# Fuzzy Vertex-Feature Clustering Framework for Predicting Survival of Health Sector Companies And Generative Explanations

Lerina Aversano
*University of Foggia*
Italy
0000-0003-2436-6835

Vincenzo Dentamaro
*University of Bari Aldo Moro*
Italy
0000-0003-1148-332X

Felice Franchini
*University of Bari Aldo Moro*
Italy
0009-0000-8887-800X

***Abstract*—Accurately forecasting the long-term viability of firms in the highly regulated and competitive health-care sector is critical for investors, managers, and policymakers. Conventional machine-learning methods often struggle to reconcile predictive power with transparency, while standard clustering techniques provide limited insight with high-dimensional, noisy business data.**

**This paper introduces *Fuzzy Vertex-Feature Clustering* (FVFC), a novel framework that couples fuzzy-set theory with a vertex–feature representation to discover interpretable survival patterns. FVFC embeds firms into a simplex, whose vertices correspond to survival classes, using Jensen–Shannon divergence to capture non-linear relationships among financial, innovation, and operational indicators. A fuzzy $k$-nearest-neighbor classifier then refines class-membership probabilities, and Kernel Principal Component Analysis maps the high-dimensional space to an interactive 3-dimensional visualization. To bridge the interpretability gap, the framework integrates a large language model that generates natural-language explanations.**

**The main contribution in AI lies in combining fuzzy-set theory, divergence-based embeddings, and automated natural-language explanations to achieve interpretable survival prediction.**

**Tests on 27,166 Italian pharmaceutical and medical-device companies, plus three public benchmarks, show FVFC attains state-of-the-art performance, with weighted F-measure of 0.969 ± 0.006 while preserving human-readable logic. Compared with tree ensembles and gradient boosting, FVFC delivers comparable or better accuracy, markedly improves minority-class recall, and provides stakeholders with 3-D plots and narrative justifications. The proposed framework offers a transparent, domain-adaptable decision-support tool that can enhance strategic planning and risk management across the health-care industry.**

***Index Terms*—Fuzzy Vertex Feature Clustering, Fuzzy k-nearest neighbors, Artificial Intelligence in survival prediction, Health sector company survival analysis, Explainable Artificial Intelligence and decision support**

## I. Introduction

The health sector represents a fundamental pillar for global economic and social stability, significantly impacting public health and technological progress. However, organizations in the health sector have to operate under challenging conditions that are characterized by high-level competitiveness, regulatory constraints, and resource scarcity. These challenges emphasize the importance of understanding survival dynamics to enable stakeholders to make informed decisions that ensure the sustainability of operations.

The prediction of health sector companies' survival requires the analysis of intricate, often nonlinear relationships among various business attributes, such as financial well-being, market stance, and technological proficiency. Conventional predictive models, although efficacious in specific scenarios, often grapple with achieving a balance between accuracy and interpretability.The development of explainable artificial intelligence (XAI) underlines the importance of developing systems that are not only accurate but also explain their predictive capabilities [1], [2].

Fuzzy prediction methods represent a potentially powerful tool for survival prediction due to their ability to naturally handle uncertainty through fuzzy membership functions, which often results in higher interpretability than traditional hard-label prediction methods. Despite their advantages, these methodologies face significant challenges in achieving high levels of interpretability and visualization when dealing with high-dimensional, nonlinear business data.

### A. Challenges

Despite recent progress, the prediction of survival probabilities for health sector enterprises still faces three operational challenges:

- **Interpretable clustering:** Existing techniques, such as k-means and hierarchical clustering, are not suitable for nonlinear and high-dimensional data of firms and often do not provide easily interpretable results [3], [4].
- **Explainability:** High-performing models are often black boxes, and this restricts the ability to trust and use the model for decision-making [5], [6].
- **Visualization:** Lack of informativeness of simple two-dimensional plots does not allow stakeholders to grasp complex structures. [7].

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: B53C22006450001, CUP: B53C22006420001.

To overcome the drawbacks of the existing predictive modeling techniques, the present work proposes a novel predictive modeling framework with the following features:

1) **Fuzzy Vertex-Feature Clustering (FVFC):** a divergence-based embedding approach that projects firms onto a simplex space and provides interpretable fuzzy class membership scores.
2) **Fuzzy K-Nearest Neighbors (Fuzzy KNN) refinement:** a neighborhood-based module that refines membership scores to improve the robustness of the prediction approach [8].
3) **Interactive 3D visualization:** a Kernel PCA-based module that projects the embedding space onto a three-dimensional space.
4) **Natural-language explanations:** an LLM-based module that provides a concise summary of the prediction rationale based on membership scores and nearest neighbors, useful for non-technical stakeholders.

To support interpretability, the proposed framework couples the FVFC prediction with an interactive 3D view (Kernel PCA) and a text explanation module. For each queried company, the system highlights its location in the embedding space, its nearest neighbors, and the resulting class-membership scores, and it generates a concise natural-language rationale conditioned on these quantities. The interface and explanation workflow are described in Section V.

This paper is organized as follows: section II analyzes the related work, section III discusses the employed methods, section IV provides and discusses the experiments and results, section V describes the graphical framework, conclusions are given in section VI.

## II. Related Work

Fuzzy systems have been able to achieve widespread acceptance in survival analysis and risk prediction due to their potential to address uncertainty through soft membership functions, which may lead to an improvement in interpretability over hard decision models. Past studies on fuzzy and hybrid models for healthcare-related prediction include fuzzy wavelet neural networks [9], fuzzy neural networks with support vector machine learning [10], and fuzzy min–max models [11]. These models were used in critical survival analysis, such as in acute decompensated heart failure patient mortality prediction [12], and were discussed in comprehensive reviews of various intelligent decision support models [13]. Fuzzy deep learning and recurrent fuzzy neural networks were also used in critical and dynamic environments, which demonstrates the flexibility of fuzzy/hybrid models in learning [14]–[16].

Interpretability is another significant factor to consider when working on high-stakes problems. Methodological work emphasizes the need to ensure human interpretable AI and argues for the use of fuzzy models to support this [17], [18]. Complementary lines of research study explainable AI techniques that help interpret complex models [7] and rule-extraction strategies that translate learned behavior into more accessible forms [19]. Interactive visualization and fuzzy AI have also been explored to support exploratory analysis and healthcare-oriented decision support where robustness and interpretability are crucial [20], [21].

Despite these advances, previous work has generally focused on fuzzy prediction models, explainable AI models, and interactive visualization models individually, while rarely integrating these models to produce a workflow that can produce both an interpretable geometric representation and stakeholder-level explanations for high-dimensional tabular data.

## III. Methods

### A. Framework Overview

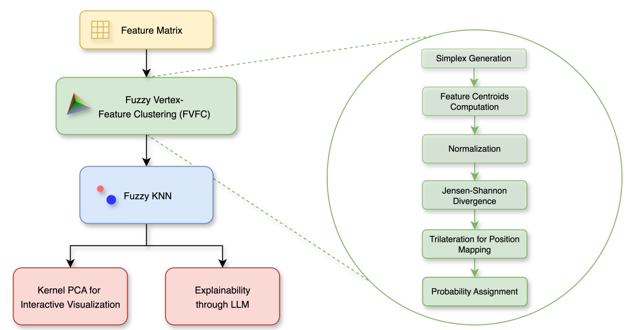


Fig. 1. End-to-end overview of the proposed FVFC framework. From an input feature matrix, FVFC computes a divergence-based simplex embedding and fuzzy memberships, which are refined by fuzzy KNN. Kernel PCA provides an interactive 3D view of the embedding, while an LLM produces a concise natural-language rationale from memberships and nearest neighbors.

Figure 1 shows the proposed workflow. The FVFC framework starts with the feature matrix and then proceeds to create an interpretable simplex representation through the following steps: building class vertices, computing centroids, and measuring dissimilarities to centroids using Jensen-Shannon divergence. Each company is then located through the trilateration step, which yields fuzzy membership scores. To make the model more robust to noisy and imbalanced data, fuzzy K-Nearest Neighbors is used to refine the fuzzy membership scores. To make the model interpretable, the high-dimensional space is reduced to three dimensions through Kernel Principal Component Analysis, allowing interactive exploration of proximity relations. Additionally, a large language model is used to create a concise textual rationale.

### B. Dataset and Preprocessing

The dataset includes 27,166 Italian companies in the pharmaceutical industry and medical device manufacturing, labeled by operational status. Features span financial, innovation, and operational indicators (e.g., profit margins, total assets, patent counts, market attractiveness, and IP efficiency).

A standard preprocessing pipeline was applied: (i) removal of records with missing/invalid labels; (ii) feature-name cleaning; (iii) separation of numerical and categorical variables;

(iv) median imputation for missing numerical values and encoding for categorical variables; (v) z-score standardization of numerical features; and (vi) consolidation into a single modeling matrix. The resulting label distribution is highly imbalanced: Issued (11,258), Lapsed (1,801), Retired (19), and Pending Transfer (18). Given the extreme scarcity of the last two categories, Retired and Pending Transfer were excluded from subsequent experiments.

### C. Fuzzy Vertex-Feature Clustering (FVFC)

The enhanced FVFC algorithm represents an extension of conventional fuzzy clustering methods. It achieves this by integrating relationships between vertices and features, and by employing the Jensen-Shannon divergence as a means to calculate distances. This methodology is designed to encapsulate the intricate relationships present in data of high dimensionality, all the while ensuring that the results remain interpretable. This approach aims to capture the complex relationships in high-dimensional data while maintaining interpretability, which is crucial for practical applications in the health sector where decisions need to be transparent and understandable.

*1) Simplex Generation:* The algorithm begins by generating simplex vertices in a $d$-dimensional space. For $C$ classes (clusters), a simplex with $C$ vertices is constructed. If $C \leq d+1$, standard simplex vertices are used; otherwise, vertices are placed uniformly (e.g., $\mathbf{v}_k = (\cos(2\pi k/C), \sin(2\pi k/C), 0)$).

*2) Feature Centroids Computation:* For each class $c \in \{1, 2, \ldots, C\}$, the feature centroid $\boldsymbol{\mu}_c$ is computed by averaging the feature vectors of all samples belonging to that class:

$$\boldsymbol{\mu}_c = \frac{1}{N_c} \sum_{i:y_i=c} \mathbf{x}_i, \tag{1}$$

where $N_c$ is the number of samples in class $c$.

*3) Normalization:* To ensure that feature vectors represent probability distributions, each feature vector $\mathbf{x}_i$ and centroid $\boldsymbol{\mu}_c$ is normalized using the $\ell_1$ norm:

$$\tilde{\mathbf{x}}_i = \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|_1}, \quad \tilde{\boldsymbol{\mu}}_c = \frac{\boldsymbol{\mu}_c}{\|\boldsymbol{\mu}_c\|_1}. \tag{2}$$

*4) Jensen-Shannon Divergence:* The Jensen-Shannon divergence (JSD) is employed as a symmetric measure of divergence between two probability distributions. For each sample $\tilde{\mathbf{x}}_i$, the JSD to each class centroid $\tilde{\boldsymbol{\mu}}_c$ is computed:

$$D_{\text{JS}}(\tilde{\mathbf{x}}_i \| \tilde{\boldsymbol{\mu}}_c) = \frac{1}{2} D_{\text{KL}}(\tilde{\mathbf{x}}_i \| \mathbf{M}_i) + \frac{1}{2} D_{\text{KL}}(\tilde{\boldsymbol{\mu}}_c \| \mathbf{M}_i), \tag{3}$$

where $\mathbf{M}_i = \frac{1}{2}(\tilde{\mathbf{x}}_i + \tilde{\boldsymbol{\mu}}_c)$ is the average distribution and $D_{\text{KL}}(\mathbf{P} \| \mathbf{Q}) = \sum_j P_j \log\left(\frac{P_j}{Q_j}\right)$. To ensure numerical stability, small values are clipped to a minimum positive value $\varepsilon$.

*5) Trilateration for Position Mapping:* The distances obtained from the JSD are used to map each sample into the simplex space by solving a trilateration problem. For each sample $i$, the position $\mathbf{p}_i$ is found by minimizing the residuals between the calculated distances and the actual distances in the simplex space:

$$\min_{\mathbf{p}_i} \sum_{c=1}^{C} \left( \|\mathbf{p}_i - \mathbf{v}_c\|_2 - D_{\text{JS}}(\tilde{\mathbf{x}}_i \| \tilde{\boldsymbol{\mu}}_c) \right)^2. \tag{4}$$

This nonlinear least squares problem is solved using optimization techniques such as the Levenberg-Marquardt algorithm.

*6) Probability Assignment:* After mapping samples to positions in the simplex space, probabilities of class membership are assigned based on the inverse of the Euclidean distance to each vertex:

$$P_{\text{FVFC}}^{(i,c)} = \frac{1}{1 + \|\mathbf{p}_i - \mathbf{v}_c\|_2}. \tag{5}$$

These probabilities are normalized to ensure they sum to one across all classes for each sample:

$$\hat{P}_{\text{FVFC}}^{(i,c)} = \frac{P_{\text{FVFC}}^{(i,c)}}{\sum_{k=1}^{C} P_{\text{FVFC}}^{(i,k)}}. \tag{6}$$

### D. Fuzzy K-Nearest Neighbors Classification

To enhance classification performance, the fuzzy KNN algorithm is integrated. [22] Fuzzy KNN extends the traditional KNN by assigning membership degrees to each class instead of hard labels.

*1) Membership Degree Calculation:* For a sample $\mathbf{x}_i$, the distances to all training samples are computed. The $K$ nearest neighbors are identified, and their distances $d_{ij}$ and class labels $y_j$ are used to compute the membership degree of $\mathbf{x}_i$ belonging to class $c$:

$$u_{ic} = \frac{\sum_{j \in \mathcal{N}_i} \delta(y_j, c) \left( \frac{1}{d_{ij}^{2/(m-1)}} \right)}{\sum_{j \in \mathcal{N}_i} \left( \frac{1}{d_{ij}^{2/(m-1)}} \right)}, \tag{7}$$

where $\mathcal{N}_i$ denotes the set of $K$ nearest neighbors of sample $i$, $\delta(y_j, c)$ is the Kronecker delta function, and $m > 1$ is the fuzziness parameter controlling the influence of the distance.

### E. Combined Classification Decision

The final class membership probabilities are obtained by combining the outputs from the FVFC and fuzzy KNN components using a weighted average:

$$P_{\text{combined}}^{(i,c)} = \alpha \hat{P}_{\text{FVFC}}^{(i,c)} + (1 - \alpha) u_{ic}, \tag{8}$$

where $\alpha \in [0, 1]$ is a weighting parameter balancing the contributions of FVFC and fuzzy KNN.

The final predicted class for sample $\mathbf{x}_i$ is determined by:

$$\hat{y}_i = \arg\max_c P_{\text{combined}}^{(i,c)}. \tag{9}$$

### F. Visualization Using Kernel Principal Component Analysis

To facilitate interpretability and exploration of complex relationships, Kernel PCA is employed to project high-dimensional data into a three-dimensional space suitable for visualization.

*1) Kernel PCA Transformation:* Kernel PCA extends traditional PCA by using kernel functions to capture nonlinear structures. The cosine kernel is selected for its ability to measure the similarity between data points based on the angle between vectors:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{x}_i^\top \mathbf{x}_j}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|}. \tag{10}$$

The kernel matrix $\mathbf{K}$ is constructed, and eigenvalue decomposition is performed to obtain the principal components. The transformed coordinates for sample $i$ are given by $\mathbf{z}_i = (z_{i1}, z_{i2}, z_{i3})^\top$, corresponding to the top three eigenvectors.

*2) Interactive Visualization:* An interactive three-dimensional graphical representation is created, wherein each data point is positioned according to its transformed coordinates, represented as $\mathbf{z}_i$. Clusters are visualized utilizing convex hulls or mesh surfaces, offering a geometric depiction of class boundaries. Interactive elements are incorporated, enabling users to hover over individual points to access detailed information, thereby enhancing interpretability.

### *G. Integration of Large Language Models for Explainability and User Interface Design*

In order to improve the interpretability of classification results for non-technical stakeholders, a large language model, namely Qwen 2.5 7B [23], is leveraged to produce a natural language explanation of the FVFC classification results.

For each classified company, feature values, predicted class probabilities, and information from the nearest neighbors are assembled into a structured prompt and sent to the LLM via the OpenAI API. The explanation process can be expressed as the mapping $\mathcal{E}(\mathbf{f}, c^*, \mathcal{N}) = \mathrm{LLM}(P(\mathbf{f}, c^*, \mathcal{N}))$, where $\mathbf{f}$ denotes the feature vector, $c^*$ the predicted class, and $\mathcal{N}$ the neighborhood information.

The user interface supports interactive inspection of each prediction by collecting company inputs through numeric fields, projecting the instance into a 3D space via FVFC and Kernel PCA, and visualizing clusters and nearest neighbors in a Plotly 3D scatter plot. The provided explanation is displayed along with the visualization, allowing users to understand the classification results in a concise narrative form based on memberships and neighborhood comparisons.

### *H. Label Mapping Using the Hungarian Algorithm*

To evaluate the clustering performance, cluster labels are mapped to true class labels using the Hungarian algorithm [24]. Given the confusion matrix $\mathbf{C} = [c_{ij}]$, where $c_{ij}$ represents the number of samples in true class $i$ assigned to cluster $j$, the optimal mapping minimizes the overall misclassification.

*1) Optimization Problem:* The assignment problem is formulated as:

$$\max_P \sum_{i=1}^{C} \sum_{j=1}^{C} p_{ij} c_{ij}, \tag{11}$$

subject to:

$$\sum_{i=1}^{C} p_{ij} = 1, \quad \sum_{j=1}^{C} p_{ij} = 1, \quad p_{ij} \in \{0, 1\}, \tag{12}$$

where $P = [p_{ij}]$ is a permutation matrix representing the assignment of clusters to true classes.

*2) Solution:* The Hungarian algorithm efficiently solves this linear assignment problem, yielding the optimal mapping between cluster labels and true class labels.

### *I. Algorithm Complexity Analysis*

The preprocessing stage runs in $\mathcal{O}(np)$ for $n$ samples and $p$ features. FVFC requires $\mathcal{O}(nCp)$ to compute Jensen–Shannon divergences to $C$ centroids, while fuzzy KNN membership refinement scales as $\mathcal{O}(nNp)$ for distance evaluation against $N$ training instances. Kernel PCA has $\mathcal{O}(n^2p)$ kernel construction and $\mathcal{O}(n^3)$ eigendecomposition cost and is applied only for visualization; similarly, the LLM module is used only for interpretability.

## IV. Experiments

### *A. Experimental Setup*

The experiments performed are multiclass classification tasks on various datasets in order to stress the capabilities of the proposed technique in various settings. Their scores are benchmarked against state-of-the-art machine learning techniques such as Decision Tree (DT), Random Forest (RF), XGBoost [25], k-Nearest Neighbors (KNN), and Linear Support Vector Machine (LSVM). For all the tests in all experiments, the model's hyperparameters reported in Table I have been used. The tests were performed using 10-Fold cross-validation to ensure models' robustness. Mean values of the weighted F1 score and their standard deviation have been reported in Table II and Tab III.

TABLE I
HYPERPARAMETER SETTINGS FOR METHODS AND MODELS

| Method/Model | Hyperparameter | Value |
|---|---|---|
| FVFC | Number of Clusters | Number of Classes |
| | Number of Neighbors | 5 |
| | Fuzziness Parameter (`m`) | 2.0 |
| | Weighting Factor (`alpha`) | 0.7 |
| Decision Tree | Criterion | Gini |
| Random Forest | Number of Estimators | 100 |
| | Criterion | Gini |
| XGBoost | Number of Classes | Unique Class Count |
| | Eval Metric | logloss |
| KNN | Number of Neighbors (`k`) | 5 |
| Linear SVM | Kernel Type | Linear |

### *B. Performance Evaluation of Proposed Framework*

The findings from the experiments illustrate the effectiveness of the improved Fuzzy Vertex-Feature Clustering (FVFC) method, which is particularly useful for forecasting the longevity of firms in the healthcare industry. Table II offers

TABLE II
CLASSIFICATION RESULTS FOR BENCHMARK MODELS

| Model | W-F1 | C0-F1 | C1-F1 |
|---|---|---|---|
| FVFC | .969 ± .006 | .969 ± .004 | .881 ± .020 |
| DT | .963 ± .005 | .979 ± .003 | .867 ± .020 |
| RF | .967 ± .004 | .982 ± .002 | .874 ± .016 |
| XGB | .971 ± .003 | .984 ± .002 | .894 ± .015 |
| KNN | .934 ± .007 | .964 ± .004 | .748 ± .023 |
| LSVM | .797 ± .016 | .925 ± .007 | .000 ± .000 |

TABLE III
CLASSIFICATION RESULTS ACROSS MODELS AND DATASETS

| Dataset | Model | Accuracy | F1 | Time(s) |
|---|---|---|---|---|
| Cervical Cancer [26] | FVFC | .978 ± .006 | .978 ± .006 | 36.52 |
| | DT | .970 ± .009 | .970 ± .009 | 00.76 |
| | RF | .976 ± .007 | .976 ± .007 | 10.55 |
| | XGB | .973 ± .007 | .973 ± .007 | 08.67 |
| | KNN | .971 ± .009 | .971 ± .009 | 00.13 |
| | LSVM | .955 ± .020 | .955 ± .020 | 00.73 |
| Breast Cancer [27] | FVFC | .969 ± .031 | .969 ± .031 | 12.59 |
| | DT | .952 ± .025 | .952 ± .025 | 00.33 |
| | RF | .962 ± .030 | .962 ± .030 | 06.10 |
| | XGB | .971 ± .032 | .971 ± .032 | 01.85 |
| | KNN | .965 ± .026 | .965 ± .026 | 00.07 |
| | LSVM | .961 ± .022 | .961 ± .022 | 00.17 |
| Iris [28] | FVFC | .953 ± .060 | .952 ± .062 | 00.08 |
| | DT | .933 ± .052 | .932 ± .052 | 04.26 |
| | RF | .947 ± .058 | .945 ± .060 | 01.14 |
| | XGB | .940 ± .047 | .939 ± .048 | 00.09 |
| | KNN | .952 ± .052 | .952 ± .054 | 00.07 |
| | LSVM | .947 ± .058 | .945 ± .060 | 01.77 |

an extensive analysis comparing multiple machine learning algorithms, emphasizing the subtle differences in performance among various classification methods.

*1) Model Performance Metrics:* The FVFC approach, as proposed, attained a weighted F1-score of $0.969 \pm 0.006$, thereby establishing its competitive standing amidst contemporary machine learning techniques. Furthermore, the model's performance was robust across various classes, with particularly strong results for the dominant class (Class 0) and consistent performance for the minority class (Class 1).

Upon conducting a comparative analysis, it becomes evident that various algorithms exhibit nuanced performance attributes. XGBoost emerged with the highest weighted F1-score of $0.971 \pm 0.003$, achieving a slight superiority over the proposed FVFC method. This marginal difference suggests that both methods are highly competitive, with XGBoost having a slight edge in this particular context. The Random Forest and Decision Tree models displayed similar performance levels, with weighted F1-scores of $0.967 \pm 0.004$ and $0.963 \pm 0.005$, respectively. Conversely, the Support Vector Machine (SVM) demonstrated the poorest performance, attaining a weighted F1-score of $0.797 \pm 0.016$. This notably lower score indicates that SVM may not be well-suited for datasets with complex, nonlinear relationships, such as the one under consideration.

### C. Class-Specific Performance

An examination of class-specific F1-scores provides valuable insights into the predictive abilities of the models under consideration. Notably, across all models evaluated, the predominant class, designated as Class 0, consistently achieved F1-scores surpassing 0.96, which attests to a strong predictive capacity for the majority class. In contrast, the performance for the minority class, Class 1, displayed more intricate variations. Specifically, the FVFC methodology attained an F1-score of 0.881 with a standard deviation of 0.020, thereby demonstrating the model's efficacy in tackling the complexities associated with class imbalance.

### D. Generalizability and Cross-Dataset Validation

Table III offers a thorough assessment of the FVFC approach across various datasets, thereby illustrating its generalizability and robustness. In the context of the Cervical Cancer dataset, the FVFC method achieved the highest accuracy rate of 0.978 with a standard deviation of 0.006, outpacing conventional machine learning models. Similarly, in the Breast Cancer dataset, the FVFC approach demonstrated competitive performance with an accuracy of 0.969 and a standard deviation of 0.031. Furthermore, consistent results were observed in the Iris dataset, where the accuracy was 0.953 with a standard deviation of 0.060.

### E. Theoretical and Practical Implications

The suggested framework overcomes key shortcomings of current predictive modeling methods by employing several innovative mechanisms. Through the integration of fuzzy clustering and interactive visualization techniques, the model offers unparalleled transparency in decision-making processes. Additionally, the inclusion of Large Language Models (LLMs) facilitates the generation of explanations in natural language, thereby effectively bridging the communication gap between technical predictions and the comprehension of stakeholders. Furthermore, the fuzzy approach exhibits remarkable proficiency in handling the inherent uncertainty that is characteristic of survival predictions in the health sector.

## V. GRAPHICAL FRAMEWORK

The proposed decision-support interface (Fig. 2) couples the FVFC embedding with interactive exploration and explanation. Companies are displayed in a 3D projection of the FVFC space (Kernel PCA), where issued and lapsed instances form a characteristic triangular structure. Transparent class surfaces approximate the regions associated with each class and typically overlap in an intersection zone that highlights higher uncertainty. Class centroids are shown as reference markers to help users contextualize where a firm lies with respect to the two survival patterns.

For a new company, the interface highlights its position (cross) together with its nearest neighbors (larger points) that drive the fuzzy membership refinement; the neighbor color intensity reflects their influence. The class membership probabilities are also provided in natural language by the LLM,

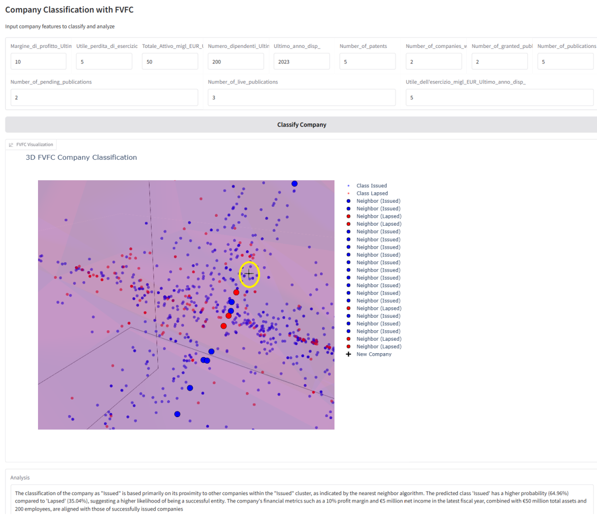


Fig. 2. Framework showing the FVFC prediction. Users enter company parameters (top), inspect the 3D embedding with the queried firm (cross) and its nearest neighbors (center), and read the LLM-generated rationale (bottom).

which offers a justified explanation of the results based on the calculated memberships.

## VI. Conclusions

The proposed Fuzzy Vertex-Feature Clustering (FVFC) framework represents a significant theoretical and practical advancement in survival prediction for health sector companies. By synthesizing fuzzy logic, advanced visualization techniques, and explainable artificial intelligence, the proposed methodology offers a comprehensive approach to addressing the complex challenges of predicting company survival under conditions of uncertainty.

The research contributes not only a novel technical solution but also provides a conceptual framework for understanding and navigating the intricate landscape of health sector company survival. Future iterations of this approach hold substantial promise for enhancing strategic decision-making processes in dynamic and complex business environments.

## References


[1] K. Cohen, L. Bokati, M. Ceberio, O. Kosheleva, and V. Kreinovich, “Why fuzzy techniques in explainable ai? which fuzzy techniques in explainable ai?” in *Springer Lecture Notes*, 2021.

[2] P. V. de Campos Souza, “Fuzzy neural networks and neuro-fuzzy networks: A review the main techniques and applications used in the literature,” *Applied soft computing*, vol. 92, p. 106275, 2020.

[3] W. Ding, M. Abdel-Basset, H. Hawash, and A. Ali, “Explainability of artificial intelligence methods, applications and challenges: A comprehensive survey,” *Information Sciences*, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S002002552201132X

[4] L. Brito, G. Susto, J. Brito, and M. Duarte, “An explainable artificial intelligence approach for unsupervised fault detection and diagnosis in rotating machinery,” *Computers in Industry*, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0888327021004891

[5] D. Minh, H. X. Wang, Y. F. Li, and T. N. Nguyen, “Explainable artificial intelligence: a comprehensive review,” *Artificial Intelligence Review*, pp. 1–66, 2022.

[6] H. Hagras, “Toward human-understandable, explainable ai,” *Computer*, vol. 51, no. 9, pp. 28–36, 2018.

[7] K. Weitz, T. Hassan, U. Schmid, and J.-U. Garbas, “Deep-learned faces of pain and emotions: Elucidating the differences of facial expressions with the help of explainable ai methods,” *tm-Technisches Messen*, vol. 86, no. 7-8, pp. 404–412, 2019.

[8] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. Wiley, 2012.

[9] Y. e. a. Zhang, “A network traffic prediction model based on quantum-behaved particle swarm optimization algorithm and fuzzy wavelet neural network,” *Discrete dynamics in nature and society*, 2016.

[10] R. Khodabandelou and M. Ebadzadeh, “Fuzzy neural network with support vector-based learning for classification and regression,” *Soft computing*, 2019.

[11] A. Mohammed and H. Lim, “An enhanced fuzzy min–max neural network for pattern classification,” *IEEE Transactions on Neural Networks and Learning Systems*, 2015.

[12] M. e. a. Yazid, “Mortality prediction for acute decompensated heart failure patient using fuzzy neural network,” *Malaysian journal of fundamental and applied sciences*, 2020.

[13] M. Ali and A. Mahmoud, “Methodologies and applications of artificial intelligence in systems engineering,” *International journal of robotics and control systems*, 2022.

[14] A. e. a. Kesse, “Development of an artificial intelligence powered tig welding algorithm for the prediction of bead geometry for tig welding processes using hybrid deep learning,” *Metals*, 2020.

[15] Y. e. a. Lin, “Improved lvrt performance of pv power plant using recurrent wavelet fuzzy neural network control for weak grid conditions,” *IEEE Access*, 2020.

[16] 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,” in *Fuzzy Logic and Technology, and Aggregation Operators (EUSFLAT 2023)*. Springer, 2023, pp. 607–619.

[17] G. Marín Díaz, R. Gómez Medina, and J. A. Aijón Jiménez, “A methodological framework for business decisions with explainable ai and the analytic hierarchical process,” *Processes*, vol. 13, no. 1, p. 102, 2025.

[18] L. Aversano, M. L. Bernardi, M. Cimitile, M. Iammarino, A. Madau, and C. Verdone, “Early diagnosis of parkinson’s disease exploting motor and non-motor symptoms: Results from the ppmi cohort,” *Procedia Computer Science*, vol. 225, pp. 2096–2105, 2023.

[19] A. Averkin and A. Yarushev, “Review of research in the field of developing methods to extract rules from artificial neural networks,” *Journal of Computer and Systems Sciences International*, 2021.

[20] J. Lee, “Application of artificial intelligence for the management of oral diseases,” *Journal of Oral Medicine and Pain*, 2022.

[21] S. Garcia-Galan, J. A. Cabrera-Rodriguez, F. J. Maldonado-Carrascosa, N. Ruiz-Reyes, M. Szczerska, P. Vera-Candeas, F. D. Gonzalez-Martinez, F. J. Canadas-Quesada, and A. J. Cruz-Lendinez, “Expert system against machine learning approaches as a virtual sensor for ventricular arrhythmia risk level estimation,” *Biomedical Signal Processing and Control*, vol. 102, p. 107255, 2025.

[22] Z. Bian, C. M. Vong, P. K. Wong, and S. Wang, “Fuzzy knn method with adaptive nearest neighbors,” *IEEE transactions on cybernetics*, vol. 52, no. 6, pp. 5380–5393, 2020.

[23] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, “Qwen technical report,” *arXiv preprint arXiv:2309.16609*, 2023.

[24] H. W. Kuhn, “The hungarian method for the assignment problem,” *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[25] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.

[26] K. Fernandes, J. S. Cardoso, and J. Fernandes, “Transfer learning with partial observability applied to cervical cancer screening,” in *Pattern Recognition and Image Analysis: 8th Iberian Conference, IbPRIA 2017, Faro, Portugal, June 20-23, 2017, Proceedings 8*. Springer, 2017, pp. 243–250.

[27] K. P. Bennett and O. L. Mangasarian, “Robust linear programming discrimination of two linearly inseparable sets,” *Optimization methods and software*, vol. 1, no. 1, pp. 23–34, 1992.

[28] R. A. Fisher, “The use of multiple measurements in taxonomic problems,” *Annals of eugenics*, vol. 7, no. 2, pp. 179–188, 1936.