An-Najah University Journal for Research – B
Humanities

مجلـــة جامعـــة النجـــاح للأبحــــــاث – ب
العلوم الإنسانية

Predicting Individual Unemployment Risk Using Explainable AI: A Two-Stage Microdata Framework from Palestine

Haitham Zeidan[1*], Adel Karaa[2], Rabeh Morrar[3]

Type: Full Article. Received: 18April 2026, Accepted: 23 June 2026

Abstract: Objective: This study develops an explainable artificial intelligence framework for predicting individual unemployment risk in Palestine using nationally representative Labour Force Survey microdata. It aims to provide a conceptually consistent and interpretable approach that can support labour-market monitoring and evidence-based policymaking. **Methodology:** A two-stage modelling framework was applied to Labour Force Survey microdata covering the period 2015–2024. The first stage predicted labour-force participation, while the second predicted unemployment conditional on labour-force participation. Logistic regression was used as the baseline model, and CatBoost was used as the primary machine-learning classifier. Survey sampling weights were incorporated into model estimation and evaluation. Model performance was assessed using weighted AUC, Brier Score, calibration analysis, and weighted accuracy. SHapley Additive exPlanations were used to interpret the contribution of individual characteristics to predicted unemployment risk**. Key Findings:** CatBoost outperformed logistic regression in predictive discrimination, improving the AUC by approximately 8–12 percentage points. The explainability analysis identified age, education, gender, region, locality type, and refugee status as important predictors of unemployment risk. Higher predicted risks were concentrated among young people, university graduates, women, residents of the Gaza Strip, and individuals living in refugee camps**. Conclusions:** The two-stage framework improves conceptual consistency by separating labour-force participation from unemployment prediction. It also combines predictive performance with transparent and policy-relevant interpretation**. Recommendations:** Official statistical institutions should consider integrating explainable machine-learning models into labour-market monitoring systems while maintaining appropriate data-governance, validation, and human-oversight procedures**.**

Keywords: Explainable Artificial Intelligence, Unemployment Risk, Labour Force Survey, CatBoost, SHAP, Official Statistics

التنبؤ بمخاطر البطالة الفردية باستخدام الذكاء الاصطناعي القابل للتفسير: إطار من مرحلتين قائم على البيانات الجزئية من فلسطين

هيثم زيدان*[1]، عادل قرعة[2]، رابح مرار[3]

تاريخ التسليم: 2026/4/18، تاريخ القبول: 2026/6/23

الملخص: الهدف: تهدف هذه الدراسة إلى تطوير إطار قائم على الذكاء الاصطناعي القابل للتفسير للتنبؤ بمخاطر البطالة على المستوى الفردي في فلسطين، باستخدام البيانات الجزئية الممثلة وطنيًا لمسح القوى العاملة. كما تسعى إلى تقديم منهج متسق من الناحية المفاهيمية وقابل للتفسير لدعم متابعة سوق العمل وصنع السياسات المستندة إلى الأدلة. **المنهج**: طُبّق إطار نمذجة من مرحلتين على البيانات الجزئية لمسح القوى العاملة خلال الفترة 2015–2024. تنبأت المرحلة الأولى بمشاركة الفرد في القوى العاملة، بينما تنبأت المرحلة الثانية باحتمال البطالة بشرط المشاركة في القوى العاملة. استُخدم الانحدار اللوجستي نموذجًا مرجعيًا، وخوارزمية CatBoost نموذجًا رئيسيًا للتعلم الآلي. وأُدرجت أوزان المعاينة في تقدير النماذج وتقييمها، باستخدام المساحة الموزونة تحت منحنى ROC، ودرجة Brier، وتحليل المعايرة، والدقة الموزونة. كما استُخدمت قيم SHAP لتفسير مساهمة خصائص الأفراد في مخاطر البطالة المتوقعة. **أهم النتائج**: تفوقت خوارزمية CatBoost على الانحدار اللوجستي في القدرة التمييزية، وحققت تحسنًا في قيمة AUC تراوح بين 8 و12 نقطة مئوية تقريبًا. وأظهر التحليل التفسيري أن العمر، والتعليم، والجنس، والمنطقة، ونوع التجمع السكاني، وصفة اللجوء من أهم العوامل المرتبطة بمخاطر البطالة. وتركزت المخاطر الأعلى بين الشباب، وخريجي الجامعات، والنساء، وسكان قطاع غزة، والمقيمين في المخيمات. **الاستنتاجات**: يعزز الإطار المقترح الاتساق المفاهيمي من خلال الفصل بين المشاركة في القوى العاملة والتنبؤ بالبطالة، كما يجمع بين الأداء التنبؤي والتفسير الشفاف ذي الصلة بالسياسات العامة. **التوصيات**: توصي الدراسة بدراسة دمج نماذج التعلم الآلي القابلة للتفسير في أنظمة متابعة سوق العمل لدى المؤسسات الإحصائية الرسمية، مع ضمان الحوكمة المناسبة للبيانات والتحقق المستمر والإشراف البشري.

الكلمات المفتاحية: الذكاء الاصطناعي القابل للتفسير، مخاطر البطالة، مسح القوى العاملة، CatBoost، SHAP، الإحصاءات الرسمية.

1 General Directorate of Information Systems, Director General of Information Systems, Palestinian Central Bureau of Statistics, Ramallah, Palestine. ORCID: 0009-0009-1803-398X
* Corresponding author: Haitham Zeidan: haitham@pcbs.gov.ps
2 Department of Computer Science, Director of the Doctoral School, Higher Institute of Management of Tunis, University of Tunis, Tunis, Tunisia.
3 Department of Economics, Faculty of Business and Communication, An-Najah National University, Nablus, Palestine

1 الإدارة العامة لأنظمة المعلومات، مدير عام أنظمة المعلومات، الجهاز المركزي للإحصاء الفلسطيني، رام الله، فلسطين.. ORCID: 0009-0009-1803-398X
* الباحث المراسل: هيثم زيدان
2 قسم علم الحاسوب، مدير مدرسة الدكتوراه، المعهد العالي للتصرف بتونس، جامعة تونس، تونس
3 قسم الاقتصاد، كلية الأعمال والاتصال، جامعة النجاح الوطنية، نابلس، فلسطين

1. Introduction

It has been widely acknowledged that unemployment is one of the key indicators of economic and social well-being. Globally, the unemployment rate was estimated at approximately 5.1% in 2023 (International Labour Organization [ILO], 2024b), and high levels of unemployment are associated with slower economic growth, increased poverty, and social instability (World Bank, 2023a). In Palestine, unemployment remains significantly higher than the global average, reaching about 24% in 2023, with youth unemployment exceeding 40% (Palestinian Central Bureau of Statistics [PCBS], 2023a). However, aggregated unemployment rates do not adequately capture the heterogeneity of unemployment risks across individuals and demographic subgroups.

In addition, the distribution of unemployment across regions and population groups is highly uneven, particularly in fragile and conflict-affected contexts. For instance, unemployment in the Gaza Strip has consistently exceeded 40%, compared to relatively lower but still substantial rates in the West Bank (PCBS, 2023a). Furthermore, unemployment is disproportionately high among women and university graduates, with female unemployment rates exceeding 40% in certain periods due to structural barriers limiting labour market participation and access to employment opportunities.

Accurate forecasting of unemployment is therefore essential for policymakers, as it enables proactive labour market interventions, supports evidence-based decision-making, and improves the efficiency of resource allocation. Reliable forecasts help governments anticipate economic downturns, design targeted employment policies, and respond effectively to labour market shocks. Recent studies highlight the increasing importance of real-time and high-frequency data in improving unemployment forecasting performance, particularly in rapidly changing economic environments (Costa et al., 2024).

Although traditional econometric models such as ARIMA and regression models are widely used due to their simplicity and interpretability (Box et al., 2015), they often fail to capture nonlinear relationships and complex temporal dynamics inherent in labour markets. In recent years, machine learning and artificial intelligence approaches have demonstrated superior performance in forecasting tasks by modelling complex patterns and interactions within large-scale datasets (Güler et al., 2024; Kim, 2025). These methods are particularly effective in handling structural breaks, seasonality, and volatility, which are common characteristics of labour market data.

Furthermore, AI-driven forecasting systems are increasingly used to support labour market analysis, workforce planning, and early detection of employment trends, including skill mismatches and sectoral shifts (Nigar, 2025). Despite these advantages, challenges related to interpretability, data quality, and integration into official statistical systems remain significant barriers to their widespread adoption.

However, despite the growing body of research on unemployment forecasting, most existing studies focus either on aggregate (macro-level) predictions or on black-box machine learning models with limited interpretability. There is a notable lack of research that integrates microdata-based modelling with explainable artificial intelligence techniques, particularly in developing and conflict-affected contexts such as Palestine.

Against this background, this study develops a micro-level explainable machine learning framework for predicting individual unemployment risk using Labour Force Survey (LFS) microdata from the Palestinian Central Bureau of Statistics (PCBS). The proposed framework combines a two-stage modelling structure with CatBoost classification and SHAP-based explainability to produce both predictive outputs and interpretable insights for labour-market monitoring.

This study contributes to the literature in three main ways. First, it introduces a two-stage modelling framework that aligns with the economic definition of unemployment by separately modelling labour force participation and unemployment status. Second, it applies explainable machine learning techniques to provide transparent and interpretable insights into the determinants of unemployment risk. Third, it demonstrates how artificial intelligence methods can be systematically integrated into official statistical processes, thereby supporting the modernization of labour market statistics and enhancing evidence-based policymaking.

2. Literature Review

In addition to methodological advancements, the selection of relevant predictors plays a crucial role in improving unemployment forecasting accuracy. Recent studies emphasize that combining macroeconomic, demographic, and socio-economic variables can enhance predictive performance, particularly when using machine learning models.

Macroeconomic indicators such as GDP growth, inflation, and economic shocks remain among the most influential predictors of unemployment. Recent empirical evidence shows that economic growth is strongly associated with labour-demand fluctuations, while inflation reflects underlying macroeconomic conditions related to employment dynamics (International Monetary Fund, 2023; World Bank, 2023a). Furthermore, recent global shocks, including the COVID-19 pandemic and geopolitical conflicts, have been associated with substantial changes in unemployment patterns, highlighting the importance of incorporating shock-related variables in forecasting models (ILO, 2023b).

Demographic and socio-economic variables also play a central role in explaining unemployment patterns at the micro level. Variables such as age, gender, education level, and region have been consistently identified as key predictors of labour-market outcomes in recent studies (Organisation for Economic Co-operation and Development [OECD], 2024a; ILO, 2022). For instance, youth unemployment remains significantly higher than adult unemployment globally, while gender disparities persist due to structural and institutional constraints in labour markets.

At the microdata level, labour-force participation should be modelled as a separate but related process in unemployment prediction because unemployment is officially defined only among individuals participating in the labour force. Survey microdata allow this distinction to be represented explicitly by separating labour-force participation from conditional unemployment status, thereby improving conceptual consistency and reducing potential classification bias. This supports the use of multi-stage modelling approaches that align with the official statistical definition of unemployment (ILO, 2023a; Eurostat, 2024).

Several recent peer-reviewed studies have applied machine learning techniques to unemployment forecasting and labour-market prediction. For example, Güler et al. (2024) compared artificial neural networks, support vector machines, and XGBoost for monthly unemployment-rate forecasting in Turkey and found that machine learning models can capture nonlinear macroeconomic relationships involving unemployment, inflation, exchange rates, and labour-force indicators. Gupta et al. (2022) used random forests to forecast changes in the United Kingdom unemployment rate and highlighted the value of nonlinear machine learning methods for modelling uncertainty-related predictors. At the micro level, Celbiş (2023) applied tree-based classification models, including random forests and gradient boosting, to predict unemployment status in rural Europe using individual-level characteristics. Similarly, Gabrikova et al. (2023) developed an ensemble model for predicting unemployment duration among jobseekers in Slovakia and highlighted the importance of age, education, and unemployment history.

These studies demonstrate the growing relevance of machine learning in labour-market prediction. However, many existing studies focus mainly on aggregate unemployment rates, unemployment duration, or specific regional contexts. In addition, although machine learning models can improve predictive performance by capturing nonlinear interactions between variables, their increasing complexity has led to growing concern about transparency and interpretability in policy-related applications (Makridakis et al., 2018; OECD, 2024b; Samek et al., 2019). The present study extends this literature by applying an explainable two-stage machine learning framework to official Labour Force Survey microdata, with a focus on individual unemployment-risk prediction in a fragile economic context.

3. Conceptual Framework

In labour economics, unemployment is generally defined conditionally upon whether or not an individual is a labour-force participant. The ILO defines an individual as unemployed only if he/she is currently a labour-force participant, i.e., available to work and actively seeking employment (ILO, 2023a). Individuals who do not participate in the labour force therefore cannot be considered unemployed. Thus, ignoring this condition could lead to classification bias when predicting unemployment directly from the full population dataset. Figure 1 depicts the logical dependency between labour-force participation and unemployment status.

Figure 1: Dependency between labour-force participation and unemployment status.

.

To account for this conditional relationship, this study uses a two-stage modelling structure to reflect the theoretical definition of unemployment in labour economics. The unemployment probability of an individual i is decomposed as follows:

$$P(U_i = 1) = P(LF_i = 1) \times P(U_i = 1 \mid LF_i = 1)$$

Where:

- P(LFi=1) represents the probability that individual i participates in the labour force.
- P(Ui=1|LFi=1) represents the probability that individual i is unemployed conditional on labour-force participation.

This formulation ensures that unemployment predictions are consistent with the conceptual definition of unemployment used in official labour statistics and helps reduce potential classification bias (Eurostat, 2024; ILO, 2023a).

Compared with a conventional one-stage unemployment prediction model, the proposed two-stage framework provides a more conceptually consistent representation of unemployment. A one-stage model predicts unemployment status directly for the full working-age population, which may mix two different labour-market processes: whether an individual participates in the labour force and whether a labour-force participant is unemployed. This may create classification ambiguity, particularly for individuals outside the labour force, who are neither employed nor unemployed according to official labour statistics. By contrast, the two-stage framework first models labour-force participation and then estimates unemployment only among labour-force participants. This structure therefore follows the official statistical definition of unemployment more closely and helps reduce conceptual classification bias. Accordingly, the value of the two-stage framework lies not only in prediction, but also in aligning the modelling process with labour-market theory and official statistical practice.

In addition, the two-stage structure provides greater clarity in predictive modelling by isolating the determinants of labour-force participation from the determinants of unemployment within the labour force. Such a separation is particularly relevant in labour markets with structural constraints and heterogeneous participation patterns, such as the Palestinian labour market (ILO, 2023a; PCBS, 2023a).

Therefore, the modelling framework used in this study consists of two sequential predictive stages, and the overall unemployment probability is obtained by combining the output of the two models using the equation above

1. **Stage 1 – Labour Force Participation Model**

This stage predicts the probability that an individual participates in the labour force based on demographic, educational, and geographic characteristics.

2. **Stage 2 – Conditional Unemployment Model**

This stage predicts the probability that an individual is unemployed conditional on being labour-force participant.

4. Data

The empirical approach to data collection used for this study was to obtain microdata from the Palestinian Labour Force Survey (LFS) conducted by the PCBS. The LFS is the central source of labour-market data for the Palestinian territory and is collected quarterly using a nationally representative sample design (PCBS, 2023b).

The time frame of this study is the period 2015–2024. Therefore, individuals aged 15 years or older were included in the study, in accordance with the ILO (2023a) international definition of the "working age." The LFS provides rich socio-economic information about individuals and their households and therefore supports examination of labour-force participation and unemployment dynamics across demographic and geographic subgroups.

As mentioned earlier, the LFS contains questions about individuals' employment status, labour-force participation, highest level of education completed, demographic characteristics, and geographical location. These variables provide the basis for developing individual-level unemployment-risk prediction models and for identifying structural factors associated with labour-market outcomes.

Therefore, the explanatory variables that were used for the development of the predictive models of this study were grouped into the following four categories:

Table 1: Variables Used in the Model

Category	Variables
Demographic	Age, Gender, Marital Status
Human Capital	Education Level
Geographic	Region (West Bank / Gaza), Locality Type (Urban, Rural, Camp)
Structural	Refugee Status

The demographic variables represent the life-cycle characteristics of individuals that are associated with labour-market behaviour, whereas the education variables serve as a proxy for an individual's human capital and employability. The geography variables represent regional labour market disparities that exist in the Palestinian territory due to the structural differences that exist between the West Bank and Gaza Strip. The refugee status variable represents a structural factor that reflects the historical and institutional factors that have limited access to labour markets in the Palestinian Territory.

To improve transparency and support interpretation of the model results, a descriptive summary of the variables used in the micro-level prediction model is presented in Table 2. The table reports the coding scheme, variable type, and descriptive distribution of the main predictors. Descriptive statistics were calculated using the Labour Force Survey person-level sampling weights to ensure that the reported percentages and averages reflect the population structure. Continuous variables are summarized using weighted means and standard deviations, while categorical variables are summarized using weighted category frequencies and percentages.

Table 2. Descriptive Summary and Coding of Variables Used in the Micro-Level Model

Variable	Description	Coding / Harmonization Rules	Use in Study
InOutlf	Labour-force participation status	1 = In labour force; 2 = Out of labour force. Records with invalid codes removed.	Used to determine labour-force population; denominator for unemployment rate.
empch	Employment status	1 = Employed, 2 = Unemployed. Non-labour-force cases filtered using *InOutlf*.	Defines target unemployment classification.
PW16	Job search activity	1 = Searched for a job during reference period; 2 = Did not search. Skip patterns validated per ILO rules.	Used to compute share searching for work.
PW21	Availability to work	1 = Available; 2 = Not available. Invalid combinations (e.g., employed persons marked “available”) corrected.	Indicator of potential labour-force attachment and underutilisation.
timerelated	Time-related underemployment	1 = Underemployed; 2 = Adequately employed or not applicable.	Used to compute share underemployed.
PW35_A	Work permit status (employed)	1 = Holds a valid work permit; 2 = Does not hold. Missing values harmonised across years.	Used to compute share of employed with work permits.
RW	Person sampling weight	Positive integer weight representing population expansion factor. Zero or negative values removed.	Used in all weighted proportions and means.
HR2	Sex	1 = Male; 2 = Female.	Used to compute female unemployment rate and female labour-force share.
HR5	Refugee status	1 = Refugee; 2 = Non-refugee.	Used to compute share of refugees.
Pr1	Age	Numeric age, harmonised for missing top-coded values.	Used for youth subgroup (15–24).
Pr3	Years of schooling / educational attainment	Continuous or categorical depending on quarter; transformed into years of schooling.	Used to compute quarterly mean education.
PW06	Weekly hours worked	Raw hours; extreme outliers winsorised at 1st and 99th percentiles.	Used for mean weekly hours and underemployment analysis.
PW09_a	Shock-related impact item (COVID-19, conflict)	Re-coded to create binary shock indicators (COVID-19 2020Q2; Conflict 2023Q4).	Basis for shock variables; used in modelling after lagging.
empch_lag1, lags of other variables	Lagged variants of labour-market indicators	Constructed automatically after aggregation; no micro-level coding needed.	Used in autoregressive modelling and PCA.

In order to ensure that the results of the empirical analyses are representative of the entire population of Palestine, sampling weights that were provided by the PCBS were used in all of the estimates. These sampling weights correct for the complex survey design used by the PCBS and allow for the results of the models to be interpreted as reflecting the population-level dynamics of the labour market.

The use of the official LFS microdata provides a reliable and relevant foundation for analyzing the risk of unemployment and allows for the incorporation of the proposed modelling framework into official statistical systems.

5. Methodology

The purpose of this study was to develop a predictive framework to estimate individual unemployment risk through the application of machine learning and explainable artificial intelligence techniques (Makridakis et al., 2018; OECD, 2024b). The methodology of this study is divided into four major areas: data preprocessing and model training, predictive models, explainability, and evaluation metrics.

5.1 Data Preprocessing and Model Training

Prior to model estimation, the Labour Force Survey microdata were reviewed and prepared to ensure consistency across survey rounds. Individuals below the working-age threshold were excluded, and the analysis was restricted to respondents aged 15 years and above. Records with missing information on the target variables, namely labour-force participation and employment status, were excluded from the modelling dataset. For explanatory variables, categorical predictors such as gender, marital status, education level, region, locality type, and refugee status were checked for consistency across years, and missing or unclear categories were treated as separate categories where applicable.

The modelling framework was implemented using a supervised learning approach. In the first stage, the target variable was labour-force participation. In the second stage, the target variable was unemployment status conditional on being in the labour force. Sampling weights provided with the Labour Force Survey were retained in the estimation and evaluation process in order to reflect the survey design and population structure.

To avoid look-ahead bias, the dataset was split using a temporal validation strategy rather than a purely random split. Survey data from 2015 to 2021 were used for model training and validation, while data from 2022 to 2024 were reserved for testing model performance. This approach ensures that the model is evaluated on observations that occur after the training period, which better reflects a realistic labour-market monitoring setting. The same temporal splitting logic was applied consistently across the labour-force participation model and the conditional unemployment model.

For logistic regression, categorical variables were transformed into dummy variables. For CatBoost, categorical variables were handled using the algorithm’s native categorical-feature processing capability. Model performance was assessed using weighted evaluation metrics, including AUC, Brier Score, calibration analysis, and weighted accuracy. Hyperparameter tuning was conducted using the validation data, with attention to controlling overfitting and maintaining stable performance across survey periods. All analyses were conducted using Python, including standard machine-learning libraries for model estimation, evaluation, and SHAP-based explainability.

5.2 Predictive Models

There were two predictive models developed to predict unemployment risk: a traditional statistical model was developed as a baseline model, and a machine learning model was developed as the primary model to predict unemployment risk.

Logistic Regression (Baseline Model)

Logistic regression was selected as the baseline model due to its widespread application in labour market research and the interpretability of its output (Molnar, 2022; James et al., 2021). Logistic regression estimates the probability of a person being unemployed given a set of explanatory variables X_i.

The logistic regression equation is:

$$P(Y_i = 1 | X_i) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 X_i)}}$$

Where

- Y_i is the unemployment status of person i,
- X_i is the vector of explanatory variables,
- β_0 is the intercept,
- β is the vector of coefficients for the predictor variables.

Although logistic regression has many benefits including transparency and interpretability, it does assume a linear relationship between the predictor variables and the dependent variable and thus may miss complex non-linear relationships between predictor variables in labour market data.

CatBoost (Machine Learning Model)

To overcome some of the shortcomings of linear models, CatBoost is a gradient boosting algorithm that effectively handles categorical variables and has demonstrated strong performance in recent machine learning applications (Dorogush et al., 2018; Hancock & Khoshgoftaar, 2020).

CatBoost has several advantages over other models when dealing with labour market microdata:

1. CatBoost can process categorical variables without having to perform extensive preprocessing prior to running the model.
2. CatBoost can perform robustly even when there is a large imbalance in the number of classes or levels in a categorical variable.
3. CatBoost can identify non-linear relationships between predictor variables and identify interactive effects between predictor variables.

Gradient boosting algorithms build an ensemble of decision trees that are combined to minimize prediction error. The general objective function that CatBoost optimizes is:

$$L = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \Omega(f)$$

Where

- l(y_i,y_i^\ ) represents the loss function,
- f represents the ensemble of decision trees,
- Ω(f) is a regularization term that controls the complexity of the model.

Therefore, CatBoost is likely to make better predictions regarding unemployment risk than the traditional statistical models.

5.3 Explainability

One of the major drawbacks of machine learning models is that they typically do not provide clear explanations of the predictions made by them. Therefore, in addition to developing a predictive model, the study also develops a model to explain the predictions made by the machine learning model. In particular, the study uses SHapley Additive exPlanations (SHAP) to interpret model outputs. SHAP values are based on cooperative game theory and measure the amount of contribution each feature makes to the final prediction for a particular instance. That is, for a given model f(x), the SHAP value for feature j measures the marginal contribution of that feature to the predicted outcome (Lundberg & Lee, 2017; Molnar, 2022).

The SHAP decomposition can be written as:

$$f(x) = \varphi_0 + \sum_{j=1}^{M} \varphi_j$$

Where

- ϕ_0 is the baseline prediction,
- ϕ_j represents the contribution of feature j.

Using SHAP analysis allows researchers to obtain both global interpretations, which identify the most important features across the entire dataset, and local interpretations, which explain why a specific instance was classified in a particular way.

The development and application of explainable AI techniques increases the level of transparency of the use of machine learning models and facilitates the acceptance of these models in official statistical agencies.

5.4 Evaluation Metrics

In order to compare the quality of the two models developed, several evaluation metrics were used. These metrics measured how well each model discriminated between unemployed and employed individuals, how accurately it estimated the probability of unemployment, and whether those probability estimates were well calibrated (Fawcett, 2006; Guo et al., 2017).

Weighted AUC

The area under the Receiver Operating Characteristic Curve (AUC) measures the degree to which a model can discriminate between unemployed and employed individuals. The weighted AUC accounts for the sampling weights in the survey data (Fawcett, 2006).

A higher AUC indicates better discrimination.

Brier Score

The Brier Score measures the accuracy of predicted probabilities. It is calculated as the average squared difference between predicted probabilities and the actual employment status of the respondent (Gneiting & Raftery, 2007; Ferro, 2007).

$$Brier = \frac{1}{N}\sum_{i=1}^{N}(p_i - y_i)^2$$

A lower Brier score indicates better calibration.

Calibration Curve

A Calibration Curve is used to examine how close predicted probabilities are to actual employment status. A good calibrated model is one whose predictions are consistent with empirical probabilities (Guo et al., 2017; Vaicenavicius et al., 2019).

The importance of calibration analysis arises from the fact that predicted probabilities are frequently used for decision-making in policy contexts.

Weighted Accuracy

Weighted accuracy measures the proportion of correctly classified instances while taking into consideration the sampling weights of the survey. By doing so, the evaluation of the model takes into consideration the population distribution in the survey data.

The use of multiple evaluation metrics provides a complete picture of how well each model performs, and whether they are able to classify respondents accurately and reliably.

6. Results

Results of the two-stage modelling framework are organized into five sub-sections. Firstly, results from the labour force participation model are reported. Secondly, the performance of the conditional unemployment model is evaluated, followed by SHAP values which highlight the importance of key features such as age, gender, education level, region, marital status, and refugee status. Fourthly, the individual-level reasons for predicted unemployment will be illustrated. Finally, the distribution of unemployment risk will be described across different population groups.

6.1 Stage 1 — Labour Force Participation

The first stage of the Two-Stage Modelling Framework, which provides the labour force participation probability for an individual, was evaluated as part of the overall model performance evaluation. The model performance was evaluated using both ROC Curve Analysis and Calibration Analysis.

Figures 2 and 3 represent the Receiver Operating Characteristic (ROC) Curve and the Calibration Curve respectively for the Labour Force Participation Model. Figure 2 indicates the capability of the labour-force participation model to differentiate between labour force participants and non-participants. Figure 3 demonstrates that the model has good calibration, where predicted probabilities of participating in the labour force are close to observed participation rates.

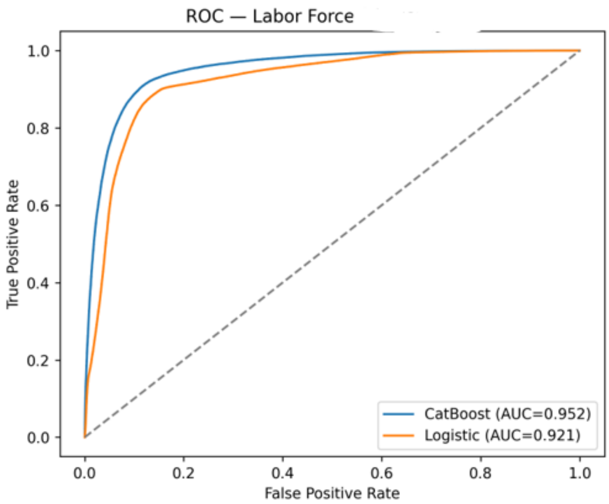

Figure 2: ROC of Labour Force Participation Models

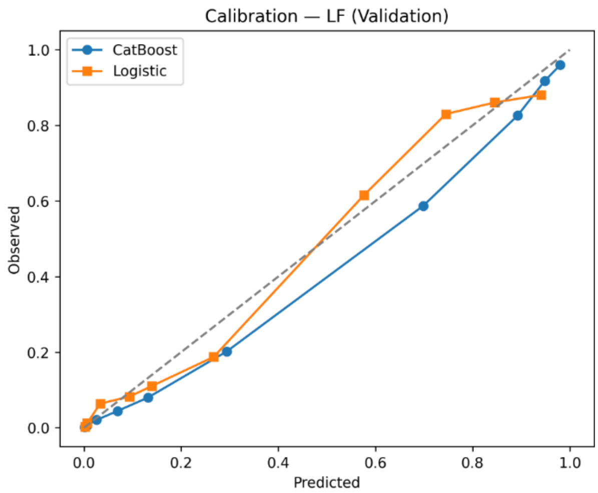

Figure 3: Calibration of Labour Force Participation Models

Results of the labour-force participation model show that the decision to participate in the labour force in Palestine is highly dependent upon gender and education barriers, similar to patterns observed in other labour markets (ILO, 2024a). Women have lower participation rates than men, and more highly educated individuals have higher participation rates.

6.2 Stage 2 — Conditional Unemployment

The second stage of the two-stage modelling framework (See table 3 below) estimates the probability of unemployment for an individual conditional on labour-force participation. The CatBoost model performs significantly better than the logistic regression model. Specifically, CatBoost increases the discriminative power of the model by around 8–12 percentage points relative to logistic regression. The reason for the improved model performance lies in the ability of machine learning models like CatBoost to identify and model non-linear relations between labour market characteristics and unemployment risk.

Table 3: Model Performance in terms of their predictive capabilities

Model	**AUC**	**Calibration**
Logistic Regression	≈0.71	Moderate
CatBoost	≈0.83	Good

Figures 4 and 5 represent the ROC Curve and the Calibration Curve respectively for the Unemployment Prediction Model. As demonstrated in Figures 4 and 5, CatBoost outperforms the logistic regression model in terms of its capability to distinguish between unemployed and employed individuals.

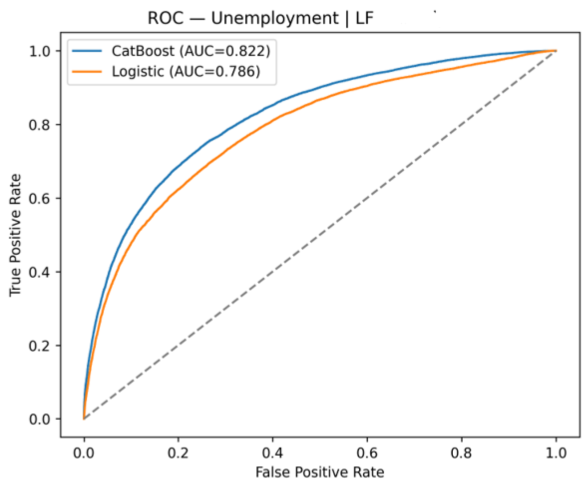

Figure 4: ROC Curves for Unemployment Given LF Participation Models

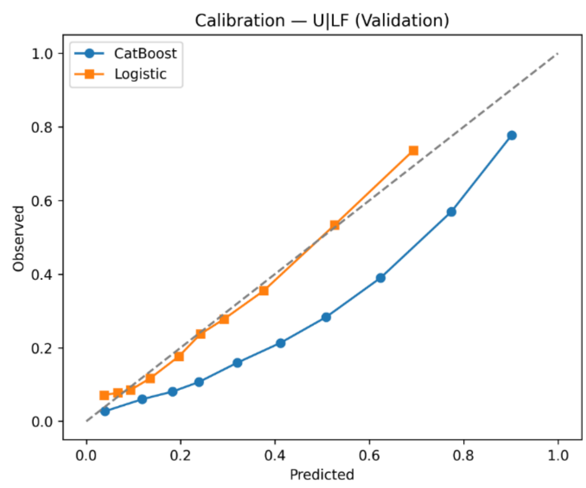

Figure 5: Calibration of Unemployment given Labour Force Participation Models

6.3 Feature Importance

SHAP values in figure 6 were used to determine the degree to which different labour market characteristics contribute to the prediction of unemployment risk.

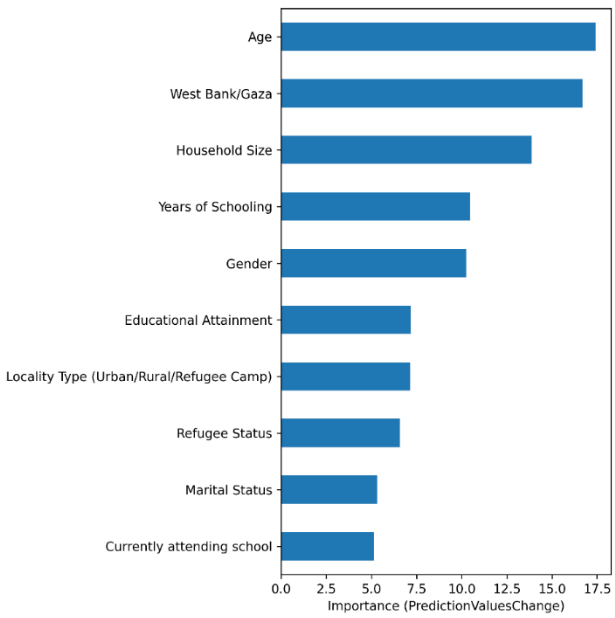

Figure 6: Ranked Importance of Features in Estimating Unemployment in the Labour Force

The ranking of features based on SHAP values identifies the top contributors to unemployment risk as follows:

1. Education Level
2. Geographic Region
3. Age
4. Gender
5. Refugee Status

These results provide further evidence that unemployment risk in Palestine is heavily driven by structural socio-economic factors.

Education is identified as one of the most significant contributing factors to unemployment risk due to the large number of unskilled workers in the Palestinian labour market, and the fact that many skilled workers leave the country in search of better economic opportunities elsewhere. The regional differences in unemployment risk can be attributed to the differences in the availability of job opportunities between the West Bank and Gaza Strip. As previously stated, there is a significant shortage of job opportunities in Gaza relative to the West Bank.

A deeper interpretation of the SHAP results suggests that unemployment risk is not shaped by a single factor, but rather by the interaction of demographic, educational, and geographic characteristics. Age is important because younger individuals, particularly those in the transition from education to work, are more exposed to entry barriers in the labour market. Education also appears as a major predictor, which may reflect the mismatch between educational attainment and available job opportunities. In this context, higher education does not necessarily protect individuals from unemployment if the labour market is unable to absorb graduates into suitable jobs. Geographic region is another central factor, as residents of Gaza face more restricted employment opportunities due to structural and institutional constraints. Gender-related differences may further interact with education and region, especially where educated women face limited employment opportunities despite their human-capital qualifications. Therefore, the SHAP results should be interpreted as evidence of overlapping risk profiles rather than isolated predictor effects.

6.4 Individual Explanation

In addition to providing information about the global feature importance, SHAP values enable the interpretation of predictions made for an individual at the individual level.

Figure 7 represents the Local SHAP Explanation for a selected individual profile. The figure illustrates how each of the labour market characteristics of the individual contribute positively or negatively to the predicted unemployment probability.

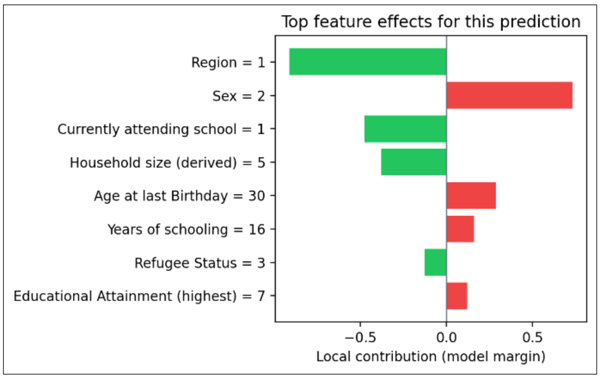

Figure 7: A SHAP bar chart illustrating the local feature contributions to the unemployment prediction for the selected profile. Positive contributions to the prediction increase the predicted unemployment risk; negative contributions decrease the predicted unemployment risk.

Local SHAP Explanations are especially useful when developing policies related to unemployment, as they provide policymaker with detailed insights into the individual-level drivers of the predictions, beyond simply looking at aggregate data. Explainable AI methods therefore improve the transparency and trustworthiness of predictive models developed for official statistical purposes (OECD, 2024b).

6.5 High-Risk Groups

To illustrate the distribution of unemployment risk across different population groups, predicted probabilities of unemployment were aggregated across demographic categories.

Table 4: Levels of Relative Unemployment Risk for Different Population Groups

Group	Risk Level
Youth (20–29)	Very High
Gaza residents	Very High
Refugee camps	High
Women graduates	High
Adults 35+	Low

There are significant differences in the risk of unemployment facing different groups within the Palestinian population as shown in the data above.

The youth are at the greatest risk of unemployment based on their transition from education to work being difficult and they have the least chance of entering into employment.

Residents of Gaza have an unemployment rate greater than other parts of Palestine, which is likely due to the structural constraints in the economy and the limited number of job opportunities available in the labour market.

In addition to youth and residents of Gaza, women who have graduated with a degree are at a similar level of relative risk as well, which suggests there are obstacles related to gender within the labour market.

Overall, the results indicate that unemployment is structurally driven in Palestine; i.e., unemployment is caused by structural factors and is not cyclically driven; however, it is focused upon several demographic and geographic groups.

7. Discussion

This study's results present significant insights into the structure and dynamics of unemployment within the Palestine labour market. Additionally, by utilizing a two-stage modelling framework, the authors provide both theoretical and practical advantages when analyzing labour markets within fragile economies.

Firstly, by using a two-stage modelling framework, the study maintains theoretical coherence with labour economics theory by explicitly modelling labour-force participation and unemployment as two sequential processes. This approach reflects the standard definition of unemployment in official statistics, where individuals are considered unemployed only if they actively participate in the labour force (ILO, 2023a; Eurostat, 2024).

By separating participation decisions from unemployment outcomes, the model reduces classification bias and improves interpretability, as each stage captures a distinct component of labour-market behaviour. This modelling strategy improves the consistency and analytical clarity of labour-market analysis, particularly when using microdata (ILO, 2023a; Eurostat, 2024).

From a labour-economics perspective, the concentration of unemployment risk among youth, university graduates, women, and Gaza residents reflects several structural characteristics of the Palestinian labour market. The high risk among university graduates may indicate a mismatch between educational outcomes and labour-market demand, where the number and type of available jobs do not fully correspond to graduates' qualifications. This suggests that unemployment among graduates is not only related to individual characteristics, but also to weak labour demand, limited private-sector absorption capacity, and possible misalignment between higher education specializations and market needs. Similarly, the elevated unemployment risk among Gaza residents reflects broader structural and institutional constraints, including restricted mobility, limited access to employment opportunities, weak business continuity, and reduced economic activity. Gender-related differences may also reflect participation barriers, limited availability of suitable jobs, and social or institutional constraints that affect women's access to employment. These findings therefore suggest that unemployment risk in Palestine should be interpreted as a structural and segmented labour-market outcome rather than only as an individual-level employment problem.

Secondly, the predictive power of machine learning models is primarily based upon an improvement in the reliability of the predicted probabilities of unemployment and not through an increase in the classification accuracy of unemployment, which supports previous studies of forecasting and predictive modelling, which have shown that machine learning algorithms generally increase the quality of probability calibration and robustness, rather than solely focusing on maximizing classification rates (Hyndman & Athanasopoulos, 2021).

Lastly, the results incorporate explainable artificial intelligence techniques to create an important link between predictive analytics and institutional responsibility and accountability. Machine learning models typically exhibit greater predictive accuracy; however, their adoption in official statistical agencies has been impeded by concerns over transparency. Using SHAP-based explainability techniques, this study demonstrates that it is feasible to explain complex machine-learning models and determine the primary drivers of unemployment risk. Incorporating explainability techniques into AI-based analytical tools can promote confidence in such systems and support their use in the production of official statistical outputs (OECD, 2024b; United Nations Economic Commission for Europe [UNECE], 2022).

It is important to emphasize that the findings of this study should be interpreted from a predictive rather than a causal perspective. The proposed framework identifies variables and population groups that are strongly associated with higher predicted unemployment risk, but it does not establish causal effects. Therefore, variables such as education, gender, region, age, and refugee status should be understood as predictors of unemployment risk within the model, rather than as causal determinants. Establishing causality would require a different research design, such as quasi-experimental methods, longitudinal causal modelling, or specific identification strategies. Accordingly, the results are intended to support labour-market monitoring, risk profiling, and evidence-informed policy discussion, rather than to provide causal estimates of the effect of each variable on unemployment.

In conclusion, this study finds that combining machine learning with explainable AI techniques will enable the creation of both more accurate predictive models, as well as models that provide meaningful insights into policy-making decisions. As such, these types of analytical techniques are particularly beneficial for the monitoring of labour markets in developing and fragile economies.

8. Policy Implications

This study presents several significant implications for labour market decision-making and for Palestinian labour market institutions. By identifying the socio-economic determinants of individual unemployment risk through the proposed framework, it is possible to develop targeted and evidence-based labour market policies.

First, the results clearly show the necessity to create more effective youth transition programs. It has been demonstrated that young people, especially those aged 20-29 years old, are at a higher unemployment risk than older age groups. Therefore, the creation of school-to-work transition programs including vocational training, internships and entrepreneurial support initiatives will likely result in a reduction in youth unemployment rates.

Second, the large regional disparities in the model results clearly show that there is a need to create regional employment investment strategies, especially in regions with structural labour market constraints. For example, targeted economic development initiatives in Gaza and other regions with lower labour market participation rates can help increase employment rates in these regions and decrease the geographic inequality in labour market outcomes.

Third, the relative high unemployment risk for university graduates also points to the need to improve job-matching mechanisms for graduates. A mismatch exists between educational outcomes and labour market requirements. Therefore, strengthening the career advice service, providing better access to labour market data and creating stronger partnerships between universities and the private sector will likely lead to an alignment of educational pathways with labour market requirements.

Fourth, the study clearly points out the need for policies to promote female labour-force participation. Gender remains the largest predictor of labour market participation and unemployment risk. Therefore, policies supporting women's participation in the labour force such as child care support, flexible work arrangements and targeted job programmes for women will contribute to decreasing the differences in labour market outcomes based on gender.

Finally, this study demonstrates the potential for using artificial intelligence-based analytical tools to support labour-market decision-making in official statistical systems. Explainable machine-learning models allow statistical institutions to produce predictive labour-market outputs while maintaining transparency and accountability. Furthermore, explainability techniques such as SHAP values enable policymakers to understand the factors driving unemployment-risk predictions, which is necessary to build trust in AI-based decision-support systems. International organisations increasingly emphasize the need to incorporate explainable artificial intelligence into official statistical systems. Transparency, interpretability, and reproducibility are core principles for adopting AI in public-sector statistical systems, as reflected in frameworks developed by the OECD and UNECE (OECD, 2024b; UNECE, 2022).

Therefore, the analytical framework presented in this report does not only contribute to labour market analysis, but also presents a practical analytical tool that can be integrated into the operational workflows of national statistical offices, e.g. the Palestinian Central Bureau of Statistics.

From an operational perspective, the proposed framework can be systematically integrated into the quarterly labour-market reporting cycle of the Palestinian Central Bureau of Statistics.

After each round of Labour Force Survey data collection, validation, and weighting, the cleaned microdata can be processed through the two-stage prediction pipeline to generate unemployment-risk indicators by age group, gender, education level, region, locality type, and refugee status. The resulting risk estimates can then be compared with official unemployment indicators and used as complementary analytical outputs in quarterly labour-market reports. In addition, SHAP-based explanations can be used to identify the main drivers of unemployment risk for each subgroup, allowing policymakers to understand not only where unemployment risk is concentrated, but also why it is concentrated in those groups. In practice, this integration could be implemented as a standard quarterly analytical workflow within PCBS, including data validation, model execution, subgroup risk profiling, technical review of model outputs, and dissemination of selected indicators through quarterly labour-market reports and internal decision-support dashboards. The framework may also support the development of early-warning tools for monitoring vulnerable groups and detecting emerging labour-market pressures before they appear clearly in aggregate indicators.

Such systems can be used to establish early warning systems for labour market developments, monitor labour markets and support evidence-based decision making.

9. Conclusion

In conclusion, the authors of this study have created an explainable artificial intelligence model to determine individual unemployment risk utilizing official Palestinian Labour Force Survey microdata. The model's two stage modelling structure utilizes the first to estimate labour force participation and the second to predict unemployment based on the probability of labour force participation. This modelling structure ensures that the model adheres to the economic definition of unemployment and reduces the likelihood of misclassification error in predictive modelling.

The empirical analysis demonstrated that machine learning models outperformed traditional statistical methods in predicting unemployment risk. The CatBoost model showed stronger discriminatory performance than the logistic regression model. Additionally, the use of explainable artificial intelligence tools, such as SHAP values, allowed for transparent explanation of the model’s predictions and the identification of the main predictors associated with unemployment risk.

Based on the empirical findings, the researchers found that unemployment risk is highly concentrated within specific demographic and geographic groups. Specifically, youth, women, residents of Gaza and residents of refugee camps exhibited significantly higher predicted unemployment rates than other groups. These findings are consistent with previous research that has indicated that unemployment in the Palestinian labour market is primarily structural (i.e., due to institutional and economic factors) and not cyclical.

In addition to providing empirical evidence regarding unemployment risk in Palestine, the authors provide evidence that microdata from official surveys can serve as an operational analytical tool for the monitoring of labour markets. By combining machine learning algorithms with explainability techniques and national level survey data, the authors propose a practical decision support system that can assist in developing evidence-based policies for labour markets.

Potential future directions of the research include expanding the framework to incorporate additional macroeconomic variables, high frequency labour market data and/or sector level employment data. Additionally, the incorporation of real-time data sources and automatic model updating processes may allow for the development of more responsive unemployment prediction systems in rapidly changing economic environments.

In addition, the post-2023 geopolitical shocks represent an important source of uncertainty for the future stability of unemployment-risk prediction models in Palestine. In fragile and conflict-affected contexts, such shocks may create structural breaks in labour-market behaviour by affecting mobility, business continuity, access to workplaces, sectoral employment opportunities, household income strategies, and labour-force participation decisions. As a result, the relationships learned by the model from historical Labour Force Survey microdata may not remain fully stable when applied to future survey rounds. This raises the possibility of model drift, where the predictive performance and the relative importance of key risk factors may change over time. Therefore, future applications of the proposed framework should not treat the model as static, but should include periodic model updating, continuous validation against newly released quarterly LFS data, and regular monitoring of prediction errors and SHAP-based explanations. Moreover, the integration of macroeconomic and shock-related indicators, such as GDP growth, inflation, restrictions on movement, conflict-related disruptions, sectoral employment changes, and emergency assistance flows, could help improve the model's robustness under rapidly changing conditions. These steps would strengthen the predictive stability, interpretability, and policy relevance of the framework in fragile and highly volatile labour-market environments.

Overall, the research makes a significant contribution to the emerging literature on the application of artificial intelligence in the analysis of labour markets and illustrates the ability of explainable machine learning models to improve labour market monitoring and policy design in fragile and developing countries.

Disclosure Statement

The authors declare that they have no relevant or material financial interests that relate to the research described in this paper

– **Ethical approval and consent to participate:** This study is based on secondary analysis of anonymized microdata from the Palestinian Labour Force Survey conducted by the Palestinian Central Bureau of Statistics (PCBS). No new data were collected directly from participants, and the researchers had no access to personally identifiable information. The data were processed in accordance with PCBS confidentiality and data-protection regulations.
– **Availability of data and materials:** The microdata used in this study are subject to the access, confidentiality, and data-dissemination regulations of the Palestinian Central Bureau of Statistics. Access may be granted through PCBS procedures and subject to institutional approval.
– **Author contribution:** Haitham Zeidan contributed to the conceptualization, methodology, data preparation, software implementation, statistical and machine-learning analysis, visualization, interpretation of results, and preparation of the original manuscript. Adel Karaa contributed to research supervision, methodological guidance, validation, and critical review of the manuscript. Rabeh Morrar contributed to the interpretation of the labour-market and policy implications and to the review and editing of the manuscript. All authors reviewed and approved the final version of the manuscript.
– **Conflict of interest:** The authors declare that they have no competing or conflicting interests related to this study.
– **Funding:** This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.
– **Acknowledgments:** The authors acknowledge the Palestinian Central Bureau of Statistics for providing the Labour Force Survey microdata used in this study. The authors also appreciate the academic guidance and institutional support provided by the University of Tunis and An-Najah National University.

References

- Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: Forecasting and control (5th ed.). Wiley.
- Celbiş, M. G. (2023). Unemployment in rural Europe: A machine learning perspective. Applied Spatial Analysis and Policy, 16, 1071-1095. https://doi.org/10.1007/s12061-022-09464-0
- Costa, E. A., Silva, M. E., & Galvão, A. B. (2024). Real-time nowcasting the monthly unemployment rates with daily Google Trends data. Socio-Economic Planning Sciences, 95, Article 101963. https://doi.org/10.1016/j.seps.2024.101963
- Dorogush, A. V., Ershov, V., & Gulin, A. (2018). CatBoost: Gradient boosting with categorical features support. arXiv. https://doi.org/10.48550/arXiv.1810.11363
- Eurostat. (2024). Methodology - Employment and unemployment (LFS). European Commission. https://ec.europa.eu/eurostat/web/lfs/methodology
- Fawcett, T. (2006). An introduction to ROC analysis. Pattern Recognition Letters, 27(8), 861-874. https://doi.org/10.1016/j.patrec.2005.10.010
- Ferro, C. A. T. (2007). Comparing probabilistic forecasting systems with the Brier score. Weather and Forecasting, 22(5), 1076-1088. https://doi.org/10.1175/WAF1034.1
- Gabrikova, B., Svabova, L., & Kramarova, K. (2023). Machine learning ensemble modelling for predicting unemployment duration. Applied Sciences, 13(18), Article 10146. https://doi.org/10.3390/app131810146
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359-378. https://doi.org/10.1198/016214506000001437
- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In D. Precup & Y. W. Teh (Eds.), Proceedings of the 34th International Conference on Machine Learning (Vol. 70, pp. 1321-1330). PMLR. https://proceedings.mlr.press/v70/guo17a.html
- Gupta, R., Pierdzioch, C., & Salisu, A. A. (2022). Oil-price uncertainty and the U.K. unemployment rate: A forecasting experiment with random forests using 150 years of data. Resources Policy, 77, Article 102662. https://doi.org/10.1016/j.resourpol.2022.102662
- Güler, M., Kabakçı, A., Koç, Ö., Eraslan, E., Derin, K. H., Güler, M., Ünlü, R., Türkan, Y. S., & Namlı, E. (2024). Forecasting of the unemployment rate in Turkey: Comparison of the machine learning models. Sustainability, 16(15), Article 6509. https://doi.org/10.3390/su16156509
- Hancock, J. T., & Khoshgoftaar, T. M. (2020). CatBoost for big data: An interdisciplinary review. Journal of Big Data, 7, Article 94. https://doi.org/10.1186/s40537-020-00369-8
- Hyndman, R. J., & Athanasopoulos, G. (2021). Forecasting: Principles and practice (3rd ed.). OTexts. https://otexts.com/fpp3/
- International Labour Organization. (2022). Global employment trends for youth 2022:

Investing in transforming futures for young people. https://doi.org/10.54394/QSMU1809
– International Labour Organization. (2023a). Resolution II: Resolution to amend the 19th ICLS resolution concerning statistics of work, employment and labour underutilization. https://www.ilo.org/resource/conference-paper/resolution-ii-resolution-amend-19th-icls-resolution-concerning-statistics
– International Labour Organization. (2023b). World employment and social outlook: Trends 2023. https://doi.org/10.54394/SNCP1637
– International Labour Organization. (2024a). The impact of care responsibilities on women's labour force participation. https://doi.org/10.54394/LPTT5569
– International Labour Organization. (2024b). World employment and social outlook: Trends 2024. https://doi.org/10.54394/HQAE1085
– International Monetary Fund. (2023). World economic outlook, October 2023: Navigating global divergences. https://doi.org/10.5089/9798400235801.081
– James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). An introduction to statistical learning: With applications in R (2nd ed.). Springer. https://doi.org/10.1007/978-1-0716-1418-1
– Kim, K. (2025). Unemployment dynamics forecasting with machine learning regression models. arXiv. https://doi.org/10.48550/arXiv.2505.01933
– Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon et al. (Eds.), Advances in Neural Information Processing Systems 30 (pp. 4765-4774). Curran Associates. https://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions
– Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 13(3), Article e0194889. https://doi.org/10.1371/journal.pone.0194889
– Molnar, C. (2022). Interpretable machine learning: A guide for making black box models explainable (2nd ed.). https://christophm.github.io/interpretable-ml-book/
– Nigar, M., Juli, J. F., Golder, U., Alam, M. J., & Hossain, M. K. (2025). Artificial intelligence and technological unemployment: Understanding trends, technology's adverse roles, and current mitigation guidelines. Journal of Open Innovation: Technology, Market, and Complexity, 11(3), Article 100607. https://doi.org/10.1016/j.joitmc.2025.100607
– Organisation for Economic Co-operation and Development. (2024a). OECD employment outlook 2024: The net-zero transition and the labour market. OECD Publishing. https://doi.org/10.1787/ac8b3538-en
– Organisation for Economic Co-operation and Development. (2024b). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449; amended 3 May 2024). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449
– Palestinian Central Bureau of Statistics. (2023a). Labour Force Survey: Third quarter 2023 [Press report].

https://www.pcbs.gov.ps/portals/_pcbs/PressRelease/Press_En_LFSQ32023E.pdf

- Palestinian Central Bureau of Statistics. (2023b). Palestinian Labour Force Survey: Annual report 2022. https://www.pcbs.gov.ps/Downloads/book2658.pdf
- Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K.-R. (Eds.). (2019). Explainable AI: Interpreting, explaining and visualizing deep learning. Springer. https://doi.org/10.1007/978-3-030-28954-6
- United Nations Economic Commission for Europe. (2022). Machine learning for official statistics. United Nations. https://unece.org/sites/default/files/2022-09/ECECESSTAT20216.pdf
- Vaicenavicius, J., Widmann, D., Andersson, C., Lindsten, F., Roll, J., & Schön, T. B. (2019). Evaluating model calibration in classification. In K. Chaudhuri & M. Sugiyama (Eds.), Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (Vol. 89, pp. 3459-3467). PMLR. https://proceedings.mlr.press/v89/vaicenavicius19a.html
- World Bank. (2023a). Racing against time: World Bank economic monitoring report to the Ad Hoc Liaison Committee. https://doi.org/10.1596/40403