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.

A Multi-Input Approach for SME Default Using Explainable Deep Learning

Belinda Laura Del Gaudio
Parthenope University of Naples
Raffaele Clemente Petrella
raffaeleclemente.petrella@phds.uniparthenop
Parthenope University of Naples
Salvatore Scognamiglio
Parthenope University of Naples
Paolo Zanetti
Parthenope University of Naples

Research Article

Keywords: Default prediction, SME, Deep Learning, Machine Learning, Explainable AI

Posted Date: October 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11279477/v1

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

Additional Declarations: No competing interests reported.

A Multi-Input Approach for SME Default Using Explainable Deep Learning

Belinda Laura Del Gaudio[1], Raffaele Clemente Petrella[2*], Salvatore Scognamiglio[1], Paolo Zanetti[1]

[1] Department of Business and Quantitative Studies, University of Naples "Parthenope", Via Generale Parisi 13, Naples, 80132, Italy .
[2] Department of Economic and Legal Studies, University of Naples "Parthenope", Via Generale Parisi 13, Naples, 80132, Italy .

*Corresponding author(s). E-mail(s): raffaeleclemente.petrella@phds.uniparthenope.it;

Abstract

This paper introduces a novel multi-input neural network model for predicting default risk among established Small and Medium-Sized Enterprises (SMEs). The proposed architecture is specifically designed to exploit heterogeneous data sources, including financial and economic time series, tabular numerical features describing firms' governance characteristics, and regional information represented by categorical variables. The model relies on a tailored Deep Neural Network architecture that integrates Feed-Forward Neural Networks for static features with Long Short-Term Memory networks for sequential financial data, and embedding layers for categorical features. This pillar-based structure enables the effective learning of complementary information from different data modalities. Using a large sample of established Italian SMEs, the proposed model achieves the best out-of sample performances, consistently outperforming logistic regression, XGBoost, and a benchmark feed-forward neural network. Furthermore, the modular design of the architecture facilitates model interpretability through the use of Shapley values, revealing that financial-economic indicators are the dominant driver of default risk, while governance characteristics and regional context provide meaningful, complementary contributions. This explainability framework mitigates the black-box nature of deep learning models and is particularly relevant for supporting their responsible adoption in credit risk assessment and regulatory-sensitive contexts.

Keywords: Default prediction, SME, Deep Learning, Machine Learning, Explainable AI

1 Introduction

The field of Artificial Intelligence (AI), particularly Machine and Deep Learning, is undergoing a rapid transformation, impacting the development of businesses and research. Corporate finance is no exception to this phenomenon (Goodell et al. 2021). In this regard, it is highlighted that the integration of AI in this field also offers new opportunities for innovation in the study of various phenomena. For example, AI and Machine Learning techniques can be used to improve corporate default forecasts. A particularly relevant application in this context concerns the prediction of default among Small and Medium Enterprises (SMEs), as they represent the backbone of national and global economies (Surya et al. 2021). Indeed, despite their central role in terms of productivity and employment, SMEs are highly dependent on bank credit and often face structural difficulties in accessing capital markets (Edmister 1972). For this reason, the use of AI and Machine Learning models to predict SME default can provide banks and institutions with more accurate and timely risk assessments, ultimately supporting financial stability and sustainable economic growth (Cheraghali and Molnár 2024). The first studies on corporate default date back to the first half of the twentieth century (Bellovary et al. 2007). These early works were predominantly based on univariate analyses and did not distinguish between different types of firms, treating SMEs and large companies according to similar analytical frameworks. A significant advance in this field was made with the work of Altman et al. (1977), who was among the first to introduce a multivariate approach to default forecasting. This contribution is particularly relevant, as it laid the foundation for subsequent developments aimed at differentiating the analysis based on the type of firm, thus paving the way for models specifically calibrated to the characteristics of SMEs. In particular, these models are based on the assumption that economic and financial data possess strong predictive power in determining the probability of default. Interest in this topic grew following two key events: the 2007–2009 financial crisis Oliveira et al. (2017) and the introduction of the Basel Accord in 2004 (Ciampi et al. 2021). The latter proposed the adoption of standardized models for assessing default forecasts. However, these models proved ineffective in capturing the specificities of SMEs, showing limited predictive capacity (Altman et al. 2008). This limitation led to growing attention towards the study and research of the determinants of SME default (Ciampi et al. 2021; Cheraghali and Molnár 2024). In particular, the resulting literature has highlighted that, despite the progress made, these models have historically relied on linear approaches, such as logistic regression and linear discriminant analysis, appreciated for their simplicity and interpretability (Altman and Sabato 2007; Calabrese and Osmetti 2013; Cheraghali and Molnár 2024). However, this reliance on linear models has raised concerns about their ability to capture the nonlinear and multidimensional dynamics that characterize the behavior of SMEs, thus paving the way for the adoption of more flexible and advanced methodologies. Indeed, despite the progress achieved, several studies have shown that many of the proposed models, based on linear approaches, continue to exhibit significant forecast errors, especially when based exclusively on balance sheet data (Altman et al. 2023; Srhoj et al. 2024; Figini and Giudici 2011). This has fueled growing interest in integrating qualitative, non-financial variables, which offer a broader and more realistic view of the functioning of

SMEs. In support of this view, numerous contributions in the literature have demonstrated that the inclusion of additional information, such as changes in management structure, staff turnover, or the geographical context in which the firm operates, significantly increases the predictive power of models compared to the exclusive use of economic and financial indicators (Algeri et al. 2021; Wang et al. 2023). At the same time, however, the difficulty of effectively incorporating a large number of predictive variables into traditional models has emerged. This finding has highlighted the need to develop more inclusive and flexible approaches capable of also meaningfully integrating non-financial information (Ciampi 2015). Given the limitations of traditional models, significant methodological advances have been made in the literature in recent years, with a growing shift toward more flexible approaches. In particular, there has been a growing adoption of Machine Learning and Deep Learning techniques, capable of handling larger datasets, integrating heterogeneous variables, and modeling nonlinear relationships (Vicario and Coleman 2020). One of the most significant contributions in this direction is the Omega Score Altman et al. (2023), which combines the robustness of Machine Learning algorithms with the joint use of financial and non-financial variables. The effectiveness of the Omega Score has also been confirmed by empirical applications, such as in the case of Croatian SMEs, where it showed superior performance compared to conventional models (Srhoj et al. 2024). Similarly, numerous studies have introduced predictive models based on advanced techniques, such as Random Forests, Support Vector Machines, Extreme Gradient Boosting (XGBoost) and Deep Neural Networks (DNNs), highlighting a significant improvement in predictive accuracy compared to traditional static linear approaches (Kou et al. 2021; Lextrait 2023; Fantazzini and Figini 2009; Sanchez-Medina et al. 2025). Among these, the work of Crosato et al. (2023) stands out, in which the authors analyze the default risk of Italian manufacturing SMEs using Machine Learning models with interpretability tools. They employ two algorithms considered "black-box", such as XGBoost and Feed-Forward Neural Networks (FFNNs) and compare them with traditional linear models to assess which best predicts default and, at the same time, which variables were most important in the default estimation process. To overcome the opacity of complex models, they apply techniques such as Shapley values, to measure the overall contribution of each variable to the probability of default. Despite the significant advances introduced by Machine Learning methods, many of these approaches are not intrinsically designed to capture the temporal structure of data, a critical feature when working with panel data. Indeed, the aforementioned models typically treat data as independent and non-sequential, thus neglecting the evolutionary patterns that often precede corporate default. In this regard, the scientific literature has proposed more sophisticated architectures that address this limitation, such as Recurrent Neural Networks and, more specifically, Long Short-Term Memory (LSTM) Hochreiter and Schmidhuber (1997). LSTMs process data sequentially by selectively retaining and releasing relevant information over time. This is particularly relevant in SME default prediction since default occurrence is a condition that evolves over time and depends on historical economic and financial data, making the integration of time-dependent models essential. Some studies have begun to explore the use of DNNs for default prediction (Giudici et al. 2020; Navarro-Galera et al. 2025; Sariev and Germano 2020), but often

with limited attention to their application to SMEs or to the management of temporal dependencies in data (Barboza et al. 2017; Altman et al. 2020).

This study extends the work of Crosato et al. (2023) in several directions. First, we integrate the classical FFNN framework with additional neural components designed to exploit information available in different formats. Specifically, three building blocks are incorporated: (i) an LSTM network to capture the temporal dynamics of financial variables, (ii) an embedding layer to model regional heterogeneity, and (iii) a feed-forward module that processes static governance attributes. According to the principles of representation learning (Bengio et al. 2013), each component is responsible for extracting a distinct type of information, and the joint architecture combines these representations to generate the final predictions. This modular design allows the model to effectively leverage both time-varying and time-invariant characteristics, as well as categorical spatial features. Second, we extend Crosato et al. (2023) analysis by broadening both the sample and the information content used for forecasting. While their study focuses exclusively on manufacturing SMEs, we consider the universe of Italian SMEs operating across all sectors over the period from 2013 to 2022. This broader perspective allows us to assess the robustness of the proposed model in a broader, more representative, and more heterogeneous context. Furthermore, building on recent advances highlighted in the literature, we complement the set of economic and financial variables with qualitative and governance-related variables, as well as with geographical information. The inclusion of these additional dimensions responds to the evidence that managerial characteristics and geographical context represent key determinants of default risk and can contribute to improved predictive performance compared to models based solely on financial data. Finally, we address the explainability criticism typically associated with deep learning models using the Shapley values. This approach allows us to quantify the marginal contribution of each subnet to the model's predictions.

The remainder of the article is organized as follows. Section 2 reviews the existing literature on default forecasting. Section 3 presents the methodological framework of the static models considered. Section 4 details the architecture of the proposed model and the approach used to ensure interpretability. Section 5 illustrates the data used and the experimental setup. Section 6 discusses the empirical results, while Section 7 concludes the article with final reflections and directions for future research.

2 Predicting the default: A Classification problem

In this work, we adopt a widely used methodology in the literature by formulating the default prediction as a supervised classification task. The goal is to develop a method capable of determining whether an SME is likely to experience default one year in advance, based on a set of observed characteristics. These features may include financial-economic indicators derived from the firm's balance sheet, attributes of its governance system, and the geographical location of its headquarters.

Let $\mathcal{K}$ denote the set of SMEs under analysis and an observation period $\mathcal{T}$. For each firm $K \in \mathcal{K}$ and time $t \in \mathcal{T}$, we define a binary response variable $Y_{T,K} \in \{0, 1\}$,

which encodes the financial status of the firm as follows:

$$Y_{T,K} = \begin{cases} 1, & \text{if company } K \text{ will experience default at time } T, \\ 0, & \text{otherwise.} \end{cases}$$

Let $\boldsymbol{x}_{T,K} \in \mathbb{R}^{q_0}$ denote the vector of explanatory variables for company K at time T, representing the firm's characteristics. We define a scoring function f as:

$$f : \mathbb{R}^{q_0} \to (0, 1),$$

where $f(\boldsymbol{x}_{T,K})$ represents the estimated probability that $Y_{T,K} = 1$, i.e., that the company will go into default given its observed characteristics.

To convert this score into a binary decision, we introduce a threshold $\nu \in (0, 1)$. The classifier $c_\nu : \mathbb{R}^{q_0} \to \{0, 1\}$ is defined as:

$$\hat{Y}_{T,K} = c_\nu(\boldsymbol{x}_{T,K}) = \mathbb{1}_{\{f(\boldsymbol{x}_{T,K}) \geq \nu\}} = \begin{cases} 1, & \text{if } f(\boldsymbol{x}_{T,K}) \geq \nu, \\ 0, & \text{if } f(\boldsymbol{x}_{T,K}) < \nu. \end{cases}$$

The threshold ν is a hyperparameter that must be appropriately calibrated.

A variety of classification algorithms have been proposed in the literature for predicting default, reflecting the evolution of machine and deep learning methods and the increasing availability of firm-level data. The primary distinction resides in the specification of the function f, which characterizes the relationship between the covariates and the probability of default. For a comprehensive overview of existing approaches, we refer the reader to relevant survey studies in the field (Ciampi et al. 2021; Cheraghali and Molnár 2024).

Traditional models such as logistic regression have long served as a standard baseline in this domain due to their simplicity, interpretability, and solid theoretical foundations. However, to capture complex, nonlinear relationships between a firm's characteristics and its likelihood of default, more advanced Machine Learning methods have increasingly been adopted. In particular, models based on DNNs have gained attention for their flexibility and predictive accuracy.

2.1 Logistic Regression

Logistic regression is a widely used statistical method for addressing binary classification problems. It models the decision function f as a composition of two functions:

$$f_{\text{LR}}(\boldsymbol{x}_{T,K}) = \sigma\left(\eta(\boldsymbol{x}_{T,K})\right) = \frac{1}{1 + e^{-\eta(\boldsymbol{x}_{T,K})}},$$

where $\sigma : \mathbb{R} \to (0, 1)$ represents the sigmoid, or logistic, function. This function maps any real-valued input to a number between 0 and 1, making it suitable for expressing probabilities associated with binary outcomes, while $\eta(\boldsymbol{x}_{T,K}) \in \mathbb{R}$, referred to as the linear predictor, is defined as:

$$\eta_{T,K}(\boldsymbol{x}_{T,K}) = \beta_0 + \boldsymbol{\beta}^\top \boldsymbol{x}_{T,K},$$

Here, $\beta_0 \in \mathbb{R}$ denotes the intercept term, $\boldsymbol{\beta} \in \mathbb{R}^{q_0}$ is the vector of regression coefficients, and $\boldsymbol{x}_{T,K} \in \mathbb{R}^{q_0}$ is the feature vector at time t for the i-th observation.

In this framework, logistic regression computes a linear combination of the input features and applies the sigmoid function to transform the result into a probability. Given a dataset $\mathcal{D} = \{(\boldsymbol{x}_{T,K}, y_{T,K})\}_{k=1,\,t=1}$ the parameters β_0 and $\boldsymbol{\beta}$ are estimated using the maximum likelihood estimator.

2.2 Deep Neural Networks

DNNs are nonlinear models inspired by the structure and functioning of the human brain. They have been extensively developed for a wide range of applications in finance (Heaton et al. 2017; Castellani et al. 2021; Giorgi et al. 2024). A DNN consists of layers of interconnected computational units, called neurons, which are trained to learn useful representations of the data. Through this training process, the model maps inputs into a high-dimensional feature space where relevant patterns are automatically extracted. This capability to discover and encode meaningful features is referred to as *representation learning* (Bengio et al. 2013).

DNNs can be viewed as a generalization of logistic regression. Although it preserves the structure of composition of function $f = \sigma(\eta(\boldsymbol{x}_{T,K}))$, the key distinction lies in the nature of the η, which is now modeled as a nonlinear function of the input covariates. This formulation allows the model to capture complex, nonlinear relationships between the covariates and the default probability. The nonlinearity is introduced by transforming the input $\boldsymbol{x}_{T,K}$ through a sequence of hidden layers that generate a new set of features, denoted by $\boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T,K})$, which are optimized to represent the data in a way that is relevant to the predictive task. Specifically, the predictor is expressed as

$$\eta(\boldsymbol{x}_{T,K}) = \beta_0 + \boldsymbol{\beta}^\top \boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T,K}).$$

The representation $\boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T,K}) \in \mathbb{R}^{q_h}$ is obtained by applying a composition of h nonlinear transformations (layers) to the original input. This composition of transformations can be written as

$$\boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T,K}) = \left(\boldsymbol{z}^{(h)} \circ \cdots \circ \boldsymbol{z}^{(1)}\right)(\boldsymbol{x}_{T,K}) \in \mathbb{R}^{q_h}, \tag{1}$$

Let $q_i \in \mathbb{N}$ denotes the number of neurons in the i-th layer, each layer $\boldsymbol{z}^{(k)}$ maps $\mathbb{R}^{q_{i-1}}$ to $\mathbb{R}^{q_i}$. The transformation carried out by the j-th neuron in layer i is defined by:

$$\boldsymbol{z}^{(i-1)} = (z_1^{(i-1)}, \ldots, z_{q_{i-1}}^{(i-1)})^\top \in \mathbb{R}^{q_{i-1}} \quad \mapsto \quad z_j^{(k)}(\boldsymbol{z}^{(i-1)}) = \phi\left(\sum_{s=1}^{q_{i-1}} w_{j,s}^{(i)}\, z_s^{(i-1)} + b_j^{(i)}\right),$$

where $\phi : \mathbb{R} \to \mathbb{R}$ is a nonlinear activation function, $w_{j,s}^{(i)} \in \mathbb{R}$ is the weight connecting the s-th neuron in the previous layer to the j-th neuron in the current layer, and $b_j^{(i)} \in \mathbb{R}$ is the corresponding bias term. By stacking multiple such transformations, DNNs are capable of modeling complex nonlinear dependencies between inputs and

outputs, significantly enhancing their expressiveness and flexibility compared to linear models.

2.3 Evaluation Metrics

There are several performance measures used to evaluate a model's ability to correctly classify observations. The choice of metric depends on the problem context and the specific objectives of the analysis. In this paper, we evaluate model performance using a set of threshold–based metrics and, more globally, the Area Under the Curve (AUC).

- Given a decision threshold ν, several metrics can be defined. Some of the most popular threshold–based metrics are:
  - Accuracy: proportion of correctly classified observations,

$$\text{Accuracy}(f,\nu) = \frac{\sum_{i=1}^{N} \mathbf{1}\{f(\mathbf{x}_{K,T}) \geq \nu\}\mathbf{1}\{Y_{K,T}=1\} + \sum_{i=1}^{N} \mathbf{1}\{f(\mathbf{x}_{K,T}) < \nu\}\mathbf{1}\{Y_{K,T}=0\}}{N}.$$

  - Sensitivity: proportion of positives correctly identified,

$$\text{Sensitivity}(f,\nu) = \frac{\sum_{i=1}^{N} \mathbf{1}\{f(\mathbf{x}_{K,T}) \geq \nu\}\mathbf{1}\{Y_{K,T}=1\}}{\sum_{i=1}^{N} \mathbf{1}\{Y_{K,T}=1\}}.$$

  - Specificity: proportion of negatives correctly identified,

$$\text{Specificity}(f,\nu) = \frac{\sum_{i=1}^{N} \mathbf{1}\{f(\mathbf{x}_{K,T}) < \nu\}\mathbf{1}\{Y_{K,T}=0\}}{\sum_{i=1}^{N} \mathbf{1}\{Y_{K,T}=0\}}.$$

- While threshold–based metrics evaluate performance at a single operating point, the AUC provides a global measure of discriminative ability across all thresholds and can be interpreted as the probability that a positive instance receives a higher score than a negative one. Using sensitivity and specificity, the ROC curve is represented by sensitivity versus $1-$specificity. Given ordered thresholds $\{\nu_i\}_{i=0}^{n}$, the AUC is approximated via

$$G(f,\nu) \approx \sum_{i=1}^{n} \frac{\text{Sensitivity}(f,\nu_i) + \text{Sensitivity}(f,\nu_{i-1})}{2} \Big[(1-\text{Specificity}(f,\nu_i)) - (1-\text{Specificity}(f,\nu_{i-1}))\Big].$$

2.4 About the limitations of conventional classifiers

Existing approaches in the literature exhibit some limitations. In many cases, the set of covariates includes only information about the firm at the same time the default occurs (Ciampi et al. 2021; Cheraghali and Molnár 2024). However, this choice has limited practical utility, especially when the objective is to anticipate default *before* it actually manifests. For instance, if we are at time $T-1$, our goal is to predict the firm's default status at time T. This implies that the probability $\mathbb{P}(Y_{T,K}=1)$ should not be estimated based on covariates observed at time T, but rather using information

available at the preceding time step $T-1$. In other word, they should use $\boldsymbol{x}_{T-1,K}$ instead of $\boldsymbol{x}_{T,K}$, and then calculate the default probability as:

$$\hat{Y}_{T,K} = c_\nu(\boldsymbol{x}_{T-1,K}) = \mathbb{1}_{\{f(\boldsymbol{x}_{T-1,K})\geq\nu\}} \tag{2}$$

In addition, historical patterns are often neglected, see for example Crosato et al. (2023)(Crosato et al. 2023). To properly assess a firm's financial condition at time T, could be suboptimal to rely solely on information observed at time $T-1$. Instead, one should consider a broader temporal window. Given a chosen lookback period $\tau \in \mathbb{N}$, covariates from previous time steps $T-1, T-2, \ldots, T-\tau$ should be incorporated into the model, extending (2) as follows:

$$\hat{Y}_{T,K} = c_\nu(\boldsymbol{x}_{T-1,K}, \boldsymbol{x}_{T-2,K}, \ldots, \boldsymbol{x}_{T-\tau,K}) = \mathbb{1}_{\{f(\boldsymbol{x}_{T-1,K}, \boldsymbol{x}_{T-2,K}, \ldots, \boldsymbol{x}_{T-\tau,K})\geq\nu\}} \tag{3}$$

These past observations can provide additional insights into the firm's trajectory and evolving risk profile. Capturing such temporal dependencies is crucial for generating more accurate and forward-looking SME default predictions.

3 A Multi-Input Network

In this section, we formally present the proposed method. We introduce a DNN architecture designed to process and exploit a large and heterogeneous set of input variables. The rationale behind this choice is that DNNs are particularly effective at modeling complex, nonlinear relationships and can integrate different sources of information. For example, we believe that enabling the model to fully leverage temporal information can significantly enhance predictive accuracy. For this reason, we allow the model to process the time series contained in the data. More precisely, the network receives three main categories of input features, each representing a different informational dimension relevant to the prediction task:

- **Financial-Economic Data**: These are balance sheet indicators that evolve over time—such as assets, liabilities, and revenues. Capturing their temporal patterns is crucial for understanding the company's trajectory and financial health. These indicators are typically reported on an annual basis, and when collected over multiple years, they naturally form time series data.
- **Governance Data**: These features encompass static or quasi-static attributes describing the corporate governance and leadership structure, such as the composition of the board of directors, the backgrounds of executive members, and the roles of other key decision makers. These features typically include metrics such as the average experience and tenure of leaders, the size of the governing body, indicators of diversity, and other governance-related attributes. As these characteristics tend to change infrequently-and their temporal evolution contributes minimally to predictive performance-only their values at the evaluation date T are used as predictors.
- **Regional Data**: These refer to contextual and geographical features such as the region's economic indicators, sectoral risks, or macroeconomic conditions in

which the company operates. They provide valuable exogenous information and are integrated using a separate processing path.

Each group of variables is passed through a dedicated subnetwork (or *subnet*), architected conforming to the specific nature of the input it handles.

The outputs of the three subnets are concatenated to form a unified latent representation, which offers a comprehensive view of the company's trajectory by integrating financial-economic performance, governance structure, and the broader economic environment. Finally, this aggregated feature vector is fed into a classification head that produces the probability score indicating the likelihood of the company experiencing default. We refer to this specialized architecture as Multi-Input Neural Network (MINET), emphasizing both its multi-input structure and its focus on predicting default. Naturally, this tri-input design can be extended to accommodate additional variable sets, should the analysis require a more complex input structure. In the following, we formally describe how the three subnets work.

3.1 Processing Financial-Economic Variables: Long Short-Term Memory Networks

Financial-Economic variables are numerical variables that evolve over time. They are available in the form of time-series. Choosen a lookback period τ, the time series of the financial-economic variables are denoted as $\left(\boldsymbol{x}_{t,K}^{(\text{financial})}\right)_{T-\tau\leq t<T}$, with $\boldsymbol{x}_{t,K}^{(\text{financial})} \in \mathbb{R}^{q^{(\text{financial})}}$. These series capture the company's trajectory over the preceding τ years, relative to the evaluation date T. To effectively model this temporal structure, it is essential to use neural architectures that can capture and learn from such dependencies across time. Recurrent Neural Networks are explicitly designed for this purpose. Unlike traditional FFNNs, Recurrent Neural Networks possess a recurrent structure that allows information to persist across time steps by maintaining a hidden state that is updated sequentially. This enables them to model complex temporal dynamics and capture the evolution of financial indicators over time. A particularly powerful variant of this class of DNNs is, as exposed, LSTM network (Hochreiter and Schmidhuber 1997), which extends the standard recurrent architecture by incorporating gating mechanisms. These gates store and release long-term information through a memory cell. The combination of these memory cells is particularly useful and interesting for various applications in finance that require sequential data. A LSTM layer

with $q_h^{(\text{financial})}$ produces the output $\boldsymbol{z}_t$ as follows:

$$\begin{aligned}
\boldsymbol{f}_t &= \boldsymbol{f}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}\right) = \sigma\left(\boldsymbol{w}_0^{(f)} + W^{(f)}\boldsymbol{x}_{t,K}^{(\text{financial})} + V^{(f)}\boldsymbol{z}_{t-1}\right) \in \mathbb{R}^{q_h^{(\text{financial})}}, \\
\boldsymbol{i}_t &= \boldsymbol{i}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}\right) = \sigma\left(\boldsymbol{w}_0^{(k)} + W^{(k)}\boldsymbol{x}_{t,K}^{(\text{financial})} + V^{(k)}\boldsymbol{z}_{t-1}\right) \in \mathbb{R}^{q_h^{(\text{financial})}} \\
\boldsymbol{o}_t &= \boldsymbol{o}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}\right) = \sigma\left(\boldsymbol{w}_0^{(o)} + W^{(o)}\boldsymbol{x}_{t,K}^{(\text{financial})} + V^{(o)}\boldsymbol{z}_{t-1}\right) \in \mathbb{R}^{q_h^{(\text{financial})}} \\
\boldsymbol{r}_t &= \boldsymbol{r}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}\right) = \tanh\left(\boldsymbol{w}_0^{(r)} + W^{(r)}\boldsymbol{x}_{t,K}^{(\text{financial})} + V^{(r)}\boldsymbol{z}_{t-1}\right) \in \mathbb{R}^{q_h^{(\text{financial})}}, \\
\boldsymbol{c}_t &= \boldsymbol{c}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}, \boldsymbol{c}_{t-1}\right) = \boldsymbol{f}_t \odot \boldsymbol{c}_{t-1} + \boldsymbol{i}_t \odot \boldsymbol{r}_t \in \mathbb{R}^{q_h^{(\text{financial})}}, \\
\boldsymbol{z}_t &= \boldsymbol{z}\left(\boldsymbol{x}_{t,K}^{(\text{financial})}, \boldsymbol{z}_{t-1}, \boldsymbol{c}_{t-1}\right) = \boldsymbol{o}_t \odot \phi\left(\boldsymbol{c}_t\right) \in \mathbb{R}^{q_h^{(\text{financial})}},
\end{aligned}$$

where $W^{(s)} \in \mathbb{R}^{q_h^{(\text{financial})} \times q_0}$, $V^{(s)} \in \mathbb{R}^{q_h^{(\text{financial})} \times q_h^{(\text{financial})}}$, and $\boldsymbol{w}_0^{(s)} \in \mathbb{R}^{q_h^{(\text{financial})}}$, for $s \in \{i, o, f, r\}$, are network coefficients to optimise. In these equations $\odot$ denotes the Hadamard product and σ represents the sigmoid activation function. The generic function ϕ, applied to the memory cell $\boldsymbol{c}_t$, in our case is the hyperbolic tangent activation, i.e., $\phi(\cdot) = \tanh(\cdot)$. The memory state at time t, denoted as $\boldsymbol{c}_t$ (with initial value $\boldsymbol{c}_0 = 0 \in \mathbb{R}^{q_1}$), is influenced by the operations of three gates: the forget gate $\boldsymbol{f}_t$, the input gate $\boldsymbol{i}_t$, and the new cell input $\boldsymbol{r}_t$. At each time step t, the update process begins with the forget gate, which removes outdated or irrelevant information from the previous cell state $\boldsymbol{c}_{t-1}$. Subsequently, the input gate $\boldsymbol{i}_t$ integrates new information derived from the current cell input $\boldsymbol{r}_t$ in order to refresh the memory state $\boldsymbol{c}_t$. The final output of the LSTM unit is then obtained by applying the output gate $\boldsymbol{o}_t$ to the updated memory cell. In typical applications, LSTM networks are trained so that the final hidden state $\boldsymbol{z}_T(\boldsymbol{x}_{T-1,K})$ represents a set of features that are optimally extracted from the input time series. These features are learned to effectively capture the dynamics of the sequence in a way that is most informative for predicting the response variable $Y_{T,K}$. In this regard, $\boldsymbol{z}(\boldsymbol{x}_{T-1,K})$ serves a role analogous to that of the feature vector $\boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T,K})$ used in classical DNNs described in Section 2.2. As noted in Section 2.4, because our goal is to anticipate default at time T, the final activation of the LSTM layer corresponds to the data available at time $T-1$. To simplify the notation in the following sections, we set

$$\boldsymbol{z}_{T-1,K}^{(\text{financial})} = \boldsymbol{z}(\boldsymbol{x}_{T-1,K}).$$

Figure 1 provides a visual representation of the internal structure of the LSTM unit.

3.2 Processing Governance Variables: Feed-Forward Neural Networks

Governance variables are numerical features that capture the structural characteristics of a company. Such variables have been shown to be relevant predictors of firm performance and stability in the literature (Ballester et al. 2020; Nwogugu 2007; Switzer

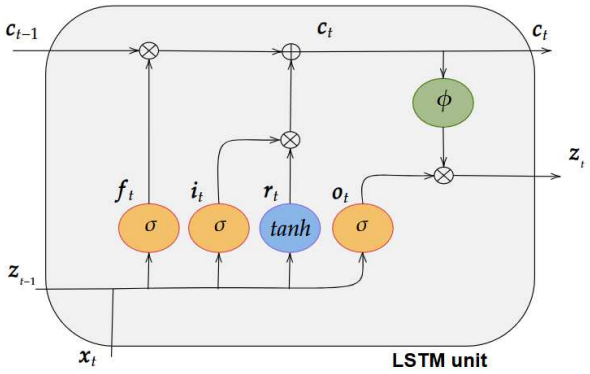

Fig. 1: Schematic representation of the LSTM unit.

et al. 2018). Given that these features tend to exhibit limited variation over time, we focus on their values at a specific time point T. We represent the governance-related features using the vector $\boldsymbol{x}_{T-1,K}^{(\text{governance})} \in \mathbb{R}^{q^{(\text{governance})}}$. This vector is processed by a sequence of FFNN layers, which learn a representation of the governance characteristics optimized with respect to the target variable. The resulting representation is consistent with the deep feature structure introduced in Equation (1):

$$\boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T-1,K}^{(\text{governance})}) = \left(\boldsymbol{z}^{(h)} \circ \cdots \circ \boldsymbol{z}^{(1)}\right)(\boldsymbol{x}_{T-1,K}^{(\text{governance})}) \in \mathbb{R}^{q_h^{(\text{governance})}}, \tag{4}$$

where each $\boldsymbol{z}^{(i)}$ denotes the transformation applied at the i-th hidden layer, and $q_h^{(\text{governance})}$ is the dimensionality of the final deep feature representation. For simplicity, we set:

$$\boldsymbol{z}^{(\text{governance})}(\boldsymbol{x}_{T-1,K}^{(\text{governance})}) = \boldsymbol{z}^{(h:1)}(\boldsymbol{x}_{T-1,K}^{(\text{governance})}).$$

3.3 Processing Regional feature: Entity Embedding Layer

The region in which a company operates is represented as a categorical variable, $r \in \mathcal{R} = \{\texttt{Abruzzo}, \texttt{Basilicata}, \ldots, \texttt{Valle D'Aosta}\}$. This variable conveys structural and economic contextual information, such as regional economic performance, labor market conditions, infrastructure quality, and the broader business environment, all of which may impact a firm's default probability .

To incorporate categorical variables like region into DNN models, Entity Embedding (EE) Networks are commonly employed. These networks map discrete categorical values into continuous, low-dimensional vector representations. Let $q_{\mathcal{L}} \in \mathbb{N}$ denote the dimensionality of the embedding space. Each category $r_k \in \mathcal{R}$ is associated with a learnable embedding vector in $\mathbb{R}^{q_{\mathcal{L}}}$, which is optimized during model training. This technique follows the approach proposed in Guo and Berkhahn (2016) for effectively incorporating categorical features into deep learning architectures.

Formally, the embedding layer for the region variable defines the following mapping:

$$\boldsymbol{z}^{(\text{region})} : \mathcal{R} \longrightarrow \mathbb{R}^{q_h^{(\text{region})}}, \quad r_K \longmapsto \boldsymbol{z}^{(\text{region})}(r_K).$$

The underlying intuition is that regions with similar influence on the target outcome, such as default a year ahead, will be mapped to nearby points in the embedding space (as measured by Euclidean distance), while dissimilar regions will be embedded farther apart. This allows the model to exploit latent relationships between categorical levels that might not be evident in one-hot or ordinal encodings.

3.4 Aggregating: Decoder

The three feature sets corresponding to financial-economic, governance, and regional covariates-extracted by their respective sub-networks-are concatenated to form a unified representation. This representation captures the temporal dynamics of balance sheet indicators, the structural attributes of the governance, and the contextual regional information:

$$\boldsymbol{z}_{T-1,K}^{(\text{full})} = \left(\boldsymbol{z}_{T-1,K}^{(\text{financial})}, \boldsymbol{z}_{T-1,K}^{(\text{governance})}, \boldsymbol{z}^{(\text{region})}\right) \in \mathbb{R}^{q_h^{(\text{financial})} + q_h^{(\text{governance})} + q_h^{(\text{region})}}.$$

It serves as the input to a final sub-network $\eta_{\text{dec}} : \mathbb{R}^{q_h^{(\text{financial})} + q_h^{(\text{governance})} + q_h^{(\text{region})}} \mapsto (0,1)$. It consists of a set of $h_d \in \mathbb{N}$ FNN layers denoted as $\boldsymbol{z}_{\text{dec}}^{j}, 1 \leq j \leq h_d$, which uses these aggregated covariates to perform the classification task:

$$\eta_{\text{decoder}}(\boldsymbol{z}_{T-1,K}^{(\text{full})}) = \left(\boldsymbol{z}_{\text{decoder}}^{(h_d)} \circ \cdots \circ \boldsymbol{z}_{\text{decoder}}^{(1)}\right)(\boldsymbol{z}_{T-1,K}^{(\text{full})})$$

This classification module follows the same structure described in Section 2.2, consisting of a sequence of FFN layers. Formally, the predicted output is given by:

$$Y_{T,K} = f_{\texttt{MI-NET}}(\left(\boldsymbol{x}_{t,K}^{(\text{financial})}\right)_{T-\tau \leq t < T}, \boldsymbol{x}_{T-1,K}^{(\text{governance})}, r_K) = \sigma(\eta_{\text{decoder}}(\boldsymbol{z}_{T-1,K}^{(\text{full})})).$$

A graphical illustration of the complete architecture, along with its constituent sub-networks, is provided in Figure 2.

3.5 Explainability

DNNs are generally perceived as *black box* since they are not fully transparent and is hard to interpret. This makes it challenging for practitioners to understand how the model derives the predictions, as well as to validate or explain those results. This issue could represent a big problem especially in the financial industry which is a highly regulated sector. The Shapley Value is a powerful tool for interpreting black box models. Originally developed in game theory, it allows for the distribution of the "payout" among players based on their individual contributions. When applied to black box models, the Shapley Value quantifies the contribution of each feature to a prediction by evaluating all possible combinations of features. Let $f(\boldsymbol{x})$ be a black

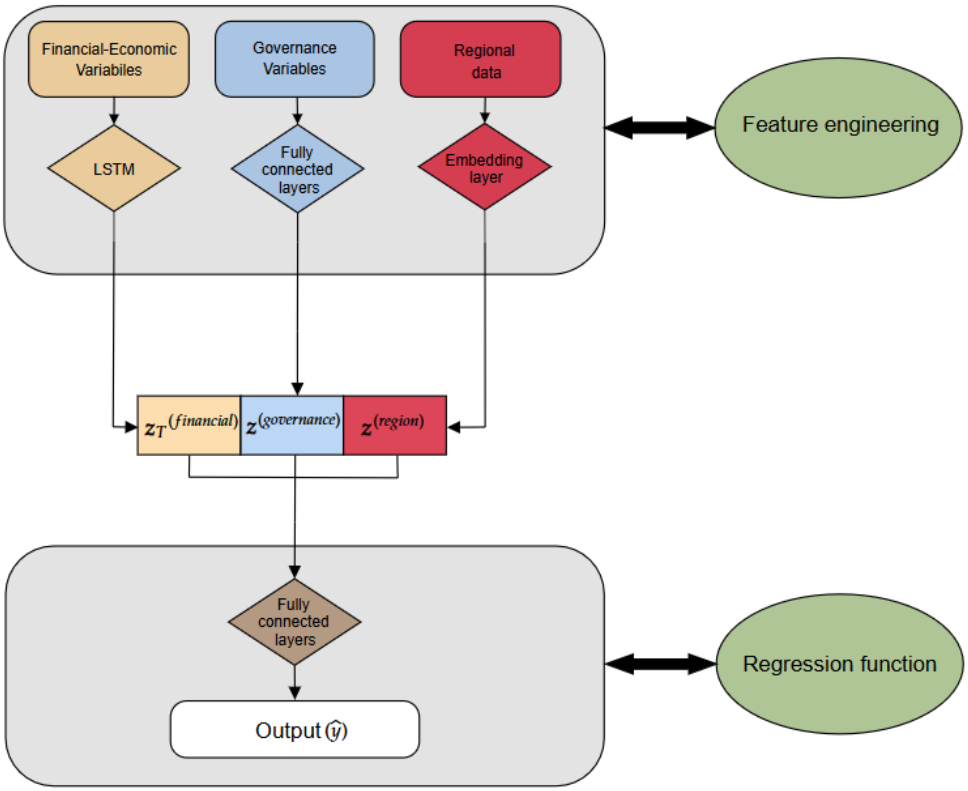

Fig. 2: Architecture of the proposed Deep Learning Model

box model. The Shapley Value for each feature x_i represents the average marginal contribution of that feature to the model's prediction, considering all possible subsets of features (Winter 2002). Let $S \subseteq \{1, 2, \ldots, q_0\} \setminus \{i\}$ be any subset of the feature set, excluding the feature x_i, the Shapley Value ϕ_i for a feature x_i is defined as:

$$\phi_i(f, \boldsymbol{x}) = \sum_{S \subseteq \{1,2,\ldots,q_0\} \setminus \{i\}} \frac{|S|!\,(q_0 - |S| - 1)!}{q_0!} \left[f(\boldsymbol{x}_{S \cup \{i\}}) - f(\boldsymbol{x}_S) \right]$$

where $f(\boldsymbol{x}_S)$ is the prediction of the model using only the features in subset S, while $f(\boldsymbol{x}_{S \cup \{i\}})$ is the prediction of the model using the features in S plus the feature x_i. The Shapley Value is the weighted average of the marginal contributions of feature x_i across all possible subsets S, ensuring a fair distribution of the "credit" for the prediction $f(\boldsymbol{x}_{T,K} =$ among all features. Given an observation $\boldsymbol{x}_j$, the Shapley Value

related to the i-th feature can be defined as:

$$\sum_{i=1}^{q_0} \phi_i(f, \boldsymbol{x_j}) = f(\boldsymbol{x_j}) - f(\emptyset)$$

where $f(\emptyset)$ represents the model's prediction with no input features, often interpreted as the baseline prediction. In the context of default prediction, this tool provides interpretability for the regression function estimating the probability of default. Since our model receives a large volume of input data, directly applying the Shapley Value to the original covariates is computationally infeasible. To address this, we apply the technique to the second stage of the network architecture, where the outputs of three sub-networks—capturing information on financial-economic history, governance composition, and regional context—are combined to generate predictions. In this setting, we can effectively measure the contribution of each of these three components to η_{decoder}. One could also derive an Importance Measure (IM) related to the three input summaring the Shapley Values (in absolute value) for all the observations $1 \leq j \leq n$:

$$\texttt{IM}^{(financial)} = \frac{\frac{1}{n}\sum_{j=1}^{n} |\phi_i(\eta_{\text{decoder}}, \boldsymbol{z}_{T-1,K}^{(financial)})|}{\texttt{IM}^{(financial)} + \texttt{IM}^{(governance)} + \texttt{IM}^{(region)}} \in (0,1)$$

$$\texttt{IM}^{(governance)} = \frac{\frac{1}{n}\sum_{j=1}^{n} |\phi_i(\eta_{\text{decoder}}, \boldsymbol{z}_{T-1,K}^{(governance)})|}{\texttt{IM}^{(financial)} + \texttt{IM}^{(governance)} + \texttt{IM}^{(region)}} \in (0,1),$$

$$\texttt{IM}^{(region)} = \frac{\frac{1}{n}\sum_{j=1}^{n} |\phi_i(\eta_{\text{decoder}}, \boldsymbol{z}_{K}^{(region)})|}{\texttt{IM}^{(financial)} + \texttt{IM}^{(governance)} + \texttt{IM}^{(region)}} \in (0,1).$$

Each importance measure quantifies the average percentage contribution of a given set of features to the model's predictions for the sample under analysis.

4 Data and numerical experiments

This study focuses on the analysis of Italian SMEs. The data were sourced from the AIDA (*Analisi Informatizzata delle Aziende Italiane*) database[1], which provides key financial statement values and other information on all companies operating within Italy. The selected reference period spans from 2013 to 2022.

Using the aforementioned data warehouse, SMEs were identified in accordance with the definition provided by the Italian Ministerial Decree of April 18, 2005[2]. This decree classifies SMEs as enterprises with fewer than 250 employees and enterprises with either an annual turnover not exceeding 50 million euros or a total annual balance sheet not exceeding 43 million euros.[3] In addition, since the analysis is intended to focus on SMEs without being influenced by short-lived companies, it was determined

[1] https://aida-r1.bvdinfo.com

[2] https://www.mimit.gov.it/images/stories/documenti/DM_18_4_2005_Definizione_PMI.pdf

[3] To align with this definition, specific filters were applied within the AIDA database to select entities meeting these criteria. Specifically, the variable *"Revenue from sales and services"* was used as a proxy for annual turnover, while the variable *"Total assets"* was employed as a substitute for the total annual balance sheet value.

that companies should survive for at least three years. This criterion was adopted to ensure that the analysis included only companies with some stability and ability to operate in the medium to long term. Microenterprises were excluded from the analysis. In line with the relevant ministerial decree, they are defined as businesses with fewer than ten employees and an annual turnover or balance sheet total not exceeding €2 million. This exclusion is motivated by several converging considerations. From a regulatory and credit-risk standpoint, although Basel capital requirement frameworks do not formally define micro-enterprises as a separate exposure class from SME borrowers, the banking literature consistently documents that these firms exhibit distinct credit-risk characteristics and monitoring requirements, owing to more limited collateral, shorter credit histories, and lower credit ratings (Belás et al. 2016). From a data-availability standpoint, a large share of the excluded micro-enterprises fall within the stricter size thresholds of Article 2435-ter of the Italian Civil Code, which permits an even more simplified "micro" financial statement format (omitting the notes to the accounts and the management report); more broadly, firms in this size range are also frequently eligible for the abbreviated statement format under Article 2435-bis, resulting in reduced balance sheet disclosure relative to standard-form filers. This limits the granularity of the financial time series required by the LSTM component of our architecture. Relatedly, the governance indicators employed in our framework presuppose a governance structure that is largely absent in micro-enterprises, which in Italy are predominantly organized as sole proprietorships or single-owner entities. For such firms, these indicators would likely be exhibit limited cross-sectional and temporal variation, thereby providing little discriminative information (e.g., a single, unchanging owner-manager), undermining the discriminative contribution of the governance subnetwork within MI-NET. From a conceptual standpoint, micro-enterprises also exhibit smaller and more heterogeneous organizational structures compared to SMEs, which risks distorting sample comparability (McKevitt and Davis 2013), and have historically shown higher and more volatile insolvency rates, complicating the analysis of long-term performance (Nartisa 2012). Furthermore, prior literature suggests that default among micro-enterprises may be more strongly influenced by factors tied to the individual entrepreneur (such as personal guarantees or household wealth) rather than by firm-level financial fundamentals (Karas and Režňáková 2021; Zheng et al. 2024), indicating that micro-enterprise default may constitute a distinct phenomenon warranting separate modeling. Taken together, these considerations motivate our focus on established SMEs, for which financial fundamentals and governance structures are expected to be well-defined and informative for default prediction, consistent with the modeling assumptions underlying MI-NET. This focus further allows us to investigate growth and sustainability dynamics that are representative of stable, consolidated business environments. SMEs, in fact, constitute an essential component of the Italian economic fabric and play a key role in terms of development and employment (Ciampi 2015). To handle missing observations, we adopted a multi-step procedure. Fields with a large missing rate were discarded prior to further processing. Firms lacking a minimum financial history (reflected by missing values in key structural variables) were excluded, as insufficient historical depth would compromise the informativeness of the LSTM component. Residual missing values were imputed using the firm-specific

median where sufficient within-firm information was available, and the year-specific cross-sectional median otherwise. Fields retaining an excessive share of missing values after this procedure were removed from the final feature set. This filtering process resulted in a total of 81186 units analyzed. These enterprises represent the SMEs, excluding microenterprises, that were active between 2013 and 2022. The analyzed group includes both companies that remain operational and those that defaulted during the study period.
Once the selection of units was made, based on the criteria defined by the Ministerial Decree of April 18, 2005, a detailed collection of different types of features was carried out (see the A for further details). First, balance sheet and income statement variables, which form the core of the economic analysis of enterprises, were considered. These variables include information on the balance sheet, revenues and liabilities. The selection of these indicators makes it possible to assess the financial-economic stability and profitability of firms. Similarly, employee-related variables such as the number of employees and productivity per worker were included in the analysis. Following the existing literature, a number of governance-related variables were included. This information provides a clear view of the corporate configuration and represents the second input of reference, presented earlier. The analysis of corporate governance and strategic decision making is critical to understanding the leadership dynamics and future direction of the SMEs. Governance-related variables were constructed by aggregating records for all individuals holding a decision-making role – board members, senior managers, and advisors – associated with each firm in a given year. For binary or categorical attributes, such as gender or whether a decision-maker also holds an equity stake, we computed the corresponding share or count among all associated individuals (e.g., the percentage of female decision-makers, or the number of advisors). For continuous attributes, such as age, we computed the average across individuals. A broader set of governance indicators was initially considered during the exploratory stage; however, some variables were subsequently discarded due to a high incidence of missing values or limited informative variability, resulting in the final subset reported in Appendix A. Finally, as outlined, the region to which each SME's operational headquarters belongs was derived. Specifically, this information was retrieved from the ZIP code of the firm's operational headquarters.

5 Results

This section presents a quantitative evaluation of the predictive performance of the models under consideration. The objective is to assess whether the proposed method allows to obtain an improvement in terms of predictive performance with respect to the other methods available in the literature.

5.1 Models Performance

We compare the performance of our MI-Net with three benchmark methods: Logistic Regression (hereafter LR), XGBoost, and the FFNN model proposed in Crosato et al. (2023). The evaluation focuses on their ability to predict default one year in advance,

using the AUC-ROC as the primary metric. Two different forecasting scenarios were considered:

- Scenario 2021: Assuming we are at 31 December 2020, the objective is to predict the probability of default for companies in 2021.
- Scenario 2022: Assuming we are at 31 December 2021, the objective is to predict the probability of default for companies in 2022.

Model	**ROC-AUC**			
	2021		**2022**	
	Train set	**Test set**	**Train set**	**Test set**
LR	0.910	0.966	0.912	0.855
XGBoost	0.924	0.973	0.926	0.920
FFNN	0.926	0.972	0.899	0.897
MI-Net	**0.938**	**0.978**	**0.941**	**0.927**

Table 1: ROC-AUC values for training and test sets across 2021 and 2022.

The performance measured in terms of AUC-ROC in the considered scenario can be seen in Table 1.

It shows how, for both scenarios, the MI-Net obtains the highest ROC-AUC, both in the training and in the test set, thus signaling the best adaptation and generalization of the data. In 2022, in particular the differences are even more pronounced, especially in the test set. Nevertheless, for a more complete understanding of the problem, we show also the out-of-sample results for some of the additional metrics introduced at the beginning of this section. These metrics are calculated after determining the optimal threshold using the Youden Point method (Youden 1950). They are reported on the Table 2.

The MI-Net model shows the highest accuracy in both 2021 and 2022, confirming its superior overall ability to correctly classify SMEs when compared to the other models under examination. In the 2021 scenario, the proposed model achieves improvements across all evaluation metrics. In particular, while MI-Net and XGBoost reach the same level of sensitivity, MI-Net provides a higher accuracy and specificity, highlighting a more balanced predictive behaviour. The FFNN model also surpasses the LR in terms of accuracy, suggesting that the adoption of non-linear architectures offers relevant advantages in this context, although the gains remain less pronounced than those obtained by MI-Net. In 2022, XGBoost achieves the highest sensitivity among all models, but this result is accompanied by a reduction in accuracy and specificity, indicating an unfavourable trade-off: the improvement in default identification comes at the cost of a higher rate of false positives. MI-Net maintains a more efficient balance

Model	Predictive Performance					
	2021			2022		
	Accuracy	Specificity	Sensitivity	Accuracy	Specificity	Sensitivity
LR	0.890	0.890	0.945	0.902	0.902	0.647
XGBoost	0.834	0.834	**0.967**	0.880	0.883	**0.824**
FFNN	0.893	0.893	0.945	0.912	0.912	0.647
MI-Net	**0.906**	**0.906**	**0.967**	**0.920**	**0.920**	0.765

Table 2: Out-of-sample predictive performance across 2021 and 2022, reported in terms of accuracy, specificity and sensitivity.

between sensitivity and specificity and therefore yields the highest overall accuracy. The FFNN and LR models display similar sensitivity levels, although FFNN obtains a higher accuracy, further reinforcing the benefits associated with non-linear representations. Overall, these results suggest that the use of a tool able to integrate non-linearity, heterogeneity and temporal sequences leads to a substantial improvement in predictive performance with respect to traditional approaches. This is particularly relevant in the context of default prediction, where misclassification may have serious consequences. Moreover, the behaviour observed on the test set highlights the greater ability of the proposed model to adapt to new data, mitigating the overfitting issues that typically affect other predictive models Figure 3 shows the out-of-sample ROC curves for 2021 and 2022, highlighting the predictive differences among the models. It is important to underline that the ROC curves have a staircase shape due to the data distribution. In fact, for both years considered, the dataset contains few positive cases and a large number of negative cases, with a further decrease in positives in 2022 compared to 2021. As a result, graphically, the true positive rate increases through discrete jumps every time a true positive is correctly classified, while between these jumps the curve remains horizontal as only false positives are added. Since the true positives are limited in number, a segmented curve is produced, which is a typical result in highly unbalanced data sets with few positive instances (Janssens and Martens 2020). It is evident that the LR model, although serving as a benchmark, is limited by its linear nature. As already highlighted, the MI-Net instead offers the best balance between sensitivity and specificity, an aspect that also clearly emerges from the analysis of the curves, which show an area under the curve greater than the other models. Overall, these results confirm the value of advanced multi-input architectures for the development of more accurate and reliable tool in the prediction of SME default.

5.2 Explainability: the contribution of the three inputs

As discussed in the its specific section, the proposed model integrates three distinct groups of variables, corresponding to as many informational blocks: financial-economic

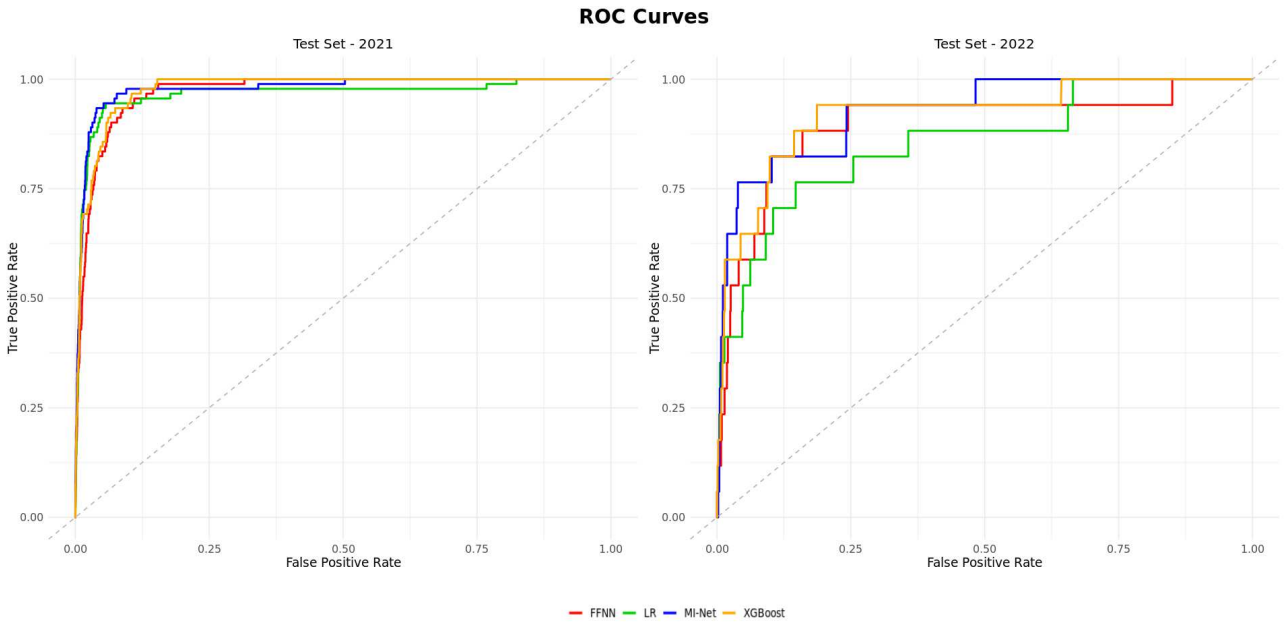

Fig. 3: ROC Curve for models out-sample for the two scenario considered

indicators, governance characteristics, and the geographic context, represented by the region. In this section, we analyze the relative impact of each block on the final default prediction, with the objective of understanding which component most influences the behavior of the model.
This objective falls within the broader goal of providing explainability to AI-based models which, although offering competitive or even superior performance compared to linear models, often lack transparency regarding the contribution of individual variables to the output.
To this end, we adopted, as exposed, the use of Shapley Values. The operation of this technique has been explained in the previous sections.

As shown in Figure 4, the dominant contribution arrives from financial-economic variables, a predictable result given their nature directly related to the economic performance of firms. However, it is interesting to note that governance-related variables (such as board composition, gender, age, or decision-making structure) also show a significant influence, despite being numerically smaller. This result reinforces the idea, which has already emerged in recent literature, that qualitative factors related to organizational structure may play a significant role in determining corporate stability.
Interestingly, geographic information—capturing the region in which the firm operates—while contributing less than 20% on average, still displays a non-negligible influence. Rather than being dismissed as marginal, this share highlights how territorial dynamics, although secondary, can offer an additional explanatory power to the overall prediction.

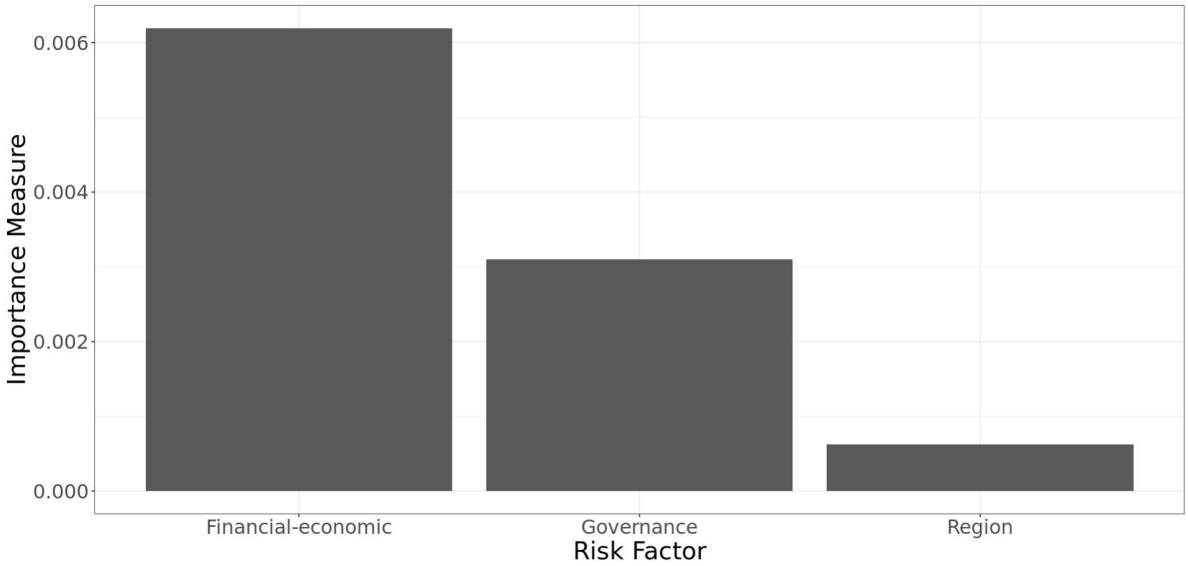

Fig. 4: Normalized average Shapley Values for the each input.

These findings suggest that while economic and governance-related features are the strongest predictors of default, the geographic context contributes relevant contextual insights that enhance the model's overall performance. Its impact, though smaller in magnitude, is nonetheless meaningful and highlights an interesting avenue for further exploration of how regional conditions may affect SMEs' default probability.

6 Conclusions

This paper proposes an explainable multi-input neural network framework for predicting default risk among SMEs, extending the line of research initiated by (Crosato et al. 2023). While prior work has shown the advantages of non-linear machine learning models over traditional linear approaches, our results indicate that further improvements can be achieved by explicitly modeling the heterogeneous and dynamic nature of SME-related information.

The MI-NET architecture integrates financial–economic time series, governance-related firm characteristics, and regional contextual information through dedicated subnetworks. In particular, the explicit treatment of temporal dynamics via an LSTM component represents a key advancement over much of the existing literature, which typically relies on static covariates observed at a single point in time. By capturing firms' financial evolution, the model provides a richer representation of default risk and offers insights into companies' future trajectories.

Empirical evidence from a large sample of Italian SMEs confirms the effectiveness of this approach. MI-Net consistently outperforms logistic regression as well as widely used non-linear benchmarks, including FFNN and XGBoost models, indicating that the structured exploitation of heterogeneous data sources leads to superior predictive

performance. These findings extend earlier results by showing that performance gains are driven not only by model non-linearity, but also by architectural choices aligned with data characteristics.

Beyond predictive accuracy, this study explicitly addresses the crucial issue of interpretability by leveraging Shapley values. The pillar-based architecture enables their efficient computation and supports a transparent decomposition of predictions, allowing for a clear identification of the relative contribution of each model component.

Future research will extend the proposed framework to younger firms and micro-enterprises. Although their exclusion in the present study was methodologically motivated by data-availability constraints and by the simplified governance structures that characterize such firms (see Section 4), these enterprises represent an important segment of the credit market. Modeling their default risk would likely require architectural adaptations to handle shorter or incomplete financial time series and less informative governance indicators. We will also evaluate Transformer-based sequence models (Vaswani et al. 2017) as a complement or alternative to LSTM networks to capture long-range and cross-input dependencies, leveraging attention-only architectures for greater parallelism on time series. We will also extend the analysis to other economic contexts, such as emerging economies, to verify robustness and generalizability. Finally, in line with recent applications in finance (Zhang et al. 2023), we will assess the adoption of transfer learning mechanisms to reduce the need for labeled data, contain overfitting on small cohorts, and shorten training times.

Acknowledgments

The authors wish to express their sincere gratitude to the Italian Ministry of University and Research for financial support within the PRIN – Progetti di Ricerca di Rilevante Interesse Nazionale 2022 programme.

Statements and Declarations

Funding

This work was supported by the Italian Ministry of University and Research under the PRIN 2022 programme.

Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

Author Contributions

Belinda Laura Del Gaudio: Funding acquisition, Conceptualization, Supervision, Writing – review and editing.
Raffaele Clemente Petrella: Data curation, Writing – original draft, Software, Investigation.
Salvatore Scognamiglio: Conceptualization, Methodology, Software, Writing – original

draft.
Paolo Zanetti: Conceptualization, Data curation, Supervision, Validation.

Data Availability

The data that support the findings of this study were obtained from the AIDA database (Bureau van Dijk). AIDA is a proprietary source accessible only through institutional subscription, and the data cannot be shared publicly. Details regarding the dataset are provided in the Data Description section of the manuscript.

Code Availability

The code used to implement the models and perform the analyses is available from the corresponding author upon reasonable request.

Ethics Approval

Not applicable.

Consent to Participate

Not applicable.

Consent for Publication

Not applicable.

Materials Availability

Not applicable.

A List of Variables

The following table reports the full list of features used in the analysis, grouped inputs.

Financial–Economic Variables

Receivables from Shareholders (thousands EUR)
Tangible Fixed Assets (thousands EUR)
Intangible Fixed Assets (thousands EUR)
Financial Fixed Assets (thousands EUR)
Current Assets (thousands EUR)
Total Inventories (thousands EUR)
Financial Current Assets (thousands EUR)
Total Receivables (thousands EUR)
Short-Term Receivables (thousands EUR)
Long-Term Receivables (thousands EUR)
Cash and Cash Equivalents (thousands EUR)
Total Assets (thousands EUR)
Total Shareholders' Equity (thousands EUR)
Share Capital (thousands EUR)
Retained Earnings (thousands EUR)
Profit or Loss for the Year (thousands EUR)
Total Liabilities (thousands EUR)
Short-Term Liabilities (thousands EUR)
Long-Term Liabilities (thousands EUR)
Liabilities Due Within the Year (thousands EUR)
Liabilities Due Beyond the Year (thousands EUR)
Net Income (thousands EUR)
Total Production Value (thousands EUR)
Revenue from Sales and Services (thousands EUR)
Inventory Variation (thousands EUR)
Work in Progress Variation (thousands EUR)
Total Changes in Inventory (thousands EUR)
Increases in Fixed Assets (thousands EUR)
Other Revenues (thousands EUR)
Operating Grants (thousands EUR)
Total Production Costs (thousands EUR)
Personnel Costs (thousands EUR)
Lease and Rental Costs (thousands EUR)
Raw Materials and Consumables (thousands EUR)
Services (thousands EUR)
Depreciation and Amortization (thousands EUR)
Social Security Contributions (thousands EUR)
Wages and Salaries (thousands EUR)
Depreciation of Intangible Fixed Assets (thousands EUR)
Depreciation of Tangible Fixed Assets (thousands EUR)
Write-down of Fixed Assets (thousands EUR)
Credit Write-downs (thousands EUR)
Provisions for Risks (thousands EUR)
Other Provisions (thousands EUR)
Other Operating Expenses (thousands EUR)
Added Value (thousands EUR)
Total Financial Income and Expenses (thousands EUR)
Other Financial Income (thousands EUR)
Financial Income from Credits (thousands EUR)
Income from Securities (thousands EUR)
Total Financial Charges (thousands EUR)
Exchange Gains and Losses (thousands EUR)
Total Adjustments to Financial Assets (thousands EUR)
Total Revaluations (thousands EUR)
Total Write-downs (thousands EUR)
Total Extraordinary Income and Expenses (thousands EUR)
Result Before Taxes (thousands EUR)
Total Income Taxes (thousands EUR)
Current Income Taxes (thousands EUR)
Deferred and Prepaid Taxes (thousands EUR)
Net Working Capital (thousands EUR)
Operating Cash Flow (thousands EUR)
Consumption Margin (thousands EUR)
Labor Cost per Employee (EUR)
Employee Productivity
Liquidity Ratio
Short-Term Debt Ratio (%)
Long-Term Debt Ratio (%)
Financial Charges on Revenue (%)
Turnover of Invested Capital
Turnover of Gross Working Capital
Return on Sales (ROS) (%)
Return on Assets (ROA) (%)
EBITDA (thousands EUR)
EBITDA on Sales (%)
Revenue per Capita (EUR)
Value Added per Capita (EUR)
Extraordinary Charges (thousands EUR)

Governance Variables

Number of Previous Managers
Percentage of Female Managers (%)
Average Age of Decision Makers (years)
Decision Maker is Also Shareholder
Number of Previous Advisors
Number of Shareholders Present

Regional Information

Region

References

Altman, E.I., Balzano, M., Giannozzi, A., Srhoj, S.: Revisiting sme default predictors: the omega score. Journal of Small Business Management **61**(6), 2383–2417 (2023)

Algeri, C., Forgione, A.F., Mucciardi, M.: Spatial dependence in small cooperative bank risk behavior and its effects on bank competitiveness and smes. Applied stochastic models in business and industry **37**(6), 1080–1096 (2021)

Altman, E.I., Haldeman, R.G., Narayanan, P.: Zetatm analysis a new model to identify bankruptcy risk of corporations. Journal of banking & finance **1**(1), 29–54 (1977)

Altman, E.I., Iwanicz-Drozdowska, M., Laitinen, E.K., Suvas, A.: A race for long horizon bankruptcy prediction. Applied Economics **52**(37), 4092–4111 (2020)

Altman, E.I., Sabato, G.: Modelling credit risk for smes: Evidence from the us market. Abacus **43**(3), 332–357 (2007)

Altman, E.I., Sabato, G., Wilson, N.: The value of non-financial information in sme risk management. Available at SSRN 1320612 (2008)

Bengio, Y., Courville, A., Vincent, P.: Representation learning: A review and new perspectives. IEEE transactions on pattern analysis and machine intelligence **35**(8), 1798–1828 (2013)

Bellovary, J.L., Giacomino, D.E., Akers, M.D.: A review of bankruptcy prediction studies: 1930 to present. Journal of Financial education, 1–42 (2007)

Ballester, L., Gonzalez-Urteaga, A., Martinez, B.: The role of internal corporate governance mechanisms on default risk: A systematic review for different institutional settings. Research in International Business and Finance **54**, 101293 (2020)

Barboza, F., Kimura, H., Altman, E.: Machine learning models and bankruptcy prediction. Expert systems with applications **83**, 405–417 (2017)

Belás, J., Vojtovič, S., Ključnikov, A.: Microenterprises and significant risk factors in loan process. Economics and Sociology (2016)

Castellani, G., Fiore, U., Marino, Z., Passalacqua, L., Perla, F., Scognamiglio, S., Zanetti, P.: Machine learning techniques in nested stochastic simulations for life insurance. Applied Stochastic Models in Business and Industry **37**(2), 159–181 (2021)

Ciampi, F., Giannozzi, A., Marzi, G., Altman, E.I.: Rethinking sme default prediction: a systematic literature review and future perspectives. Scientometrics **126**, 2141–2188 (2021)

Ciampi, F.: Corporate governance characteristics and default prediction modeling for

small enterprises. an empirical analysis of italian firms. Journal of Business Research **68**(5), 1012–1025 (2015)

Crosato, L., Liberati, C., Repetto, M.: Lost in a black-box? interpretable machine learning for assessing italian smes default. Applied Stochastic Models in Business and Industry **39**(6), 829–846 (2023)

Cheraghali, H., Molnár, P.: Sme default prediction: A systematic methodology-focused review. Journal of Small Business Management **62**(6), 2847–2905 (2024)

Calabrese, R., Osmetti, S.A.: Modelling small and medium enterprise loan defaults as rare events: the generalized extreme value regression model. Journal of Applied Statistics **40**(6), 1172–1188 (2013)

Edmister, R.O.: An empirical test of financial ratio analysis for small business failure prediction. Journal of Financial and Quantitative analysis **7**(2), 1477–1493 (1972)

Fantazzini, D., Figini, S.: Random survival forests models for sme credit risk measurement. Methodology and computing in applied probability **11**(1), 29–45 (2009)

Figini, S., Giudici, P.: Statistical merging of rating models. Journal of the operational research society **62**(6), 1067–1074 (2011)

Guo, C., Berkhahn, F.: Entity embeddings of categorical variables. arXiv (2016) https://doi.org/10.48550/ARXIV.1604.06737

Giudici, P., Hadji-Misheva, B., Spelta, A.: Network based credit risk models. Quality Engineering **32**(2), 199–211 (2020)

Giorgi, F., Herzel, S., Pigato, P.: A reinforcement learning algorithm for trading commodities. Applied Stochastic Models in Business and Industry **40**(2), 373–388 (2024)

Goodell, J.W., Kumar, S., Lim, W.M., Pattnaik, D.: Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. Journal of Behavioral and Experimental Finance **32**, 100577 (2021)

Heaton, J.B., Polson, N.G., Witte, J.H.: Deep learning for finance: deep portfolios. Applied Stochastic Models in Business and Industry **33**(1), 3–12 (2017)

Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)

Janssens, A.C.J., Martens, F.K.: Reflection on modern methods: Revisiting the area under the roc curve. International journal of epidemiology **49**(4), 1397–1403 (2020)

Karas, M., Režňáková, M.: The role of financial constraint factors in predicting sme default. Equilibrium. Quarterly Journal of Economics and Economic Policy **16**(4), 859–883 (2021)

Kou, G., Xu, Y., Peng, Y., Shen, F., Chen, Y., Chang, K., Kou, S.: Bankruptcy prediction for smes using transactional data and two-stage multiobjective feature selection. Decision Support Systems **140**, 113429 (2021)

Lextrait, B.: Scaling up smes' credit scoring scope with lightgbm. Applied Economics **55**(9), 925–943 (2023)

McKevitt, D., Davis, P.: Microenterprises: how they interact with public procurement processes. International Journal of Public Sector Management **26**(6), 469–480 (2013)

Nartisa, I.: Openness and knowledge as leading tendencies in development of micro enterprises. Economics and management **17**(4), 1579–1584 (2012)

Navarro-Galera, A., Lara-Rubio, J., Novoa-Hernández, P., Cruz Corona, C.A.: Using decision trees to predict insolvency in spanish smes: Is early warning possible? Computational Economics **65**(1), 91–116 (2025)

Nwogugu, M.: Decision-making, risk and corporate governance: A critique of methodological issues in bankruptcy/recovery prediction models. Applied Mathematics and Computation **185**(1), 178–196 (2007)

Oliveira, M.D., Ferreira, F.A., Pérez-Bustamante Ilander, G.O., Jalali, M.S.: Integrating cognitive mapping and mcda for bankruptcy prediction in small-and medium-sized enterprises. Journal of the Operational Research Society **68**(9), 985–997 (2017)

Sariev, E., Germano, G.: Bayesian regularized artificial neural networks for the estimation of the probability of default. Quantitative finance **20**(2), 311–328 (2020)

Sanchez-Medina, A.J., Blázquez-Santana, F., Cerviño-Cortínez, D.L., Pellejero, M.: Ensemble methods for bankruptcy resolution prediction: A new approach. Computational Economics, 1–36 (2025)

Surya, B., Menne, F., Sabhan, H., Suriani, S., Abubakar, H., Idris, M.: Economic growth, increasing productivity of smes, and open innovation. Journal of Open Innovation: Technology, Market, and Complexity **7**(1), 20 (2021)

Switzer, L.N., Tu, Q., Wang, J.: Corporate governance and default risk in financial firms over the post-financial crisis period: International evidence. Journal of International Financial Markets, Institutions and Money **52**, 196–210 (2018)

Srhoj, S., Vitezić, V., Giannozzi, A., Mikulić, J.: Tourism sme default: A note on

predictors. Tourism management **103**, 104910 (2024)

Vicario, G., Coleman, S.: A review of data science in business and industry and a future view. Applied Stochastic Models in Business and Industry **36**(1), 6–18 (2020)

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)

Winter, E.: The shapley value. Handbook of game theory with economic applications **3**, 2025–2054 (2002)

Wang, L., Jia, F., Chen, L., Xu, Q.: Forecasting smes' credit risk in supply chain finance with a sampling strategy based on machine learning techniques. Annals of Operations Research **331**(1), 1–33 (2023)

Youden, W.J.: Index for rating diagnostic tests. Cancer **3**(1), 32–35 (1950)

Zhang, D., Lin, R., Wei, T., Ling, L., Huang, J.: A novel deep transfer learning framework with adversarial domain adaptation: application to financial time-series forecasting. Neural Computing and Applications **35**(34), 24037–24054 (2023)

Zheng, C., Weng, F., Luo, Y., Yang, C.: Micro and small enterprises default risk portrait: evidence from explainable machine learning method. Journal of Ambient Intelligence and Humanized Computing **15**(1), 661–671 (2024)