Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Cost-Sensitive Feature Selection in Medical Diagnostics: Integrating Fuzzy Logic for Optimal Accuracy and Resource Efficiency

Felice Franchini
University of Studies of Bari “Aldo Moro”
Giuseppe Pirlo
University of Studies of Bari “Aldo Moro”
Lucia Sarcinella
University of Studies of Bari “Aldo Moro”
Gianfranco Semeraro
gianfranco.semeraro@uniba.it

University of Studies of Bari “Aldo Moro”

Research Article

Keywords: Value Based Procurement (VBP), Public Healthcare System, Interpretable Machine Learning Models, Cost-Accuracy Trade-Off

Posted Date: May 7th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8268455/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Cost-Sensitive Feature Selection in Medical Diagnostics: Integrating Fuzzy Logic for Optimal Accuracy and Resource Efficiency

Felice Franchini[1](Orcid: 0009-0000-8887-800X), Giuseppe Pirlo [1](Orcid: 0000-0002-7305-2210), Lucia Sarcinella[1](Orcid: 0000-0002-8550-8588), Gianfranco Semeraro [1,*] (Orcid: 0000-0003-1666-8323)

[1] Department of Computer Science, University of Studies of Bari "Aldo Moro", Via Edoardo Orabona, 4, 70125 Bari, BA, Italy

* Correspondence: gianfranco.semeraro@uniba.it

Abstract: The increasing pressures on healthcare systems highlight the need for resource-efficient and cost-effective decision-making frameworks. This study investigates cost-sensitive feature selection for medical diagnostics, focusing on frailty prediction as a case study. Frailty, characterized by gradual and imprecise clinical changes, is well-suited for fuzzy logic integration, which manages uncertainties in medical data and enhances decision-making. Three methodologies (Combinatorial Optimization, Multi-Objective Optimization, and Mono-Objective Optimization) are evaluated to balance classification accuracy and resource costs. Combinatorial Optimization achieves high accuracy (>94%) by exhaustively exploring feature subsets, though at significant computational expense. Multi-Objective Optimization constructs a Pareto-Optimal Frontier, identifying feature subsets with high accuracy (91.9–95%) and varied costs (€120–396.16), offering a versatile decision-making framework for balancing priorities. Mono-Objective Optimization emphasizes computational efficiency and cost minimization, identifying the most affordable subset (€2.3) but with lower accuracy (63.98%). Results underscore the utility of fuzzy logic in defining and optimizing cost-accuracy trade-offs, enabling interpretable and resource-efficient diagnostic tools. Combinatorial and Multi-Objective approaches are suited for scenarios requiring high diagnostic precision, while Mono-Objective methods are appropriate where cost constraints dominate. This work demonstrates how integrating fuzzy principles into AI-driven frameworks enhances decision-making in cost-sensitive healthcare applications. Future research should incorporate temporal costs and explore broader applications in cost-sensitive domains. Refining these methods with advanced fuzzy reasoning can further optimize adaptability and applicability, supporting the development of sustainable, accurate, and interpretable healthcare technologies.

Keywords: Value Based Procurement (VBP); Public Healthcare System; Interpretable Machine Learning Models; Cost-Accuracy Trade-Off

Competing Interests: The authors declare that they have no direct or indirect financial or non-financial conflicts of interest that could influence the work presented in this manuscript.

1. Introduction

Pressures on healthcare systems, vividly exposed by the strains the National Health Service (NHS) put under during the COVID-19 pandemic, have placed immense emphasis on the need for a solid distribution of resources and an efficient decision-making process. Healthcare expenditure in the EU27 region has been increasing gradually since 2008, with a more marked acceleration between 2013 and 2019, which can be seen

in Figure 1 [1]. This situation calls for a major re-evaluation of current frameworks, incorporating fuzzy logic approaches, to ensure sustainable and efficient healthcare provisioning.

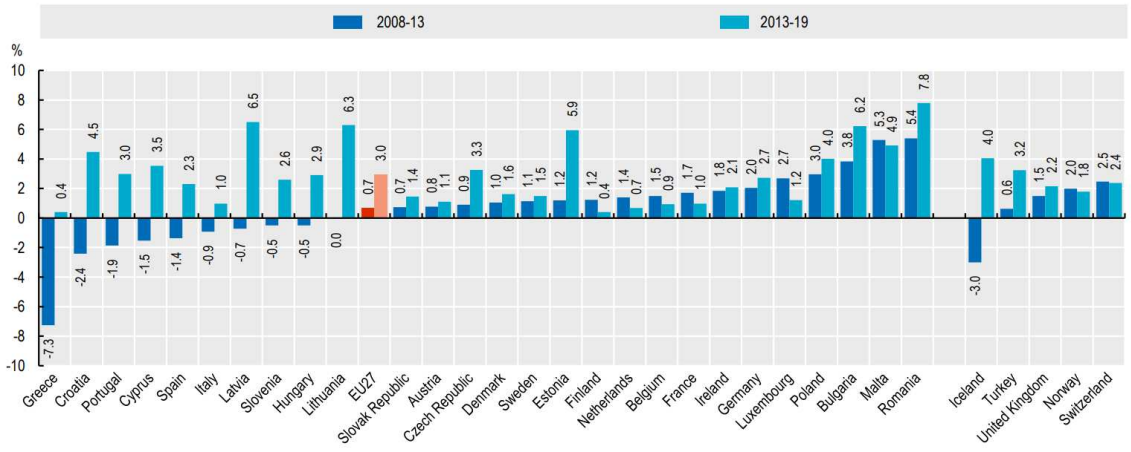


Figure 1 Annual average growth rate (real terms) in per capita health spending, 2008-2019(or nearest year)[1]

The rise in healthcare spending is attributed to a complex interplay of demographic, social, and economic factors, as elaborated in the fifth chapter of "Health at a Glance: Europe 2020" [1]. In reaction to these diverse challenges, value-based procurement (VBP) has surfaced as a novel method for healthcare acquisition, emphasizing quality and cost-effectiveness with a focus on patient results [2], [3]. The term "value" can in itself be imprecise, incorporating various measures that are both subjective and context dependent. Such a change in basic assumptions would essentially demand that the structure of health care be reexamined to meet access and quality-of-care standards, to be sure, while balancing within inevitably uncertain boundaries that relate to cost-effectiveness and diagnostic accuracy.

The present landscape is characterized by the exponential growth in medical data, driven by the continuous generation of digital records containing laboratory results, diagnostic reports, and medical diagnoses [4], [5] . This abundance of data opens up possibilities for the application of machine learning algorithms—possibly combined with fuzzy logic methods—in predictive modeling, especially within a medical diagnosis framework [6], [7], [8], [9]. The diagnostic problem is usually cast as a binary classification problem; whereby trained classifiers make predictions based on a predefined set of input features that describe medical procedures or tests. The features, however, come at a cost: monetary, temporal, or in patient burden. Furthermore, the "relevance" or "importance" of a particular characteristic can occur in different extents, leading to the application of fuzzy membership functions. Therefore, expert decision-making, possibly aided by fuzzy decision support systems, lies at the heart of determining the best feature set which optimizes the trade-off between acquisition cost and the inherently imprecise notion of predictive accuracy [10].

This research reviews the literature on such an optimization challenge in the medical domain with an especial focus on how fuzzy logic methods can be integrated in that optimization process. It is proposed that this problem has not been given ample attention with an overwhelming concern on classical disease diagnosis that does not explicitly take account of the medical costs involved or the intrinsic imprecision in the medical data. Consequently, the study utilizes a dataset from the IRCSS "De Bellis" in Castellana Grotte (BA), Italy, which includes laboratory tests, diagnostic tests, and diagnoses, to evaluate three optimization methodologies aimed at identifying feature subsets that achieve a favorable balance between accuracy and cost. The concept of a "favorable balance" itself can be formalized using fuzzy rules and membership functions. Specifically, Extreme Gradient Boosting (XGBoost), optionally combined with fuzzy feature selection methods, is used for cost-accuracy optimization.

The class variable is selected to be the patient's frailty status based on the potential advantages of predicting it accurately, such as lowering the risk of hospitalization, public expenditure, and increasing the survival rate among the elderly. The clinical concept of frailty usually shows gradual changes and does not have sharp limits, so it can be modeled by fuzzy logic. These considerations deal with accuracy and the monetary cost of medical tests, in line with which the performance measure and cost type fall. The intrinsic interpretability of the

XGBoost algorithm, improved by fuzzy rule extraction methods, is used in combination with the feature selection methods to decrease the dimensionality of the dataset and find good feature subsets for further investigation.

Although the size of the obtained feature subsets varies, as do the processing times, the results show the efficiency of the applied methods in finding a good trade-off between cost and accuracy. This paper presents a novel approach by applying the principles of Value-Based Payment (VBP) within the framework of interpretable machine learning models, specifically XGBoost. The interpretability of XGBoost can be further improved by representing the feature importance or decision boundaries using fuzzy sets. This allows for refinement of the search space for the three feature selection optimization techniques in order to achieve an optimal cost-accuracy trade-off within a healthcare context. While machine learning is typically applied to disease diagnosis [11], [12], [13], [14], this work extends the scope to explicitly consider the consideration of medical costs, thus filling a gap in existing literature. The choice of frailty prediction as the target variable further differentiates this study and points out the potential for a huge healthcare benefit. Moreover, the vagueness inherent in the definition and assessment of frailty points to the potential for fuzzy logic to make a more subtle and powerful contribution to its prediction.

The importance of this work extends beyond the specific healthcare application. Methodologies for cost-sensitive feature selection and optimization on a high-dimensional medical dataset provide insights valuable for AI-driven technology transfer. The problem of choosing the most informative yet cost-effective features in medical diagnosis has much in common with determining important factors for successful technology transfer, including relevant patents, market needs, and potential risks. The very assessment of "relevance" and "risk" in technology transfer is very often subjective and imprecise, the very conditions under which fuzzy logic comes into play as an applicable tool. Interpretable machine learning techniques, applied via XGBoost (possibly augmented with fuzzy rule-based systems) may provide an increase in the transparency and trust of technology transfer decisions, showing the factors that drive success or failure. Similarly, value-based procurement in healthcare runs parallel with a need for value-driven technology transfer, where the focus has shifted away from volumes of transactions to what truly counts—the real impact and societal benefit conveyed by the technology. Inherent in this, the definitions of "impact" and "benefit" are likely to be fuzzy and context dependent.

This work contributes to the special issue by showing, as a practical application of AI, incorporating fuzzy logic principles, in resource allocation and decision optimization, principles that can be applied directly to the effective and efficient transfer of technology. The methodologies explored in balancing cost and accuracy in feature selection can be adapted for the identification of key indicators and predictors of successful technology transfer outcomes. Fuzzy logic may be especially useful in dealing with the uncertainty and vagueness surrounding predictions of technology transfer success. The use of interpretable models, including fuzzy rule extraction, has to be in line with the need for transparency and explainability of AI-driven technology transfer processes for generating trust and facilitating adoption. Fuzzy rules can represent, for example, expert knowledge concerning conditions under which a certain technology transfer will succeed.

The following three sections outline the state-of-the-art in cost-sensitive learning, cost-sensitive feature selection, and feature selection categorization in order to provide the theoretical framework necessary for the applied methodologies. Section 3 subsequently describes the three optimization methodologies applied in this work (integrating fuzzy logic techniques into these methodologies) and Section 3.4 presents the used dataset. The results, discussion, and conclusions are given in Sections 4, 5, and 6, respectively, pointing out the implications for AI-driven technology transfer and explicitly dealing with the role of fuzzy logic in enhancing these processes. The ability to identify a minimal but most informative set of features using feature selection methods, comes to address the challenge directly for efficient knowledge representation and decision-making within the domain of technology transfer.

2. State of art review

2.1 Cost-sensitive learning

In the machine learning domain, the main goal is minimizing the error rate of predictions, with usually small emphasis on the difference between types of errors. In statistics, a type I error is a false positive while a type II error is a false negative thus, the basic application of machine learning algorithms, usually, assumes that all the errors have the same costs but, in real life the two types of error could have - and in many scenarios they have - a different cost. This consideration is easily express in the work of Sammut C. and Webb G. I. [2].

Therefore, when facing the costs of errors, this type of analysis is called cost-sensitive learning [3]. On the other hand, when the same cost for the two types of error is assumed, the analysis is called cost-insensitive learning.

In the context of cost-sensitive learning for binary classification, a cost matrix, denoted as C, plays a pivotal role in weighting errors and correct classifications. This matrix, characterized as a 2x2 matrix, is designed such that its rows represent the actual (ground-truth) labels, while its columns correspond to the predicted classification labels. The element C(i,j) within this matrix signifies the cost associated with an error where the classifier erroneously assigns label j when the actual label is i. Conversely, the diagonal elements C(i,i) represent the correct classifications, essentially quantifying the "benefit" of accurate predictions. This binary classification framework can be extrapolated to multi-class classification scenarios involving N classes, necessitating the expansion of the cost matrix to an N x N structure. The employment of such a cost matrix enables the execution of classification tasks, wherein samples are categorized into classes that yield the minimum associated cost, aligning with the principles of cost-sensitive learning [3].

One of the firsts works about cost-sensitive learning was made by Turney P.D. [4] where the authors introduced the “Inexpensive Classification with Expensive Tests” ICET, a cost-sensitive classifier. In that work, a genetic algorithm is used in combination with a decision tree. The fitness function is defined as the average cost of classification. To compute the average the tests’ costs and the classification errors’ costs were used [4].

Looking to recent results, Lijue Liu, Xiaoyu Wu et al [10] applied the cost-sensitive learning to face the challenge of class imbalance in a medical dataset. Lijue Liu, Xiaoyu et al [10] have shown that cost-sensitive learning can overcome the imbalanced dataset problem.

A similar focus is on the work of Fen Liu and Quan Qian [15] that propose the use of cost-sensitivity factors so that the overall misclassification cost is used as the optimization target with the aim of improving of the classification accuracy with respect to the minority class. The proposed method uses the domain knowledge to build the cost matrix and modify it by refining the cost-sensitive factors.

2.2. Cost-sensitive feature selection

Feature selection plays a crucial role in supervised learning across various applications, particularly in addressing Bellman's "curse of dimensionality." It facilitates the exploration of relationships between features and between features and classes, enhancing comprehension of the outcomes and the final model. Typically, machine learning algorithms operate under the presumption that the cost of generating or extracting each feature is uniform. However, this assumption may not hold true in certain domains, such as healthcare, where the cost of acquiring specific features can vary significantly, irrespective of their predictive capabilities.

Pawel Teisseyre and Tomasz Klonecki [16] proposed a new feature selection method based on the use of criteria to describe features relevance and the penalty cost. Moreover, they proposed a method to choose the optimal value of a cost factor that controls the trade-off between the features relevance and the penalty cost. The presented results in [16] show that their proposed approach is superior to existing traditional feature selection methods when there is a limited budget.

Thus, cost-sensitive feature selection is more about the selection of sub-set of features that provides the best trade-off between costs and accuracy. To cope with this goal, Hong Zhao et al. [17] tackled the cost-sensitive feature selection problem applying a new data model using backtracking and heuristics algorithms.

With the same perspective Arun Verma et al. [18] presented the use of Thompson Sampling (TS) and Upper Confidence Bounding (UCB) index-based algorithms to solve the combinatorial problem of subset selection that minimizes the total cost.

Mai Abdulla and Mohammad T. Khasawneh presented in their work [19] a new ensemble method to perform the cost-sensitive feature selection. In particular, they presented the G-Forest algorithm to face the high dimensionality, the rare variance and the costs to extract genomes features. In this way they built an algorithm that embedded the cost of experimentation (to obtain genomes features) in the process of optimization. The results presented show higher accuracy and lower costs for the selected sub-set.

Syed Imran Ali et al. applied the cost-sensitive feature selection to chronic kidney disease (CDK) diagnosis [20] since it is an expensive condition requiring special treatments. The authors have proposed two distinct methodologies for cost-sensitive feature ranking, both of which rely on the implementation of an ensemble of decision tree models. Additionally, the utilization of an automated threshold selection heuristic, which is predicated on balancing the value of the features against their associated costs, is a notable aspect of this approach. The outcome of this process is the derivation of a more compact subset of features, which contributes to a reduction in costs. Consequently, these subsets are deemed both practical and cost-efficient.

2.3. Feature selection categorization

The rapid growth of data in some domains, including healthcare, involves huge amounts of high-dimensional data. This proposes challenges for their effective and efficient management, involving the application of data mining and machine learning techniques to automatically discover knowledge.

When data mining and machine learning algorithms are applied to high-dimensional data, a critical issue is known as the Bellman's curse of dimensionality [21]. In addition, with a large number of features, learning models tend to overfit, which can cause performance degradation on data not yet seen. In addition to this, highly dimensional data can significantly increase memory requirements and computational costs for data analysis.[22]

Dimensionality reduction methodologies, encompassing Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), and Independent Component Analysis (ICA), present robust solutions for addressing the outlined challenges. An alternative approach involves conducting feature selection, which, by maintaining a subset of original features, enhances the models' interpretability and readability. Given the significance of these attributes in assessing the cost-accuracy balance, feature selection is the focal technique in this research.

Feature selection algorithms are conventionally classified into two primary categories based on their selection strategies: filter, wrapper, and embedded methods. Filter methods function independently of learning algorithms, relying solely on data characteristics to assess feature significance. Despite their computational efficiency, these methods may not yield optimal features for the intended learning algorithm due to the absence of algorithm-specific guidance in the selection process. Typically, filter methods involve a two-stage procedure: initially, features are ranked based on specific evaluation criteria, which can be univariate (individual assessment of each feature) or multivariate (collective assessment). Subsequently, features at the lower end of the ranking are eliminated. Over time, various criteria for filter methods have emerged, including feature correlation, mutual information, and the capability of features to reconstruct original data.[23], [24], [25]

Wrapper methods rely on a chosen learning algorithm's performance or predictive performance to assess the quality of selected features. Given a specific learning algorithm, a typical wrapper method performs two steps: it searches for a subset of features and evaluates the selected ones. These steps are repeated until some stopping criteria is satisfied. Thus, the search algorithm first generates a subset of features, then the machine learning model acts as a black box to evaluate the quality of these features based on their performance.

For example, the whole process works iteratively until the highest learning performance is achieved or the desired number of selected features is obtained. Then, the subset of features that provides the highest learning performance is returned as a single subset. A known problem with wrapper methods is that the search space of a d-dimensional task is 2d, thus a very large number for high dimensional data. Therefore, various search strategies such as sequential search, hill-climbing search, best-first search, branch-and-bound search and genetic

algorithms are proposed to produce optimal local learning performance. However, the search space is still extremely large for high-dimensional datasets, and therefore wrapper methods are not often used in practice [25], [26], [27].

Hybrid methods, sometimes confused with embedded methods or considered as one of the possible categories of embedded methods, can be viewed as a combination of multiple feature selection algorithms. The main goal is to address the instability problems of many feature selection algorithms. For example, for highly dimensional data, but with a limited number of instances, a small perturbation of the training data can lead to totally different feature selection results. By aggregating multiple subsets of features selected by different methods, the results are more robust and thus the credibility of the selected features is better [28], [29].

Features may also exhibit certain types of structures, examples include group structures, tree structures, and graph structures. Incorporating prior knowledge of feature structures can help select the most relevant ones to greatly improve learning performance. When structural information is not taken into account during feature selection, the inherent dependencies may not be captured and thus the selected features may not be suitable for the target application [30].

The categorization provided by Li et al. in [31] provides a decline of feature selection algorithms based on the data perspective. This taxonomy is presented in Figure 2.

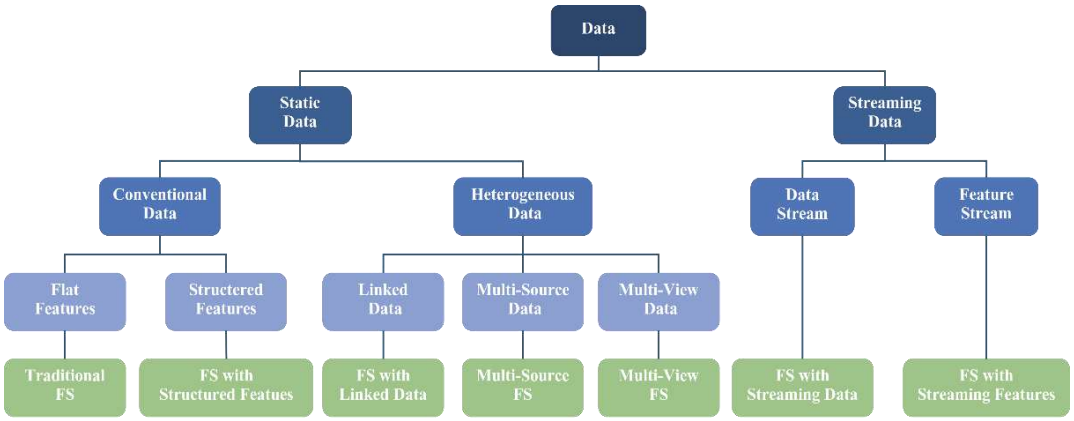


Figure 2 Feature selection algorithms from the data perspective

3. Methods and Materials

This section presents three novel fuzzy feature selection methods integrated with cost-sensitivity analysis. The proposed framework aims to optimize the transfer of healthcare technologies by providing an effective trade-off between classification performance and resource costs, crucial factors in successful technology transfer initiatives. The core classification algorithm employs Extreme Gradient Boosting (XGBoost) enhanced with fuzzy rule extraction capabilities. This choice is motivated by XGBoost's superior classification accuracy, inherent feature ranking capabilities, and interpretability confirmed by the state of art . Fuzzy's logic was implemented ensuring an enhancement for better handling of uncertainties inherent in features selection of medical data [32].

The target variable in our healthcare technology transfer study focuses on frailty level prediction, represented as a fuzzy binary classification problem (frail versus non-frail subjects). This application demonstrates significant technology transfer value, as preventing frailty conditions increases survival rates threefold for patients over 65 years and improves Quality Adjusted Life Years (QALY) by 0.3. QALY serves as a fuzzy weighted metric in evaluating life expectancy improvements related to health interventions. The successful transfer of this technology can substantially reduce healthcare system costs and decrease hospitalization risks and severe disability occurrence by up to six times [27].

In the technology transfer context, two critical cost dimensions are considered: monetary and temporal. The

temporal dimension represents access time to specific treatments or diagnostics. For instance, in neurodegenerative disease diagnosis, the timely implementation of diagnostic technologies significantly impacts patient hospital journey and treatment outcomes.

Following Turney's definition [4], each medical test (feature) in our framework carries different monetary costs while maintaining cost consistency across patients. This standardization facilitates technology transfer across different healthcare institutions.

The following sections present three hybrid methods for fuzzy cost-sensitive feature selection, specifically designed for healthcare technology transfer applications.

3.1 First method: Combinatorial Optimization

The first method introduces a fuzzy combinatorial optimization approach, merging sequential methods with exponential complexity algorithms. This method adapts the Sequential Backward Selection (SBS) algorithm to construct an exponential complexity algorithm [28], incorporating fuzzy decision boundaries for improved technology transfer applicability.

The base implementation of the SBS algorithm is similar to that presented in the book of Raschka S and Mirjalili V [33].The base algorithm consists of 4 steps:

- d is the dimensionality of the feature space Xd and k it's initialized as k=d;
- determine the feature x that maximize the formula $\bar{x} = argmax(J(X_k - x))$ where x ∈ X_k
- once the x is found, x is removed from the features set, so $X_{(k-1)} = X^k - \bar{x}$and
- If k is equal to the desired number of features, then the algorithm stops; otherwise, it returns to step 2.

It's clear that the algorithm needs, before the beginning of the computation, the definition of the criterion to maximize the desired number of features.

This combinatorial method does not choose the best subsets but computes all the subsets for each dimension k in order to compute the performance and the costs of each subset. The computation of each subset is stopped when k reaches the desired minimum number of features for each subset. The pseudocode of the first method is shown in Table 1.

The outputs of the algorithm are the subsets and their relative scores of performances and costs. In this way, this method allows the user to choose among several subsets and let the decision maker choose the one that better fits his/her needs.

Table 1 First Method Pseudocode

First method: Combinatorial Optimization
Input: *TRD* – Training Dataset; *TED* – Test Dataset; *classifies* – Choosed classifier; *measure* – Choosed performance measure; *k_features*- minimum number of features for a subset; *costs* – features costs.
Output: *all_subsets* – all the subsets; *scores_subsets* – all the computed scores by the performance measures for each subsets; *costs_subsest* – all the total costs for each subsets
COMBINATORIAL_OPT(TRD, TED, classifier, measure, k_features, costs):
1. *dim* := number of features of *TRD*
2. *indices* := indices of the features of *TRD*
3. *all_subsets* <- [*indices*]
4. *score* <- score_calculation(*TRD,TED, classifier, measure, indices*)
5. *scores_subsets* <- [score]
6. *cost* <- cost_calculation(*indices*, *costs*)
7. *costs_subsets* <- [*cost*]
8. **while** (*dim* > *k_features*) **do**

```
        for (c in combinantions(indices, r = dim-1)) do
            score <- score_calculation(TRD,TED, classifier, measure, c)
            cost <- cost_calculation(c, costs)
            scores_subsets <- scores_subsets + score
            costs_subsets <- costs_subsets + cost
            all_subsets <- all_subsets + c
        end for
        dim = dim – 1
    end while
    return all_subsets, scores_subsets, costs_subsets
```

3.2 Second method: Multi-Objective Optimization

In the realm of technology transfer, the evolution of methodologies for selecting features necessitates the incorporation of sophisticated techniques capable of addressing the inherent ambiguities and uncertainties present in practical applications. This discourse introduces an augmented iteration of the methodology initially proposed by McCrae [30] which integrates the principles of fuzzy logic to navigate the intricate complexities associated with contemporary technology transfer endeavors. The advanced methodology, referred as "Fuzzy-Enhanced McCrae Method", is based on four foundational postulates, which have been expanded to assimilate the principles of fuzzy set theory, thereby enhancing the adaptability of the methodology within the technology transfer domain.

The initial postulate introduces a fuzzy threshold, denoted by λ, which is conceptualized as a fuzzy numeric entity accompanied by a membership function. The threshold serves to highlight the model's performance into area of acceptability and non-acceptability. Performance metric, often quantified as classification accuracy, is usually compared to this fuzzy threshold, and the degree of its acceptability is measured by the previously mentioned membership function.

The second postulate confirms McCrae's original hypothesis regarding the selection of subsets of characteristics, but extends this concept through the use of fuzzy subset theory. It suggests that the pinnacle of performance is approached through the utilization of the comprehensive feature set, with a fuzzy degree of optimality. As informative features are excised, the resultant performance diminution is articulated through a fuzzy membership function, offering a more granular insight into the significance of features within technology transfer scenarios.

The third postulate codifies the delineation of a classifier. Consider a classifier D that employs a machine learning method m on a feature set F. The criterion for acceptability is encapsulated by the following expression:

$$\mu_{acceptable}(D(F,m)) = \mu_{\lambda}(A(D(F,m))) \quad \textbf{(1)}$$

In this formulation, $\mu_{acceptable}$ is the membership function for acceptability, while μ_{λ} corresponds to the membership function linked to the fuzzy threshold λ. This fuzzy construct provides a more elastic framework for technology transfer applications, where traditional binary acceptability criteria may be unduly restrictive.

The fourth postulate addresses the ranking of features via a fuzzy-weight adaptation of the Extreme Gradient Boosting technique. Here, the conventional Gini index is refined with the introduction of fuzzy weights, thereby providing a superior mechanism for capturing the intrinsic uncertainty surrounding the importance of features in technology transfer settings. This refinement bolsters the method's capacity to contend with the nuanced demands of knowledge transfer within specific domains.

These postulates culminate in the conceptualization of the Fuzzy Acceptable Feature Set (FAFS), encapsulated by the set {F*, $\mu_{acceptable}$ (DF ∗), A(DF*), C(DF*)}. In this representation $\mu_{acceptable}$ (DF ∗) signifies the degree of acceptability, A(DF*) indicates the accuracy, and C(DF*) denotes the associated costs.

The original objectives set by McCrae are maintained but further refined through the application of fuzzy logic. The main objective is to identify all FAFS whose acceptability rating exceeds a predefined α level. This fuzzy logic-based approach allows greater freedom in selecting characteristics relevant to technology transfer applications.

The second objective addresses computational efficiency by using fuzzy prioritization of feature ordering. Three distinct ordering strategies—descending, random, and ascending—are evaluated in relation to fuzzy measures of feature importance. Empirical data suggest that descending order, when combined with fuzzy weights, produces an optimal balance between computational efficiency and maintenance of solution quality.

The third objective highlights the pragmatic advantages of delineating acceptable fuzzy models in the context of technology transfer. The integration of fuzzy set theory facilitates the creation of a more sophisticated cost accuracy landscape, leading to the formation of an optimal fuzzy Pareto frontier. This frontier embodies a continuum of non-dominated solutions, offering those engaged in technology transfer projects a wider range of selection options [31], [32].

The algorithmic framework, as initially outlined in Table 2, remains intact, although with modifications aimed at integrating fuzzy logic operations. The fuzzy augmentation primarily influences the evaluation of acceptability and the computation of feature importance, while the overarching computational structure is preserved.

Empirical studies within the domain of technology transfer have demonstrated that the fuzzy-augmented method confers several benefits:

- Enhanced resilience to the ambiguity inherent in the assessment of feature importance
- More pliable acceptability criteria tailored to the diverse needs of technology transfer scenarios
- Improved capacity to encapsulate the expertise of domain authorities through fuzzy rules
- Superior adaptation to the fluctuating cost-performance dynamics presents in various transfer environments

The fuzzy-augmented approach is particularly efficacious in technology transfer scenarios where conventional binary decision-making frameworks are found wanting. The method's proficiency in managing the transition between acceptable and non-acceptable solutions is well-aligned with the practical exigencies of technology transfer projects, which often involve stakeholders with disparate acceptability standards.

The relationship between the costs of characteristics and their manifestation on the fuzzy Pareto-Optimal frontier retains the characteristics outlined by McCrae [30], while exhibiting heightened responsiveness to incremental changes. The fuzzy paradigm offers a more authentic representation of the impact that cost variations exert on decisions related to feature selection in technology transfer contexts.

This refinement of McCrae's methodology through the incorporation of fuzzy set theory signifies a substantial leap forward in the domain of feature selection for technology transfer applications. It offers heightened flexibility and resilience, while simultaneously upholding computational efficiency. This approach equips practitioners with a sophisticated instrument for harmonizing considerations of performance and cost in the tangible world of technology transfer.

Table 2 Second Method Pseudocode

Second method: Multi-Objective optimization
Input: *TRD* – Training Dataset; *TED* – Test Dataset; *classifies* – Choosed classifier; *measure* – Choosed performance measure; *t*- value between 0 and 1 used to scale the accuracy computed on the entire feature set in order to define the threshold of acceptability; *costs* – features costs.
Output: *s_pareto_frontier* – subsets on the Pareto-Optimal Frontier
MULTI_OBJECTIVE_OPT(*TRD,TED,classifier, measure, t, costs*):
1. *subsets_a*, *scores_a* <- search(*TRD,TED,classifier,measure,t*)
2. *s_pareto_frontier* <- pareto_front(*subsets_a*, *scores_a*, *costs*)

```
3.  return s_pareto_frontier
search(TRD,TED,classifier, measure,t):
1.  subsets_acceptable <- []
2.  scores_acceptbale <- []
3.  subsets_analized <- []
4.  n_trained <- 0
5.  dim := number of features
6.  indices := indices of the features
7.  priorityQ <- [[indices,'']]
8.  all_features <- priorityQ[0]
9.  tot_score <- pattern_score(TRD,TED, classifier, measure, all_features[0])
10. min_score <- t*tot_score
11. while (priorityQ ≠ ) do
12.         n_trained = n_trained + 1
13.         this_subset <- priorityQ[0]
14.         subsets_analized <- subsets_analized + this_subset[0]
15.         priorityQ <- priorityQ – priorityQ[0]
16.         d_pattern := dimension(this_subset[0])
17.         if (d_pattern ≠ 0) then
18.             score_result <- pattern_score(TRD,TED,classifier, measure, this_subset[0])
19.             if (score_result ≥ min_score) then
20.                 subsets_acceptable <- subsets_acceptable + this_subset[0]
21.                 scores_acceptable <- scores_acceptable + score_result
22.                 priorityQ <- priorityQ + (successors (this_subset, dim))
23.             end if
24.         end if
25. end while
26. return subsets_acceptable, scores_acceptable
successors(this_subset, dim):
1.  Fs <- this_subset[0]
2.  fp <- this_subset[1]
3.  subsets <- []
4.  d_fs := dimension(Fs)
5.  if (d_fs = dim) then
6.          for (f in Fs) do
7.              app <- []
8.              for (i in Fs) do
9.                  if (i ≠ f) then
10.                     app <- app + i
11.                 end if
12.             end for
13.             subsets <- subsets + [app, f]
14.         end for
```

```
15. else
16.         for (f in Fs) do
17.                 if (f > fp) then
18.                         app <- []
19.                         for (i in Fs) do
20.                                 if (i ≠ f) then
21.                                         app <- app +i
22.                                 end if
23.                         end for
24.                         subsets <- subsets + [app, f]
25.                 end if
26.         end for
27. end if
28. return subsets
```

3.3 Third method: Mono-Objective Optimization

The third methodological approach under consideration employs a fuzzy-integrated strategy, expanding upon the pioneering research of Jiang, Kong, and Li [33], which tackles the challenge of multi-objective optimization within the context of Cost-Sensitive Feature Selection (CSFS). This strategy is of particular significance in the realm of technology transfer, where the concurrent objectives of enhancing classification precision and reducing the expenses associated with feature acquisition are paramount for the effective deployment of artificial intelligence (AI) systems.

The decision to adopt a hybrid methodological framework is informed by the seminal studies of Kohavi and John [34], which have evidenced considerable improvements in accuracy compared to conventional filter methodologies. This hybrid paradigm is especially pertinent in technology transfer contexts, as the strength of feature selection has a direct bearing on the successful cross-domain implementation of AI technologies.

In their innovative reformulation, Jiang et al. [33] recast the multi-objective optimization problem into a singular objective framework by integrating a fuzzy cost assessment component. This recasting simplifies the decision-making process in feature selection, thereby enhancing the efficacy of technology transfer. The proposed evaluation function is articulated as follows:

$$Merit(model, S) = ModelAccuracy - \sum(i = 1)^{k\,Ci} \quad (2)$$

Here, the Merit function quantifies the composite value of the feature subset S that has been selected. It encapsulates the disparity between the predictive accuracy of the model and the aggregate cost of the features included in S, with S encompassing k elements. The equation encapsulates the core principle that AI systems intended for real-world technology transfer must achieve a judicious equilibrium between superior classification accuracy and feasible cost parameters.

The framework introduced by Jiang and colleagues accommodates fuzzy decision thresholds within the feature selection process, thereby enabling a more sophisticated management of ambiguity inherent in practical technology transfer settings. This method is particularly advantageous when confronted with uncertain or incomplete knowledge regarding the costs associated with features, a common occurrence in the deployment of AI systems. Empirical evidence presented by Jiang et al. [33] substantiates the efficacy of this framework in assisting professionals to discern feature subsets that optimize the trade-off between classification efficacy and cost minimization.

Within the implementation framework, the "Max_stale" parameter is introduced to regulate the extent of backtracking, specifically the quantity of suboptimal nodes to be examined prior to reversion to prior states.

This parameter is critical in maintaining a balance between exploration and exploitation throughout the feature selection algorithm, a balance that is particularly crucial in technology transfer applications where computational efficiency is a priority.

The outcome of this methodology is the identification of an optimal set of features, thereby providing practitioners in the field of technology transfer with a solution that is both clear and actionable, one that carefully weighs the imperatives of accuracy and cost. This approach, therefore, significantly contributes to the seamless deployment of AI systems across various sectors, thereby advancing the overarching objectives of successful technology transfer.

3.4 Dataset

The used dataset is composed of clinical data related to patients of a public health structure. The dataset is composed of 83 features and 1929 instances (in this case patients). The features can be divided into 3 macro-categories: laboratory tests, diagnostic tests and diagnosis.

3.4.1 Laboratory tests

The "laboratory tests" are tests needed to draw a general overview on the health status of a patient. These types of tests are related to the analysis of biochemical components analyzable from the patients' blood.

The laboratory tests used were:

- Glucose tests[S/P/U/DU/LA]: these tests consist of the measurement of the glucose level in the blood in order to monitor glycemia and verify the chance that the patients have developed a form of diabetes or prediabetes.
- HDL cholesterol: this test is prescribed primarily for the purpose of assessing the risk of developing cardiovascular diseases or monitoring the effectiveness of any therapeutic treatments for an already established cardiovascular disease.

Table 3 Third Method Pseudocode

Third Method: Mono-Objective Optimization
Input: *TRD* – Training Dataset; *TED* – Test Dataset; *classifies* – Choosed classifier; *measure* – Choosed performance measure; *costs* – features costs; *max_stale* – number of maximum node, with no enhancement, to consider before to stop the research
Output: *best_subset* – best subset
MONO_OBJECTIVE_OPT(*TRD,TED,classifier, measure, costs, max_stale*):
1. *open* := []
2. *scores_o* := []
3. *closed* := []
4. *scores_c* := []
5. *best_subsest* := []
6. *best_merit* <-∞
7. *k* <- 0
8. *indices* := indices of the features
9. *dim* := number of features
10. *i* <- 0
11. **while** (*i* < *dim*) **do**
12. *f* <- [*indices*[*i*]]
13. *temp_merit* <- merit(*TRD*, *TED*, *classifier*, *measure*, *f*, *costs*)

```
        open <- open + {f}
        scores_o <- scores_o + temp_merit
        if (temp_merit > best_merit) then
            best_merit <- temp_merit
            best_subset <- f
        end inf
        i <- i+1
end while
while (k < max_stale) do
        add_judge <- false
        v <- arg_max(scores_o)
        v_o <- open[v]
        closed <- closed + v_o
        scores_c <- scores_c + scores_o[v]
        scores_o <- scores_o – scores_o[v]
        open <- open – v_o
        i <- 0
        while( i < dim) do
            f <- {[indices[i]]}
            if ((f  v_o)  ((v_o  f)  open)  ((v_o  f)  closed)) then
                temp_merit <- merit(TRD,TED,classifier, measure, [v_o  f], costs)
                open <- open + (v_o  f)
                scores_o <- scores_o + temp_merit
                if (temp_merit > best_merit) then
                    k <- 0
                    add_judge <- true
                    best_merit <- temp_merit
                    best_subset <- (v_o  f)
                end if
            end if
        i <- i + 1
        end while
        if (! add_judge) then
            k <- k +1
        end if
end while
return best_subset
```

- LDL cholesterol: high amounts of this cholesterol are responsible for plaques formation and the possibility of developing cardiovascular disease.
- Total cholesterol: this value belongs to a range in which it's considered fundamental to some physiological processes, but when the total cholesterol is higher than a maximum value, its presence can develop cardiovascular problems
- Triglycerides: monitoring this value is useful to assess the chance that patients can develop cardiovascular disease

- Blood count: it's a test related to the blood and with some components in the blood. In particular, in this test are calculated: the concentration of hemoglobin (HB), red blood cells(GR), white blood cells (GB), hematocrit (HCT), platelets (PLT), leukocyte formula (FL).
- Vitamin D: optimal value of Vitamin D is useful to prevent rickets and osteoporosis in elderly. But a higher value can cause headache, nausea, vomiting, hypertension, and other more serious diseases, including growth retardation.
- Hemoglobin (HB): this test is divided into [SG] and [LA] to detect serious diseases such as anemia, jaundice and risk of infection etc.
- Glycated Hemoglobin (HbA1c): this test is performed for retrospective evaluation of glycemic control and diagnosis of diabetes.
- Aspartate aminotransferase (AST or GOT) [S]: is performed to detect the presence of a liver lesion and/or as an aid to diagnosis, screening, and monitoring of hepatopathy.
- Alanine aminotransferase (ALT or GPT) [S/U]: is performed to assess liver damage and/or as an aid in the diagnosis, screening, and monitoring of hepatopathy.
- Gamma-glutamyl transferase (GGT): this test is performed to determine possible liver or biliary tract disease, to differentiate between liver and bone disease, and for screening and monitoring of alcohol abuse.
- Insulin [S]: is performed to assess pancreatic function, as an aid to the diagnosis of pancreatic cancer, and for the evaluation and optimization of insulin therapy in a patient with type 2 diabetes
- PCR: identifies the presence of inflammation due to pathogens and tissue damage, monitors acute attacks of chronic autoimmune diseases and response to therapeutic treatments
- Tumor Necrosis Factor (TNF); increased values can be found in all situations of inflammation, infectious diseases, sepsis, autoimmune diseases, and in organ transplant rejection
- Inflammatory Markers (IL6, TNFALfa); whose blood concentrations imply the correct or incorrect immune response in inflammatory diseases, infections, autoimmune disorders, cardiovascular diseases, and some cancers; as an aid to diagnosis and monitoring of patients with diabetes

3.4.2 Diagnostic tests

The "diagnostic tests" are tests needed to be able to distinguish diseases that have similar symptoms. With this type of test, a clinician can orient its diagnosis in a more accurate way.

The used diagnostic tests were:

- PTA (Audiometric Examination): this is the most frequent examination performed by an otolaryngologist, through which hearing ability is tested for the main frequencies perceptible by the human ear. Based on the speed and accuracy of the patient's responses, the audiogram is processed to assess the hearing threshold and, more specifically, the threshold of minimal audibility
- SSI/Vocal (Audiometric Examination): speech audiometry allows the so-called threshold of perception to be established. In this case, the patient must recognize words and phrases recorded or spoken by the audio metrist with different levels of intensity
- SBP/DBP (General Assessment): to obtain a diagnosis of hypertension or hypotension, it is necessary to measure blood pressure defined by systolic blood pressure (SBP) and diastolic blood pressure (DBP) values
- BMI - Body Mass Index (General Assessment): Body Mass Index (BMI) is a parameter used to obtain a general assessment of one's body weight, not distinguishing, however, weight changes due to the amount of muscle and bone mass, water content or fat accumulation
- MMSE/FAB/Apathy (Clinical Neuropsychological Examination); the Mini Mental State Evaluation (MMSE) is a frequently used test for an initial assessment of cognitive function, and

the Frontal Assessment Battery (FAB) is a screening battery of global executive function using a series of cognitive and behavioral tests

- Apathy is a test designed to assess an emotional-affective condition characterized by an attitude of indifference toward people and reality
- MNA/SPPB/Other Specific Tests (Geriatric Examination); the Mini Nutrional Assessment (MNA) is a monitoring and assessment tool designed for the purpose of identifying elderly patients at risk of malnutrition, while the Short Physical Performance Battery (SPPB) is a short battery of tests created to assess lower limb function in elderly patients. Other tests are specific for the assessment of balance, mobility and postural stability
- Fatty Liver Index (FLI): test that helps predict the histologic severity of chronic hepatopathy by providing information on the likelihood of having steatosis
- Fibrois-4 (FBI-4): is a noninvasive marker to diagnose the evolution of liver fibrosis into cirrhosis

3.3.3 Diagnosis

The "diagnosis" is a medical judgment recognizing a chronic condition regarding laboratory and instrumental research.

The diagnoses used were:

- Mild Cognitive Impairment (Neurological Assessment); Mild Cognitive Impairment (MCI) represents the transitional state between normal aging and Alzheimer's disease and is characterized by mild acquired cognitive deficits without significant decline in the activities of the subject's working, social, and daily life
- Tinnitus/Age Related Hearing Loss/Central Auditory Processing Disorder (ENT Evaluation); Tinnitus is a phantom noise that the patient hears but is not present in the external environment. Age-Related Hearing Loss (ARHL) or Presbycusis is a deficit related to the irreversible aging process that recognizes multifactorial pathogenesis. Central Auditory Processing Disorders (CAPD) are a complex group of symptoms that describe difficulty in integrating auditory information as in using language and following or remembering instructions
- Depression: is a term used to refer to the presence of a sad, empty, or irritable mood accompanied by physical, physiological, and cognitive changes that significantly affect the individual's ability to function
- Diabetes: a condition that detects an increase in blood glucose levels due to a deficit in the amount and, often, in the biological effectiveness of insulin, the hormone that controls blood glucose and is produced by the pancreas
- Hypertension: assessment of cardiovascular disease through identification of risk factors, early diagnosis, treatment of the acute phase of disease, and secondary prevention, also making use of advanced methods
- Dementia: is characterized by the presence of a memory deficit that is associated with disturbances in other cognitive areas and causes a significant reduction in the patient's ability to live daily life
- APOe: support used in the diagnosis of late-onset Alzheimer's disease (AD) in symptomatic adults with signs and symptoms of progressive dementia with suspected AD
- BCPO; chronic obstructive pulmonary disease is a narrowing, occlusion, or obstruction, persistent airway, due to emphysema, chronic bronchitis, or both conditions
- Frail: a general frail condition of the elderly patient, detected by the physician after a thorough geriatric evaluation.

3.4.4 Costs

Costs associated with each feature were determined by querying the public healthcare database system of the Apulia Region and are expressed in euros. Costs are intended for the public administration and not for the patient.

4. Experimental setup and results

For preprocessing, it has been decided to eliminate the features whose cost was not communicated. In this way the number of features changed from 83 to 67.

In order to determine the true dimensionality of the feature set on which apply the three optimization methods, the technique presented by Minka in [34] has been used. This technique allows to choose a value in range [0,1] and return the features set whose sum of cumulative importance is within an established range. For this work a threshold value of 0.55 was chosen and it has been obtained 10 features instead of 67. Then features were ordered in descending order.

After this preprocessing step, the final feature set belongs to two categories: "laboratory tests" and "diagnostic test". Thus, no features from the "diagnosis" category have been automatically selected for the specific problem of frailty prediction. The selected features are presented in Table 4.

Table 4 Selected features from implemented method

Importance Order	"t._sedia" - "gait_speed_5mt" - "salita_gradino" - "dbp" - "piastrine.y" - "SSI_mean" - "mini nutritional assessment deficit" - "GPT+BI1 glutammato Triticum" - "Vocale60dx" - "fattore necrosi tumorale TNFa"	
Category	**Features Names**	**Costs**
Laboratory test: Blood count (Platelets - PLT)	"piastrine.y" -	1,17 €
Laboratory test: Tumor Necrosis Factor (TNF)	"fattore necrosi tumorale TNFa"	198 €
Laboratory Test: - Alanine aminotransferase (ALT or **GPT**) [S/U]	"GPT+BI1 glutammato Triticum"	1.04 €
Diagnostic test: **MNA**/SPPB/Other Specific Tests (Geriatric Examination)	"mini nutritional assessment deficit"	65 €
Diagnostic test: SSI/Vocal (Audiometric Examination)	"SSI_mean" – "Vocale60dx" –	38 €
Diagnostic test: SBP/**DBP(diastolic blood pressure)** (General Assessment)	"dbp"	2,30 €
Diagnostic test: MNA/**SPPB**/**Other Specific Tests** (Geriatric Examination)	"t._sedia" - "gait_speed_5mt" - "salita_gradino"	60€

4.1 Results of First Method: Combinatorial Optimization

The experiment was performed with the parameter "k_features" set to 1 being the minimum number of selectable features. With this value, this method returns 2n-1 subsets of features, the set of parts of the 10 features excluding the empty set. Thus, 2n-1 subsets mean in our case 1023 subsets and for each one the cost and the classification accuracy must be computed. In this way the user can choose the right trade-off. The top 19 subsets with best accuracy are showed in Table 5.

The method took a processing time of 1h 32min 32s.

Table 5 Top 19 - Subsets with best accuracy found by method 1.

Features Name	Costs	Accuracy
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "mini nutritional assessment deficit" - "Vitamina D”	502,33 €	0,95077720 2
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "Vitamina D”	437,33 €	0,95077720 2
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "Vitamina D”	396,16 €	0,95077720 2
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "Vocale60sx" - "Vitamina D”	475,33 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "SSI_dx_0" - "Vocale60sx" - "mini nutritional assessment deficit" - "Vitamina D”	537,16 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "Vocale60sx”	459,47 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0”	421,47 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "Vitamina D”	399,33 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y”	383,47 €	0,94818652 8
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "mini nutritional assessment deficit”	486,47 €	0,94559585 5
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "SSI_dx_0" - "Vocale60sx" - "Vitamina D”	472,16 €	0,94559585 5
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "piastrine,y" - "SSI_dx_0" - "Vocale60sx" - "Vitamina D”	277,33 €	0,94559585 5
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "SSI_dx_0" - "Vocale60sx”	456,3 €	0,94559585 5

“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "fattore necrosi tumorale TNFa" - "Vocale60sx" - "Vitamina D”	434,16 €	0,945595855
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "piastrine,y" - "SSI_dx_0" - "Vitamina D”	239,33 €	0,945595855
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "fattore necrosi tumorale TNFa" - "piastrine,y" - "SSI_dx_0" - "mini nutritional assessment deficit”	484,17 €	0,945595855
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp" - "Vocale60sx" - "mini nutritional assessment deficit”	285,3 €	0,945595855
“gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "piastrine,y" - "SSI_dx_0" - "Vocale60sx”	259,17 €	0,945595855
gait_speed_5mt" - "t,_sedia" - "salita_gradino" - "dbp”	182,3 €	0,945595855

4.2 Results of Second Method: Multi-Objective Optimization

According to the explanation of the method provided in Section 3.2, the costs and accuracy values are combined to compute a Pareto-Optimal Frontier in order to exploit the combination of high accuracy and low cost. The method took a processing time of 1h 11min 46s. In total there are 4 optimal subsets as presented in Table 6. The appropriate subset is selected by the expert based on their experience and specific requirements.

Table 6 Optimal Subsets found by method 2.

Features Name	Accuracy	Costs
“t._sedia” - “gait_speed_5mt”	0.91968911917	120 €
“t._sedia” - “gait_speed_5mt” – “Vitamine D”	0.92746113989	135.86 €
“t._sedia” - “gait_speed_5mt” – “Vitamine D” – “dbp” – “Vocale60dx”	0.93782383419	176.16 €
“t._sedia” - “gait_speed_5mt” – “salita_gradino” – “dbp”	0.94559585492	182.3 €
“t._sedia” - “gait_speed_5mt” – “salita_gradino” – “dbp” – “fattore necrosi tumorale TNFa” – “piastrine.y”	0.94818552849	383.47 €
“t._sedia” - “gait_speed_5mt” – “salita_gradino” – “dbp” –	0.95077720207	396.16 €

"fattore necrosi tumorale TNFa"

– "Vitamina D"

4.3 Results of Third Method: Mono-Objective Optimization

For this method the parameter "max_stale" was set to 5, it means that after 5 iterations with no improvement, the search strategy stops. This value is considered from Jiang L. et al [35] as the standard threshold.

The output of this method is the subset with the higher accuracy and the lower cost. The subset is show in Table 7. With this method, the user is not asked to choose, because the algorithm just picks the subset having the best cost-accuracy trade off. The method took a processing time of 4min 35s.

Table 7 Best subset from the third method

Feature Name	Accuracy	Cost
"dbp"	0.639896373	2.3 €

5. Discussion

This study analyzes the application of three distinct methodologies for feature selection (Combinatorial, Multi-Objective, and Mono-Objective Optimization) within a paradigm that is acutely sensitive to the costs associated with healthcare operations. The central objective of this research is to discern an optimal balance between the efficacy of classifying patient data and the financial and temporal expenses incurred by the chosen feature subsets. An exhaustive comparison of these strategies is presented, elucidating their respective merits and demerits, and exploring their potential impact on the translation of artificial intelligence (AI) technologies into fuzzy systems, with a particular focus on medical applications. The implementation of the Combinatorial Optimization method has led to the delineation of the foremost 19 feature subsets, each demonstrating classification accuracies that surpass 94%. This method, predicated on the Sequential Backward Selection (SBS) heuristic, rigorously assesses the entire spectrum of feature combinations by iteratively excising the least influential feature, as dictated by a specified performance criterion. The consistent emergence of specific attributes, namely "gait_speed_5mt," "t._sedia," "salita_gradino," "dbp," "tumor necrosis factor TNFa," and "Vitamin D", within these elite subsets underscores their indispensable contribution to the prediction of frailty. Despite the high levels of accuracy achieved, this approach necessitates a substantial investment of computational resources, as evidenced by the protracted processing duration averaging one hour and thirty-two minutes, which may pose challenges in environments characterized by expansive datasets or the need for expeditious data processing. The Multi-Objective Optimization framework, drawing upon the theoretical underpinnings of McCrae's model, seeks to construct the Pareto-Optimal Frontier as a means to equilibrate the demands of accuracy with those of cost. Through this lens, six optimal feature subsets were identified, each embodying a unique balance between superior accuracy, ranging from roughly 91.9% to 95%, and a spectrum of associated costs, from €120 to €396.16. The delineation of the Pareto-Optimal Frontier serves as a critical decision-making tool, empowering stakeholders to select feature subsets that best align with their specific preferences at the intersection of cost and performance. In particular, the processing time for this method was observed to be relatively expeditious, with an average of approximately one hour and eleven minutes. The strategic selection of subsets, including those such as "t._sedia" and "gait_speed_5mt," that maintain high accuracy while reducing costs, attests to the method's proficiency in resource optimization without a significant diminution of predictive capacity. The Mono-Objective Optimization strategy, informed by the research of Jiang et al., employs a composite merit function that concurrently evaluates the accuracy of classification and

the attendant costs of feature inclusion. This tactic yielded the selection of a solitary feature subset, “dbp,” which achieved a classification accuracy of 63.98% alongside a minimal cost outlay of €2.3. Although the computational efficiency of this approach is laudable, with a processing time of just four minutes and thirty-five seconds, the stark reduction in accuracy, when juxtaposed with alternative methodologies, is conspicuous. This discrepancy points to a critical trade-off inherent in the Mono-Objective Optimization method: while it excels in cost minimization, it may inadvertently compromise the fidelity of the predictive models, rendering it less optimal for contexts where precision in diagnosis is of paramount importance. The principal advantage of the Combinatorial Optimization method resides in its exhaustive analysis of feature subsets, which ensures the discovery of models with exceptional accuracy. However, this comprehensive search incurs a significant computational burden, which may impede its application in large-scale or time-constrained environments. Additionally, the method's reliance on a priori defined performance metrics necessitates careful calibration to avert the inadvertent omission of potentially valuable variables. The Multi-Objective Optimization approach adeptly manages the tension between classification accuracy and cost considerations, offering a suite of feature subsets that accommodate a diversity of stakeholder priorities. The utilization of the Pareto-Optimal Frontier is particularly beneficial in the healthcare domain, where the allocation of resources must judiciously balance fiscal prudence with the imperative for accurate patient assessments. However, the efficacy of this method is contingent upon the robustness and clinical relevance of the feature importance rankings generated by algorithms such as Extreme Gradient Boosting (XGBoost), which, if flawed, could introduce biases into the selection process. In contrast, the Mono-Objective Optimization method distinguishes itself through its computational speed, making it an attractive option for scenarios demanding swift decision-making. Nevertheless, the pronounced decrease in classification accuracy exposes a significant limitation: the potential for an oversimplified approach to cost-benefit analysis, which may be ill-suited for medical applications where the integrity of diagnostic accuracy cannot be compromised. The integration of these feature selection methodologies into frameworks for AI-driven technology transfer is poised to significantly enhance the operational efficiency and clinical effectiveness of healthcare systems. The identification of high-accuracy feature subsets through Combinatorial and Multi-Objective Optimization can inform the creation of diagnostic tools that are both robust and resource-efficient. The Multi-Objective Optimization method's capacity to articulate the Pareto-Optimal Frontier confers a tactical advantage in technology transfer by enabling informed choices that are in harmony with the unique constraints and objectives of individual healthcare providers. This approach resonates with the complex, multi-faceted decision-making processes inherent in healthcare systems, where the optimization of outcomes, costs, and practical implementation is paramount. Despite its limitations in terms of accuracy, the Mono-Objective Optimization method serves as a reminder of the importance of computational efficiency, particularly in settings where rapid turnaround times are essential. Its application may be most appropriate in scenarios where cost constraints are particularly stringent, provided that a modest relaxation of accuracy standards is tenable.

6. Conclusion

This investigation scrutinizes the effectiveness of three discrete feature selection strategies, Combinatorial, Multi-Objective, and Mono-Objective Optimization, in the realm of healthcare that is sensitive to both financial and temporal costs, utilizing the principles of fuzzy artificial intelligence. The central aim was to achieve a harmonious equilibrium between the precision of classification tasks and the associated fiscal and time-related expenses that are intrinsic to the chosen sets of features.

The Combinatorial Optimization technique, predicated on the Sequential Backward Selection (SBS) heuristic, meticulously examines the universe of potential feature combinations by progressively eliminating the least influential feature based on a specified metric of performance. This meticulous procedure isolated the nineteen most effective feature combinations, each yielding classification accuracies in excess of 94%. Pivotal attributes such as "gait_speed_5mt," "t._sedia," "salita_gradino," "dbp," "tumor necrosis factor TNFa," and "Vitamin D" were consistently represented within these elite combinations, underscoring their pivotal

importance in the prediction of frailty. Nonetheless, the comprehensive nature of this method engendered considerable computational demands, with an average processing duration of approximately one hour and thirty-two minutes, which could potentially impede its applicability in expansive feature spaces or time-sensitive scenarios.

Contrastingly, the Multi-Objective Optimization paradigm, drawing inspiration from McCrae's conceptual framework, prioritizes the equilibrium between accuracy and cost by delineating the Pareto-Optimal Frontier. This approach discerned six superior feature subsets that epitomize various compromises between commendable accuracy (varying from around 91.9% to 95%) and disparate cost profiles (ranging from 120 € to 396.16 €). The Pareto-Optimal Frontier serves as a decision-making aid, enabling stakeholders to opt for feature subsets that correspond to their distinct cost-effectiveness priorities. This method exhibited a processing time of roughly one hour and eleven minutes, rendering it a practical alternative for the judicious allocation of resources without forfeiting a significant degree of predictive efficacy.

The Mono-Objective Optimization framework, traceable to the contributions of Jiang et al., utilizes a merit function that simultaneously accounts for the accuracy of classification and the costs associated with feature inclusion, thus transforming the multi-objective challenge into a single-objective one. This method culminated in the isolation of a solitary feature subset, "dbp", which attained an accuracy of 63.98% at a nominal cost of 2.3 €. Although this technique boasts exceptional computational expediency, with processing times of a mere four minutes and thirty-five seconds, the accuracy achieved is markedly lower than that of the other methodologies. This suggests that while Mono-Objective Optimization excels in cost minimization, it may not be the method of choice for scenarios where the fidelity of predictions is paramount.

The Combinatorial Optimization technique's comprehensive exploration guarantees the discovery of models with exceptional accuracy, which is particularly beneficial in scenarios where the utmost precision in classification is indispensable. However, the need for substantial computational resources represents a significant limitation for applications involving large datasets or requiring real-time processing. The Multi-Objective Optimization strategy adeptly navigates the trade-off between accuracy and cost, offering a continuum of optimal solutions tailored to the diverse preferences of various stakeholders. Its reliance on sophisticated feature ranking methodologies, such as Extreme Gradient Boosting (XGBoost), is instrumental in mitigating biases and ensuring the thorough identification of feature subsets. On the other hand, the Mono-Objective Optimization approach, despite its streamlined efficiency, presents a compromise between cost containment and accuracy that may be untenable in healthcare contexts where precision is of the essence.

The incorporation of these feature selection methodologies into AI-driven technology transfer frameworks presents an opportunities for augmenting the efficiency and efficacy of healthcare systems. The high-fidelity feature subsets pinpointed by Combinatorial Optimization and Multi-Objective Optimization can inform the creation of resilient diagnostic instruments that optimize both patient outcomes and the utilization of healthcare resources. The Multi-Objective Optimization method's Pareto-Optimal Frontier confers a strategic edge by facilitating informed decision-making that aligns with specific operational and budgetary parameters. This method's capacity to furnish a multitude of optimal solutions empowers healthcare practitioners to select feature subsets that best align with their particular needs, thereby enhancing the adaptability and responsiveness of AI-driven diagnostic tools.

Subsequent inquiries should consider the amalgamation of these feature selection methodologies with supplementary fuzzy AI principles to further refine their applicability and resilience. The integration of linguistic variables and fuzzy reasoning systems could enhance the interpretability and adaptability of the feature selection process, rendering it more suitable for the complexities of clinical data. Furthermore, broadening the research scope to encompass temporal cost considerations—such as the duration necessary to secure diagnostic findings—would yield a more exhaustive assessment of feature subsets, mirroring the practical constraints of healthcare delivery more faithfully. Applying these methodologies to other sectors where cost sensitivity is pivotal, such as finance or ecological surveillance, could showcase their adaptability and wider utility. Lastly, validating the efficacy of these methods in scenarios characterized by non-constant

costs, where the penalties for misclassification vary, would offer deeper insights into their performance under a spectrum of operational conditions.

This study, thus, delineates the proficiency and relevance of three feature selection methodologies within a cost-conscious healthcare milieu, harnessing fuzzy AI principles to calibrate the equilibrium between classification accuracy and the prudent expenditure of resources. The Combinatorial Optimization and Multi-Objective Optimization methodologies display considerable promise in extracting feature subsets that are both highly accurate and cost-efficient, thereby facilitating the seamless integration of AI-driven technologies into the healthcare domain. Conversely, the Mono-Objective Optimization method underscores the pivotal trade-offs between cost reduction and the maintenance of accuracy, highlighting the necessity for context-sensitive feature selection tactics. Collectively, these insights contribute to the evolution of fuzzy systems within healthcare, fostering the development of AI models that are clinically effective and financially sustainable.

Acknowledgements:

DARE BA : This research was co-funded by the Italian Complementary National Plan PNC-I.1 "Research initiatives for innovative technologies and pathways in the health and welfare sector" D.D. 931 of 06/06/2022, "DARE - DigitAl lifelong pRevEntion" initiative, code PNC0000002, CUP: B53C22006420001.

EDIH4DT : This research was co-funded by the Project "European Digital Innovation Hub for Secure Digital Transformation EDIH4DT" - Next Generation EU (PNRR) - M4C2I2.3) - CUP B97H22004920001.

References

[1] "Health at a Glance: Europe 2020," Nov. 2020, doi: 10.1787/82129230-EN.

[2] "Encyclopedia of Machine Learning," *Encyclopedia of Machine Learning*, 2010, doi: 10.1007/978-0-387-30164-8.

[3] T. R. Shultz *et al.*, "Cost-Sensitive Learning," *Encyclopedia of Machine Learning*, pp. 231–235, 2011, doi: 10.1007/978-0-387-30164-8_181.

[4] P. D. Turney, "Cost-Sensitive Classification: Empirical Evaluation of a Hybrid Genetic Decision Tree Induction Algorithm," *Journal of Artificial Intelligence Research*, vol. 2, pp. 369–409, Mar. 1994, doi: 10.1613/JAIR.120.

[5] C. Esposito *et al.*, "Assessment of Primary Human Liver Cancer Cells by Artificial Intelligence-Assisted Raman Spectroscopy," *Cells*, vol. 12, no. 22, p. 2645, Nov. 2023, doi: 10.3390/CELLS12222645/S1.

[6] S. Bello, A. Monaco, L. Musti, G. Pirlo, and G. Semeraro, "Filtered Random Hybrid Strokes (Frhs): Filtering Time-Series Considerding Velocity Profile Filtered Random Hybrid Strokes (FRHS): Filtering Time-Series Considering Velocity Profile," *Workshop on Medical Condition Assessment Using Pattern Recognition: Progress in Neurodegenerative Disease and Beyond*, 2024, doi: 10.5220/0012567500003654.

[7] G. F. P. Cuccovillo, D. Impedovo, A. Monaco, G. Pirlo, G. Semeraro, and D. Veneto, "Handwriting Detection Test (HWDT): Android Application for the Recognition of Neurodegenerative Diseases," *Workshop on Medical Condition Assessment Using Pattern Recognition: Progress in Neurodegenerative Disease and Beyond*, 2024, doi: 10.5220/0012593600003654.

[8] V. Dentamaro, D. Impedovo, G. Pirlo, and G. Semeraro, "Decision-Level Fusion for the Early Dementia Identification by Hand-Writing and -Drawing," *2024 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering, MetroXRAINE 2024 - Proceedings*, pp. 843–848, 2024, doi: 10.1109/METROXRAINE62247.2024.10796523.

[9] L. Aversano, M. Iammarino, A. Madau, G. Pirlo, and G. Semeraro, "Process mining applications in healthcare: a systematic literature review," *PeerJ Comput Sci*, vol. 11, p. e2613, Jan. 2025, doi: 10.7717/PEERJ-CS.2613/FIG-15.

[10] L. Liu, X. Wu, S. Li, Y. Li, S. Tan, and Y. Bai, "Solving the class imbalance problem using ensemble algorithm: application of screening for aortic dissection," *BMC Med Inform Decis Mak*, vol. 22, no. 1, Dec. 2022, doi: 10.1186/S12911-022-01821-W.

[11] AversanoLerina, IammarinoMartina, MadauAntonella, MontanoDebora, and VerdoneChiara, "A Hybrid Approach integrating Clinical Data and Tomography to improve diagnosis of Parkinson's Disease," *ACM Trans Comput Healthc*, Apr. 2024, doi: 10.1145/3735660.

[12] L. Aversano, M. L. Bernardi, M. Cimitile, M. Iammarino, A. Madau, and C. Verdone, "An Explainable Approach for Early Parkinson Disease Detection Using Deep Learning," *Communications in Computer and Information Science*, vol. 1875 CCIS, pp. 326–339, 2023, doi: 10.1007/978-3-031-39059-3_22.

[13] L. Aversano, M. L. Bernardi, M. Cimitile, M. Iammarino, D. Montano, and C. Verdone, "Early Parkinson's Disease Detection from EEG Traces Using Machine Learning Techniques," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 14069 LNCS, pp. 607–619, 2023, doi: 10.1007/978-3-031-39965-7_50.

[14] L. Aversano *et al.*, "Using Machine Learning for Classification of Cancer Cells from Raman Spectroscopy," pp. 15–24, Jul. 2024, doi: 10.5220/0011142600003277.

[15] F. Liu and Q. Qian, "Cost-Sensitive Variational Autoencoding Classifier for Imbalanced Data Classification," *Algorithms*, vol. 15, no. 5, May 2022, doi: 10.3390/A15050139.

[16] P. Teisseyre and T. Klonecki, "Controlling Costs in Feature Selection: Information Theoretic Approach," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12743 LNCS, pp. 483–496, 2021, doi: 10.1007/978-3-030-77964-1_37/COVER/.

[17] H. Zhao, F. Min, and W. Zhu, "Cost-sensitive feature selection of numeric data with measurement errors," *J Appl Math*, vol. 2013, 2013, doi: 10.1155/2013/754698.

[18] A. Verma, M. K. Hanawal, and N. Hemachandra, "Unsupervised Online Feature Selection for Cost-Sensitive Medical Diagnosis," *2020 International Conference on COMmunication Systems and NETworkS, COMSNETS 2020*, pp. 1–6, Dec. 2019, doi: 10.48550/arxiv.2001.00626.

[19] M. Abdulla and M. T. Khasawneh, "G-Forest: An ensemble method for cost-sensitive feature selection in gene expression microarrays," *Artif Intell Med*, vol. 108, p. 101941, Aug. 2020, doi: 10.1016/J.ARTMED.2020.101941.

[20] S. I. Ali *et al.*, "Cost-Sensitive Ensemble Feature Ranking and Automatic Threshold Selection for Chronic Kidney Disease Diagnosis," *Applied Sciences 2020, Vol. 10, Page 5663*, vol. 10, no. 16, p. 5663, Aug. 2020, doi: 10.3390/APP10165663.

[21] L. Chen, "Curse of Dimensionality," *Encyclopedia of Database Systems*, pp. 545–546, 2009, doi: 10.1007/978-0-387-39940-9_133.

[22] X. Ying, "An Overview of Overfitting and its Solutions," *J Phys Conf Ser*, vol. 1168, no. 2, p. 022022, Feb. 2019, doi: 10.1088/1742-6596/1168/2/022022.

[23] D. Effrosynidis and A. Arampatzis, "An evaluation of feature selection methods for environmental data," *Ecol Inform*, vol. 61, p. 101224, Mar. 2021, doi: 10.1016/J.ECOINF.2021.101224.

[24] A. O. Balogun *et al.*, "Empirical Analysis of Rank Aggregation-Based Multi-Filter Feature Selection Methods in Software Defect Prediction," *Electronics 2021, Vol. 10, Page 179*, vol. 10, no. 2, p. 179, Jan. 2021, doi: 10.3390/ELECTRONICS10020179.

[25] H. Sun, J. Jin, R. Xu, and A. Cichocki, "Feature Selection Combining Filter and Wrapper Methods for Motor-Imagery Based Brain–Computer Interfaces," *https://doi.org/10.1142/S0129065721500404*, vol. 31, no. 9, Aug. 2021, doi: 10.1142/S0129065721500404.

[26] M. Mandal, P. K. Singh, M. F. Ijaz, J. Shafi, and R. Sarkar, "A Tri-Stage Wrapper-Filter Feature Selection Framework for Disease Classification," *Sensors 2021, Vol. 21, Page 5571*, vol. 21, no. 16, p. 5571, Aug. 2021, doi: 10.3390/S21165571.

[27] J. Chen, S. Yuan, D. Lv, and Y. Xiang, "A novel self-learning feature selection approach based on feature attributions," *Expert Syst Appl*, vol. 183, p. 115219, Nov. 2021, doi: 10.1016/J.ESWA.2021.115219.

[28] X. Li, J. Zhang, and F. Safara, "Improving the Accuracy of Diabetes Diagnosis Applications through a Hybrid Feature Selection Algorithm," *Neural Process Lett*, pp. 1–17, Mar. 2021, doi: 10.1007/S11063-021-10491-0/TABLES/5.

[29] X. F. Song, Y. Zhang, D. W. Gong, and X. Z. Gao, "A Fast Hybrid Feature Selection Based on Correlation-Guided Clustering and Particle Swarm Optimization for High-Dimensional Data," *IEEE Trans Cybern*, 2021, doi: 10.1109/TCYB.2021.3061152.

[30] S. Lv, S. Shi, H. Wang, and F. Li, "Semi-supervised multi-label feature selection with adaptive structure learning and manifold learning," *Knowl Based Syst*, vol. 214, p. 106757, Feb. 2021, doi: 10.1016/J.KNOSYS.2021.106757.

[31] J. Li *et al.*, "Feature selection: A data perspective," *ACM Comput Surv*, vol. 50, no. 6, Dec. 2017, doi: 10.1145/3136625.

[32] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, vol. 13-17-August-2016, pp. 785–794, Aug. 2016, doi: 10.1145/2939672.2939785;CSUBTYPE:STRING:CONFERENCE.

[33] S. Raschka and V. Mirjalili, *Python machine learning : machine learning and deep learning with python, scikit-learn, and tensorflow 2*, vol. 11, no. 1. 2021. Accessed: Jun. 30, 2022. [Online]. Available: https://www.libreriauniversitaria.it/python-machine-learning-third-edition/book/9781789955750

[34] T. P. Minka, "Automatic choice of dimensionality f o r PCA," 2000.

[35] L. Jiang, G. Kong, and C. Li, "Wrapper Framework for Test-Cost-Sensitive Feature Selection," *IEEE Trans Syst Man Cybern Syst*, vol. 51, no. 3, pp. 1747–1756, Mar. 2021, doi: 10.1109/TSMC.2019.2904662.