalvin, Peter Kazanzides, Robert Van Vorhis, Jianhua Yao, Rajesh Kumar, Andrew Bzostek, Alind Sahay, Martin Borner and Armin Lahmer, “Computer-Integrated Revision Total Hip Replacement Surgery: Concept and Preliminary Results,” Medical Image Analysis, Vol 3., No. 3, pp 301-319, 1999.

[51] R. Teale, “The Concept of Specific Energy in Rock Drilling,” International Journal of Rock Mechanics and Mining Science, Vol. 2, pp57-73, 1965.

[52] A. Thangaraj and P. K. Wright, "Computer-assisted Prediction of Drill-failure Using In-process Measurements of Thrust Force," *Transactions of the ASME, Journal of Engineering for Industry*, Vol. 110, p. 192, May, 1988.

[53] Richard E. Thill, "Coal and Rock Properties for Premine Planning and Mine Design," *Proceedings Bureau of Mines Technology Transfer Seminars*, U. S. Department of the Interior, Pittsburgh, PA December 6-7, 1983, Denver, CO, December 8-9, 1983.

[54] C. E. Tsoutrelis, "Determination of the Compressive Strength of Rock In Situ or in Test Blocks Using a Diamond Drill,' *International Journal of Rock Mechanics and Mining Science*, Vol. 6, pp. 311-321, 1969.

[55] Walter Utt, "Neural Network Technology for Strata Strength Characterization," Spokane Research Center, National Institute for Occupational Safety and Health, Spokane, WA, 1999.

[56] G. Wijk, "Rotary Drilling Prediction," *International Journal of Rock Mechanics, Mining Science and Geomechanical Abstracts*, Vol. 28, No. 1, pp. 35-42, 1991.

[57] Patrick M. Wong, Tamas D. Gedeon, and Ian J. Taggart, "An Improved Technique in Porosity Prediction: A Neural Network Approach," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 33, No. 4, July, 1995

[58] X. L. Zhao and J.-C. Roegiers, "A Fracture Mechanics Approach to Rock Drilling," *Rock Mechanics*, Daeman & Schultz (eds), 1995 Balkema, Rotterdam.