

Empirical Analysis of Machine Learning Configurations for Prediction of Multiple Organ Failure in Trauma Patients

Yuqing Wang^{1*}, Yun Zhao^{1*}, Rachael Callcut², and Linda Petzold¹

¹ Department of Computer Science, University of California, Santa Barbara

² UC, Davis Health

wang603@ucsb.edu, yunzhao@cs.ucsb.edu

Abstract. Multiple organ failure (MOF) is a life-threatening condition. Due to its urgency and high mortality rate, early detection is critical for clinicians to provide appropriate treatment. In this paper, we perform quantitative analysis on early MOF prediction with comprehensive machine learning (ML) configurations, including data preprocessing (missing value treatment, label balancing, feature scaling), feature selection, classifier choice, and hyperparameter tuning. Results show that classifier choice impacts both the performance improvement and variation most among all the configurations. In general, complex classifiers including ensemble methods can provide better performance than simple classifiers. However, blindly pursuing complex classifiers is unwise as it also brings the risk of greater performance variation.

Keywords: Machine learning · Multiple organ failure · Trauma.

1 Introduction

Multiple organ failure (MOF) is a clinical syndrome with variable causes including pathogens [1], complicated pathogenesis [2], and a major cause of mortality and morbidity for trauma patients who are admitted to Intensive Care Units (ICU) [3]. Based on recent studies on ICU trauma patients, up to 47% have developed MOF, and MOF increased the overall risk of death 6 times compared to patients without MOF [4]. To prevent the development of MOF for trauma patients from progression to an irreversible stage, it is essential to diagnose MOF early and effectively. Many scoring systems have been proposed to predict MOF [5–8] and researchers have attempted to predict MOF on trauma patients using predictive models in an early phase [9, 10].

The rapid growth of data availability in clinical medicine requires doctors to handle extensive amounts of data. As medical technologies become more complicated, technological advances like machine learning (ML) are increasingly needed to improve real-time analysis and interpretation of the results [11]. In recent years, practical uses of ML in healthcare have grown tremendously, including

* These two authors contributed equally to this paper.

cancer diagnosis and prediction [12–14], tumor detection [15, 16], medical image analysis [17], and health monitoring [18, 19].

Compared to traditional medical care, ML-assisted clinical decision support enables a more standardized process for interpreting complex multi-modality data. In the long term, ML can provide an objective viewpoint for clinical practitioners to improve performance and efficiency [20]. ML is often referred to as a black box: explicit input data and output decisions, but opaque at intermediate learning process. Additionally, in medical domains, there is no universal rule for selecting the best configuration to achieve the optimal outcome. Moreover, medical data has its own challenges such as numerous missing values [21] and colinear variables [22]. Thus it is difficult to process the data and choose the proper model and corresponding parameters, even for a ML expert. Furthermore, detailed quantitative analysis of the potential impacts of different settings of ML systems on MOF has been missing.

In this paper, we experiment with comprehensive ML settings for prediction of MOF, considering 6 different dimensions from data preprocessing (missing value treatment, label balancing, feature scaling), feature selection, classifier choice, to hyperparameter tuning. To predict MOF for trauma patients at an early stage, we use only initial time measurements (hour 0) as inputs. We mainly use area under the receiver operating characteristic curve (AUC) to evaluate MOF prediction outcomes. We focus on analyzing the relationships among configuration complexity, predicted performance, and performance variation. Additionally, we quantify the relative impacts of each dimension.

The main contributions of this paper include:

- (1) To the best of our knowledge, this is the first paper to conduct a thorough empirical analysis quantifying the predictive performance with exhaustive ML configurations for MOF prediction.
- (2) We provide general guidance for ML practitioners in healthcare and medical fields through quantitative analysis of different dimensions commonly used in ML tasks.
- (3) Experimental results indicate that classifier choice contributes most to both performance improvement and variation. Complex classifiers including ensemble methods bring higher default/optimized performance, along with a higher risk of inferior performance compared to simple ones on average.

The remainder of this paper is organized as follows. Section 2 describes the dataset and features we use. All of the ML configurations are available in Section 3. Experimental results are discussed in Section 4. Finally, our conclusions are presented in Section 5.

2 Dataset

Our dataset, collected from the San Francisco General Hospital and Trauma Center, contains 2190 highest level trauma activation patients evaluated at the

level I trauma center. Due to the urgency of medical treatment, there are numerous missing values for time-dependent measurements. Thus we have chosen to consider only those features with a maximum missing value percentage of 30% over all patients. To obtain a timely prediction, early lab measurements (hour 0) as well as patients’ demographic and illness information were extracted as the set of features. Detailed feature statistics are available in Table 1.

Feature type	# of extracted features	Features
Demographic	5	<i>gender, age, weight, race, blood type</i>
Illness	2	<i>comorbidities, drug usage</i>
Injury factors	4	<i>blunt/penetrating trauma, # of rib fractures, orthopedic injury, traumatic brain injury</i>
Injury scores	8	injury severity score, 6 abbreviated injury scale (head, face, chest, abdomen, extremity, skin), Glasgow coma scale score
Vital sign measurements	4	heart rate, respiratory rate, systolic blood pressure, mean arterial pressure
Blood-related measurements	13	white blood cell count, hemoglobin, hematocrit, serum CO ₂ , prothrombin time, international normalized ratio, partial thromboplastin time, blood urine nitrogen, creatinine, blood pH, platelets, base deficit, <i>factor VII</i>

Table 1: MOF dataset statistics. Italicized features are categorical.

Our target variable consists of binary class labels (0 for no MOF and 1 for MOF). Then, the data with feature and target variables is randomly split into training and testing sets at the ratio of 7 : 3.

3 Methods

Based on ML pipelines and special characteristics of our data such as large number of missing values and varying scales in feature values, we consider comprehensive ML configurations from the following 6 dimensions: data preprocessing (missing value treatment (MV), label balancing (LB), feature scaling (SCALE)), feature selection (FS), classifier choice (CC), and hyperparameter tuning (HT). In the remainder of the paper, we will interchangeably use the full name and corresponding abbreviations shown in parentheses. Further details on each dimension are described below.

3.1 Data Preprocessing

Methods to handle the dataset with missing values, imbalanced labels, and unscaled variables are essential for the data preprocessing process. We use several different methods to deal with each of these problems.

Missing Value Treatment In our dataset, numerous time-dependent features cannot be recorded on a timely basis, and missing data is a serious issue. We consider three different ways to deal with missing values, where the first method serves as the baseline setting for MV, and the latter two methods are common techniques of missing value imputation in ML.

1. Remove all patients with any missing values for the features listed in Section 2.
2. Replace missing values with mean for numerical features and mode for categorical features over all patients.
3. Impute missing values by finding the k -nearest neighbors with the Euclidean distance metric for each feature respectively.

Label Balancing Our dataset is imbalanced as the sample class ratio between class 0 and class 1 is 11 : 1. Keeping imbalanced class labels serves as the baseline setting for LB. Three different ways are considered to resample the training set.

1. Oversampling the minority class (label 1)
 - 1.1 Method: SMOTE (synthetic minority over-sampling technique) [23].
 - 1.2 Explanation: choose k -nearest neighbors for every minority sample and then create new samples halfway between the original sample and its neighbors.
2. Undersampling the majority class (label 0)
 - 2.1 Method: NearMiss [24].
 - 2.2 Explanation: when samples of both classes are close to each other, remove the samples of the majority class to provide more space for both classes.
3. Combination of oversampling and undersampling
 - 3.1 Method: SMOTE & Tomek link [25].
 - 3.2 Tomek link: two samples are k -nearest neighbors to each other but come from different classes.
 - 3.3 Explanation: first create new samples for the minority class and then remove the majority class sample in any Tomek link.

Feature Scaling Since the range of feature values in our dataset varies widely, we perform feature scaling. No scaling on any feature serves as the baseline setting for SCALE. Two common scaling techniques are used for numerical features.

1. Normalization: rescale values to range between 0 and 1.
2. Standardization: rescale values with mean 0 and standard deviation 1.

3.2 Feature Selection

In medical datasets, there usually exist many highly correlated features, and some features that are weakly correlated to the target [22, 26]. Thus it is essential to identify the most relevant features that may help to improve the outcome of the analysis. Using all of the features described in Section 2 serves as the baseline setting for FS. We consider two main feature selection techniques: filter and wrapper methods.

1. Filter-based methods (independent of classifiers):
 - 1.1 Use correlation between features and the target to select features which are highly dependent on the target.
 - 1.2 Filter out numerical features using ANOVA F -test and categorical features using χ^2 test.
2. Wrapper-based methods (dependent on classifiers):
 - 2.1 Method: RFE (recursive feature elimination) in random forest.
 - 2.2 Explanation: perform RFE repeatedly such that features are ranked by importance, and the least important features are disregarded until a specific number of features remains.

3.3 Classifier Choice

We experimented with 15 classifiers on the dataset. In general, these classifiers can be divided into two main categories: single and ensemble. Lists of all classifiers are available in Table 2. For ensemble classifiers (combination of individual classifiers), we tried bagging (BAG, RF, ET), boosting (GB, ABC, XGB, LGBM), voting (VOTE) and stacking (STACK). In bagging, DT is a homogeneous weak learner. Multiple DTs learn the dataset independently from each other in parallel and the final outcome is obtained by averaging the results of each DT. In boosting, DT also serves as a homogeneous weak learner. However, DTs learn the dataset sequentially in an adaptive way (new learner depends on previous learners' success), and the final outcome is determined by weighted sum of previous learners. In voting, heterogeneous base estimators (LR, RF, SVM, MLP, ET) are considered, where each estimator learns the original dataset and the final prediction is determined by majority voting. In stacking, several heterogeneous base learners (RF, KNN, SVM) learn the dataset in parallel, and there exists a meta learner (LR) that combines the predictions of the weak learners. Abbreviations of classifiers shown in parentheses for voting and stacking are the ones we use.

Single classifiers	Ensemble classifiers
Logistic Regression (LR)	Bagged Trees (BAG)
Support Vector Machine (SVM)	Random Forest (RF)
Naive Bayes (NB)	Extra Trees (ET)
K-nearest Neighbors (KNN)	Gradient Boosting (GB)
Decision Tree (DT)	Adaptive Boosting (ABC)
Multi-layer Perceptron (MLP)	Extreme Gradient Boosting (XGB)
	Light Gradient Boosting Machine (LGBM)
	Voting (VOTE)
	Stacking (STACK)

Table 2: List of 6 single classifiers and 9 ensemble classifiers. Corresponding abbreviations of each classifier are shown in parentheses.

3.4 Hyperparameter Tuning

Hyperparameters are crucial for controlling the overall behavior of classifiers. Default hyperparameters of classifiers serve as the baseline setting for HT. We apply grid search to perform hyperparameter tuning for all classifiers. Detailed information about tuned hyperparameters is available in Table 3.

Classifiers	# of tuned hyperparameters	Hyperparameter lists
LR	3	C, class_weight, penalty
SVM	4	C, gamma, kernel, class_weight
KNN	3	n_neighbors, weights, algorithm
NB	1	var_smoothing
DT	5	min_samples_split, max_depth, min_samples, leaf_max_features, class_weight
MLP	3	activation, solver, alpha
BAG	2	base_estimator, n_estimators
RF	2	n_estimators, max_features
ET	2	n_estimators, max_features
GB	2	n_estimators, max_depth
ABC	3	base_estimator, n_estimators, learning_rate
XGB	2	min_child_weight, max_depth
LGBM	4	num_leaves, colsample_bytree, subsample, max_depth
VOTE	2	C (SVM), n_estimators (ET)
STACK	2	C (SVM), n_neighbors (KNN)

Table 3: Detailed configurations of tuned hyperparameters for all classifiers. All of the hyperparameter names come from *scikit-learn* [27].

4 Experiments and Results

We formulated MOF prediction as a binary classification task. All of the experiments in this paper were implemented using *scikit-learn* [27]. As mentioned in Section 2, our training and testing dataset is randomly split with a ratio of 7 : 3. One-hot encoding is applied to all categorical features. For each classifier, we use the same training and testing dataset. We use AUC as our main performance metric, as it is commonly used for MOF prediction in the literature [6, 28, 29]. It provides a “summary” of classifier performance compared to single metrics such as precision and recall. AUC represents the probability that a classifier ranks a randomly chosen positive sample (class 1) higher than a randomly chosen negative sample (class 0), and thus useful for imbalanced datasets. In this section, we quantify the impacts (improvement and variation) of each dimension on the predicted performance over our testing dataset.

4.1 Influence of Individual Dimensions

First, we evaluate how much each dimension contributes to the AUC score improvement and variation respectively, and find the correlation between performance improvement and variation over all dimensions.

Performance Improvement across Dimensions For HT, MV, LB, SCALE, and FS, we define the *baseline* as default hyperparameter choices, using no missing value imputation, no label balancing, no feature scaling, and no feature selection, respectively. For CC, we choose SVM, which achieves the median score among all classifiers, as the *baseline*. Then we quantify the performance improvement of each dimension. Fig. 1 shows the percentage that each dimension contributes to the improvement in the AUC score over baseline by tuning only one dimension at a time while leaving others at baseline settings. We observe that CC contributes most to the performance improvement (15.00%) for MOF prediction. After CC, LB (10.81%), FS (10.09%), MV (7.90%), HT (6.94%), and FS (2.45%) bring decreasing degrees of performance improvement in the AUC score.

Table 4 shows the improvement of every single dimension on each classifier over the baseline. In general, MV and LB tend to provide the greatest performance improvement for most classifiers. For RF, ET, and LGBM, FS contributes the most to improvement in performance since these classifiers require feature importance ranking intrinsically, and external FS improves their prediction outcomes to a large extent. Note that the classifier for which SCALE has the largest impact is KNN, as it is a distance-based classifier which is sensitive to the range of feature values. Also, due to instability and tendency to overfit, HT is the most critical for DT improvement.




Fig. 1: Performance improvement in the AUC score of each dimension over the baseline when tuning only one dimension at a time while leaving others at baseline settings. CC brings the greatest performance improvement, followed by LB, FS, MV, HT, and SCALE in decreasing order of improvement.

Classifier	MV (%)	LB (%)	SCALE (%)	FS (%)	HT (%)
LR	2.78	11.48	0.30	5.50	3.03
SVM	3.37	26.83	2.38	20.95	3.21
KNN	13.60	11.85	17.68	13.12	15.60
NB	0.60	38.90	0.17	4.12	2.84
DT	12.87	16.22	0.42	15.34	38.85
BAG	2.94	8.91	0.28	7.05	5.43
RF	4.13	5.34	0.28	5.85	1.04
ET	3.82	7.87	0.00	18.96	1.33
ABC	19.33	7.02	0.00	16.99	12.99
GB	12.44	3.81	0.02	6.63	4.08
LGBM	7.03	1.85	2.75	10.39	3.13
XGB	11.46	3.97	0.02	7.47	4.27
MLP	10.78	5.08	6.05	7.53	5.69
STACK	6.94	8.94	4.32	5.48	1.82
VOTE	6.38	4.00	2.11	6.04	0.85

Table 4: Column 1 shows a total of 15 classifiers. Columns 2 to 6 represent the percentage (two decimal places accuracy) of AUC score improvement when tuning each individual dimension while leaving other dimensions at baseline settings for each classifier. Bold entries represent the dimension that contributes to the largest improvement for the specific classifier. MV and LB tend to dominate in performance improvement for most classifiers.

In addition to AUC, 6 other performance metrics are used to measure the performance improvement degree of each dimension. The results in Table 5 reveal that CC brings the greatest improvement regardless of the metrics we use. Contributions from HT and SCALE are relatively small compared to other dimensions.

	AUC	F-score	G-mean	Precision	Sensitivity/ Recall	Specificity	Accuracy
CC (%)	15.00	15.58	10.50	16.41	10.50	11.86	10.50
LB (%)	10.81	11.34	9.33	13.27	9.33	10.72	9.34
FS (%)	10.09	7.33	6.30	10.61	6.30	6.94	6.30
MV (%)	7.90	5.30	4.60	5.83	4.59	4.95	4.59
HT (%)	7.46	2.11	3.21	3.41	3.20	4.64	3.20
SCALE (%)	2.45	1.04	0.65	3.03	0.65	0.48	0.65

Table 5: Performance improvement in different metrics of each dimension. The performance improvement of each dimension on other metrics displays an order consistent with that of the AUC score.

Performance Variation across Dimensions For all of the ML configurations, we further investigate how much each dimension contributes to the performance variation in the AUC score. By tuning only one dimension at a time while leaving other dimensions at baseline settings, we obtain a range of AUC scores. Performance variation is the difference between the maximum and the minimum score of each dimension. Fig. 2 shows the proportion of each dimension that brings the performance variation in the AUC score. Based on Fig. 2, we notice that CC, which brings the largest performance improvement, also brings the largest performance variation (10.98 %). After CC, LB (7.00 %), FS (6.93 %), MV (5.64 %), HT (4.97 %), and SCALE (1.66 %) bring decreasing degrees of performance variation in the AUC score.

Table 6 shows the variation of every single dimension on each classifier over the baseline. We observe that for each classifier, if one dimension brings a larger performance improvement, it also results in a larger performance variation. For our assessment of performance variation, the same metrics as above are used for evaluation on each dimension. Using the same metrics as above, Table 7 shows that the proportion of performance variation in different metrics from each dimension follows an order that is consistent with the performance improvement in Table 5. Thus, for different metrics, greater improvement brings greater variation of each dimension. For every step that researchers take when predicting MOF using ML, they should always be aware of the trade-off between benefits (improvement in performance) and risks (variation in performance) when adjusting each dimension.




Fig. 2: Performance variation in the AUC score when tuning only one dimension at a time while leaving others at baseline settings. CC brings the greatest performance variation, followed by LB, FS, MV, HT, and SCALE in decreasing order of variation. Larger improvement also brings the risk of larger variation for each dimension.

Classifier	MV (%)	LB (%)	SCALE (%)	FS (%)	HT (%)
LR	2.28	8.44	0.25	4.27	2.49
SVM	2.45	16.87	1.79	13.04	2.42
KNN	7.97	6.95	10.36	7.69	9.14
NB	0.48	22.22	0.13	3.25	2.25
DT	7.13	8.99	0.23	8.50	21.53
BAG	2.22	6.17	0.21	5.26	4.10
RF	3.35	4.14	0.23	4.49	0.84
ET	3.22	6.14	0.00	13.41	1.12
ABC	13.09	4.61	0.00	11.51	9.40
GB	9.59	2.83	0.02	5.11	3.14
LGBM	5.54	1.43	2.16	7.76	2.47
XGB	8.70	3.01	0.02	5.59	3.24
MLP	7.89	3.55	4.43	5.30	4.16
STACK	5.47	6.47	3.41	4.20	1.43
VOTE	5.19	3.13	1.71	4.63	0.69

Table 6: Columns 2 to 6 represent the proportion (two decimal places accuracy) of each dimension that contributes to the performance variation in the AUC score. Bold entries represent the dimension that contributes to the largest variation for the specific classifier. MV and LB tend to result in larger performance variation for most classifiers.

	AUC	F-score	G-mean	Precision	Sensitivity/ Recall	Specificity	Accuracy
CC (%)	10.98	11.87	8.57	12.86	8.57	10.60	8.57
LB (%)	7.00	7.29	6.83	9.02	6.83	10.14	6.82
FS (%)	6.93	5.27	4.38	7.62	4.37	4.70	4.38
MV (%)	5.64	4.36	3.69	4.77	3.69	2.98	3.68
HT (%)	4.87	2.55	3.52	1.54	3.52	2.72	3.53
SCALE (%)	1.66	0.88	0.57	1.47	0.57	0.46	0.57

Table 7: Performance variations in different metrics of each dimension. The performance variation of each dimension on other metrics displays an order that is consistent with that of the AUC score.

4.2 Performance Comparison across Classifiers

We have shown that classifier choice is the largest contributor to both performance improvement and variation in the AUC score. Hence, we further investigate the performance differences among classifiers. Specifically, we investigate the relationships among classifier complexity, performance, and performance variation.

Default versus Optimized Performance *Default* classifiers are defined as classifiers with default parameters, while *optimized* classifiers are those for which hyperparameter tuning with 10-fold cross validation is applied using grid search. We compare the performance of default and optimized classifiers in consideration of all other dimensions, i.e., MV, LB, SCALE, and FS. The average AUC scores of all classifiers with default and optimized settings are shown in Fig. 3. In general, ensemble classifiers perform better than single classifiers regardless of default or optimized performance.

In addition to AUC, 6 other performance metrics are used to evaluate the performance of all classifiers. We use the median score to rank classifiers with both default and optimized settings. Then, NDCG (normalized discounted cumulative gain), one of the most prevalent measures of ranking quality [30], is used to compare classifier rankings between each of these metrics and the AUC score. Detailed relevance scores are shown in Table 8. The result indicates that the median performance of each classifier is similar no matter which metric is used. This also suggests that the AUC score can represent classifiers’ overall performance well.

Based on the above experiments, ensemble classifiers should be prioritized in MOF prediction since they usually bring better predictive performance than single classifiers.




Fig. 3: Comparison of default and optimized performance over all classifiers. Classifiers listed on the left-hand side of BAG are single while the ones on the right-hand side are ensemble and MLP. Overall, ensemble methods have better default and optimized performance compared with single classifiers.

	Default (%)	Optimized (%)
F-score	96.92	97.92
G-mean	96.46	97.63
Precision	95.49	90.01
Sensitivity/Recall	98.42	97.59
Specificity	95.35	97.46
Accuracy	96.46	97.59

Table 8: Column 1 represents 6 other performance metrics. Columns 2 and 3 show the NDCG score between each of these metrics and the AUC score when ranking 15 classifiers by their median performance in default and optimized settings, respectively. Median performance of classifiers is similar regardless of which metric to use.

Performance Variation across Classifiers We measure the performance variation for each classifier in consideration of all other dimensions, i.e., MV, LB, SCALE, FS, and HT. For each classifier, we get a range of AUC scores. The size of the range determines the extent of performance variation. Fig. 4 shows the performance variation in the AUC score of all classifiers. The order of listed classifiers on the x -axis is based on increasing model complexity, which is measured by classifier training time with default settings. The complexity of classifiers and performance variation demonstrates an evident ‘U-shaped’ relationship. When the classifier is ‘too simple’, its performance variation is relatively large. When the complexity of the classifier is ‘appropriate’, the performance variation is relatively small. If the classifier becomes ‘too complex’, it is also at the risk of larger performance variation. Therefore, classifiers with ‘appropriate’ complexity are more stable, with smaller changes in performance, while ‘too simple’ or ‘too complex’ classifiers are relatively unstable with larger changes in performance in general.

In addition to AUC, the same metrics as above were used to validate the performance variation of all of the classifiers. We use the range (difference between maximum and minimum scores) to rank classifiers in consideration of MV, LB, SCALE, FS, and HT. Then, NDCG is used to compare classifier rankings between each of these metrics and the AUC score. Table 9 displays detailed relevance scores. The result suggests that other metrics show a similar ‘U-shaped’ relationship between classifier complexity and performance variation as the AUC score. When predicting MOF, it is inappropriate for clinical practitioners to choose ‘too simple’ and ‘too complex’ classifiers since they may run the risk of underfitting and overfitting, respectively.




Fig. 4: Performance variation comparison over all classifiers. The order of classifiers listed on the x -axis is based on increasing model complexity. ‘Too simple’ and ‘too complex’ classifiers result in larger performance variation. The performance variation of classifiers with ‘appropriate’ complexity is relatively small.

	F-score	G-mean	Precision	Sensitivity/ Recall	Specificity	Accuracy
Relevance (%)	93.15	94.98	94.37	93.24	93.13	93.77

Table 9: NDCG score between each of 6 other performance metrics and the AUC score in terms of classifier complexity and performance variation. Different metrics show a similar ‘U-shaped’ relationship.

5 Discussion

We have provided a timely MOF prediction using early lab measurements (hour 0), patients’ demographic and illness information. Our study quantitatively analyzes the performance via the AUC score in consideration of a wide range of ML configurations for MOF prediction, with a focus on the correlations among configuration complexity, predicted performance, and performance variation. Our results indicate that choosing the correct classifier is the most crucial step that has the largest impact (performance and variation) on the outcome. More complex classifiers including ensemble methods can provide better default/optimized

performance, but may also lead to larger performance degradation, without careful selection. Clearly, more MOF data is needed to provide a more general conclusion. Our work can potentially serve as a practical guide for ML practitioners whenever they conduct data analysis in healthcare and medical fields.

6 Acknowledgments

This work was funded by the National Institutes for Health (NIH) grant NIH R01 - HL149670.

References

1. Harjola, V.P., Mullens, W., Banaszewski, M., et al.: Organ dysfunction, injury and failure in acute heart failure: from pathophysiology to diagnosis and management. A review on behalf of the Acute Heart Failure Committee of the Heart Failure Association (HFA) of the European Society of Cardiology (ESC). *European Journal of Heart Failure*, vol. 19, pp. 821-836 (2017).
2. Wang, Z.K., Chen, R.J., Wang, S.L., et al.: Clinical application of a novel diagnostic scheme including pancreatic β -cell dysfunction for traumatic multiple organ dysfunction syndrome. *Molecular Medicine Reports*, vol. 17, pp. 683-693 (2018).
3. Durham, R. M., Moran, J. J., Mazuski, J. E., et al.: Multiple organ failure in trauma patients. *Journal of Trauma and Acute Care Surgery*, vol. 55, pp. 608-616 (2003).
4. Ulvik, A., Kvåle, R., Wentzel-Larsen, T., et al.: Multiple organ failure after trauma affects even long-term survival and functional status. *Critical Care*, vol. 11, pp. R95 (2007).
5. Barie, P.S., Hydo, L.J., Fischer, E.: A prospective comparison of two multiple organ dysfunction/failure scoring systems for prediction of mortality in critical surgical illness. *The Journal of Trauma*, vol. 37, pp. 660-666 (1994).
6. Bota, D.P., Melot, C., Ferreira, F.L., et al.: The Multiple Organ Dysfunction Score (MODS) versus the Sequential Organ Failure Assessment (SOFA) score in outcome prediction. *Intensive Care Medicine*, vol. 28, pp. 1619-1624 (2002).
7. Dewar, D.C., White, A., Attia, J., et al.: Comparison of postinjury multiple-organ failure scoring systems. *Journal of Trauma and Acute Care Surgery*, vol. 77, pp. 624-629 (2014).
8. Hutchings, L., Watkinson, P., Young, J.D., et al.: Defining multiple organ failure after major trauma. *Journal of Trauma and Acute Care Surgery*, vol. 82, pp. 534-541 (2017).
9. Sauer, A., Moore, F.A., Moore, E.E., et al.: Multiple Organ Failure Can Be Predicted as Early as 12 Hours after Injury. *Journal of Trauma and Acute Care Surgery*, vol. 45, pp. 291-303 (1998).
10. Vogel, J.A., Liao, M.M., Hopkins, E., et al.: Prediction of postinjury multiple-organ failure in the emergency department. *Journal of Trauma and Acute Care Surgery*, vol. 76, pp. 140-145 (2014).
11. Obermeyer, Z., Emanuel, E.J.: Predicting the Future — Big Data, Machine Learning, and Clinical Medicine. *New England Journal of Medicine*, vol. 375, pp. 1216-1219 (2016).
12. Cruz, J.A., Wishart, D.S.: Applications of Machine Learning in Cancer Prediction and Prognosis. *Cancer Informatics*, vol. 2, pp. 117693510600200030 (2006).

13. Kourou, K., Exarchos, T.P., Exarchos, K.P., et al.: Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal*, vol. 13, pp. 8-17 (2015).
14. Asri, H., Mousannif, H., Al Moatassime, H., et al.: Using Machine Learning Algorithms for Breast Cancer Risk Prediction and Diagnosis. *Procedia Computer Science*, vol. 83, pp. 1064-1069 (2016).
15. Sharma, K., Kaur, A., Gujral, S.: Brain Tumor Detection based on Machine Learning Algorithms. *International Journal of Computer Applications*, vol. 103, pp. 7-11 (2014).
16. Wang, Z., Yu, G., Kang, Y., et al.: Breast tumor detection in digital mammography based on extreme learning machine. *Neurocomputing*, vol. 128, pp. 175-184 (2014).
17. De Bruijne, M.: Machine learning approaches in medical image analysis: From detection to diagnosis. *Medical Image Analysis*, vol. 33, pp. 94-97 (2016).
18. Farrar, C.R., Worden, K.: Structural health monitoring: a machine learning perspective, <https://onlinelibrary.wiley.com/doi/book/10.1002/9781118443118> (2013).
19. Worden, K., Manson, G.: The application of machine learning to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 365, pp. 515-537 (2007).
20. Ahmed, Z., Mohamed, K., Zeeshan, S., et al.: Artificial intelligence with multi-functional machine learning platform development for better healthcare and precision medicine. *Database*, vol. 2020 (2020).
21. Janssen, K.J., Donders, A.R.T., Harrell Jr, F.E., et al.: Missing covariate data in medical research: To impute is better than to ignore. *Journal of Clinical Epidemiology*, vol. 63, pp.721-727 (2010).
22. Tuba, E., Strumberger, I., Bezdán, T., et al.: Classification and Feature Selection Method for Medical Datasets by Brain Storm Optimization Algorithm and Support Vector Machine. *Procedia Computer Science*, vol. 162, pp. 307-315 (2019).
23. Chawla, N.V., Bowyer, K.W., Hall, L.O., et al.: SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, vol. 16, pp. 321-357 (2002).
24. Mani, I., Zhang, I.: kNN approach to unbalanced data distributions: a case study involving information extraction. In: *Proceedings of workshop on learning from imbalanced datasets*, vol. 126 (2003).
25. Batista, G.E., Bazzan, A.L., Monard, M.C.: Balancing Training Data for Automated Annotation of Keywords: a Case Study. In: *WOB*, pp. 10-18 (2003).
26. Dağ, H., Sayin, K.E., Yenidoğan, I., et al.: Comparison of feature selection algorithms for medical data. In: *2012 International Symposium on Innovations in Intelligent Systems and Applications*, pp. 1-5 (2012).
27. Pedregosa, F., Varoquaux, G., Gramfort, A., et al.: Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, vol. 12, pp. 2825-2830 (2011).
28. Bakker, J., Gris, P., Coffernils, M., et al.: Serial blood lactate levels can predict the development of multiple organ failure following septic shock. *The American Journal of Surgery*, vol. 171, pp. 221-226 (1996).
29. Papachristou, G.I., Muddana, V., Yadav, D., et al.: Comparison of BISAP, Ranson's, APACHE-II, and CTSI Scores in Predicting Organ Failure, Complications, and Mortality in Acute Pancreatitis. *American Journal of Gastroenterology*, vol. 105, pp. 435-441 (2010).
30. Chen, W., Liu, T.Y., Lan, Y., et al.: Ranking measures and loss functions in learning to rank. In: *Advances in Neural Information Processing Systems*, vol. 22, pp. 315-323 (2009).