

The Application of Machine Learning Techniques in Prediction of Quality of Life Features for Cancer Patients*

Miloš Savić¹, Vladimir Kurbalija¹, Mihailo Ilić¹, Mirjana Ivanović¹, Dušan Jakovetić¹, Antonios Valachis², Serge Autexier³, Johannes Rust³, and Thanos Kosmidis⁴

¹ Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad
Trg D. Obradovića 3, Novi Sad, Serbia
{svc, kurba, milic, mira, dusan.jakovetic}@dmi.uns.ac.rs

² Department of Oncology, Faculty of Medicine and Health, Örebro University
SE 70182 Örebro, Sweden
antonios.valachis@oru.se

³ German Research Center for Artificial Intelligence GmbH, Cyber-Physical Systems
Bremen, Germany
{serge.autexier, johannes.rust}@dfki.de

⁴ Care Across Ltd
London, England

thanos.kosmidis@careacross.com

Abstract. Quality of life (QoL) is one of the major issues for cancer patients. With the advent of medical databases containing large amounts of relevant QoL information it becomes possible to train predictive QoL models by machine learning (ML) techniques. However, the training of predictive QoL models poses several challenges mostly due to data privacy concerns and missing values in patient data. In this paper, we analyze several classification and regression ML models predicting QoL indicators for breast and prostate cancer patients. Three different approaches are employed for imputing missing values, and several settings for data privacy preserving are tested. The examined ML models are trained on datasets formed from two databases containing a large number of anonymized medical records of cancer patients from Sweden. Two learning scenarios are considered: centralized and federated learning. In the centralized learning scenario all patient data coming from different data sources is collected at a central location prior to model training. On the other hand, federated learning enables collective training of machine learning models without data sharing. The results of our experimental evaluation show that the predictive power of federated models is comparable to that of centrally trained models for short-term QoL predictions, whereas for long-term periods centralized models provide more accurate QoL predictions. Furthermore, we provide insights into the quality of data preprocessing tasks (missing value imputation and differential privacy).

Keywords: Quality of Life, Cancer Patients, Predictive Models, Federated Learning, Breast Cancer, Prostate Cancer

* This is an extended version of a MEDES 2021 conference paper [1].

1. Introduction

Nowadays in modern society people of all ages suffer from different chronic and critical diseases. The one of important trends in medical domains is to find out adequate services that will help patients to keep, or even if it is possible to increase, their quality of life and health parameters.. Information and Communication Technologies offers powerful tools, devices and services for collection of wide range of patient's data: digital health records, wearable devices, dashboards, unobtrusive sensors and other types of smart devices. Using advanced and effective techniques of Artificial intelligence (AI) and Machine learning (ML) [2] such data are processed and can be used to improve patient's health status and different quality of life (QoL) aspects [3, 4].

Predictive analytics in medicine and healthcare plays essential role and recently is more concentrated on obtaining personalized interventions, treatment and improvement of QoL issues. However, health QoL issues depend on type of disease and patients status (beginning disease phase, after surgery, during treatment, follow-ups, etc.) and in this paper we concentrate of various QoL issues for breast and prostate cancer such as depression, anxiety, insomnia, etc.

Research results presented in the paper are achieved within ASCAPE project: Artificial intelligence supporting cancer patients across Europe (<https://ascape-project.eu/>) where two the most prevalent types of cancer are considered: breast and prostate cancer. One of the main purposes of the project is to employ powerful AI/ML mechanisms to support cancer patients' health status and QoL.

Apart from patients' data collected from different sources like health records, wearable devices, environmental data, within ASCAPE also different validated questionnaires will be used to capture the particular QoL issues depending on type of cancer. Such questionnaires are good instruments to minimize the risk for measurement bias, and to enable the better ASCAPE ability for suitable predictions and interventions. Clinical ASCAPE partners identified 15 QoL issues for breast and 12 for prostate cancer that will be predicted through AI-based models.

To the best of our knowledge, ASCAPE is a unique research project that prospectively investigate an AI-based approach, towards a personalised follow-up strategy for cancer patients focusing on their QoL issues. The aim of the project is to offer substantial benefits for after-treatment health-related QoL improvements.

Importance of applying adequate approaches to properly process huge amounts of sensitive patients' data is evident in different medical disciplines and areas [5–8]. Two approaches to train predictive models are prevalent: centralized and distributed. In the centralized approach patients' data collected from different sources are stored locally in clinics and model training is performed locally as well. However, recently popular decentralized ML technique is Federated learning (FL) which enables the quality use and learning from decentralized data. Patients' data collected by clinics could be adequately processed by applying FL approach as it enables the training of shared global models with a central server, while keeping all the sensitive data in local. Moreover, raw data could not be used directly for ML model training since it needs some additional activities for curation and pre-processing. In this paper we will examine the importance and the impact of two pre-processing techniques which are required with existing form of data: missing value imputation (MVI) and differential privacy (DP). The primary purpose of

FL in ASCAPE is to enable democratized access to ML models promoting patient QoL without revealing private or sensitive patient data.

Presented experimental results are based on already existing retrospective datasets of an ASCAPE clinical partner. The aim of conducted experiments was to identify appropriate ML models for breast and prostate cancer patients. However, numerous ML techniques which are utilized during project showed to be very promising and we will continue our research activities on prospectively collected data expecting accurate and reliable behavior of developed models. The experimental evaluation of the ML models presented in this paper is based on datasets containing QoL indicators derived from prescribed medications and LISAT-11 questionnaires. As far as we are aware, this is the first analytical study of various ML models predicting such QoL features. Another important novelty of our study is that the selected ML models are not examined only in a centralized learning scenario, but also in two different federated learning settings. Medical datasets usually contain missing values, while ensuring privacy during predictive model training is of the uttermost importance. Thus, as the last contribution of our paper, we study the impact of appropriate pre-processing techniques (MVI and DP) to the quality of investigated ML models.

The rest of paper is organized as follows. The second Section is devoted to related work. The central, third Section, is focused on application of ML techniques for predictive QoL models. Section four is devoted to the description of experimental datasets. Experimental results are discussed in Section five. The last Section brings concluding remarks.

2. Related Work

The tracking and monitoring of QoL parameters has attracted many research attention recently. This has a particular importance in chronic diseases for the prevention and early detection of symptoms and signs. The proper monitoring should provide the positive effect on patient's quality of life, economic impact and resource management. As expected, substantial amounts of data on quality of life are combined into clinical trials using a variety of instruments [9].

Regular assessment of QoL parameters of chronic patients has an impact on physician-patient communication and result in benefits for some patients, who had a better QoL [5]. Furthermore, continuous collection of medical attributes and growth of the database could allow the selection of proper medical variables and the selection of adequate models for a more accurate prediction of QoL [9].

Several recent studies emphasized the importance of applying AI and data mining (DM) techniques in successful prediction of QoL issues. For example, in [10] authors used QoL issues to monitor patients' clinical condition and health status for the patients with ultra-rare autosomal recessive disease Alkaptonuria. One of the main problems in the follow-up of patients with ultra rare diseases is the lack of a standardized methodology to assess disease severity or response to treatment and QoL scores could be a successful way to monitor such patients. Also, different machine learning approaches (Linear Regression, Neural networks and k-nearest-neighbor) were implemented with the aim to perform a prediction of QoL scores based on clinical data.

The authors in [3] present the basis of the DIAL system which represents an early warning system for the early detection of a deteriorating QoL score in the hemodialysis population using machine learning algorithms. Here two models (classification tree and Naïve Bayes) were generated to predict an increase or decrease of 5% in a patient's QoL score over one month. The classification tree was selected as the better model with an area under curve (AUC) of 83.3% and accuracy of 81.9%. The authors concluded that their system DIAL, if implemented on a larger scale, is expected to help patients in terms of ensuring a better QoL and a reduction in the financial burden in the long term.

The predictor of quality of life called the Better Life Index (BLI) was used for QoL assessment in [11]. It is based on the measurement of different aspects of human life in the whole population: environment, jobs, health, civic engagement, governance, education, access to services, housing, community, and income. The paper presents a supervised machine-learning analytical model that predicts the life satisfaction score using several DM models like: decision tree, elastic net, neural network, random forest, support vector machine, etc. The results showed that the ensemble model based on the stacked generalization framework is a significantly better predictor of the life satisfaction of a nation, compared to base models.

QoL are studied in different cancer related diseases since they present a common chronic conditions nowadays. For example, in [12] authors presented the project which tries to design a new patient reported outcome measure to assess QoL issues for patients with locally recurrent rectal cancer. The authors identified the fact that it is very tedious to administer simultaneously several questionnaires required for QoL assessment especially in patients with major limitations caused by the disease. So it is essential to find alternatives which largely replace questionnaires, and their approach involves use of biometric devices with the help of DM techniques.

Some successful applications of QoL prediction in cancer diseases also includes: the comparison of the pre and the post treatment quality of life in cervical cancer patients [13], construction of machine learning and statistical models for prediction of gastro-urinary symptoms and QoL issues following prostate radiation treatment [14], and the prediction of 5-year lung cancer survival on the basis of QoL issues [4].

To the best of our knowledge, this study is the first one which tries to develop a QoL prediction model for breast and prostate cancer patients. The study [14] also investigates prostate cancer patients but focus only on one type of intervention – radiation treatment. Here we will try to encompass both types of cancer patients (breast and prostate) with as many as possible common attributes, with all types of common interventions, and with a set of relevant QoL attributes.

3. Machine Learning Techniques for Predictive QoL Models

The QoL indicator of a cancer patient can be predicted either by a classification or a regression machine learning model depending on the type of the indicator. In this work we consider two types of QoL indicators: (1) binary indicators indicating whether the patient will experience QoL related symptoms after diagnosis (e.g., anxiety and depression) and (2) numeric indicators indicating overall QoL of the patient reported by filling an appropriate questionnaire (e.g., the QoL score of the LISAT-11 questionnaire [15]). Binary

classification models are the most appropriate for the first type of QoL indicators, while for the second type of QoL indicators predictions can be obtained by regression models.

For predictive QoL models (as well as for medical predictive models in general) it is important to distinguish two practical scenarios: centralized and decentralized (federated). In the centralized scenario all anonymized training data coming from different medical organizations is collected to a central place where machine learning models are trained. However, in majority of cases medical organizations are not willing to share anonymized patient data, or it is not even allowed for them to do so according to governmental regulations and laws. One solution in this case is to apply federated learning [16] of predictive QoL models in which several edge nodes collectively train predictive QoL models by exchanging locally updated machine learning models instead of sharing training datasets. Edge nodes are computational devices owned by medical organizations and deployed within their computational infrastructure, so training datasets never leave the boundaries of data owners.

Before training AI models, the available datasets must be pre-processed. Some pre-processing steps are mandatory, while the others only improve the predictive power or security/privacy of trained models. In this paper we will consider two pre-processing approaches: missing values imputation (MVI) and differential privacy (DP). MVI is obligatory step since the model training can not be performed with incomplete data. DP is optional since its role is to protect the privacy of patients data.

3.1. Centralized QoL Predictive Models

The following machine learning algorithms are examined for training centralized binary classification models predicting binary QoL indicators: NB (Naive Bayes), kNN (k nearest neighbors), SVM (support vector machines), DT (decision trees) and RF (random forests).

NB is a probabilistic classification algorithm learning a predictive model giving the most probable class (positive or negative in the case of binary classification) for a given data instance (in our case a patient described by a set of features). Class probabilities are computed using conditional probability estimates learned from a training dataset under the assumption that features describing data instances are conditionally independent.

kNN is a lazy learning algorithm. The class for a given data instance is predicted by majority voting from the classes of the k closest data instances belonging the training dataset according to some distance functions (e.g., the Euclidean or Manhattan distance).

SVM classifiers are based on the idea of using linear models to identify non-linear boundaries of classes. This is achieved by transforming data instances into a new higher dimensional space using a non-linear mapping. Quadratic programming algorithms are then employed in the higher dimensional space to determine the maximum margin hyper-plane separating instances from different classes.

Decision-tree based classifiers make predictions according to decision trees constructed from the training dataset by a recursive divide-and-conquer algorithm utilizing some information theoretic measure (e.g, information gain or the Gini impurity). This means that decision trees are formed in the divisive manner from the root of the tree to its leafs. The underlying information theoretic measure is used to find and select the most discriminative feature to form a corresponding node in the tree and split training data to recursively form its subtrees.

A random forest is an ensemble of decision trees learned from bootstrapped samples of the training dataset. The RF algorithm employs a feature bagging procedure to determine a random subset of features for learning individual decision trees. The class for a given input data instance is then determined as the most frequent class predicted considering all decision trees in the ensemble.

For predicting numeric QoL indicators we examine the following algorithms for learning regression models: LINEAR (linear regression), RIDGE (ridge regression), LASSO (lasso regression), ELASTICN (elastic net regression), KRIDGE (kernel ridge regression), SVM (regression by support vector machines), RF (regression by random forests), and kNN (k-nearest neighbours regression).

The linear regression algorithms determine coefficients of a linear model by minimizing the residual sum of squares (RSS) between real values of the target variable and predictions derived from the model. Ridge, Lasso and Elastic Net find linear models by minimizing RSS with incorporated regularization penalties: Ridge incorporates the L2 regularization penalty, Lasso is based on the L1 regularization penalty, while Elastic Net uses both previously mentioned penalties. Kernel Ridge regression performs Ridge regression in a space obtained by a non-linear mapping of the training dataset. SVM, RF and kNN are adaptations of the corresponding classification algorithms for regression tasks.

In our experiments we have used the Scikit-learn machine learning library [17] to develop a set of Python modules for training centralized QoL predictive models.

3.2. Federated QoL Predictive Models

A federated model is a machine learning model collectively trained by several edge nodes running federated learning clients. Each federated learning client has its own dataset for training the model and those local training datasets are never exchanged among federated learning clients participating in federated learning. The federated learning process is coordinated by a federated learning server. The main purpose of the federated learning server is to enable the exchange of locally updated federated models among federated learning clients.

Two basic federated learning schemas are incremental and concurrent. Let C_1, C_2, \dots, C_k denote k federated learning clients each having its own training dataset D_i ($i \in [1..k]$). In the incremental federated learning scheme federated learning clients incrementally build a machine learning model from the first to the last client. This means that C_1 creates M on D_1 and sends it to the federated learning server. Then, C_2 retrieves M from the federated learning server, updates it on D_2 and returns the updated model back to the server. Each next federated learning client does exactly the same until the last client C_k .

In the concurrent federated learning scheme M is collectively trained in parallel. In the first step each federated learning client C_i creates its own model M_i on D_i . All models are then sent to the federated learning server which averages M_1 to M_k into a single model M . Once all edge nodes have submitted their local models, the global model is updated using the Federated Averaging [18] approach. The federated learning server then sends M to all federated learning clients which update M on their local datasets and the updated models are returned back to the federated learning server for the second averaging. The previous operation is repeated for an arbitrary number of learning rounds and the averaged model after the last learning round is the final federated model.

Neural networks are the most natural model choice for federated learning for the following two reasons: (1) neural networks can be incrementally updated, and (2) neural networks can be easily averaged by averaging edge weights and biases. Since we deal with two types of predictive problems we also have two types of federated neural networks:

1. Federated neural networks for regression. The last layer of such neural networks contains exactly one node activated by the linear function. In our work we use the mean squared error (MSE) as the loss function when training regression neural networks.
2. Federated neural networks for binary classification. In this case, one node contained in the last layer is activated by the sigmoid function. Output values higher than 0.5 indicate the positive class, while values lower than 0.5 correspond to the negative class. The binary cross-entropy function is used at the loss function when optimizing parameters of binary classification neural networks.

Nodes in hidden layers of both types of federated neural networks are activated by the ReLU activation function. We also consider two mechanisms to prevent overfitting: dropout and regularization strategies (the kernel, bias and activation regularization).

To compare federated QoL models to centralized QoL models we have developed a federated learning simulator based on the Tensorflow machine learning library [19]. The realized simulator supports both incremental and concurrent federated learning mode for an arbitrary number of simulated edge nodes (federated learning clients). The architecture of a federated neural network can be specified by providing its type (regression or binary classification), the number of hidden layers and the number of nodes per hidden layer. The user can also specify the number of epoch (learning rounds) and the batch size (the number of training instances propagated through the network when updating model parameters). At the beginning, the simulator divides training data (training folds when the k -fold cross-validation is applied to evaluate models) into p stratified parts, where p is the number of simulated edge nodes. Then, simulated edge nodes use their part of training data when creating or updating Tensorflow neural networks by the Adam optimization algorithm [20]. In the case of concurrent federated learning, the averaging of Tensorflow neural networks formed by simulated edge nodes is performed after each epoch. More precisely, we utilized the callback mechanism provided by Tensorflow to implement the federated client-server communication. A socket connection is open between an edge node and the coordinating server once a client wishes to update a global model, and at key training steps such as the end of each epoch, the edge node sends its model updates to the server.

3.3. Missing Value Imputation (MVI)

Datasets for training machine learning models often contain missing values. This is also the case with the experimental datasets used in this work. To train machine learning models it is necessary to infer and fill missing values of predictor features or, alternatively, to remove data instances containing missing values. For the MVI process we use three methods provided by the Scikit-learn library: (1) simple MVI, (2) iterative MVI and (3) nearest neighbours imputation.

The simple MVI fills missing values using a simple approach: all missing values for a feature f are filled with the mean of existing values in f .

The iterative MVI is based on the idea to train a regression model for each feature containing missing values [21]. The regression model for feature f is trained based on values of other predictor features. Then, missing values for f are filled based on predictions of its regression model. After obtaining predictions for all missing values, the iterative MVI repeats the whole procedure for predefined number of times in a round-robin fashion (i.e., predicted missing values of f together with known values are then used to retrain the regression model for other features also containing missing values). The predictions of the final round are then used to fill missing values. It is important to emphasize here that we do not use a datasets' target attribute for the training of regression model. If the target attribute was used, the regression models of the iterative imputer would be fit with target values, which must not happen since it would introduce erroneous dependence between predictors and target attribute.

In nearest neighbours imputation a Euclidean distance metric is used to find the nearest neighbours. Each missing feature is imputed using values from k (parameter of the method) nearest neighbours that have a value for the feature. The features of the neighbours could be averaged uniformly or weighted by distance to each neighbour.

3.4. Differential Privacy (DP)

The task of DP [22] is to preserve the privacy of patients data and to prevent model inversion attacks [23]. Using these attacks an attacker can reconstruct some of the training data from AI models (under particular circumstances). Differential privacy is based on the idea that the outcome of the query posed to protected database (or machine learning model) is essentially equally likely independent of whether any individual joins or refrains from joining the database. In such a way the private data about a particular patient is protected since the system returns the result with the same probability whether a particular patient was involved in the analysis or not.

There are several mechanisms to implement DP, but the most common are: Laplace mechanism, Gaussian mechanism and Exponential mechanism. Laplace mechanism is most commonly used for numeric types of data, and it consists of adding Laplacian noise (a noise which follows Laplace distribution) to the data model. Furthermore, in machine learning algorithms noise could be added in different ways: (1) to the training dataset, (2) to the prediction model itself (for example in edge weights in neural network), and (3) to the predicted results of the model. Here, first option will be applied as the simplest one, but still sufficiently general and robust.

As expected, more noise means more privacy, but a lot of noise could lead to the reduction of model performance/accuracy. So, the main challenge in introduction of privacy preserving techniques in machine learning models is the optimal balance between privacy and the utility of models. The amount of privacy is controlled through parameter ϵ . Smaller value of ϵ means stronger privacy, which is achieved by adding more noise to the training data. The value 0 represents total privacy, but the usability of such an algorithm is none since in that case training data represents pure randomness.

3.5. Model Evaluation

To estimate errors of the examined regression models we use the mean absolute error (MAE):

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad (1)$$

where n is the number of data instances in a test dataset (or a test fold), y_i the real value of the target variable for i -th instance and \hat{y}_i is the predicted value of the target value according to a regression model.

In our experimental evaluation of QoL regression-based predictive models, MAE estimates of the examined regression models are also compared to MAE estimates of the so-called DUMMY regression algorithm. The DUMMY regression algorithm always predict the same value for a target feature: the mean of the target feature from a training dataset (or training folds).

Binary classifiers predicting QoL issues are evaluated using accuracy, precision, recall and F_1 scores. The accuracy of a binary classifier is equal to the number of correctly classified instances divided by the total number of instances in a test dataset (or a test fold). Precision and recall metrics are defined per class. The precision for a class c (c is the positive or negative class) is the number of instances correctly classified to c divided by the total number of instances classified to c . On the other hand, the recall for c is the number of instances correctly classified to c divided by the total number of instances that belong to c . Precision and recall score for the whole binary classifier are obtained by averaging precision and recall scores per class. Since precision and recall measure two different aspects of classifier's performance it is useful to aggregate them into a single score. The usual way to aggregate precision and recall is to compute the F_1 score which is the harmonic mean of precision (P) and recall (R):

$$F_1 = \frac{2 \cdot P \cdot R}{P + R}. \quad (2)$$

All above-mentioned model evaluation metrics are estimated by the 10-fold cross validation procedure. In the 10-fold cross-validation, an input dataset is split into 10 folds. Then, 10 iterations of training and evaluation are performed. In the i -th iteration, i -th fold is used to compute evaluation metrics of models trained on the remaining folds. In this way we obtain 10 estimates of evaluation metrics that are then averaged into final estimates.

The effects of missing value imputation to the performance of predictive QoL models is investigated by comparing models trained on datasets to which the simple MVI is applied, to models trained on datasets obtained after the iterative MVI and after the MVI based on kNN rule. The impact of differential privacy is measured by comparing models trained without and with applied differential privacy with different values of parameter ϵ .

The evaluation methodology described in this section may also be utilized by practitioners for selecting the best predictive ML model for any other medical dataset with numeric or discrete target variables. After applying MVI and DP techniques to a dataset, each ML model from a list of candidates is validated by the 10-fold cross validation that results with a metric indicating the overall predictive ability of the model. In case of numeric target variables, the regression-based model with the lowest MAE is the best one.

On the other hand, the model with the largest F_1 is the best performing model in the case of discrete target variables.

4. Experimental Datasets

Previously described machine learning models are evaluated on datasets formed from two databases: ORB and BcBase. BcBase is a population-based research database containing data about early breast cancer patients from three healthcare regions in Sweden (accounting for nearly 60% of the total population) in terms of patient and tumor characteristics, treatment strategies and prescribed medications. The database does not include QoL indicators captured through questionnaires, but from prescribed medications it is possible to derive the presence of certain QoL issues. From BcBase we formed 4 datasets for training and evaluating binary classifiers deciding whether a patient will suffer from anxiety (BcBase-Anxiety), depression (BcBase-Depression), insomnia (BcBase-Insomnia) and pain (BcBase-Pain) after breast cancer treatment. All BcBase datasets contain 18988 data instances (patients) described by 97 predictor features.

ORB is a database containing data about patients with localised prostate cancer treated with radiotherapy in the Örebro healthcare region. It includes data on patient and tumour characteristics, treatment approaches, dosimetric parameters regarding radiotherapy, side effects based on direct questions or validated questionnaires (IPSS, IIEF-5) and QoL issues based on a validated questionnaire (LISAT-11). ORB contains 2466 health records with follow-up examinations repeated at six months intervals and LISAT-11 QoL scores at the time of the diagnosis and three different times relative to the date of diagnosis at months 36, 60 and 120. From ORB data we created six datasets with the naming scheme ORB- n - m for training regression models predicting the LISAT-11 QoL score at month m considering all patient data collected up to month n . The created datasets are: ORB-30-36, ORB-30-60, ORB-30-120, ORB-54-60, ORB-54-120 and ORB-108-120. For example, ORB-30-36 is used to train regression models predicting the LISAT-11 QoL score in month 36 considering data collected up to month 30 as predictor variables. The number of instances (patients) in ORB datasets ranges from 1138 in ORB-30-36 to 610 in ORB-108-120, while the number of predictor variables ranges from 96 in ORB-30-36 to 158 in ORB-108-120 (the number of patients decreases due to recoveries, dropouts and deaths, while the number of predictors increases due to longer time intervals. Also some patients got their diagnosis less than 10 years ago. They haven't had their follow-up yet).

5. Results and Discussion

In this section we first present the evaluation of centralized machine learning models trained on experimental datasets described in previous section. The influence of missing value imputation and differential privacy to the performance of those models is then discussed. Finally, we examine federated QoL machine learning models and compare them to their centralized counterparts.

5.1. Evaluation of Centralized QoL Models

The performance of binary classification models trained on BcBase datasets obtained by the 10-fold cross-validation are summarized in Tables 1, 2, 3, and 4. All classifiers were

trained and evaluated on BcBase datasets obtained after missing value imputation by the iterative MVI method.

Table 1. Evaluation of binary classification models on BcBase-Anxiety

Classifier	Accuracy	Precision	Recall	F_1
RF	0.673	0.484	0.535	0.514
SVM	0.698	0.411	0.349	0.5
NB	0.629	0.552	0.553	0.552
KNN	0.682	0.458	0.529	0.507
DT	0.583	0.511	0.511	0.511

Table 2. Evaluation of binary classification models on BcBase-Depression

Classifier	Accuracy	Precision	Recall	F_1
RF	0.677	0.473	0.524	0.509
SVM	0.702	0.413	0.351	0.5
NB	0.566	0.534	0.543	0.551
KNN	0.688	0.462	0.536	0.509
DT	0.589	0.515	0.515	0.515

Table 3. Evaluation of binary classification models on BcBase-Insomnia

Classifier	Accuracy	Precision	Recall	F_1
RF	0.538	0.526	0.535	0.532
SVM	0.541	0.533	0.539	0.537
NB	0.555	0.554	0.555	0.554
KNN	0.521	0.502	0.516	0.514
DT	0.516	0.515	0.515	0.515

The SVM classifier exhibits the highest accuracy on three out of four BcBase datasets (BcBase-Anxiety, BcBase-Depression and BcBase-Pain). However, precision and recall scores of SVM on those three BcBase datasets are significantly lower compared to other classifiers. A more detailed examination of precision and recall per class revealed that SVM has zero precision and zero recall for the positive class (the presence of anxiety,

Table 4. Evaluation of binary classification models on BcBase-Pain

Classifier	Accuracy	Precision	Recall	F_1
RF	0.698	0.486	0.554	0.518
SVM	0.714	0.417	0.357	0.5
NB	0.53	0.517	0.553	0.564
KNN	0.699	0.457	0.528	0.506
DT	0.604	0.522	0.522	0.522

depression and pain after treatment) and that it dominantly predicts the negative class (no negative QoL related symptoms). The highest accuracy of SVM on those three dataset is a consequence of its bias towards the negative class on class imbalanced datasets in which approximately 70% of the patients belong to the negative class and 30% to the positive class. Therefore, it can be concluded that accuracy is not an appropriate measure for comparing binary classification models trained on BcBase datasets. F_1 score is more adequate measure since it takes into account precision and recall of both classes. It can be observed that SVM exhibits the lowest F_1 score on those datasets where it has the highest accuracy. KNN has the lowest F_1 score on BcBase-Insomnia and the second lowest on other BcBase datasets. Thus, it can be concluded that those two methods are the worst performing binary classification models for BcBase datasets.

The largest F_1 score on three BcBase datasets (anxiety, depression and insomnia) is achieved by the NB classifier. The best model for the fourth dataset (pain) is DT, but its F_1 score is very close to the F_1 score of NB. Consequently, it can be concluded that NB is the best choice to train centralized QoL predictive models for BcBase datasets.

The results of the evaluation of centrally trained regression models on the ORB datasets are shown in Table 5 including the DUMMY regressor as the baseline. The best model (the lowest MAE) for ORB-30-36 is RF. For the rest of ORB datasets, the best performing model is LASSO. KNN is the worst performing regression algorithm on all ORB datasets: predictions made by this model are even more erroneous than predictions made by DUMMY. Excluding kNN, all others considered models exhibit smaller prediction errors compared to DUMMY except in one case: DUMMY is better than linear regression on ORB-30-120. The prediction errors of the best performing model are in the range [4.84, 6.47], which is an acceptable level of prediction errors taking into account that the target variable (the LISAT QoL index) is in the range [11, 66].

The improvement of LASSO (the best performing regression model) over DUMMY are significant for short term QoL predictions (30-36, 54-60, 108-120) when the reduction of the MAE score ranges from 20% to 30%. For medium term QoL predictions (30-60, 54-120) the improvements are between 10% and 15%. As expected, the lowest improvement is for long term QoL predictions on ORB-30-120 where the reduction of MAE scores is slightly higher than 5%.

Table 5. MAE scores of regression models on ORB datasets (best value per column is bolded)

Regressor	30-36	30-60	30-120	54-60	54-120	108-120
DUMMY	6.541	6.89	6.909	6.89	6.909	6.909
LINEAR	5.311	6.129	7.003	5.238	6.899	6.524
RIDGE	5.1	5.925	6.652	5.07	6.356	5.977
LASSO	5.089	5.886	6.478	4.84	6.18	5.437
ELASTICN	5.126	5.913	6.504	4.859	6.216	5.448
KRIDGE	5.147	5.958	6.75	5.115	6.492	6.155
SVM	6.519	6.773	6.871	6.772	6.859	6.875
RF	5.051	6.015	6.685	5.009	6.357	5.635
KNN	6.72	6.968	7.133	6.906	7.128	7.033

5.2. Influence of Missing Value Imputation

To evaluate the impact of MVI and DP to the accuracy of the examined ML models we focus on one dataset – ORB-30-36⁵. The following MVI techniques are applied to ORB-30-36 in order to obtain a complete dataset with missing values: (1) Simple imputer, (2) Iterative imputer, and (3) Imputer base on kNN rule. The value of parameter k in kNN imputer is varied from 1 to 10.

After imputation of missing values all regression models described in Section 3.1 are trained using datasets with imputed values. The performance of these models is measured using MAE regression evaluation metrics.

Table 6 shows the MAE values for all 3 imputer approaches, and for 10 values of the parameter k in kNN approach. Since there are a lot of regression methods, table is split on two sub-tables.

In order to better visualize the impact of MVI techniques, the values of MAE of two most promising regressors (LASSO and ElasticNet) are plotted in Figure 1 and Figure 2. The presented behaviour is similar to all regressors, although it is not the same in all circumstances. Considering kNN approach, better values (smaller MAE) are achieved for higher values of k . Simple and iterative imputers are usually in the range of lower MAE values of kNN approach. In some cases, simple imputer is better, while in some other cases iterative imputer shows better performance.

Since the differences are quite small, and since the simple imputer is significantly faster than other two, it can be concluded that the choice of the MVI method does not make a significant influence to the performance of examined models and that the simple MVI can be as equally effective as the iterative or kNN MVI.

5.3. Influence of Differential Privacy

Our experimental datasets contain missing values. Consequently, it is first necessary to impute missing values before applying differential privacy. The experimental results presented in the previous section shows that the choice of MVI imputer does not significantly

⁵ The results of MVI and EP evaluation for other datasets are highly similar to the results obtained on ORB-30-36.

Table 6. MAE values for different MVI techniques and for different regressors on ORB-30-36 dataset

MVI	Linear	Ridge	Lasso	ElasticNet	KernelRidge
simple	5.3040	5.2366	5.0814	5.1084	5.2619
iterative	5.2494	5.1141	5.1344	5.1912	5.1501
knn_1	5.4895	5.4254	5.3072	5.3204	5.4575
knn_2	5.3644	5.3084	5.1617	5.1774	5.3374
knn_3	5.2343	5.1727	5.1313	5.1435	5.1998
knn_4	5.2735	5.2117	5.1567	5.1619	5.2364
knn_5	5.2565	5.1927	5.1525	5.1695	5.2207
knn_6	5.2625	5.1948	5.1038	5.1134	5.2306
knn_7	5.2732	5.2021	5.1201	5.1132	5.2353
knn_8	5.2619	5.1885	5.0952	5.0958	5.2258
knn_9	5.2584	5.1773	5.0950	5.1004	5.2179
knn_10	5.2682	5.1836	5.0807	5.0873	5.2240
MVI	SVM	RF	KNN	AdaB	TFNN
simple	6.5197	5.1142	6.4550	5.5006	6.0614
iterative	6.5000	5.1079	6.7064	5.3787	6.0978
knn_1	6.5110	5.2839	6.6407	5.6586	6.3338
knn_2	6.5164	5.1627	6.7020	5.4869	6.0745
knn_3	6.5148	5.1237	6.7146	5.4681	6.0500
knn_4	6.5200	5.1076	6.7535	5.5114	6.1146
knn_5	6.5208	5.1107	6.7520	5.4943	5.9661
knn_6	6.5148	5.1442	6.7349	5.4876	6.6542
knn_7	6.5183	5.0847	6.7675	5.4033	5.9290
knn_8	6.5201	5.1352	6.7311	5.4286	5.8820
knn_9	6.5205	5.1653	6.7735	5.3508	5.7765
knn_10	6.5204	5.1455	6.7483	5.3892	6.0941

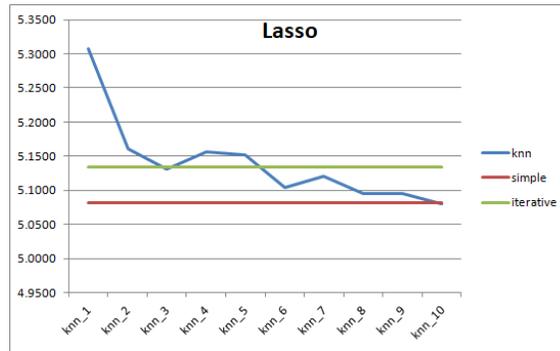


Fig. 1. MAE values for Lasso regressor for different MVI techniques

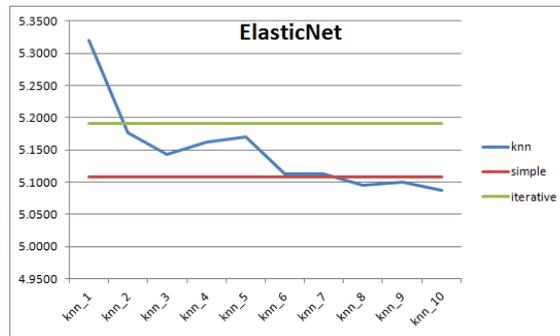


Fig. 2. MAE values for ElasticNet regressor for different MVI techniques

influence the predictive ability of the ML models. Thus, in the experimental evaluation of DP, we impute missing values in the ORB-30-36 dataset by the simple imputer. Then, the dataset without missing values is treated with the DP component which adds a specified amount of noise. The amount of noise is expressed through the ϵ parameter. The noise is added only on numerical features while the categorical features, Boolean features, one-hot-encoding features and class features stayed untouched. The following values of DP parameter (ϵ) were used for initial experiments: {0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50, 100}. These initial experiments should provide an insight in the behaviour of ML models with regard to the amount of added noise.

All obtained results for all ML regression models behave in a similar manner. These results could be summarized as shown in Figure 3 and Figure 4 (here the results for the most prominent regressors LASSO and RF are shown). For a very small value of ϵ (e.g., 0.1), the MAE is notably high. After that, MAE significantly drops with an increase of ϵ up to the value of approximately 1. For the values of $\epsilon > 1$, MAE remains approximately constant on the level of MAE of models trained with no DP data (baseline level).

Since the lower ϵ parameter means more privacy, the goal here was to find the lowest value of ϵ with satisfactory MAE value. The conclusion from this preliminary set of experiments is that the value of ϵ should be in the neighborhood of value 1. For the values of

$\epsilon > 1$ the models do not have significantly higher performance (do not have lower MAE), and the level of privacy is decreased.

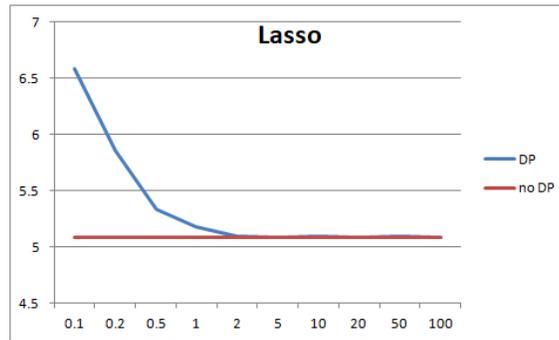


Fig. 3. MAE values for Lasso regressor for different values of DP parameter.

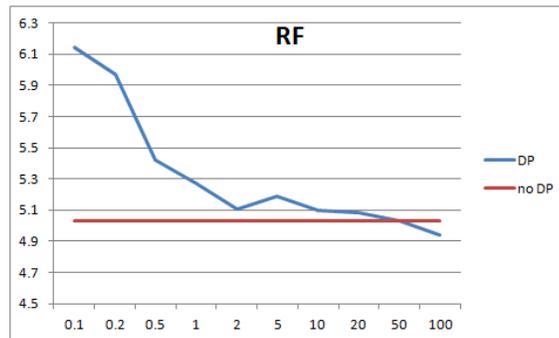


Fig. 4. MAE values for RF regressor for different values of DP parameter.

Since the privacy of the patient data is particularly important we performed additional experiments to find an optimal value of ϵ (lowest value of ϵ with satisfactory MAE value of regressors). We tested the performance of all regressors with the following values of ϵ (with finer granularity around value 1): $\{0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 2\}$. The results are shown in Table 7.

The MAE values of the most prominent regressors are plotted in Figure 5 and Figure 6 (LASSO and RF). It is evident that in this interval of ϵ values there are no significant changes in trend of regressors' performance. Same behaviour is also observed in case of other regressors. However, as already mentioned, we are interested in the smallest value of ϵ with satisfactory MAE value. From previous experiments we know that the MAE value steadily drops for the small values of ϵ . And now it is evident from these results that this constant drop-down is present from $\epsilon = 0.1$ to $\epsilon = 0.7$ which is applicable for all

Table 7. MAE values for different values of DP parameter and for different regressors on ORB-30-36 dataset.

EPS	Linear	Ridge	Lasso	Elastic Net	Kernel Ridge
0.5	5.4861	5.4378	5.3345	5.3501	5.4648
0.6	5.4210	5.3804	5.2231	5.2518	5.4001
0.7	5.3296	5.2879	5.1340	5.1646	5.3153
0.8	5.3002	5.2633	5.2317	5.2624	5.2781
0.9	5.3366	5.2832	5.1960	5.2313	5.3093
1	5.2954	5.2535	5.1809	5.2119	5.2568
1.1	5.5842	5.5404	5.4678	5.4899	5.5640
1.2	5.2935	5.2500	5.1456	5.1799	5.2737
1.3	5.5551	5.5114	5.4355	5.4589	5.5350
1.4	5.2051	5.1577	5.0898	5.1276	5.1847
1.5	5.3807	5.3366	5.2439	5.2728	5.3582
1.6	5.3220	5.2722	5.1698	5.2036	5.2958
1.7	5.3602	5.3156	5.2295	5.2603	5.3371
1.8	5.3633	5.3184	5.2390	5.2699	5.3392
1.9	5.3696	5.3239	5.2398	5.2706	5.3454
2	5.2583	5.2039	5.0817	5.1219	5.2276
NO DP	5.3109	5.1000	5.0895	5.1259	5.1474
EPS	SVM	RF	KNN	AdaB	TFNN
0.5	6.5186	5.4221	6.7054	5.5476	6.3555
0.6	6.5186	5.2985	6.7123	5.4608	6.3131
0.7	6.5186	5.2177	6.7159	5.4039	6.7360
0.8	6.5186	5.3210	6.7152	5.5178	6.2085
0.9	6.5186	5.1730	6.7250	5.4338	6.0307
1	6.5186	5.2733	6.7239	5.4399	6.1016
1.1	6.5186	5.4779	6.7139	5.6432	6.3989
1.2	6.5186	5.2679	6.7204	5.4214	6.2083
1.3	6.5186	5.4569	6.7145	5.6210	6.3799
1.4	6.5186	5.1505	6.7270	5.3515	5.8816
1.5	6.5186	5.3059	6.7173	5.4841	6.2614
1.6	6.5186	5.2953	6.7212	5.4282	5.8575
1.7	6.5186	5.2939	6.7194	5.4745	6.2064
1.8	6.5186	5.3016	6.7198	5.4815	6.1535
1.9	6.5186	5.2996	6.7202	5.4779	6.1480
2	6.5186	5.1055	6.7194	5.4234	5.7977
NO DP	6.5186	5.0332	6.7200	5.3623	6.0351

regressors. After value $\epsilon = 0.7$ the MAE value starts to fluctuate without no significant trend of decrease. Therefore, the value 0.7 is imposed as the appropriate choice for DP parameter ϵ .

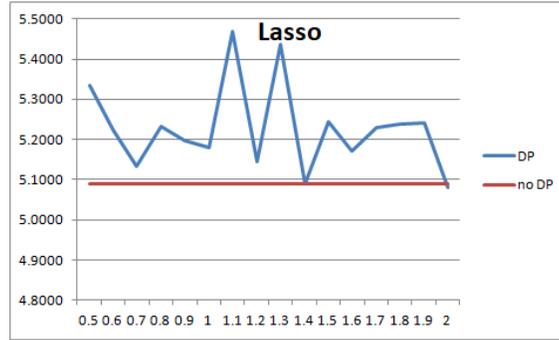


Fig. 5. MAE values for Lasso regressor for different values of DP parameter (with finer granularity).

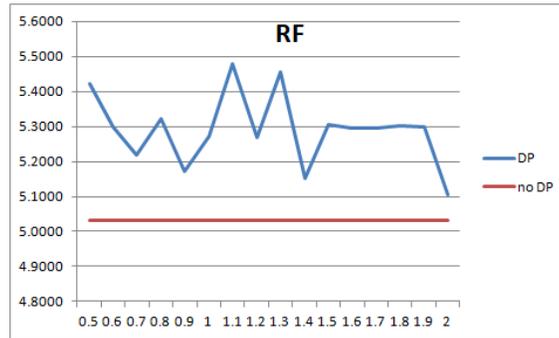


Fig. 6. MAE values for RF regressor for different values of DP parameter (with finer granularity).

5.4. Evaluation of Federated QoL Models

In experiments with simulated federated models, we have used different neural network architectures for BcBase and ORB datasets. A preliminary investigation, in which we have varied the number of hidden neural network layers between 1 and 10 and the batch size in the set $\{16, 64, 128, 256, 512\}$, showed that shallow neural networks (a small number of hidden layers) trained with a large batch size are more suitable for BcBase datasets, while deeper neural networks (a larger number of hidden layers) trained with a small batch size result with better predictive models for ORB datasets. We have simulated from 2 to 4 edge nodes training models in both incremental and concurrent federated learning mode.

Performance metrics for simulated federated models were obtained in the same way as for centrally-trained predictive models (by the 10-fold cross validation procedure).

The architecture of neural networks for federated binary classification models on the BcBase datasets consists of 4 hidden layers each having 20 nodes. Federated neural networks were trained in 200 epochs per simulated edge node with batch size equal to 512. The comparison of F_1 scores of centrally trained and simulated federated binary classification models is presented in Table 8. As baselines we use a centrally trained TensorFlow-based neural network (TFNN), the best centrally trained non-neural network classifier (Best C) and the worst centrally trained non-neural network classifier (Worst C). INC- k and CON- k denote federated neural network binary classification models trained in the incremental (INC) and concurrent (CON) learning mode for k simulated edge nodes.

Table 8. F_1 scores of federated binary classification models on BcBase datasets.

	Anxiety	Depression	Insomnia	Pain
INC-2	0.536	0.512	0.546	0.542
INC-3	0.542	0.507	0.529	0.542
INC-4	0.539	0.515	0.538	0.532
CON-2	0.522	0.504	0.542	0.548
CON-3	0.512	0.519	0.55	0.534
CON-4	0.53	0.509	0.542	0.545
Best C	0.552	0.534	0.554	0.522
Worst C	0.411	0.413	0.502	0.457
TFNN	0.438	0.53	0.54	0.542

For BcBase-Anxiety, Depression and Insomnia datasets, we have observed that federated models are significantly better than the worst performing local model (SVM and kNN depending on the dataset). F_1 scores of federated models are close to F_1 scores of NB which is the best performing centrally trained model. Federated models trained on BcBase-Pain have higher F_1 scores compared to the best performing centrally trained model on that dataset (DT). It is also important to emphasize that there are no significant differences in F_1 scores of incremental models and concurrent federated models. Additionally, the performance of federated models does not tend to significantly drop with the number of simulated edge nodes.

For federated regression models trained on the ORB dataset, we have used neural networks with 10 hidden layers each with 40 neurons. The training was performed in 200 epochs per simulated edge node with the batch size equal to 32. The obtained MAE scores are summarized in Table 9. For all six datasets, the best local model (LASSO) has lower prediction errors than simulated federated models. Large differences between federated models trained in different federated learning modes are absent. In contrast to federated models trained on BcBase datasets, here we can observe a tendency of increasing errors with the number of simulated edge nodes. Federated models are better than the DUMMY baseline for ORB-30-36, ORB-30-60, ORB-54-60 and ORB-108-120, but

worse than DUMMY for ORB-30-120 and ORB-54-120. This result implies that different neural network architectures should be employed for short term and long term QoL predictions. Therefore, our subsequent work will be to examine a wider range of neural network architectures for federated regression and determine architectures providing satisfactory results for long-term QoL predictions.

Table 9. MAE scores of federated regression models on ORB datasets.

	30-36	30-60	30-120	54-60	54-120	108-120
INC-2	6.012	6.775	7.488	5.931	7.188	6.625
INC-3	6.472	6.751	7.226	5.867	7.169	6.43
INC-4	6.595	7.042	7.463	6.22	7.206	6.484
CON-2	5.881	6.705	7.444	5.904	7.193	6.463
CON-3	6.404	6.826	7.427	5.986	7.098	6.327
CON-4	6.534	6.883	7.538	6.269	7.221	6.652
DUMMY	6.541	6.89	6.909	6.89	6.909	6.909
LASSO	5.089	5.886	6.478	4.84	6.18	5.437
TFNN	5.783	6.572	7.323	5.811	7.206	6.562

6. Conclusions and Future Work

In this paper we have examined several classification and regression machine learning algorithms for training models predicting binary and numeric QoL indicators, respectively, for breast and prostate cancer patients within ASCAPE project. The focus of our experimental evaluation was on two types of predictive models: (1) centrally-trained QoL models that are relevant either for individual data owners or for multiple data owners when it is allowed to collect training data in a central location, and (2) federated QoL models trained in distributed environments encompassing multiple data owners without data sharing. We also examined the influence of three different MVI algorithms and the privacy-accuracy trade-off parameter in DP approach on the performance of examined predictive models.

In the MVI pre-processing task we tested 3 approaches: simple imputer, iterative imputer and imputer based on kNN rule. The kNN imputer is further tested with 10 values of the parameter k , from 1 to 10. It is notable that kNN imputer has almost steady down trend (in the MAE values) while the k is increasing, obtaining the best values for $k = 10$. However, these differences are not particularly important, and the simple and iterative imputers obtain similar results as kNN for higher values of k . Our analysis of simple, iterative and kNN MVI algorithms showed that the choice of the MVI does not make a significant influence to the performance of examined models. Therefore, our choice for future analysis will be the simple imputer since it achieves comparable results as the other two and is much faster than the other two.

In the differential privacy task the goal was to add a controlled amount of noise to the training data in order to preserve the privacy of the patients. More noise means more

privacy, but lower performance of AI models (measured with MAE and other evaluation metrics values). The amount of noise is controlled through parameter ϵ , and the goal here was to find the value of ϵ with highest privacy and with satisfying regression performance. After extensive set of experiments our conclusion is to further utilize value of DP parameter $\epsilon = 0.7$.

Our experimental evaluation on real datasets showed that numeric QoL indicators can be accurately predicted by both centrally-trained and federated models for short term future periods. Centrally-trained regression models provide also accurate long term predictions. On the other hand, federated regression models exhibit prediction errors close to the dummy regression model indicating that different neural network architectures should be employed for learning regression-based federated models providing short term and long term QoL predictions.

For classification-based models predicting binary QoL indicator it was observed that centrally-trained and federated models have comparable prediction performances. However, for both types of model we noticed that they achieve relatively low precision and recall scores for the minority class due to imbalanced training datasets indicating that appropriate data sampling techniques should be examined to form more class-balanced training datasets prior to model training.

In our future work, we will also examine various feature selection techniques to identify the most relevant features for making QoL predictions and examine the performance of predictive QoL models trained on selected features. Furthermore, it would be interesting to examine federated learning approaches which are not based on neural networks like: decision trees [24] or logistic regression [25].

Acknowledgments. This research was supported by the ASCAPE project. The ASCAPE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 875351. The authors would also like to thank the anonymous reviewers for their insightful suggestions and comments that helped improve the quality of the paper.

References

1. Savić, M., Kurbalija, V., Ilić, M., Ivanović, M., Jakovetić, D., Valachis, A., Autexier, S., Rust, J., Kosmidis, T.: Analysis of Machine Learning Models Predicting Quality of Life for Cancer Patients, p. 35–42. Association for Computing Machinery, New York, NY, USA (2021), <https://doi.org/10.1145/3444757.3485103>
2. Sidey-Gibbons, J., Sidey-Gibbons, C.: Machine learning in medicine: a practical introduction. *BMC Medical Research Methodology* 19 (03 2019)
3. Saadat, S., Aziz, A., Ahmad, H., Imtiaz, H., Sohail, Z., Kazmi, A., Aslam, S., Naqvi, N., Saadat, S.: Predicting quality of life changes in hemodialysis patients using machine learning: Generation of an early warning system. *Cureus* 9 (09 2017)
4. Sim, J., Kim, Y., Kim, J., Lee, J., Kim, M.S., Shim, Y., Zo, J., Yun, Y.H.: The major effects of health-related quality of life on 5-year survival prediction among lung cancer survivors: applications of machine learning. *Scientific Reports* 10, 10693 (07 2020)
5. Velikova, G., Booth, L., Smith, A., Brown, P., Lynch, P., Brown, J., Selby, P.: Measuring quality of life in routine oncology practice improves communication and patient well-being: A randomized controlled trial. *Journal of clinical oncology : official journal of the American Society of Clinical Oncology* 22, 714–24 (03 2004)

6. Singh, A., Pannu, H.S., Malhi, A.: Explainable information retrieval using deep learning for medical images. *Computer Science and Information Systems* 19(1), 277–307 (01 2022)
7. Šušteršič, T., Peulić, M., Peulić, A.: FPGA Implementation of Fuzzy Medical Decision Support System for Disc Hernia Diagnosis. *Computer Science and Information Systems* 18(3), 619–640 (06 2021)
8. Bratić, B., Kurbalija, V., Ivanović, M., Oder, I., Bosnić, Z.: Machine learning for predicting cognitive diseases: Methods, data sources and risk factors. *J. Med. Syst.* 42(12) (oct 2018), <https://doi.org/10.1007/s10916-018-1071-x>
9. Sinha, R., Heuvel, W.: A systematic literature review of quality of life in lower limb amputees. *Disability and rehabilitation* 33, 883–99 (06 2011)
10. Spiga, O., Cicaloni, V., Fiorini, C., Trezza, A., Visibelli, A., Millucci, L., Bernardini, G., Bernini, A., Marzocchi, B., Braconi, D., Prischi, F., Santucci, A.: Machine learning application for development of a data-driven predictive model able to investigate quality of life scores in a rare disease. *Orphanet Journal of Rare Diseases* 15 (12 2020)
11. Kaur, M., Dhalaria, M., Sharma, P., Park, J.: Supervised machine-learning predictive analytics for national quality of life scoring. *Applied Sciences* 9, 1613 (04 2019)
12. Gonçalves, J., Faria, B.M., Reis, L.P., Carvalho, V., Rocha, A.: Data mining and electronic devices applied to quality of life related to health data. In: 2015 10th Iberian Conference on Information Systems and Technologies (CISTI). pp. 1–4 (2015)
13. Kumar, S., Rana, M., Verma, K., Singh, N., Sharma, A., Maria, A., Singh, G., Khaira, H., Saini, S.: Prediqt-cx: Post treatment health related quality of life prediction model for cervical cancer patients. *PloS one* 9, e89851 (02 2014)
14. Yang, Z., Olszewski, D., He, C., Pinteá, G., Lian, J., Chou, T., Chen, R.C., Shtylla, B.: Machine learning and statistical prediction of patient quality-of-life after prostate radiation therapy. *Computers in Biology and Medicine* 129, 104127 (2021), <https://www.sciencedirect.com/science/article/pii/S0010482520304583>
15. Melin, R., Fugl-Meyer, K., Fugl-Meyer, A.: Life satisfaction in 18-to 64-year-old swedes: In relation to education, employment situation, health and physical activity. *Journal of rehabilitation medicine : official journal of the UEMS European Board of Physical and Rehabilitation Medicine* 35, 84–90 (04 2003)
16. Yang, Q., Liu, Y., Chen, T., Tong, Y.: Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.* 10(2) (Jan 2019), <https://doi.org/10.1145/3298981>
17. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830 (2011)
18. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D’Oliveira, R.G.L., Eichner, H., Rouayheb, S.E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P.B., Gruteser, M., Harchaoui, Z., He, C., He, L., Huo, Z., Hutchinson, B., Hsu, J., Jaggi, M., Javidi, T., Joshi, G., Khodak, M., Konečný, J., Korolova, A., Koushanfar, F., Koyejo, S., Lepoint, T., Liu, Y., Mittal, P., Mohri, M., Nock, R., Özgür, A., Pagh, R., Raykova, M., Qi, H., Ramage, D., Raskar, R., Song, D., Song, W., Stich, S.U., Sun, Z., Suresh, A.T., Tramèr, F., Vepakomma, P., Wang, J., Xiong, L., Xu, Z., Yang, Q., Yu, F.X., Yu, H., Zhao, S.: Advances and open problems in federated learning (2021)
19. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*. p. 265–283. OSDI’16, USENIX Association, USA (2016)

20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015)
21. Buuren, S., Groothuis-Oudshoorn, C.: MICE: Multivariate imputation by chained equations in R. *Journal of Statistical Software* 45 (12 2011)
22. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) *Theory of Cryptography*. pp. 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
23. Fredrikson, M., Jha, S., Ristenpart, T.: Model inversion attacks that exploit confidence information and basic countermeasures. In: *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. p. 1322–1333. CCS '15, Association for Computing Machinery, New York, NY, USA (2015), <https://doi.org/10.1145/2810103.2813677>
24. Li, Q., Wen, Z., He, B.: Practical federated gradient boosting decision trees (2019)
25. Yang, S., Ren, B., Zhou, X., Liu, L.: Parallel distributed logistic regression for vertical federated learning without third-party coordinator (2019)

Miloš Savić is an Associate Professor with the Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, where he received his B.Sc., M.Sc., and Ph.D. degrees in computer science in 2010, 2011 and 2015, respectively. His research interests are related to complex network analysis and machine learning. He is/was a member of several international and bilateral research projects including three H2020 projects. He published more than 40 papers in international journals and proceedings of international conferences.

Vladimir Kurbalija holds the position of Full Professor from 2021 at the Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Serbia, where he received his B.Sc., M.Sc. and Ph.D. degrees. He was/is a member of several international projects supported by DAAD, TEMPUS, Horizon, and bilateral programs. Vladimir (co)authored over 40 papers in Case-Based Reasoning, Time-Series Analysis, Medical decision-support systems, and related fields. He was a member of Program Committees of several international conferences, and a reviewer in more than 20 international journals.

Mihailo Ilić currently works as a Teaching Assistant at the Faculty of Sciences, University of Novi Sad, Serbia, where he is also pursuing his PhD. His main fields of interests are agent technologies and machine learning, with a focus on federated learning techniques. As a member of the ASCAPE H2020 project, his work is focused on the implementation of federated learning in the medical domain. In addition, Mihailo has notable experience from working in industry as both a software developer responsible for system design and a data scientist in the field of smart agriculture.

Mirjana Ivanović holds the position of Full Professor at Faculty of Sciences, University of Novi Sad, Serbia. She is a member of National Scientific Committee for Electronics, Telecommunication and Informatics within Ministry of Education, Science and Technological Development, Republic of Serbia. Prof. Ivanovic is author or co-author of 14 textbooks, several international monographs and more than 450 research papers in

areas: agent technologies, intelligent techniques, applications of machine learning techniques in medical domains and technology enhanced learning. She is member of Program Committees of more than 300 international conferences and Program/General Chair of several international conferences. Mirjana Ivanovic delivered several keynote speeches at international conferences and visited numerous academic institutions all over the world as visiting researcher. Currently she is Editor-in-Chief of the Computer Science and Information Systems.

Dušan Jakovetić received the Dipl. Ing. Degree from the School of Electrical Engineering, University of Belgrade, Belgrade, Serbia, in 2007, and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, and the Instituto de Sistemas e Robótica (ISR), Instituto Superior Técnico (IST), Lisbon, Portugal, in 2013. Dr. Jakovetic is an associate professor at the Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Serbia. Prior to that, he worked as an assistant professor with the same faculty, a research fellow at the BioSense Institute, Novi Sad, Serbia, and a postdoctoral researcher at IST. Dr. Jakovetic's research interests include distributed inference, learning, and optimization.

Antonios Valachis is Associate Professor of Oncology since 2018 and senior consultant at the Department of Oncology, Örebro University Hospital. His clinical work and research is focused on breast cancer and melanoma.

Serge Autexier studied and received his doctorate (2003) in computer science with a focus on AI. He is a principal researcher at the German Research Center for Artificial Intelligence (DFKI), which he joined in 2002, and his research areas are formal methods, cyber-physical systems, human-environment/-robot interaction, and active and healthy living. Since 2014 he is head of the Bremen Ambient Assisted Living Lab and DFKI Research Fellow since 2017 and was and engaged as PI or Co-PI in more than 15 national projects and EU projects.

Johannes Rust received the Master of Science degree in Systems Engineering from the University of Bremen in 2020. Since then he worked in the department for Cyber-Physical Systems of the German Research Center for Artificial Intelligence (DFKI) at the site Bremen. His main research topics are explainable AI and computer vision.

Thanos Kosmidis is cofounder and CEO of CareAcross, a digital health company focusing on cancer, providing personalised support with evidence of improving patients' quality of life. Thanos has been at Imperial College, Carnegie Mellon University, Massachusetts Institute of Technology. Before CareAcross he worked in leading technology companies in the USA and Europe.

Received: February 27, 2022; Accepted: August 22, 2022.