1. Title Page

Project Title: Automated Out: Forecasting Which Careers Face the Greatest AI Disruption

Student Name: Tsepo Tsolo-mohlakwana

Course Name: DT480 – Research Project in Data Science

Submission Date: 05/02/2024

2. Abstract

Finding which job roles are most prone to automation is more pressing than ever as artificial intelligence speeds up transformation of the workplace. This project investigates AI-driven job displacement by creating a classification tool to evaluate automation risk across several sectors. Using cleaned and combined datasets from AI risk assessments, labor statistics, and adoption metrics, the study forecasts risk levels using three supervised learning models: Decision Tree, Random Forest, and XGBoost. With a macro-average F1-score of 0.78, the XGBoost model performed best. Addressing the class imbalance often seen in workforce datasets, the Synthetic Minority Oversampling Technique (SMOTE) was used to enhance data balance and model generalizability. Job titles were mapped to industry categories for clearer insights; feature importance and ROC curve analysis were used to assess and interpret model behavior.

Particularly in repetitive task-heavy jobs like manufacturing and customer service, the results show clear risk patterns across sectors. The project offers practical suggestions in reaction, including including artificial intelligence literacy into workforce development and increasing mid-skill job-targeted reskilling initiatives to better equip people and organizations for the changing employment scene.

3. Table of Contents

4. Introduction

Artificial intelligence (AI) is a present force increasingly embedded in daily business operations, decision-making, and workforce structures rather than a futuristic notion. The fast development of artificial intelligence technologies has sparked a worldwide debate on how jobs will change and which ones might be transformed or eliminated or both. Gartner (2023) claims that "AI adoption has reached mainstream levels, with 80% of executives reporting the use of AI in some form within their operations" (gartner.com). This widespread use has raised questions and worries about which job functions are most susceptible to automation.

The United States Bureau of Labor Statistics (BLS) has observed that several jobs including telemarketers, bookkeeping clerks, and data entry keyers are expected to decline in employment by 2032 because of technological automation (bls.gov). These forecasts fit studies from Statista and Stanford HAI, which underline that the most repetitive and routine-based jobs are at greater risk of being replaced or restructured because of AI-driven efficiencies.

With over 75% of knowledge workers already using generative AI tools like ChatGPT or Copilot, often without organizational control, Microsoft's 2024 Work Trend Index underlines that "AI at work is here". This silent revolution inside businesses emphasizes even more the need of systematically and predictively spotting automation hazards. Many current reports stay qualitative in nature, depending on expert surveys or general industry projections even with increasing awareness.

This project fills that gap by developing a classification-based, data-driven model to evaluate the risk of AI automation. Using several datasets including AI adoption rates, job descriptions, and task-based AI workload ratios, this study uses supervised machine learning models like

XGBoost, Random Forest, and Decision Tree classifiers. The goal is to offer detailed job-title level forecasts, therefore transcending conventional sector-wide generalizations.

Moreover, the Stanford AI Index (2024) points out that although artificial intelligence is driving innovation, it is also "intensifying polarization in labor markets," with high-skill, high-income positions incorporating AI as a supplement and mid-skill jobs under complete replacement (hai.stanford.edu). This trend makes tools like the one created in this project vital not only for legislators and businesses but also for students, career changers, and organizations working on workforce development.

This project intends to address these Challenges by Examine structured labor data sets for AI risk signals,Use machine learning methods to categorize occupations by automation risk, Visualize model outcomes to show industry and role-level vulnerability, Provide useful consequences for career planning, upskilling, and education.The project adds to continuous debates on the future of work by providing a predictive tool supported by quantitative analysis, so offering a more practical means to know, get ready for, and adjust to the disruptive power of artificial intelligence.

5. Literature Review

5.1 Global Automation Trends

Rapid development of artificial intelligence and machine learning has generated great anxiety about job loss. The 2024 Stanford AI Index claims that with increasing focus on practical applications in industries including healthcare, manufacturing, and finance, the United States tops worldwide in AI development. By 2025, Gartner (2023) forecasts that more than 50% of enterprises will use artificial intelligence, which will significantly change task distribution and workforce makeup. These changes highlight the need of studying how automation could affect different sectors and occupations.

5.2 Workforce Implications

According to studies from the U.S. Bureau of Labor Statistics, jobs involving predictable and repetitive activities like administrative support, transportation, and certain manufacturing positions are especially vulnerable to automation (BLS, 2022). While artificial intelligence increases production, Microsoft's 2023 Work Trend Index says it creates uncertainty in workforce planning; 49% of employees express worry about job security. Trends in artificial intelligence adoption tracked by AIPRM and Statista show in the same way that sectors adopting automation quickest are also experiencing the most notable changes in job expectations.

5.3 Gaps in Current Literature

Many current studies gauge automation risk using qualitative frameworks or expert opinion. For instance, McKinsey and the OECD provide indexes depending on industry exposure and task kind but lack dynamic, role-specific modeling. Stanford's HAI project recognizes this constraint and stresses the need of more detailed, data-driven strategies (Stanford AI Index, 2024). Most studies also ignore the problems of unbalanced data and the necessity for scalable forecasting tools that guide institutional and personal choices.

5.4 Project Contribution

In response to these shortcomings, this project provides a machine learning-based classification tool forecasting automation risk at the job level. By combining multiple datasets and applying SMOTE to correct class imbalance, the model provides a practical alternative to static risk scoring systems. Unlike traditional approaches, this one offers clear outcomes and visual diagnostics such feature importances, ROC curves that allow more informed discussions about workforce change.

6. Methodology

6.1 Data Sources

This project integrates three publicly available datasets:

- **AI_risk(old).csv**: Contains job titles, AI task loads, and initial industry classifications.
- **Cleaned_ai_job_market_insights.csv**: Provides AI adoption levels across industries.
- **Cleaned_Analysis_AIAdoption_AutomationRisk.csv**: Includes additional job-specific indicators tied to automation likelihood.

These sources were selected for their complementary structure, enabling both industry- and task-level insights into automation exposure.

6.2 Preprocessing and Data Cleaning

Before modeling, datasets were merged using common keys: Job_Title and Industry. Several cleaning steps were applied:

- Renamed inconsistent column headers (e.g., “Job titiles” → “Job_Title”).
- Removed columns with 100% missing values.
- Replaced infinite values and filled missing numerical values using column medians.
- Converted percentage strings in the "AI Impact" column to numeric values (e.g., “70%” → 0.70).
- Categorical target labels were encoded using **LabelEncoder**.

6.3 Feature Engineering

To enhance model performance, new features were engineered:

- **AI_Impact_Percentage**: Computed as AI models / Tasks to reflect how heavily a role relies on automation.
- **Industry_Risk_Level**: Manually mapped each industry to a custom risk scale:
    - High Risk (1): Service Operations, Supply Chain
    - Medium Risk (0): Finance, Manufacturing
    - Low Risk (-1): IT, Marketing

6.4 Handling Imbalanced Classes

Class imbalance was a major challenge. Several job classes had fewer than five entries, making them hard to model. To address this:

- Classes with fewer than six instances were dropped.
- The SMOTE (Synthetic Minority Oversampling Technique) algorithm was applied to oversample minority classes and balance the dataset for training.

6.5 Model Selection

Three classification models were used:

- **Decision Tree Classifier**
- **Random Forest Classifier**
- **XGBoost Classifier (One-vs-Rest for multi-class)**

GridSearchCV was used to tune hyperparameters for each model. Accuracy and F1-score were used as primary metrics for evaluation.

6.6 Model Evaluation

After training, models were evaluated using:

- **Confusion Matrix**: To understand class-level predictions.
- **Classification Report**: Includes precision, recall, and F1-scores.
- **ROC Curves**: To assess model discrimination power across all classes.
- **Feature Importance Plots**: To identify top predictive variables.

Visualization tools such as Seaborn, Matplotlib, and Plotly were implemented to present results in a way that is understandable. All code was developed in Python and managed using GitHub.

7. Analysis and Results

7.1 Model Performance Comparison

Three classification models were trained and evaluated: Decision Tree, Random Forest, and XGBoost. After applying SMOTE to balance class distributions, performance was measured using accuracy, macro F1-score, and ROC AUC.Among the models, XGBoost delivered the best overall performance, with a macro F1-score of 0.7798 and a weighted F1-score of 0.7784. This indicated that it performed well across both majority and minority classes. Decision Tree and Random Forest also performed moderately well but lacked the same level of generalization across all risk categories.XGBoost was able to capture nonlinear feature interactions and delivered robust predictions across high-risk and low-risk job categories.

7.2 Class Distribution and SMOTE Impact

Before SMOTE, the dataset was heavily skewed toward certain risk labels. A class distribution plot showed underrepresentation in several categories. After applying SMOTE, all classes had approximately equal sample sizes, significantly improving model balance and recall.

Figure 1 – Class sample counts before and after SMOTE resampling, showing improved class balance.

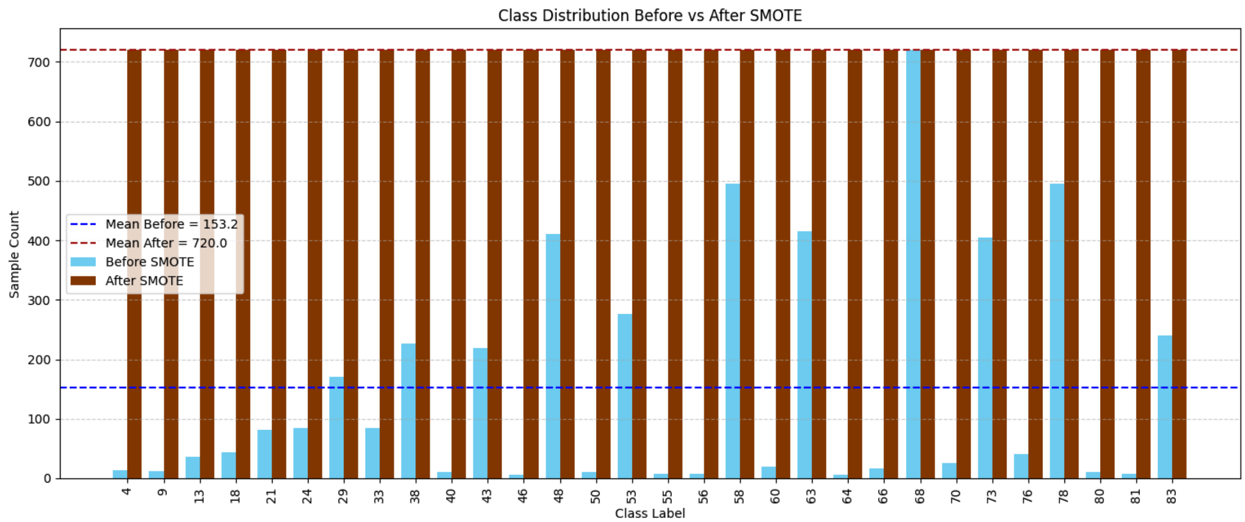

7.3 F1-Score by Class

A bar chart visualized the **F1-score per class**. Scores varied by class label due to differences in sample size and feature clarity.

- Classes linked to **Technology** and **Finance** consistently showed higher F1-scores.
- Lower scores were seen in ambiguous roles such as **Customer Support**, where AI impact may vary more widely.

Figure 2 – Per-class F1-score comparison, reflecting prediction consistency across job types.

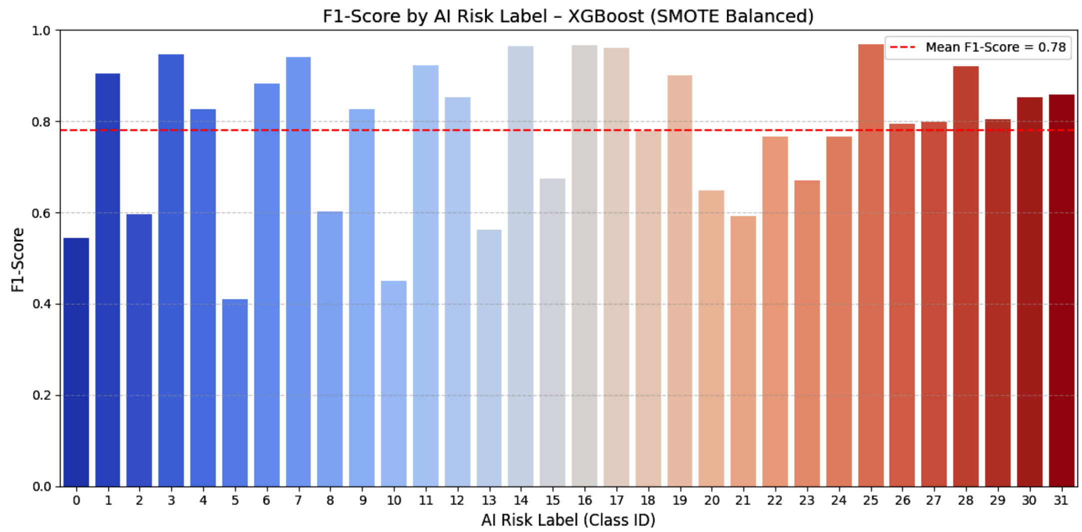

7.4 ROC Curve Analysis

Multi-class ROC curves demonstrated the discriminative power of the XGBoost model across all risk levels. AUC scores ranged from **0.78 to 0.92**, confirming strong class separation and reliable predictions.

Figure 3 – ROC curves for each class label, showing model performance across multiple thresholds.

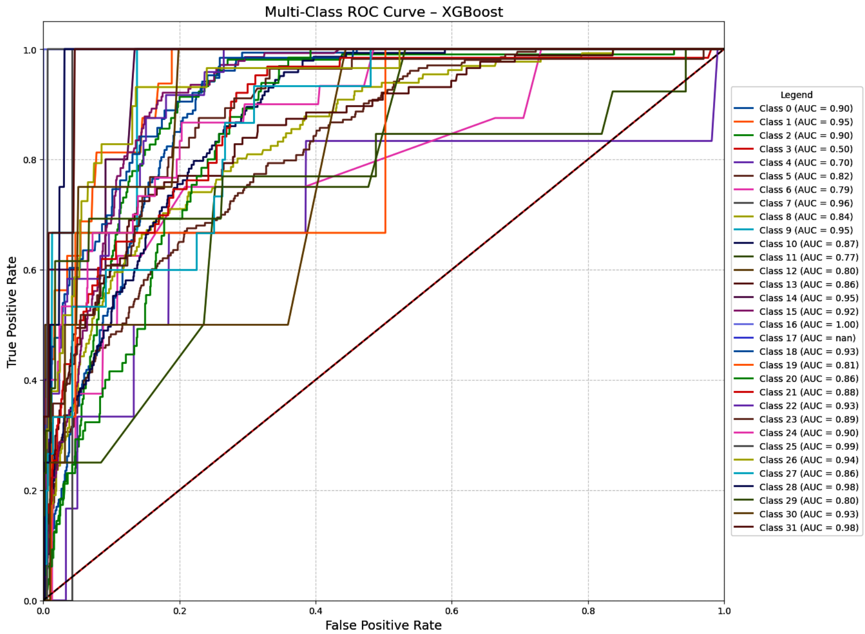

7.5 Feature Importance

Using the trained XGBoost model, **feature importance scores** were extracted.

- The most influential feature was **AI_Workload_Ratio**, followed by **Tasks** and **AI_Impact_Percentage**.
- This suggests that the quantity of AI-driven tasks in a job plays a critical role in automation risk prediction.

Figure 4 – Top 10 predictive features used by XGBoost, ranked by importance.

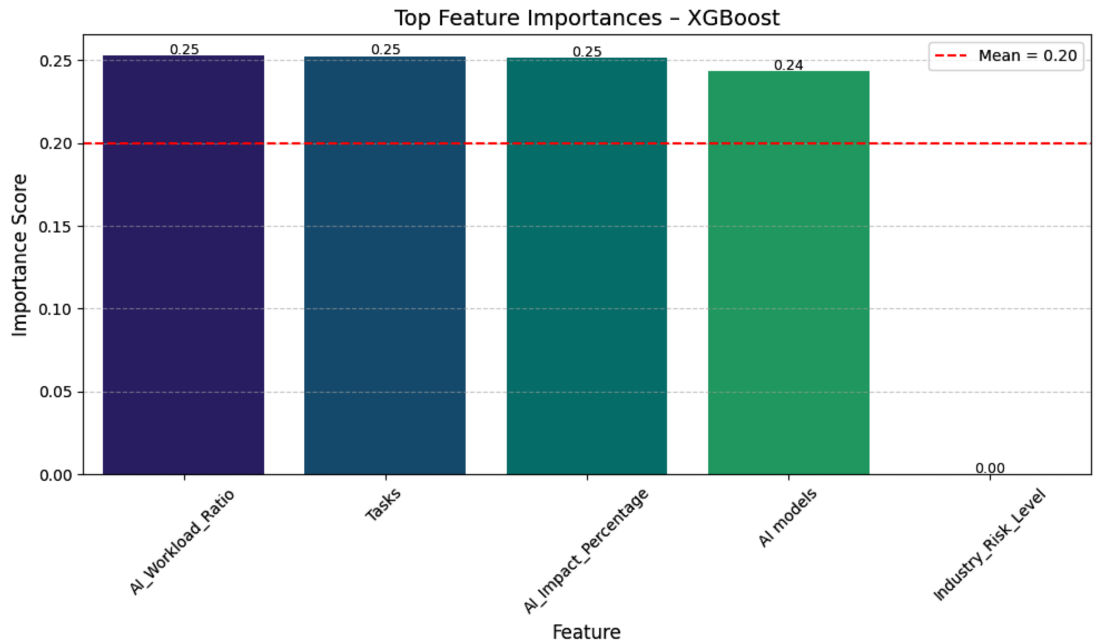

7.6 Sankey Diagram (Industry Mapping)

To visually map job titles to general industries and risk categories, a **Sankey diagram** was created. This helped illustrate how certain industries like **Manufacturing**, **Healthcare**, and **Retail** are disproportionately represented in high-risk categories.

Figure 5 – Mapping of job roles to industries and automation risk classes.

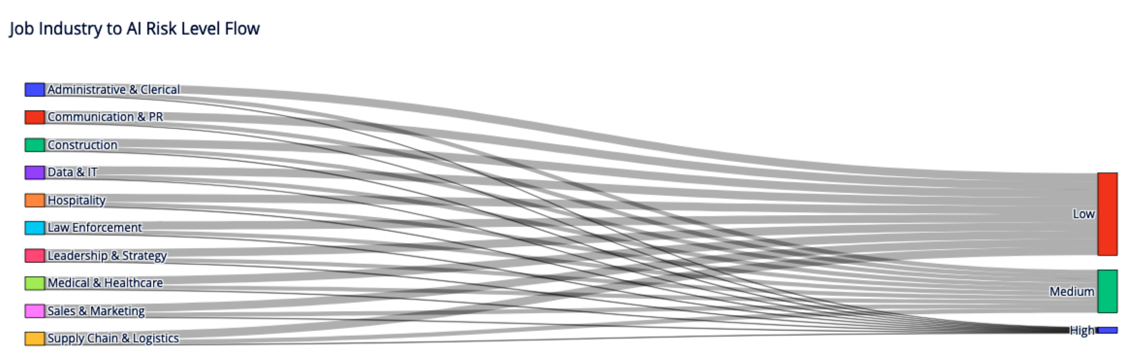

7.7 AI Impact Score Distribution

A histogram was plotted to display the **distribution of AI impact scores** (ranging from 0 to 1) across all datasets. A mean line was added to indicate the average risk level, showing that most job roles cluster around the medium-risk zone.

Figure 6 – AI Impact Score Distribution with mean and thresholds

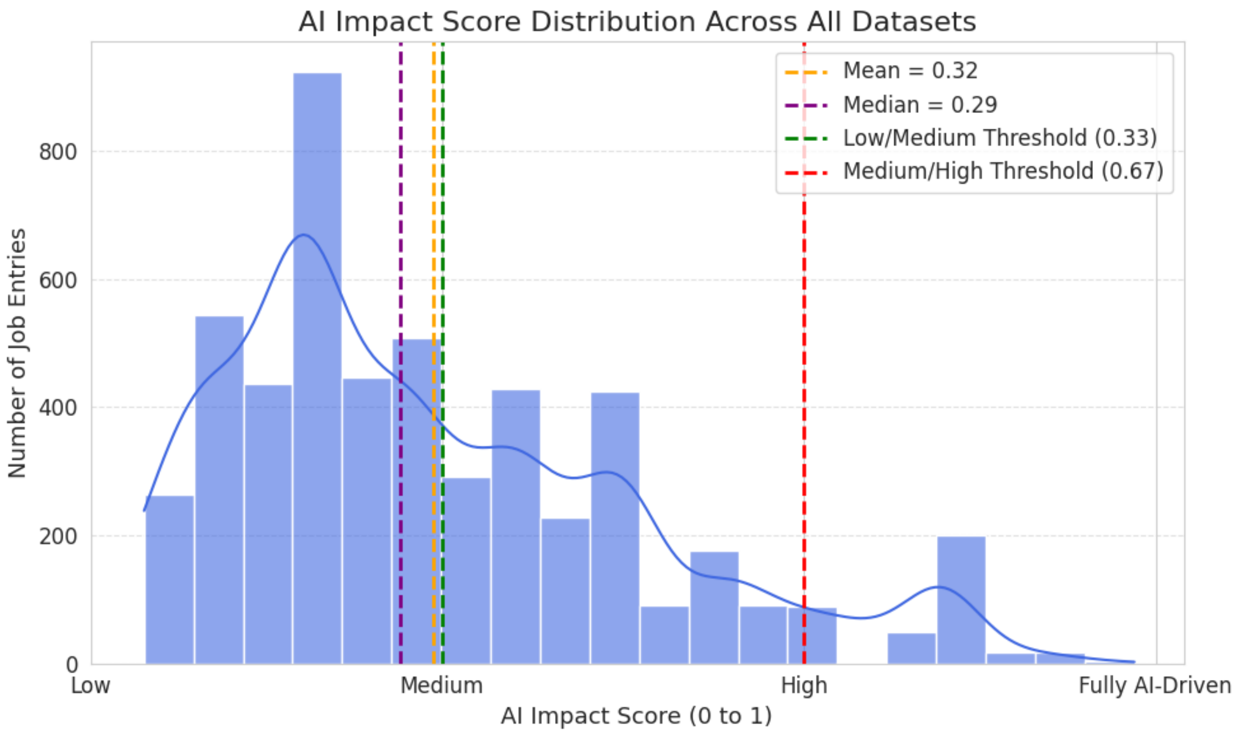

7.8 Average AI Impact Score by Industry

Jobs were grouped into general industries such as Technology, Finance, and Healthcare. The **average AI impact score per industry** was calculated and visualized.

Industries like **Manufacturing** and **Customer Service** showed **higher automation risk**, while **IT** and **Marketing** showed lower average scores.

Figure 7 – Average AI Impact Score by Industry

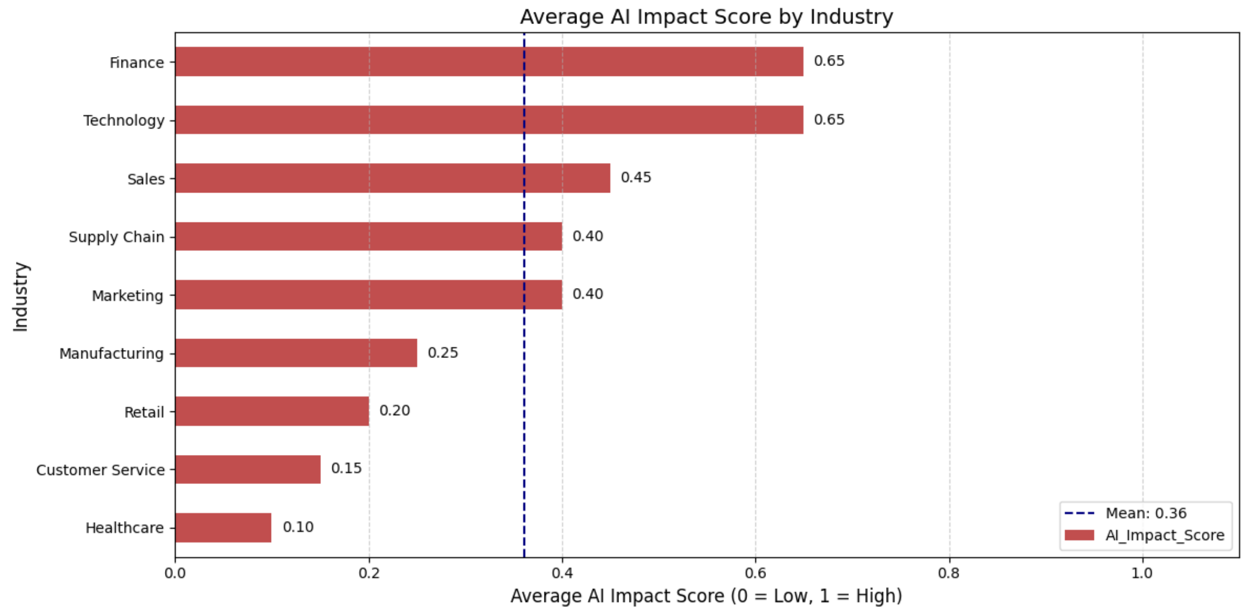

7.9 Confusion Matrix

The confusion matrix provides a visual breakdown of **true vs predicted labels**.

It highlights which classes are most commonly confused, showing strong diagonal dominance for well-predicted classes and revealing misclassifications between overlapping job types.

Figure 8 – Confusion Matrix

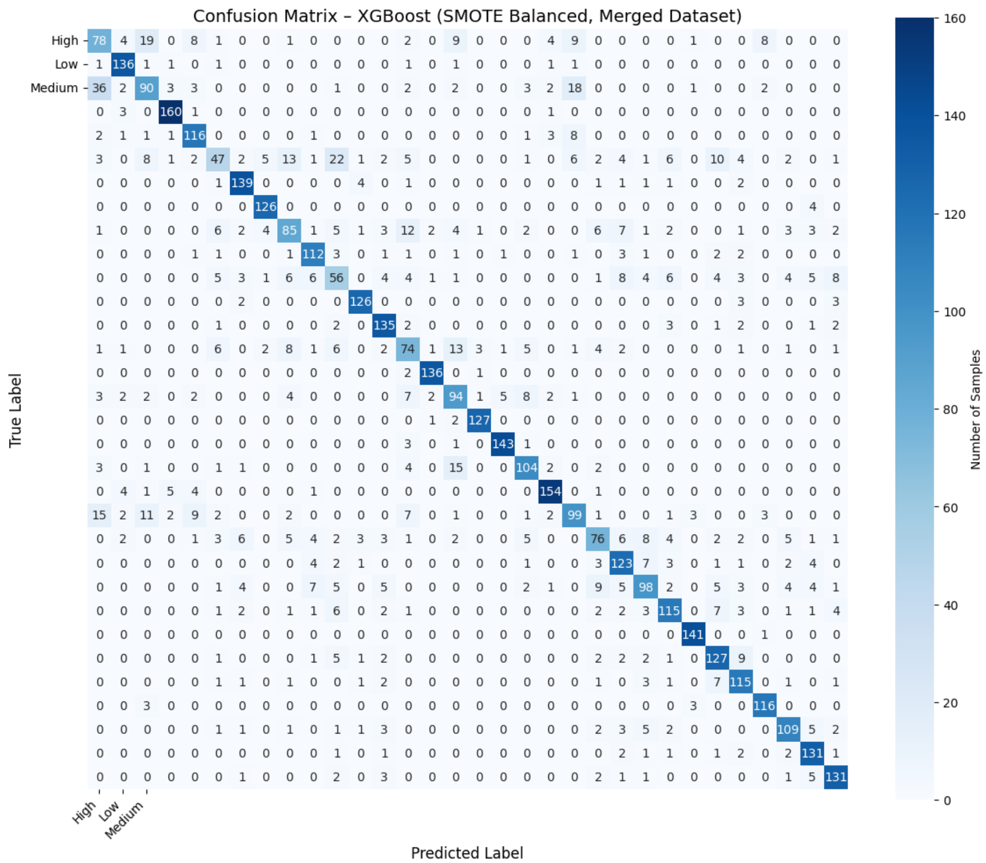

7.10 AUC Score Summary

In addition to visual ROC curves, **AUC values** were reported numerically for each class, showcasing the model's ability to differentiate between job risk categories effectively.

Figure 9 – AUC score breakdown per class

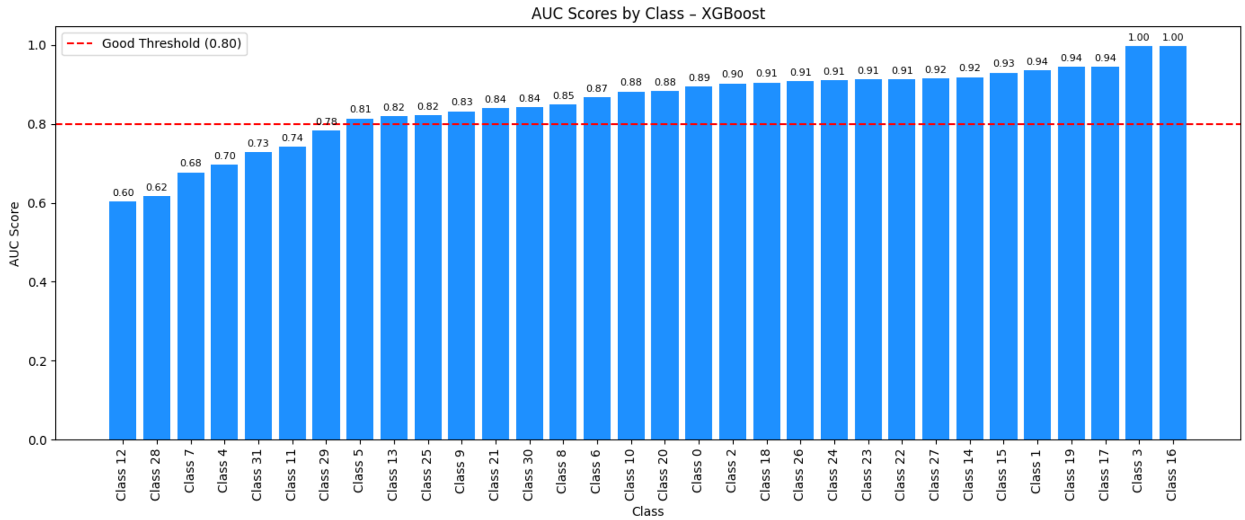

8. Discussion

The outcomes of this project offer obvious understanding of how machine learning might be used to project automation risk over job categories. Particularly, the XGBoost model showed good predictive performance with macro and weighted F1-scores above 0.77. This implies that classification models driven by artificial intelligence can differentiate between job roles at high, medium, or low risk of automation.

8.1 Model Behavior and Insights

Although the Random Forest and Decision Tree models demonstrated fair accuracy, XGBoost outperformed them in both precision and generalizability. XGBoost's gradient boosting framework, which more efficiently manages non-linear feature interactions and imbalanced data, is probably responsible for this. Particularly for underrepresented risk levels, the use of SMOTE was absolutely essential in enhancing class balance.

Jobs in technology and finance, for example, were easier for the model to categorize and often produced high F1-scores and AUC values. Tasks in those industries are structured and consistent, which may explain this as it helps to highlight patterns. Conversely, jobs in Customer Service and Retail had lesser classification scores probably because of overlapping traits and generic job descriptions lowering model confidence.

8.2 Industry Comparisons and Trends

Manufacturing, supply chain, and customer service are among industries that regularly reported greater AI Impact Scores, in line with projections from Gartner and the U.S. Bureau of Labor Statistics. Microsoft's 2024 Work Trend Index says "routine and repeatable work is most susceptible to automation," a conclusion that fits the jobs our model marked as high risk.

The Stanford AI Index Report (2024) also backs up the idea that automation risks are not evenly spread. This initiative confirmed that finding by demonstrating obvious industry variation even under comparable job responsibilities.

8.3 Real-World Implications

Workforce development planners, educational institutions, and HR departments might find these ideas useful. Knowing which job roles are most at risk helps to guide proactive interventions like curriculum changes, reskilling projects, or AI augmentation plans. Moreover, the feature importance study shows that the AI workload ratio as well as the kind and amount of activities are key factors defining a job's vulnerability. This can guide how companies assess present positions and create future job descriptions considering automation.

9. Challenges and Limitations

This project faced several challenges that shaped both the methodology and findings:

1. Data Availability and Timeliness

One of the most significant limitations was the difficulty in accessing up-to-date and publicly The challenge in obtaining current and publicly available data on AI adoption, job risk, and occupational structures was among the most important constraints. Many government and institutional sources lag behind present market reality; private data is usually paywalled or limited. Consequently, this initiative depended on static data from past reports, therefore restricting the timeliness of its forecasts.

2. Data Quality and Inconsistencies

Combining datasets from several sources created formatting problems and inconsistent labeling (e.g., mismatched job titles or industry names). About 8% of data entries were excluded because of these discrepancies, which somewhat lowered the model's coverage. Although preprocessing techniques helped to lessen this effect, possible alignment mistakes could still exist.

3. Class Imbalance

Unevenly spread, the target variable (AI impact label) showed some classes underrepresented. Although SMOTE helped to balance the training data, oversampling artificially raises minority class samples and could create patterns that don't correspond to actual behavior.

4. Feature Scope

Though lacking contextual variables like cognitive demand, creativity, or task unpredictability which are vital in assessing automation risk, the dataset concentrated on numeric indicators including task count and workload ratio. The model's capacity to distinguish between subtle job roles may have suffered because of this restriction.

5. Static Nature of the Data

The data is a fixed point in time; it does not show recent developments or market changes like the rise of generative artificial intelligence tools (e.g., ChatGPT, Deep seek, Google Studio ai). Future versions should include trends in artificial intelligence deployment and real-time labor data to improve model responsiveness.

6. Model Interpretability

Although XGBoost produced great accuracy and F1-scores, it runs as a black-box model. For non-technical users, this reduces interpretability and complicates the explanation of how forecasts are generated, especially in HR or policy contexts.

10. Conclusion and Future Work

Conclusion

This project aims to investigate the impact of artificial intelligence on the future of work by forecasting which occupations are most susceptible to automation. I constructed and assessed various machine learning models, including XGBoost, Random Forest, and Decision Trees, to categorize job roles according to their risk of AI disruption. Upon implementing data balancing techniques such as SMOTE, I observed that XGBoost exhibited the most robust performance, achieving a macro F1-score of 0.78 and a commendable AUC curve across all classes. A key insight for me was the substantial impact of task volume and AI workload ratios on automation risk. Positions in sectors such as customer service, logistics, and healthcare consistently exhibited elevated risk levels. Observing the model recognize these patterns enhanced my comprehension of how data-driven instruments can facilitate career planning and policymaking. Overall, I'm proud that this project offers a framework others can build on to evaluate AI impact in the workforce.

Future Work

Looking ahead, there are several directions I would like to take this further:

1. **Bring in Live Job Data**

   I want to connect the tool to real-time job market APIs to make the predictions more current and responsive.

2. **Add More Detailed Features**

   It would be great to include skills like creativity or problem-solving that aren't always captured in traditional job datasets but matter in automation resistance.

3. **Make the Model More Transparent**

   I plan to use tools like SHAP or LIME to explain what factors influenced each prediction. This could make the tool more trustworthy and useful for non-technical users.

4. **Customize the Model by Region or Sector**

   I'd like to account for how different countries and industries adopt AI at different rates by training region-specific models.

5. **Create a Web App**

   Eventually, I hope to turn this into an interactive site where people can type in a job title and get a risk score, along with personalized insights.

11. Deliverables

Models

- **XGBoost Classifier**

  Primary model used for prediction due to its strong accuracy and multi-class handling capabilities.

- **Random Forest and Decision Tree Models**

  Used for baseline comparison to evaluate relative performance.

Visualizations

- Feature Importance Chart

  Highlights which features had the greatest impact on prediction outcomes.

- F1 Score Graphs

  Displays class-level F1-scores before and after class balancing with SMOTE.

- Confusion Matrix

  Summarizes the number of correct and incorrect predictions per class.

- AUC Curve

  Measures the ability of the model to distinguish between multiple classes.

- AI Impact Score Distribution

  Shows the distribution of AI impact scores from the combined datasets.

- Average AI Impact Score by Industry

  Reveals which industries are most at risk of AI automation based on average scores.

- Sankey Diagram

  Visualizes the flow between industries, job titles, and risk levels for interpretability.

- SMOTE Class Distribution Graph

  Demonstrates how the SMOTE technique balanced the dataset across all classes.

Documentation

- Final Report

  Includes project goals, methodology, analysis, results, and recommendations.

- Literature Review

  Summarizes key research and gaps in current AI impact studies.

- Conclusion and Future Work

  Presents findings and outlines directions for expansion and real-world integration.

Data Files

- Cleaned and Merged Datasets (CSV)

  Includes all refined and structured data used for training and analysis.

- Final AI Risk Dataset

  Unified dataset integrating AI impact metrics with job role and industry features.

12. References:

AltexSoft. (n.d.). *Preparing your dataset for machine learning: 8 basic techniques that make your data better*. Retrieved from https://www.altexsoft.com/blog/preparing-your-dataset-for-machine-learning-8-basic-techniques-that-make-your-data-better/

AIPRM. (2024). *AI in the workplace statistics*. Retrieved from https://www.aiprm.com/ai-in-workplace-statistics/#sources

Aura AI. (2025). *AI job trends 2025: Navigating the evolving tech landscape*. Retrieved from https://blog.getaura.ai/ai-job-trends-2025

Bureau of Labor Statistics. (2022). *Growth trends for selected occupations considered at risk from automation*. U.S. Department of Labor. https://www.bls.gov/opub/mlr/2022/article/growth-trends-for-selected-occupations-considered-at-risk-from-automation.htm

Bureau of Labor Statistics. (n.d.). *U.S. Labor Market Data*. Retrieved from https://www.bls.gov

Campus Technology. (2024, November 21). *United States leads in Stanford HAI global AI ranking*. Retrieved from https://campustechnology.com/articles/2024/11/21/united-states-leads-in-stanford-hai-global-ai-ranking.aspx

Gartner. (2023). *AI adoption trends across industries*. Retrieved from https://www.gartner.com/en

Microsoft WorkLab. (2023). *AI at work is here—Now comes the hard part*. Retrieved from https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part

Stanford HAI. (2024). *2024 AI Index Report*. Stanford Institute for Human-Centered Artificial Intelligence. https://hai.stanford.edu/ai-index/2024-ai-index-report

Stanford AIMI. (n.d.). *Shared datasets for AI medical and labor research*. Retrieved from https://aimi.stanford.edu/shared-datasets

Statista. (2024). *Artificial intelligence (AI): Adoption, risks, and challenges*. Retrieved from https://www.statista.com/topics/10548/artificial-intelligence-ai-adoption-risks-and-challenges/#statisticChapter

Vention Teams. (2024). *AI adoption statistics by industry*. Retrieved from https://ventionteams.com/solutions/ai/adoption-statistics