RESEARCH ARTICLE　　　　　　　　　　　　　　**OPEN ACCESS**

ARTIFICIAL INTELLIGENCE FOR SUPPORT OF THE BIOETHICAL DECISION-MAKING OF NURSES RELATED TO END OF LIFE ISSUES

*[1]Gabriel Resende Machado, [2,3]Oswaldo Jesus R. da Motta, [4,5]Eugênio Silva, [2]Luciene Muniz Braga, [6]Lumaira M. Nascimento Silva da Rocha Marques, [7]Selma Vaz Vidal, [1]Ronaldo Ribeiro Goldschmidt and [2,8]Rodrigo Siqueira-Batista**

[1]Computer Engineering Departament (SE/9), Military Institute of Engineering (IME), Rio de Janeiro, RJ – Brazil; [2]Department of Medicine and Nursing, Universidade Federal de Viçosa (UFV), Viçosa, MG – Brazil; [3]Departament of Public Health Sciences and Pediatrics, Università degli Studi di Torino, Torino (UNITO); [4]School of Computer Science, Centro Universitário Serra dos Órgãos (UNIFESO), Teresópolis, RJ – Brazil; [5]Computing Unit, Fundação Centro Universitário da Zona Oeste (UEZO), Rio de Janeiro, RJ – Brazil; [6]Postgraduate Program of Bioethics, Applied Ethics and Collective Health, Universidade Federal Fluminense (UFF), Niterói, RJ – Brazil; [7]School of Nursing, Centro Universitário Serra dos Órgãos (UNIFESO), Teresópolis, RJ – Brazil; [8]School of Medicine, Faculdade Dinâmica do Vale do Piranga (FADIP), Ponte Nova, MG – Brazil

---

ARTICLE INFO

ABSTRACT

Health professionals, especially nurses, are often compelled to make decisions in face of bioethical issues commonly related to intensive care environments. Due to their complexities and particularities, these issues end up generating great physical and emotional strain on the professional. Therefore, this study raised the hypothesis that artificial intelligence techniques could assist nurses in making bioethical decisions. For this, a group of four specialists in the field of nursing prepared and answered a questionnaire based on a hypothetical ICU scenario. The data collected from the responses were then used in the training and evaluation of five machine learning (ML) algorithms. The results culminated from a comparative study between the algorithms point to the veracity of the hypothesis raised in this research, by showing that it is possible to adopt ML algorithms to support bioethical decision making.

INTRODUCTION

The intensive care unit (ICU) aims to offer care to clinically unstable patients, who could possibility recover, in a way that they are constantly monitored by an interdisciplinary team of specialists. The ICU environment, when compared to other health care environments, has several distinct aspects that bring technical, scientific and ethical issues to discussion. Situations like the manipulation of critical patients in borderline situations (*i.e.*, characterized by a high risk of death), for example, often force nurses and other health professionals to make complex decisions in a short period of time.

These decisions take into account, mainly, the balance between logistical and bioethical factors, which causes the professional great tension, often leading to physical and emotional stress [SIQUEIRA-BATISTA *et al.*, 2019]. The difficulties encountered in similar palliative situations make artificial intelligence (AI) techniques promising alternatives to assist health professionals in making bioethical decisions, since machine learning (ML) algorithms have already been used for diagnosing illnesses and administrating treatments [SIQUEIRA-BATISTA *et al*., 2014]. Therefore, this work raises the hypothesis that the adoption of AI algorithms can assist health professionals in making bioethical decisions.

To verify such conjecture, a questionnaire has been elaborated, based on a hypothetical scenario related to the research topic, and answered afterwards by four experts. Based on the data gathered from the four questionnaires, a comparative study has been carried out among five ML algorithms, namely: (i) Multilayer Perceptron Neural Network, (ii) Decision Tree, (iii) Random Forest, (iv) Support Vector Machine with RBF function and (v) Logistic Regression. The performance of each model has been evaluated from the perspective of four different metrics: (a) accuracy, (b) confusion matrix, (c) F1 metric and (d) ROC/AUC analysis. The results obtained from the comparative study corroborate the hypothesis raised by this work, by showing the feasibility of ML algorithms to support the decision making process of health professionals in similar bioethical conflicts. The rest of this study is organized as follows: Section 2 details the materials and methods used in the design of the experiments. Section 3 describes the experiments that have been carried out and shows the obtained results. Finally, Section 4 presents the final considerations.

MATERIALS AND METHODS

Specialists: Four specialists, with degrees in Nursing and PhDs in Bioethics/Public Health or Nursing, participated in the elaboration and filling out of the questionnaire (as described in Section 2.2.) A summary of these specialists' training and experience is described in Table 1.

Table 1. Graduation and experience of the specialists that assisted the construction of the database

Specialists	Graduation	Post-Graduation	Time since Graduation
1	Nursing	Doctorate – Bioethics	12 years
2	Nursing	Doctorate – Bioethics	27 years
3	Nursing	Doctorate – Nursing	22 years
4	Nursing	Masters – Bioethics	8 years

Source: Research data

(Bio)ethical Problem: The bioethical problem addressed in the questionnaire, which has been later distributed to the four specialists, is described by Situation 1:

Situation 1*: An adult cancer patient, with no therapeutic possibility of cure, needs sedation due to severe pain. The drug is properly prescribed. The nurse and doctor, both members of the multiprofessional care team, have a dialogue and come to the conclusion that the dose to be administered, within the proposed time interval, may result in respiratory depression, with the possible risk of death.*

Below, each question elaborated by the experts is listed and explained in more detail.

Question 1: Is the administration of the medication the patient's wish? The main problem concerns the patient's own desire, who will exercise his autonomy either by verbalizing his will or through his advance directives of will (ADW), which, according to Dadalto [2014, p. 411], can be understood as: "the kind of desire for future medical care that a person wants to receive". The ADW is an "optional document and can be prepared, modified or revoked at any time in life" [CFM, 2012, p. 2], in which the patient is treated according to his wishes in relation to the treatment he wants to receive.

Question 2: Is the administration of the medication the wish of the patient's family? Regarding this aspect, the need

to appoint a representative to decide whether or not the patient is unable to answer for his/herself is inferred. Criteria such as "the patient's best interests" and "the patient's will must be really known or inferred from their values and concepts", such as "respect for the patient's previously expressed wishes, characterizing an important element" [BEAUCHAMP & CHILDRESS, 2013]. Such representatives can use three models: "(i) they substitute judgment when decisions are taken on behalf of the patient, thus assuming their will; (ii) pure autonomy, when the patient that previously expressed their will and in the given circumstances must be fully respected (previous guidelines); (iii) in the patient's best interest, when it is for the good or the charity of the subject [BEAUCHAMP & CHILDRESS, 2013].

Question 3: Is the administration of the medication the doctor's exclusive decision? Often, the medical opinion differs from that of other health professionals involved in patient care. Although there are consultations with nurses and other professionals in the face of certain clinical situations, the final decision is always made by the doctor.

Question 4: Is the administration of the medication the nurse's exclusive decision? Decision making is exclusive to nurses, because, although there is a medical prescription, nurses will be the one to decide on the patient's circumstances, not only because they are the professionals who accompany the patient's care, but also because they are the ones who administer medication.

Question 5: Is the administration of the medication necessary to relieve the patient's pain and suffering? This question assesses whether the needs of a patient, who is out of the therapeutic possibility of cure, are similar to those of a critically ill patient. According to Gomes [GOMES, 2010], "in general it is a progressive failure of the different systems, as they become weaker. The treatment in this phase of illness is aimed at the immediate control of the multiple symptoms presented by the patient that deteriorate their quality of life". Therefore, the proposal of comfort in palliative care is, in its essence, ethical "especially when regarding the terminality of life. It trails in conventional treatment, transgresses the aegis of the disease and offers itself to the transformation of assistance to people with illnesses", according to Burlá and Py [BURLÁ & PY, 2014].

Question 6: Is the administration of no medication due to the nurse's conscientious objection (that is, the professional refuses to administer the medication because of the possible risk of shortening the process of dying)? Professional autonomy, underlying consciousness, is understood as "a normative tool of professional codes and public policies that aims to protect the integrity of those involved in a situation of moral conflict" [DINIZ, 2007, p. 982]. As a result, a conscientious objection refers to a normative field of professions that a doctor or other health professional can use, on behalf of individual beliefs, to perform (or not perform) a specific procedure based on their moral principles.

Question 7: Will not the administration of any medication cause damage or harm the patient (non-maleficence)? Non-maleficence is understood as the duty of not intentionally causing harm to an individual, a maxim also attributed to Hippocrates [REZENDE, 2009]. In this way, we can interpret that harming someone's interests is not admitted. This principle

is closely related to discussions about the end-of-life process, regarding the distinction between killing and letting die and the moral relevance of this distinction [BEAUCHAMP & CHILDRESS, 2013]. Among the discussions on non-maleficence, we have the one related to abstention and withdrawal of life support.

Question 8: **Is there a judicial decision that makes the administration of the medication mandatory for maintaining the patient's life?** Calling upon the Judiciary Power for health demands has become common in recent years, this being a valid conduct in a democratic state of law. The fact that people seek justice to invoke their constitutional right to health (for medical treatments, for example) may indicate that the population is more aware of their rights and that the Judiciary Power is aware of the people's social rights, especially those of the most vulnerable groups [REZENDE, 2009]. Thus, the existence of a court order that makes the administration of certain medication necessary for maintaining the patient's life must be considered.

Ethical Aspects: The study has not been sent to the Research Ethics Committee (REC) for consideration due to the fact that it did not involve human beings as research subjects, in accordance with Resolution No. 466 of the National Health Council, of December 10, 2012, and since all the research steps have been carried out by the authors of the article.

Dataset: The dataset used in this study has been generated with the help of the four health professionals described in Table 1, who were given the task of deciding whether or not to administer medication to the hypothetical patient considered in the study. Thus, each professional has been provided with an electronic spreadsheet, where each question q_i, such that $1 \leq i \leq 8$, has the possibility of answering "Yes" or "No" (*i.e.*, q_i {"Yes", "No"}), thus totaling $2^8 = 256$ possible scenarios. Each C_n scenario is defined as $C_n = (q_{1n}, q_{2n}, ..., q_{8n})$, such that $1 \leq n \leq 256$. The 256 scenario combinations have been then presented to the volunteer professionals, who have been responsible for providing an answer r_n (such that r_n {'Yes', 'No'}), for each C_n scenario, where "Yes" indicates that the professionals would administer medication and "No", otherwise. The union of a combination C_n and its corresponding response r_n forms up a register R_n, *i.e.* $R_n = (C_n, r_n)$. As mentioned earlier, the questionnaire has been answered by four specialists, thus totaling 4 x 256 = 1024 entries. For this reason, all 1024 entries have been finally brought together into a single dataset, called R, such that $R = \{R_1, R_2, ..., R_m\}$ and $1 \leq m \leq 1024$. This has been done in order to facilitate the training processes and the evaluation of the ML algorithms, which are described in more detail in Section 3.

Data Preprocessing: After all of the entries have been brought together in a single data file, a categorical-numerical transformation has been carried out in R, where all occurrences with the value "Yes" of each R_m record have been replaced by the value "1", that is, $q_{im} \leftarrow$ "1" and $r_m \leftarrow$ "1" $q_{im} =$ "Yes" and $r_m =$ "Yes", respectively. In contrast, occurrences with a value of "No" received a value of "0". This transformation has been made necessary since the ML algorithms evaluated in this work require numerical input values during the learning and evaluation processes. Then, the C_m scenarios have been subjected to z-score normalization, in order to avoid saturation of the ML algorithms during training. The z-score

normalization is described by Equation 1, where μ represents the population mean and σ the standard deviation.

$$Z = \frac{X - \mu}{\sigma} \qquad (1)$$

After normalization, the next step in the preprocessing phase consisted of applying Principal Component Analysis (PCA) in the C_m scenarios.

PCA is a dimensionality reduction technique that aims to transform a set of possibly correlated variables (represented here by the eight questions) into an equal or lesser set of non-correlated variables, called main components [FACELLI *et al.*, 2011, p. 47]. Therefore, the objective of the PCA in this study was to try to eliminate possible redundancies between the eight questions in the questionnaire in order to simplify the process of identifying the patterns present in the data by the ML algorithms. With the application of the PCA in each C_m scenario, the preprocessed set of records $R = \{R_1, R_2, ..., R_m\}$ has been finally divided into two subsets: (i) a training subset T_r, randomly formed by 80% of the entries (*i.e.* 819 entries) which have been used to feed the training process with samples in order to adjust the internal parameters of the ML algorithms and calibration of hyperparameters; as well as (ii) the test subset Te, formed by the remaining 20% of entries (205 entries) which have been used to evaluate the generalization ability of the algorithms.

ML Algorithms: After completing the preprocessing step, the next step has focused on defining the ML algorithms that would be used in the experiments. For this, five classification algorithms have been chosen, namely: (i) Multilayer Perceptron Neural Network (MLP), (ii) Decision Tree (DT), (iii) Random Forest (RF), (iv) Support Vector Machines with RBF function (SVM-RBF) and (v) Logistic Regression (LR). It is important to mention two observations here: (i) in this work, a classification algorithm seeks, after the learning process, to simulate health professionals by associating a given scenario with an answer r {"0", "1"}, such that "0" means the class corresponding to the non-administration of the medication, and "1" the class corresponding to the administration of the medication; (ii) all five ML algorithms have been chosen because they are widely used in related works [FERNANDES et al;. 2017; SIQUEIRA-BATISTA & SILVA, 2019]. A brief explanation of each ML algorithm is detailed below. MLPs have an architecture inspired by the functionality of the biological brain and are usually used to solve more complex classification problems, *i.e.* composed of several attributes. A MLP model generally consists of: (i) an input layer made up of k attributes (where k represents the number of attributes), (ii) one or more inner layers (also known as hidden layers) and (iii) an output layer, usually formed by n neurons, where n represents the number of classes in the classification problem. An activation function is applied to each neuron, whose objective is to produce non-linear outputs and that are within a specific range [FACELI et al., 2011, p. 111]. MLP's learning is represented by its set of weights, which is adjusted with the aid of an optimization algorithm called backpropagation [RUMELHART et al., 1986]. This algorithm uses the computation of the gradient descent of the difference between the values predicted by the model with the expected values as responses to adjust the weights and to propagate the error to the inner layers. This adjustment process is repeated until a certain number of iterations or until the error reaches a minimum tolerance,

defined *a priori*. Decision trees, like the popular C4.5 trees, are classification algorithms where, unlike MLPs, it is possible to access the decision rules found by the model from the learning process. A decision tree usually uses the divide-and-conquer strategy to divide a complex decision problem into simpler ones, where the same strategy is recursively applied. In the end, the solutions of the subproblems are combined, in the form of a tree, to produce the solution of the original problem [FACELI et al., 2011, p. 83].

The random forests, in turn, are a classification method based on ensembles, where the results of several decision trees are combined and the final answer is usually obtained by majority vote. One of the premises of random forests is to fix the bad habit of overfitting (that is, overfitting the model to the training data, which does not allow for the generalization of the knowledge acquired for new samples) of the decision trees in training sets [HASTIE et al., 2009, p, 587].SVM's approach consists of establishing, for a linearly separable problem in the multidimensional space, a hyperplane with maximum margin between the border samples of each class as a way to prevent overfitting [FACELI et al., 2011, p. 122-123]. In cases where the data is not linearly separable, kernel functions, such as the Radial Basis Function (RBF), are used to map the data in a larger dimension in space, in order to make it separable from a hyperplane in a linear form. Finally, logistic regression is a statistical model that, unlike linear regression, uses the logistic function to calculate a probability. In classification problems, this probability is compared to a threshold defined *a priori* in order to assign a label that will represent the class in which it belongs to the sample provided.

Experiments and Results

ML Algorithm Training: With the definition of the ML algorithms, the learning process, also called adjustment or training, has been initialized, with the help of the *Tr* subset. In this subset, each ML algorithm has been adjusted using a process known as cross-validation with $k = 10$ partitions. In this approach, the subset *Tr* has been iterated 10 times; in each iteration, *Tr* has been divided into 10 partitions: one partition has been used for validation while the 9 remaining partitions have been unified into a single partition for training the ML algorithm. The performance of the algorithm in each iteration has been stored and, at the end of the 10 iterations, the average of the 10 performances has been computed. Figure 1 illustrates a hypothetical example of cross-validation with $k = 5$ partitions.
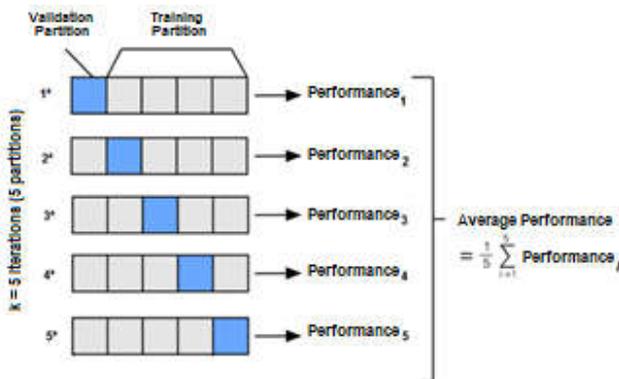

Figure 1. Illustration of the cross-validation process with $k = 5$ partitions.[1]

[1] Adapted from http://ethen8181.github.io/machine-

In order to evaluate the performance of each ML algorithm, in each of the 10 iterations, the accuracy metric (ACC), described by Equation 2, has been used.

$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \qquad (2)$$

In Equation 2, *TP* and *TN* mean, respectively, the number of true positives (entries of class '1' that the ML algorithm correctly classified as '1') and the number of true negatives (entries of class '0' that the ML algorithm correctly classified as '0'). On the other hand, *FP* and *FN* represent, respectively, the number of false positives (entries of class '0' that the ML algorithm wrongly classified as '1') and the number of false negatives (entries of class '1' that the ML algorithm wrongly classified as '0'). The ACC value calculated by this metric belongs to the range [0.1], where the closer to 1, the better. Table 2 shows the average accuracies of the five ML algorithms in 10 iterations of cross-validation in the *Tr* subset, accompanied by their respective standard deviations . It is worth mentioning that all the source code of the experiments, along with the values of the hyperparameters defined empirically for each learning algorithm, are available online.

Table 2. Performance of the ML algorithms in cross-validation with 10 iterations in subset *Tr*

ML Algorithm	$\overline{ACC} \pm \sigma$ (10 iterations)
MLP	0.843 ±0.025
DT	0.768 ± 0.048
RF	0.788 ± 0.040
SVM-RBF	0.856 ± 0.045
LR	0.848 ± 0.039

Source: Research data

By analyzing the results in Table 2, resulting from the cross-validation process in subset *Tr*, it is observed that all ML models had an average accuracy of above 76% with low instabilities, due to the small values of the standard deviations. It is also worth mentioning that the MLP, SVM-RBF and LR models achieved a performance rate of above 84%, with emphasis on the SVM-RBF, with the highest average accuracy of the ML algorithms: 85.6%.

Assessment of the Learning Algorithms: After the learning stage, each ML algorithm has been evaluated on the assessment stage. The assessment stage consisted of verifying whether the knowledge gained by the ML algorithms during training has been able to generalize, *i.e.* to show a good performance when classifying new entries. Therefore, the *Te* set has been used in the assessment stage, since it is formed by the entries that have not been observed by the ML algorithm during the learning stage. In the assessment stage, in addition to Accuracy, three other performance metrics have been used: (i) confusion matrix, (ii) F1 metric and (iii) ROC/AUC analysis. Each performance metric is explained below in more detail. The confusion matrix provides an overview of the performances of the ML algorithms by reporting the values of *TP, FP, FN* and *TN* in a matrix, organized according to Figure 2. From the confusion matrix of a binary problem, illustrated by Figure 2, it is possible to observe two metrics, in addition to Accuracy: Precision and Revocation. Precision can be seen as a measure of model accuracy, while Revocation can be seen as

learning/model_selection/model_selection.html. Access on July 14 of 2020.

Figure 2. Binary Confusion Matrix with some classification metrics.[2]
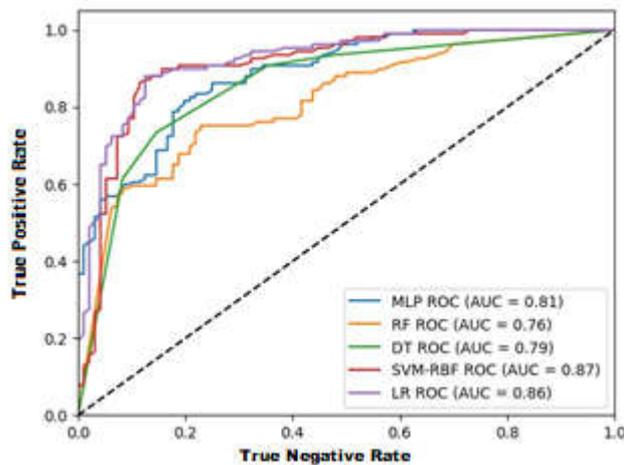

Figure 3. ROC graphs of the five ML algorithms in the *Te* subset

a measure of its completeness. An accuracy of 1.0 for a given class *C* means that each entry labeled by the ML algorithm as belonging to class *C*, actually belongs to class *C*. However, no information is provided regarding the number of class C entries that have not been correctly classified by the ML algorithm.

Table 3. Performance of the five ML algorithms in the *Te* subset

Model	TP	FP	FN	TN	ACC	F1	AUC
MLP	77	20	19	89	0,810	0,820	0,809
DT	82	29	14	80	0,790	0,788	0,794
RF	74	28	22	81	0,756	0,764	0,757
SVM-RBF	83	13	13	96	0,873	0,881	0,873
LR	78	11	18	98	0,859	0,871	0,856

Source: Research data

In contrast, a Revocation of 1.0 means that each class *C* entry has been correctly classified by the ML algorithm as belonging to class *C*, but nothing is reported about how many other entries have been incorrectly labeled by the ML algorithm for class *C* [FACELI et al., 2011, p. 165]. In this way, Precision and Revocation are usually combined and analyzed with the help of the F1 metric, computed according to Equation 3. It is worth mentioning that, similarly to the accuracy, the closer to 1.0 a F1 metric is, computed for a ML algorithm, the better is its performance.

$$F1 = \frac{2 * Precision * Revocation}{Precision + Revocation}$$

To complete the analysis, a ROC/AUC analysis has been performed. ROC (Receiving Operating Characteristics) is an

[2]Adapted from https://manisha-sirsat.blogspot.com/2019/04/confusion-matrix.html. Access on July 15, 2020.

analysis of binary classification problems using a two-dimensional graph, where the X and Y axes represent, respectively, the false positive rates (or the specificity complement) and the true positive rate (or revocation). ROC graphs are usually represented by a diagonal straight line, which represents random predictions, and by the ROC curve, which is formed from a set of *p* points that represent the specificity/revocation ratio with *p* threshold values, which can be used by ML algorithms in the classification of entries. The closer the ROC curve is to the point (0.1), the better the performance of the ML algorithm. In turn, the metric AUC (Area Under the Curve) corresponds to the area under the ROC curve. Similar to the F1 metric, the AUC metric provides an overview of the model's performance. The closer the AUC metric is to 1, the better the performance of the algorithm is considered. Next, Table 3 organizes the performances of the five ML algorithms in the *Te* subset according to the four metrics discussed: (i) confusion matrix (represented by the T*P, FP, FN,* and *TN* columns), (ii) accuracy, (iii) metric F1 and (iv) AUC metric.

Afterwards, the ROC graphs of the respective models are shown in Figure 3 and, finally, the results obtained will be discussed in more detail. When analyzing the results presented in Table 3, it is possible to observe that the SVM-RBF (line highlighted in blue) confirmed the partial numbers in the *Tr* set and has been, again, the best algorithm among the models tested in this work, presenting an accuracy and AUC of 0.873. It is also worth mentioning that the SVM-RBF model has shown the best result for the F1 metric among all five tested models. These promising results justify somehow the growing use of this ML learning technique in countless related works, where results are often comparable and even superior (like the results obtained in this work) to the results of MLPs [FACELI et al., 2011, p. 122]. Another important point to be mentioned here is the performance of both the MLP and RF algorithms. Due to the fact that these ML algorithms have more elaborate architectures than the other algorithms, they were expected to present the best results. However, the MLP algorithm achieved only the third best performance, while the RF model (line highlighted in red) presented the worst performance among the five algorithms, with a lower accuracy than the average accuracy presented in the *Tr* set (see the Table 2).

Two factors might have contributed to the results of both algorithms: (i) the calibration of the models might have empirically selected a non-optimal set of values for both algorithms, given the fact that a brute force scan, *i.e.* that assessed all possible combinations of hyperparameters for MLP and RF, would take too long to complete; (ii) the complex architecture of both algorithms might also have contributed to some kind of saturation, caused by some transformation in the values of the records used during training, negatively affecting the learning process of the models.

FINAL CONSIDERATIONS

Health professionals are often faced with sensitive and stressful issues when caring for critical or end-of-life patients in the ICU setting. Therefore, this study raised the hypothesis that artificial intelligence techniques could assist health professionals in making bioethical decisions. For this, a group of four specialists with solid training in the area of nursing has prepared and answered a questionnaire based on a hypothetical

ICU scenario. The data collected from the responses of the four specialists have been then used in the training and in the evaluation of five ML algorithms. The results obtained from the experiments, especially those regarding the SVM-RBF and LR algorithms, corroborated the hypothesis raised in this study, by showing that it is possible to adopt ML algorithms to support the decision making process of health professionals where bioethical issues are involved.

REFERENCES

BEAUCHAMP, Tom L.; CHILDRESS, James F. Principles of Biomedical Ethics. 7ª ed. New York: Oxford University Press; 2013.

BURLÁ, Claudia; PY, Lygia. Cuidados paliativos: ciência e proteção ao fim de vida. Caderno de Saúde Pública, Rio de Janeiro, 30(6):1-3, jun, 2014.

CONSELHO FEDERAL DE MEDICINA. Resolução CFM no 1.995, de 31 de agosto de 2012. Dispõe sobre as diretivas antecipadas de vontade dos pacientes. [Internet]. 2012 (acesso 20 mai. 2017). Disponível em: http://www.portalmedico.org.br/resolucoes/CFM/2012/1995_2012.pdf. Acesso em: 01 jul 2020

DADALTO, Luciana. Testamento vital. 3ª ed. São Paulo: Atlas; 2014

DINIZ, Debora. Objeção de consciência e aborto: direitos e deveres dos médicos na saúde pública. Revista de Saúde Pública. 2011 e esclarecido. Jornal do CRM-PB Nº 72, abr./jun. 2007.

FACELI, K., CARVALHO, A.., LORENA, A., GAMA, J. Inteligência Artificial: uma Abordagem de Aprendizado de Máquina. Rio de Janeiro: LTC, 2011.

FERNANDES, Daniel L.; SIQUEIRA-BATISTA, Rodrigo; GOMES, Andréia P. ; SOUZA, Camila R. ; COSTA, Israel T. ; CARDOSO, Felippe; ASSIS, João V. ; CAETANO, Gustavo H. L. ; CERQUEIRA, Fabio R. Investigation of the Visual Attention Role in Clinical Bioethics Decision-Making Using Machine Learning Algorithms. Procedia Computer Science, v. 108, p. 1165-1174, 2017

GOMES, Ana Margarida Rodrigues. O cuidador e o doente em fim de vida - família e/ou pessoa significativa. Rev Elet cuatrimestral de Enfermería. Nº 18, Fev, 2010.

HASTIE, Trevor; FRIEDMAN, Jerome; TIBSHIRANI, Robert. The Elements of Statistical Learning. 2ª Ed. New York: Springer Series in Statistics, 2009.

REZENDE, Joffre M. À sombra do plátano: crônicas de história da medicina [online]. São Paulo: Editora Unifesp, 2009. O juramento de Hipócrates. p. 31-48.

RUMELHART, David E.; HINTON, Geoffrey E.; WILLIAMS, Ronald J. Learning Representations by Back-Propagating Errors. Nature, v. 323, n. 6088, p. 533-536, 1986.

SIQUEIRA-BATISTA, Rodrigo. A Ressurreição de Frankenstein: Uma Metáfora das Unidades de Terapia Intensiva contemporâneas. In: SCHRAMM, F. R.; REGO, S.; BRAZ, M.; PALÁCIOS, M. Bioética, Riscos e Proteção. Rio de Janeiro: Editora UFRJ; Editora FIOCRUZ, 2009, v., p. 148-163

SIQUEIRA-BATISTA, Rodrigo; GOMES, Andréia Patrícia; MAIA, Polyana Mendes; COSTA, Israel Teoldo da; PAIVA, Alcione Oliveira de; CERQUEIRA, Fábio Ribeiro. Modelos de tomada de decisão em bioética clínica: apontamentos para a abordagem computacional. Revista Bioética, v. 22, n. 3, p. 456-61, 2014.

SIQUEIRA-BATISTA, Rodrigo; SILVA, Eugênio. Notas sobre os fundamentos matemáticos da Inteligência Artificial. Revista de Ciência, Tecnologia e Inovação, v. 4, n. 6, p. 44-54, 2019.

****