

Early Warning System for Seismic Events in Coal Mines Using Machine Learning

Robert Bogucki, Jan Lasek, Jan Kanty Milczek, Michał Tadeusiak
 deepsense.io,

Email: {robert.bogucki, jan.lasek, jan.milczek, michal.tadeusiak}@deepsense.io

Abstract—This document describes an approach to the problem of predicting dangerous seismic events in active coal mines up to 8 hours in advance. It was developed as a part of the AAIA'16 Data Mining Challenge: Predicting Dangerous Seismic Events in Active Coal Mines. The solutions presented consist of ensembles of various predictive models trained on different sets of features. The best one achieved a winning score of 0.939 AUC.

I. INTRODUCTION

IN 2015, the mining industry in Poland reported 2158 dangerous incidents with 19 casualties and 12 severe injuries [1]. Underground mining work poses a number of threats including fires, methane outbreaks or seismic tremors and bumps. Monitoring and decision support systems might play an essential role in limiting the number of incidents and their prevention. Such systems, often based on machine learning or data mining techniques, can be effectively applied to lessen the danger to employees and prevent potential losses arising from lost and damaged equipment, see, e.g., [2], [3], [4].

In this paper, we present a model for predicting dangerous seismic events in coal mines. Using different machine learning models, we address the classification problem of whether the total seismic energy in the upcoming few hours is going to reach a warning level. The model was developed for the AAIA'16 Data Mining Challenge: Predicting Dangerous Seismic Events in Active Coal Mines and proved to be the most successful approach among the 203 teams participating in the challenge [3].

The paper is structured as follows: the first section outlines the problem and describes the main challenges. Next, we describe our approach, focusing on feature engineering, model optimization and evaluation. Finally, in the last section we conclude the work.

II. THE CHALLENGE

In this section we introduce the problem and describe the provided data. We also make some preliminary remarks about the data and its nature.

A. Problem statement

The given problem is a classification task. The goal is to develop a prediction model that, based on the recordings from a 24-hour long period, predicts whether an energy warning level is going to be reached in the upcoming 8 hours. The warning is reached when the total energy of seismic bumps exceeds 50 000 J = 50 kJ. The accuracy of a model is

determined with respect to the Area Under ROC curve (AUC) metric. This accuracy measure is defined as follows. Let $(\mathbf{x}_i, y_i) \in \mathbf{X}$ denote an instance from the dataset \mathbf{X} , i.e., \mathbf{x}_i stands for the feature vector associated with a single measurement and $y_i \in \{0, 1\}$ stands for its label. Let f be a model that maps each instance to probability that it belongs to class '1' (or, more generally, a real-valued risk score). Then AUC is derived as

$$AUC(f, \mathbf{X}) = \frac{\sum_{i:y_i=0} \sum_{j:y_j=1} \mathbb{1}(f(\mathbf{x}_i) < f(\mathbf{x}_j))}{|\{y_i : y_i = 0\}| \cdot |\{y_j : y_j = 1\}|} \quad (1)$$

where $\mathbb{1}(\cdot)$ denotes an indicator function that returns 1 if a given condition is satisfied or 0 otherwise, and $|S|$ denotes the cardinality of set S . This accuracy measure returns values in the range range $[0, 1]$, where 1 is achieved by a perfect predictor. A random predictor yields values around 0.5.

B. Data

Two sets of observations were provided: training dataset with accompanying labels and the test set without them. The former was provided so that the problem could be approached from a machine learning angle, the latter serves for evaluation purposes. The competitors were asked to submit the likelihood of the label 'warning' for each record in the test set.

In total, the training set consists of 133 151 observations. Each observation (instance) is described by a set of 541 numbers. Below, we briefly introduce the data provided. For a more thorough description of the dataset please refer to the competition website [5].

The instances are described by a set of 13 features of different type and 22 time series over last 24 hours prior to the forecasting period. The time series' names are followed by 1, 2, ..., 24, indicating consecutive hours of measurements (with the most recent hour prior to forecasting period being 24). Possible suffixes `_eξ`, $\xi \in \{2, 3, 4, 5, 6plus\}$ refer to orders of magnitude of a given time series in a certain range, e.g. `sum_e3.5` stands for sum of energies within range $(10^2, 10^3]$ in the 5th hour of the time series. The series are listed below:

- `count_e2, ..., count_e6plus` - number of registered seismic bumps;
- `sum_e2, ..., sum_e6plus` - sum of energy of registered seismic bumps;
- `total_number_of_bumps`;
- `number_of_rock_bursts`;

- `number_of_destressing_blasts`;
- `highest_bump_energy`.

Additionally, for the most active geophones, the following series are provided:

- `max_gactivity`;
- `max_genergy`;
- `avg_gactivity`;
- `avg_genergy`;
- `max_difference_in_gactivity`;
- `max_difference_in_genergy`;
- `avg_difference_in_gactivity`;
- `avg_difference_in_genergy`.

There are also 4 *assessments* provided by experts. They are provided as categorical variables with four levels ranging from 'a' (the lowest risk) to 'd' (the highest risk):

- `latest_seismic_assessment`;
- `latest_seismoacoustic_assessment`;
- `latest_comprehensive_assessment`;
- `latest_hazards_assessment`.

Finally, several features which we will refer to as *general* are provided:

- `total_bumps_energy`;
- `total_tremors_energy`;
- `total_destressing_blasts_energy`;
- `total_seismic_energy`;
- `latest_progress_estimation_l`;
- `latest_progress_estimation_r`;
- `latest_maximum_yield`;
- `latest_maximum_meter`.

Metadata: For each observation we are given its location, i.e., a *longwall* in a particular coal mine that the measurement comes from. Each location is accompanied with additional information (metadata included in a separate file):

- `main_working_id` - ID of the main working site (at a longwall);
- `main_working_name` - name of the main working site;
- `region_name` - name of region where the main working site is located;
- `bed_name` - name of coal bed;
- `main_working_type` - type of the main working site;
- `main_working_height` - height of the main working site;
- `geological_assessment` - geological assessment of the main working site made by experts before the beginning of exploration (ordered categorical variable ranging from 'a' (lower risk) to 'c' (higher risk)).

Most of metadata were unique to the working sites, therefore were discarded early due to the reasons discussed later. The only information of potential use were `main_working_height` and `geological_assessment`, however they still had to be treated with caution:

- `geological_assessment`: A closer insight revealed that there is none mine assessed as 'd' and only one

marked as 'c'. It was replaced by 'b'. Moreover, the proportion of 'a' assessments for longwalls in the training and test dataset varied significantly, 25% to 48%, respectively.

- `main_working_height`: many working sites had unique working heights - this posed a danger that the feature would be used by a model as a proxy for particular location rather than a potentially valuable information about the height. One solution, discussed later, could be to add extra noise, to diminish the relations between the mines and their heights.

The test set consists of 3 860 unlabeled observations. Approximately 25% of them were used for evaluation on the preliminary leaderboard, which was updated throughout the contest when participants submitted their solutions. The remaining observations were used for selection of the best solutions at the end of the competition.

We should also note that the observations in the test set were randomly selected events rather than time series as provided within the training set. More precisely, given a series consecutive observations, samples were uniformly drawn from them to form a test set. If two samples collected laid within the same window of 32 hours (for 24-hour long time series describing seismic activity plus 8 hour window for prediction), one of them was dropped so as to assure that the samples were approximately independent. This procedure removes a significant amount of observations hence the size of the competition test set was relatively modest in comparison to the amount of training data available. This resulted in a very unreliable leaderboard evaluation during the competition that was based on ca. 1 000 observations. Therefore, we put great emphasis and efforts to develop reliable evaluation methods given the available training data as discussed in the next section.

C. Initial remarks

When we approached the problem we quickly realized that the main challenge was to develop a prediction model that generalizes well to new locations. Table I presents the warning frequencies per location in the dataset. We observe that first of all, different locations vary considerably in terms of the frequency of warnings. Secondly, the sets of locations differ between the training and test dataset. Additionally, the test set in the competition originated from future recordings with respect to the training data available. This is the root of the problem. Hence a proposed model should be both location and time independent in the sense that it yields unbiased predictions for instances with no regard to their origin and time they are collected. We also see that the number of instances originating from different locations varies considerably. These preliminary observations should be carefully considered during model building and evaluation steps. We elaborate on this in the next section.

III. THE SOLUTION

In this section we describe in detail our solution to the given problem. We discuss different sets of features that were proposed, various evaluation methods, models and their set up.

A. Feature engineering

In our experiments we created several feature sets for model training. For the sake of simplicity and completeness, we describe them under consecutive headers and denote as FS_n which stands for the n -th *feature set* we proposed. These feature sets were developed independently by members of our team. Inevitably, there are significant overlaps between them. FS_1 : The processing of the data focused mainly on aggregation, aiming to reduce the number of the hourly measurements as the majority of them were just zeros (for the training set, about 66% of all numbers were 0). The feature extraction step ended up with 133 features, over 4 times less than the original set. From the original features we kept:

- all *general* features;
- all seismic assessments converted to consecutive integers and their average;
- number of bumps (`count_e*`) and their energies (`sum_e*`) summed over all 24 hours, together with mean energies resulting from division (if `count_e*` was 0, then we were substituting the result by 0);
- number of bursts and the highest bump energies were just summed.

We also aggregated the remaining time series related to most active geophones (8 time series), however this time we introduced some aggregations over subsets of hourly measures based on their relative importance. The process is described below.

In order to assess the impact of features we used a functionality provided by the implementation of Gradient Boosting Trees available in the **XGBoost** [6] package. The library allows building a tree classifier and assessing the importance of particular features by providing the number of times the feature was used in a split. The more often a feature is used, the more separation gain it offers and therefore the more important it is. We used an XGBoost classifier with 150 trees (other parameters were default). Fig. 1 presents an example of such feature importance analysis for `avg_genergy`. It seems that features are gaining importance towards the end of the time-series - it agrees with the intuition that the measurements closer to the forecasting period are more informative. Therefore in this case, apart from the entire time-series statistics, we are also interested in statistics based on the last five hours (they stand out from the preceding hours). Also, we keep the measurements from the very last hour as a separate feature. Having applied analogous analysis to the above feature groups, we selectively compute statistics such as:

- average and average over absolute values;
- standard deviation;
- max and max over absolute values;

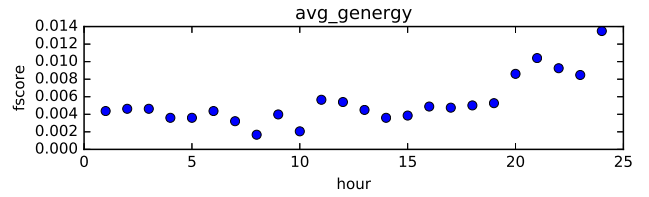


Fig. 1. Importance of hourly measurements of `avg_genergy`.

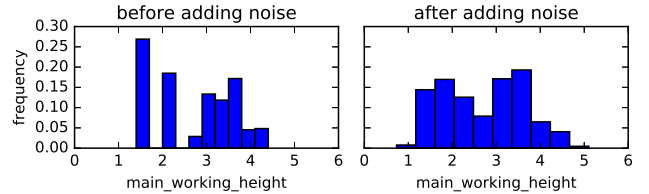


Fig. 2. Distribution of `main_working_height`, before and after adding noise.

- average over last γ hours (where γ varies from 1 to 6);
- standard deviation over last γ hours;
- slope of a linear regression over last 5 hours with respect to time.

As mentioned above, competitors were also provided with the metadata describing specific mine sites. Most of them were discarded. The only metadata used here were the *main working height*, but only after adding Gaussian noise ($\sigma = 0.2$) resulting in more even distribution, see Fig. 2. This step was performed to prevent a model from recognizing a particular location by its height.

In addition to the above features, we produced a vast set of more than 6 000 interactions between them, i.e., pairwise products of features. This is obviously an exhaustive number and we were not planning to use all of them. However, some interactions proved to be valuable. We applied an iterative process of selecting the most promising subsets of features and their interactions. We will come back to them when describing the final model (Section III-C Model₁).

FS_2 : In constructing this set, at first we decided to drop time series describing maximum statistics (`max*`) since they were highly correlated with corresponding average records. For the series `count_e*`, `sum_e*`, `number_of_rock_bursts`, `number_of_destressing_blasts`, `avg_gactivity`, `avg_genergy` we extracted the following features:

- minimum,
- maximum,
- standard deviation,
- indicator variable, if there is a non-zero value in the series,
- hours elapsed from the last non-zero observation.

Moreover, these statistics were derived over the window of the last 2, 4, 8 and 24 hours prior to the forecasting period in order to describe the most recent data in greater

detail. These features were appended to data with the original time series that they were computed for. Additionally, for series `avg_difference_in_gactivity` and `avg_difference_in_genergy` maximal absolute value was derived over the last 2 and 24 hours. Finally, the categorical variables were converted to binary features using one-hot encoding, i.e., for each possible value of a categorical a separate column was created which indicates that a given observation has this particular category. These operations produced a feature set with a total number of 700 features.

FS₃: In this set of features, we first derived several new series based on the original ones

- 1) `log_max_avg_diff_genergy` which was derived as difference in `max_genergy` and `avg_genergy`, with application of logarithm hereafter. Analogous operation was performed for other series `avg` and `max_difference_in_genergy` and a corresponding series for `gactivity`
- 2) `log_ave_energy` series was produced by computing average energy based on `sum` and `count` series.

In addition to features enumerated for **FS₂**, we derived the following statistics:

- 0.25, 0.5 and 0.75-quantiles,
- number of times a series increased in comparison to the previous hour's recording,
- number of positive values in a series,
- indicator variable, if there is a non-zero value in the series.

The statistics were computed on 4, 8 and 24 hours window. Furthermore, we computed the coefficient, intercept and R^2 statistic for a fit of linear model of series `avg_difference_in_gactivity`, `avg_gactivity`, `log_ave_energy` to an independent hourly temporal variable (1, 2, ..., 24). Finally, we computed correlations between `avg_difference_in_gactivity` and `avg_difference_in_genergy` as well as `avg_gactivity` with `avg_genergy`. After extracting features, constant features were dropped from the feature set. Also, if there were features that were correlated over 0.99 (according to Pearson correlation coefficient), one of them was removed. For categorical variables, they were one-hot encoded for logistic regression model or converted to integers (with higher risk categories being assigned a higher integer) for tree-based models.

These steps produced a training set of 426 features (for the integer encoding of categorical features).

FS₄: The feature set that was created with the goal of being simple and as such leaving little room for overfitting. Out of the basic (not time-based) features, `main_working_id` was dropped. Out of the metadata, only `geological_assessment` was used. The time-based features were ran through maximum and standard deviation functions on 8-hour time periods with 4 hour increments, only the features concerning quantities and maximums were used (features listing averages and sums were left out).

B. Evaluation procedures

Evaluation methodology is a crucial part of creating a successful application of a model. Below we list different validation techniques that we employed to assess the accuracy of a model. Here, the issue of overfitting a model to particular locations and time-frames of samples is considered in detail.

k-fold cross validation (k-CV): This is one of the basic validation procedures. It is performed by assigning each example in the training set randomly to one of k folds (in our application we used $k = 10$ or 20). Note that due to temporal alignment of instances in the training data, this evaluation procedure tends to produce overly optimistic evaluation scores (we observed that during the contest by, e.g., large discrepancy in local evaluation and leaderboard scores). This is because consecutive instances are likely to share the same label. If some of them pertain to a training fold and the others to test fold, then a classifier has a relatively easy task to assign this instance to the proper label.

Leave one location out (LOLO): This evaluation method was chosen to estimate the model's performance on mining sites not included in the training data (see Table I). It was supposed to promote models that overfit less and filter out those whose good performance was actually based on data leaks. We have decided to not use the three largest locations (with IDs 264, 373 and 437) for testing. These three locations constitute to a large portion of the total training data (48%) and were not appearing in the test set. Locations that had no 'warning' labels were also not used as validation data, as AUC could not be computed for them. This approach resulted in a 8-fold cross-validation that gave much lower scores than the other ones (not even the best models could exceed 0.9 AUC), and the scores varied between folds (from as low as 0.6 to as high as 0.999) but it should not be perceived as a flaw — it was intended behavior.

Train and test split #1 (TrTs₁): This evaluation methodology was devised to reflect the way the leaderboard was constructed. It is based on multiple train and test splits of the data. It proceeds in two steps:

- 1) 5 series are chosen at random and included in the validation set,
- 2) among series that have not been selected in 1) we include the first 70% observations in the training sample and the other 30% in the validation set.

Moreover, in each of those 70%-30% splits, 32 observations between the split point were removed to assert approximate independence between the training and validation set (by introducing a gap of 32 hours between them). Again, data from locations with IDs 264, 373, 437 where included only in the training set.

In order to arrive at a reliable error estimate, this evaluation was repeated 25 times and consecutive measurements were

averaged. With that many iterations we arrived at stable results for mean AUC value.

Train and test split #2 (TrTs₂): The evaluation was based on multiple train and test splits (20 in the final model) with some restrictions. By comparing the total seismic energy (TSE) of mines (which turned out to be linearly correlated with the frequency of appearances of warnings) we tried to make the split, so the TSE in the inferred test sets resembled the level of energies in the private test set.

TABLE I

NUMBER OF INSTANCES ORIGINATING FROM DIFFERENT LOCATIONS IN THE TRAINING AND TEST SET ALONG WITH AVERAGED TOTAL SEISMIC ENERGIES (TSE) AND FREQUENCY OF WARNINGS (NOT AVAILABLE FOR TEST SET CASES).

Mine ID	Instances	Train set		Test set	
		Mean TSE [J]	Warnings Frequency	Instances	Mean TSE [J]
373	31236	81002	1.1%	-	-
264	20533	7563	0.4%	-	-
725	14777	190232	9.4%	330	106741
777	13437	0	0.0%	330	29061
437	11682	4727	0.4%	-	-
541	6429	9397	0.9%	5	324
146	5591	678	0.1%	98	1
575	4891	9775	0.5%	253	7503
765	4578	136	0.0%	329	51265
149	4248	48357	7.3%	98	72749
155	3839	322021	17.2%	98	527229
583	3552	2595	0.2%	215	73302
479	2488	5548	0.0%	35	102
793	2346	0	0.0%	330	11547
607	2328	6027	0.0%	209	9470
599	1196	29932	1.9%	363	39962
171	-	-	-	49	33
470	-	-	-	258	10701
490	-	-	-	160	13698
508	-	-	-	58	32183
641	-	-	-	97	10672
689	-	-	-	83	63889
703	-	-	-	145	44031
799	-	-	-	317	8

Table I presents averaged TSE for each mine grouped over train and test datasets, together with the frequencies of warnings in the training dataset. It is worth to point out significant discrepancies between the activity levels of mines in both sets. For mine 765, the activity in the training set is mere 136 J, with no warnings. In the test set, the average activity is above 50 kJ, so there must have been several warnings emitted. A closer look reveals that there are some abnormalities in the training set. Fig. 3 presents the TSE of mines 155 and 765. While the activity of the former looks realistic, 765 is mute for majority of the time, only to exhibit a few spikes towards the end of the time series. On the other hand, its activity in the test set greatly increased. Some mines do not exhibit any activity in the training set, i.e. TSE equals zero (mines 777, 793). This is one of the reasons we have to avoid producing models that would be able to recognize the mines, the classifiers should generalize correctly from the activity records, regardless any behavior specific to certain mines. Also, it poses a problem - whether to consider the suspicious

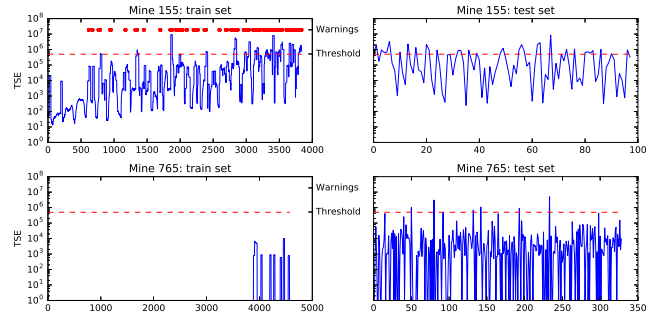


Fig. 3. Total seismic energy over time.

mines during the training or not. It is rather unusual for a mine to have a zero seismic activity and supposedly the data in these periods might be corrupted.

The final train and test splits were based on the above TSE analysis and were produced in the following way:

- 1) mines 777, 793 were excluded due to their suspicious lack of any activity in the training set (although they represented a significant amount of data);
- 2) in every split, five randomly selected mines were left only for testing (to evaluate the generalization properties of classifiers);
- 3) from the remaining sites we were taking 20% of samples for testing. For mines 146 and 599 samples were drawn from the beginning (due to corresponding energy levels in the test set), for the remaining - from the end;
- 4) in some cases (mines 373, 437) only samples where TSE were nonzero were taken into account.

The process was repeated several times to obtain multiple train/test splits. The final evaluation was based on the average score over 20 splits.

C. Model training and optimization

There are several models that we employed in creating the final solution to the problem. We used the implementation of models available in Python’s **scikit-learn** package for machine learning (ver. 0.17.1) [7], [8] and **XGBoost** package for tree boosting models (ver. 0.4) [6]. Throughout the paper, for brevity, we use the following abbreviations for the model names: Linear Discriminant Analysis - LDA, Logistic Regression - LR, Extra Trees Classifier - ETC (all from **scikit-learn** library) and Extreme Gradient Boosting Classifier - XGB (from **XGBoost** library).

Model₁: The first model was built using **FS₁** and TrTs₂ evaluation method. Several models were considered, apart from XGB and ETC, also logistic regressions and neural networks, finally only the first two were used in the final blend. They were performing particularly well in spite of rather large number of features.

First, we ran learning on all the features and interactions we produced. Based on the importance scores provided by XGB (described in Section III-A **FS₁**) we kept the first 982

interactions and all *individual* features. Then, a grid search returned sets of parameters scoring the highest:

The optimal parameters for XGB model were (otherwise default):

- `n_estimators = 100`
- `max_depth = 2`
- `learning_rate = 0.08`

The optimal parameters for ETC were (otherwise default):

- `n_estimators = 1000`
- `max_depth = 7`
- `criterion = entropy`

It is worth to note, that trees, by their design, are relatively powerful in discovering interactions between features. However, in their case the interactions are not discovered concurrently, but rather in a multilevel manner, between consecutive splits. By explicitly using interactions as features, they can be made use of directly.

Having obtained well performing hyperparameters, we ran a randomized search for best features' subsets. In each iteration we were randomly selecting from 20 to 40 individual features (out of 133) and additionally up to 10 interactions (out of 982). We ran several thousands evaluations on XGB and several hundreds on ETC, tracking their validation scores.

The idea was to produce many models built only on subsets of features and to take advantage of assembling them which reduces variance of predictions and minimize the risk of overfitting to anomalies in particular features. This is a powerful method for increasing the performance of the model [9].

The final blend was composed of:

- single ETC of 10 000 trees using all 133 features and 20 best interactions;
- single XGB using the same features;
- a blend of 20 ETCs built on 20 best subsets of features;
- a blend of 20 XGBs built on 20 best subsets of features;

The final submission scored 0.9199 on the public leaderboard. The score in the final evaluation reached 0.9393 and turned out to be the best in the competition.

Model₂: The second model involved the following classifiers: two linear models (LDA and LR) as well as the tree-based ETC model.

The first part of the solution was the LDA model trained on **FS₂** using *k*-CV evaluation procedure. The regularization shrinkage parameter selection was done in an automated way (i.e., the parameter `shrinkage` set to "auto") in **scikit-learn**'s LDA implementation. The other models were LR and ETC trained on **FS₃** using TrTs_1 evaluation method. The parameter values were set using grid search. The optimal values for LR model were:

- `penalty = l1`
- `C = 0.003`

The optimal parameters for ETC were:

- `n_estimators = 1000` (number of trees)
- `max_depth = 3`

- `max_features = 200`
- `min_samples_split = 3`
- `class_weight = 10` (for label '1').

The three models were blended by averaging their predictions with equal weights to produce a solution. Prior to averaging, the model predictions were standardised so that their standard deviations would equal 1. This step aims to convert the probabilities yielded by individual models to the same scale. Note that the mean values of predictions are irrelevant since AUC is invariant to monotonic transformations of output, see Equation 1. On the competition test set, the model yielded 0.9385 and 0.9340 of preliminary and final evaluation score, respectively.

Model₃: This model used only **FS₄** and was meant to be more universal than the other models and thus was tuned on LOLO validation. The algorithms used were ETC, XGB and logistic regression. For each algorithm, many sets of predictions were generated (using the top results from a grid search). This model achieved 0.928 and 0.933 on preliminary and final evaluations, respectively.

Below we list the best parameters found for each algorithm:

ETC

- `min_samples_leaf = 5`
- `n_estimators = 40 000`

XGB

- `subsample = 1.0`
- `num_round = 200`
- `max_depth = 10`
- `objective = binary:logistic`
- `base_score = 0.05`
- `eta = 0.04`
- `colsample_bytree = 0.8`

LTR

- `solver = sag` (Stochastic Average Gradient, just for speed)
- `C = 1.0`

D. Model ensemble

We have decided to use sorted order position averaging (as the AUC assessment method considers only the rank of predicted likelihoods and not the values) of the three presented models' predictions with the final weights being 1, 3 and 2 for models 1, 2 and 3 respectively. The averaging was employed in order to leverage various approaches and come up with yet a better predictor for the given task. The weights for the ensemble were chosen basing on individual model's scores on the preliminary leaderboard. The ensemble produced a model scoring 0.933 and 0.938 on preliminary and final leaderboard, respectively. All in all, it turned out that model 1 outperformed the full ensemble by a small margin (0.939 to 0.938). However, it might be caused by the relatively small test set size.

E. Things we tried that did not work

Throughout the process of creating the most successful model we tested a couple of ideas that turned out not to

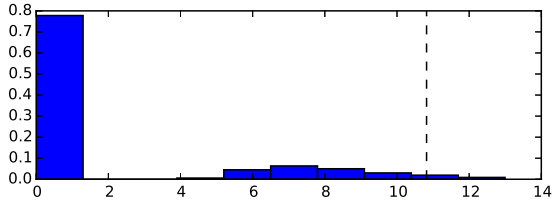


Fig. 4. Histogram of the total energy within 8 next hours (after with application of $\log(x + 1)$ transformation). The dashed line indicates the warning level of $\log(5 \cdot 10^4 + 1)$ J.

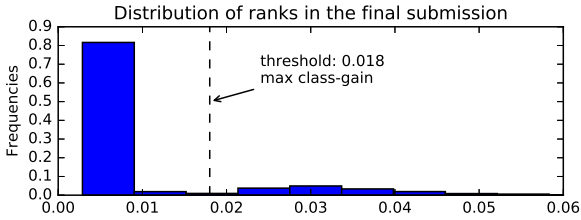


Fig. 5. Histogram of likelihoods returned by the winning model.

improve our results. First of all, we framed the problem given in the competition as a regression task. Because the observations in the training data were given as time series, it was possible to retrieve the energy level for the next hours, see Fig. 4. This allowed us to forecast energy levels within the target window of 8 hours. Note that predictions from a regression model may be directly evaluated using AUC accuracy measure as it can be considered as a *risk scoring* model for high values of energy. For the original classification problem, we tried to modify the energy levels and train the models on an enriched set of labels, e.g., we assumed 30 kJ (and a couple of other values) as the warning level and estimate the model.

We also experimented with undersampling of training instances pertaining to class 0 so that the proportion of ‘1’ in the training data increases. We also tried to reduce samples from locations 264, 373 and 437 in the training set by undersampling or assigning them a lower sample weight (in, e.g., LR model).

However, in this particular application, our efforts were not successful as the performance of the models (in terms of evaluation scores) was not improving.

F. Model performance on the final test set

After the competition we were provided with the true labels used during the final evaluation. We were able to compute different metrics than AUC. The winning model’s predictions had a strongly skewed distribution (Fig. 5), corresponding to total energies seen in Fig. 4. The distribution has two modes, however of a much lower mode related to predicted warnings - this is due to imbalance of classes. Depending on the threshold beyond which we consider predictions as warnings we can derive the confusion matrix:

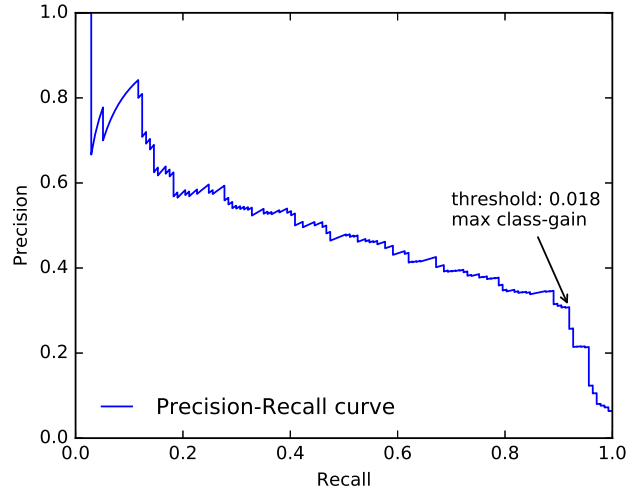


Fig. 6. Precision-recall curve.

		True warning	
		1	0
Predicted warning	1	$TP = 126$	$FP = 284$
	0	$FN = 11$	$TN = 2390$

Based on the matrix we can compute several useful accuracy measures of the model:

$$\text{precision} = TP / (TP + FP) \tag{2}$$

$$\text{sensitivity (recall)} = TP / (TP + FN) \tag{3}$$

$$\text{specificity} = TN / (TN + FP) \tag{4}$$

$$F1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \tag{5}$$

$$\text{Class-gain} = \text{specificity} + \text{recall} - 1 \tag{6}$$

The threshold maximizing the class-gain score is 0.018 (see Fig. 5) and yields the following accuracy on the final test set:

precision	recall	specificity	F1	class-gain
0.31	0.92	0.89	0.46	0.81

The entire precision-recall curve can be seen in Fig. 6.

In order to assess our results we looked for research addressing similar problems as the one considered here. In [4] we found an approach to solve the same problem, however the results are not directly comparable since they are based on different datasets. Our results prove to outperform results reported there for all presented classifiers (the highest class-gain reported is 0.75). Also, in the cited paper there was no inter-coal-mine validation, while the models described in our work were cross-validated on separate coal mines. Therefore, the models proposed are designed to generalize well and should be applicable also to working sites with no historical data available.

IV. SUMMARY

Given that the dataset originates from working mine sites, with the entire measurement infrastructure already installed,

we hope that the approach presented in this paper could be implemented and serve as a valuable tool for alerting about dangerous seismic events early. This hopefully should result in preventing possible accidents which pose a threat to employees and generate losses from damaged coal mine infrastructure and machinery.

Even though the models presented here have outperformed the other models in the competition, we recommend they be ensembled with other high-scoring models, because properly combined efforts of multiple participants are expected to yield better results than individual solutions.

Lastly, we would like to thank the organizers for the opportunity to solve a real-life problem and the contestants for creating such a competitive environment.

REFERENCES

- [1] Wyższy Urząd Górniczy, "Wypadkowość w górnictwie od 1 stycznia 2015 do 31 grudnia 2015," 2015, in Polish, last accessed 18 April 2016. [Online]. Available: http://www.wug.gov.pl/bhp/Statystyki_archiwalne_2015
- [2] A. Zagorecki, *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing: 15th International Conference, RSFDGrC 2015, Tianjin, China, November 20-23, 2015, Proceedings.* Cham: Springer International Publishing, 2015, ch. Prediction of Methane Outbreaks in Coal Mines from Multivariate Time Series Using Random Forest, pp. 494–500. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-25783-9_44
- [3] A. Janusz, M. Sikora, Ł. Wróbel, and D. Ślęzak, "Predicting Dangerous Seismic Events: AIAI16 Data Mining Challenge," in *Proceedings of FedCSIS 2016*. IEEE, 2016, in print September 2016.
- [4] M. Sikora, "Induction and pruning of classification rules for prediction of microseismic hazards in coal mines," *Expert Systems with Applications*, vol. 38, no. 6, pp. 6748–6758, 2011.
- [5] Knowledge Pit, a host platform for data challenges, 2016, last accessed 18 April 2016. [Online]. Available: <https://knowledgepit.fedcsis.org/>
- [6] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," *arXiv preprint arXiv:1603.02754*, 2016.
- [7] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [8] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and G. Varoquaux, "API design for machine learning software: experiences from the scikit-learn project," in *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, 2013, pp. 108–122.
- [9] T. K. Ho, "The random subspace method for constructing decision forests," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 20, no. 8, pp. 832–844, 1998.