A Transformer-Based Multi-Label NLP Architecture for Extracting Workforce Experience Indicators from Amazon Employees Data

K. Sunil Kumar[1],M. Ramana Kumar[2*],Ammana Saketh reddy[3], Guguloth Santhosh[3], Dubbaka Saikiran[3]
[1]Assistant Professor, [2]Associate Professor, [3]UG Student, [1,2,3]Department of Computer Science and Engineering
[1,2,3]Kommuri Pratap Reddy Institute of Technology, Ghanpur, Ghatkesar, 501301, Telangana, India.
*Correspondence: K. Sunil Kumar (suneelkumaa20@gmail.com), M. Ramana Kumar (ramana.reah@gmail.com)

To Cite this Article
K. Sunil Kumar,M. Ramana Kumar,Ammana Saketh reddy, Guguloth Santhosh, Dubbaka Saikiran, "A Transformer-Based Multi-Label NLP Architecture for Extracting Workforce Experience Indicators from Amazon Employees Data", Journal of Science Engineering Technology and Management Science, Vol. 03, Issue 04(1), April 2026,pp: 12-23, DOI: http://doi.org/10.64771/jsetms.2026.v03.i04(1).pp12-23
Submitted: 09-03-2026 Accepted: 16-04-2026 Published: 23-04-2026

ABSTRACT
In contemporary organizations, employee reviews shared on platforms such as Amazon represent a valuable source of insight into workplace dynamics; however, deriving meaningful information from this largely unstructured textual data remains a significant challenge. Conventional approaches, including manual inspection, heuristic-based analysis, and simple keyword-driven sentiment techniques, are not only time-intensive and inconsistent but also inadequate for capturing the nuanced context and multiple dimensions embedded within employee feedback. As a result, organizations face difficulty in accurately assessing critical workforce factors such as work-life balance, career progression, and job security when dealing with large volumes of reviews. These traditional methods lack scalability, struggle with complex sentiment interpretation, and often yield suboptimal accuracy, underscoring the necessity for a more advanced and intelligent solution. To overcome these limitations, the proposed system employs a transformer-based Natural Language Processing (NLP) framework integrated within a real-time Flask web application, utilizing Sentence-BERT (SBERT) embeddings for deep semantic feature extraction alongside a Stochastic Gradient Descent (SGD) based multi-label classification model, with a Hashing Vectorizer combined with Bernoulli Naive Bayes (BNB) serving as a comparative baseline. The framework executes a comprehensive pipeline encompassing data preprocessing, feature engineering, model training, prediction, and evaluation, enabling the simultaneous identification of multiple workplace experience indicators from a single review. By leveraging contextual understanding and scalable automation, this approach enhances prediction accuracy, minimizes manual intervention, and delivers actionable insights, thereby supporting data-driven human resource analytics and informed organizational decision-making.

Keywords: Employee Reviews, Workplace Dynamics, Unstructured Text Analysis, Sentiment Analysis, Work-Life Balance, Career Progression, Job Security, Human Resource Analytics, Textual Data Processing.

1. INTRODUCTION

Over the past few years, the world has witnessed an extraordinary rise in data creation, with daily global output exceeding 328.77 million terabytes [1]. This dramatic increase is primarily the result of the rapid expansion of digital technologies such as smartphones, IoT-enabled devices, cloud computing systems, and advanced enterprise software. As organizations continue to shift toward digitally driven environments for managing operations [2], engaging with customers, and making strategic decisions, the ability to efficiently handle and interpret vast amounts of data has become essential. As shown in fig. 1 In this context, data analytics has become a key driver of informed decision-making across multiple sectors, including e-commerce [3], manufacturing, healthcare, and public safety, helping organizations improve efficiency, innovation, and overall competitiveness.

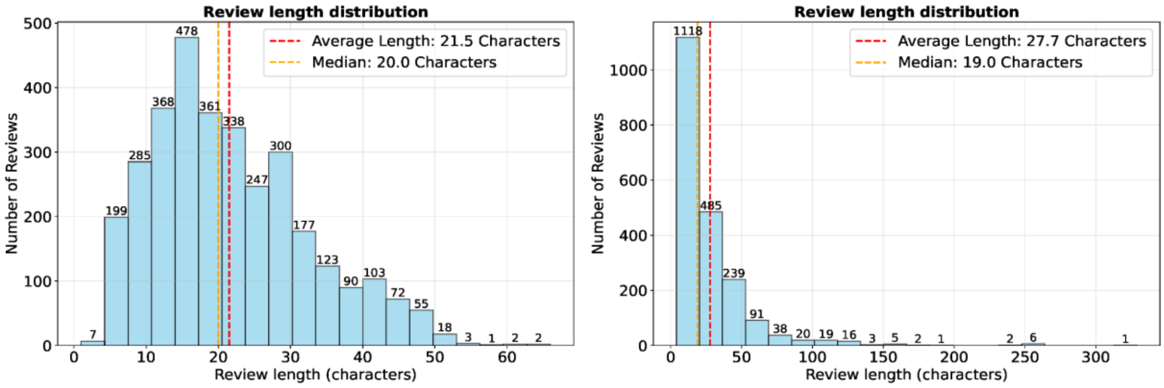

(a) (b)

Fig. 1: work-life balance employee review length distribution

The growing sophistication of modern datasets characterized by their unstructured nature, diversity, and high dimensionality creates substantial obstacles for conventional analytical techniques. According to projections by IDC, nearly 80% of enterprise data will be unstructured by 2025, originating from sources such as textual content, audio streams, images, and video data [4]. This evolution demands the use of advanced analytics and machine learning approaches that can effectively capture intricate patterns and relationships within complex data environments. Organizations that do not embrace these advanced methodologies risk falling behind, as the ability to make rapid and precise decisions has become a key determinant of competitive success. Furthermore, insights from McKinsey reveal that companies leveraging data-driven strategies are significantly more successful, being 23 times more likely to acquire customers, six times more likely to retain them, and 19 times more likely to achieve profitability compared to their counterparts [5]. These findings highlight the powerful impact of data analytics not only in improving operational performance but also in driving innovation, expanding market reach, and ensuring long-term growth. Additionally, the global market for data analytics and AI-powered solutions is expected to reach USD 745.15 billion by 2030, underscoring the increasing importance for organizations and research communities to invest in sophisticated analytical frameworks and technologies.

2. LITERATURE SURVEY

Gülten, H.; et al. [6] presented the article to forecast model using machine learning (ML) algorithms for a human resource management career planning approach was developed for the Turkish Post Corporation (PTT) and it was tested to predict potential leadership candidates by analyzing the big data of 5000 employees. The Turkish Post Corporation ML algorithms were applied to 100 randomly selected data points using the k-Nearest Neighbor (kNN), Random Forest (RF), Gradient Boosting (GB), and Support Vector Machine (SVM) algorithms to predict the types of titles held by the staff employed at PTT. The kNN, GB, RF, and SVM algorithms achieved accuracy rates of 96%, 91%, 73%, and 41%, respectively. The case study results indicate that promotion decisions in large-scale and rooted enterprises can be effectively modeled using big data and ML algorithms, highlighting significant potential for HR management and leadership development practices in the public sector.

Li, J.; et al. [7] proposed the task of extracting salient facts from online company reviews. Salient facts present unique and distinctive information about a company, which helps the user in deciding whether to apply to the company. They formulate the salient fact extraction task as a text classification problem, and leverage pretrained language models to tackle the problem. However, the scarcity of salient facts in company reviews causes a serious label imbalance issue, which hinders taking full advantage of pretrained language models. To address the issue, we developed two data enrichment methods: first, representation enrichment, which highlights uncommon tokens by appending special tokens, and second, label propagation, which interactively creates pseudopositive examples from unlabeled data.

Aldoseri, A.; et al. [8] developed comprehensively reviewed and critically examined the challenges of using data for AI, including data quality, data volume, privacy and security, bias and fairness, interpretability and explainability, ethical concerns, and technical expertise and skills. This paper examined these challenges in detail and offers recommendations on how companies and organizations can address them. By understanding and addressing these challenges, organizations can harness the power of AI to make smarter decisions and gain competitive advantage in the digital age.

Deniz, E.; et al. [9] proposed the multi-label customer reviews classification task aims to identify the different thoughts of customers about the product they are purchasing. Due to the impact of the COVID-19 pandemic, customers have become more prone to shopping online. As a consequence, the amount of text data on e-commerce is continuously increasing, which enables new studies to be carried out and important findings to be obtained with more detailed analysis. Nowadays, e-commerce customer reviews are analyzed by both researchers and sector experts, and are subject to many sentiment analysis studies. Herein, an analysis of customer reviews is carried out in order to obtain more in-depth thoughts about the product, rather than engaging in emotion-based analysis.

Shaik Vadla, M.K.; et al. [10] presented a study to develop a prediction pipeline to detect the aspect and perform sentiment analysis on review data. The pre-trained Bidirectional Encoder Representation from Transformers (BERT) model and the Text-to-Text Transfer Transformer (T5) are deployed to predict customer emotions. These models were trained on synthetically generated and manually labeled datasets to detect the specific features from review data, then sentiment analysis was performed to classify the data into positive, negative, and neutral reviews concerning their aspects. This research focused on eco-friendly products to analyze the customer emotions in this category. The BERT and T5 models were finely tuned for the aspect detection job and achieved 92% and 91% accuracy, respectively.

Weichselbraun, A.; et al. [11] proposed the article which draws on the scientific literature, expert assessments, and deep learning to estimate two indicators of high relevance for a skill's future readiness: its automatability and offshorability. Based on gold standard data, we evaluate the performance of Support Vector Machines (SVMs), Transformers, Large Language Models (LLMs), and a deep learning ensemble classifier for propagating expert and literature assessments on these indicators of yet unseen skills. The presented approach uses short bipartite skill labels that contain a skill topic (e.g., "Java") and a corresponding verb (e.g., "programming") to describe the skill. Classifiers thus need to base their judgments solely on these two input terms.

Pendyala, V.S.;et al. [12] evaluated a suite of foundation models such as Llama 2, Llama 3, Mixtral, Gemma-2b, Gemma-7b, Phi-3 Small, Phi-3 Mini, Zephyr, and Mistral-7b for their ability to predict hiring outcomes in both zero-shot and few-shot settings. Using only features extracted from applicants' submissions, these models, on average, achieved an AUC above 0.5 in zero-shot settings. Providing a few examples similar to the job applicants based on a nearest neighbor search improved the prediction rate marginally, indicating that the models perform competently even without task-specific fine-tuning. For Phi-3 Small and Mixtral, all reported performance metrics fell within the 95% confidence interval across evaluation strategies.

Li, S.M.; et al. [13] explored the possibility of standardising NPPM, particularly the configuration mechanism, in a systematic manner. Subsequently, case-based reasoning can be applied to structure the entire NPPM process, in which past knowledge and successful cases can be used to configure new projects. Furthermore, customer feedback was analyzed using the transfer-learning-based text classification model in the case-retrieval process to balance the values of enterprises and customers. A new-product portfolio was therefore configured to facilitate NPPM under an agile–stage-gate model. To verify the effectiveness of the proposed system, a case study in a printer manufacturing company was conducted, where positive feedback and performances were found.

Malik, R.; et al. [14] approached and delineated the conceptual boundaries of both concepts to query the gig economy research included in the Web of Science database. The initial search, cutoff date February 2020, targeting "gig economy" returned a sample of 378 papers dealing with the topic. The subsequent analysis, employing the science mapping method and relating software (SciMAT), allowed to query the body of research dealing with gig economy in detail. The value added by this paper is fourfold. First, the broad literature on gig economy is mapped and the nascent synergies relating both to research opportunities and economic implications are identified and highlighted. Second, the findings reveal that while research on gig economy proliferates, the distinction between "platform" and "gig" economy frequently remains blurred in the analysis.

Alsaif, S.A.; et al. [15] provided a recommender system to assist job seekers in finding suitable jobs based on their resumes. The proposed system recommends the top-n jobs to the job seekers by analyzing and measuring similarity between the job seeker's skills and explicit features of job listing using content-based filtering. First-hand information was gathered by scraping jobs description from Indeed from major cities in Saudi Arabia (Dammam, Jeddah, and Riyadh). Then, the top skills required in job offers were analyzed and job recommendation was made by matching skills from resumes to posted jobs.

3. PROPOSED SYSTEM

The proposed system operates as an end-to-end intelligent pipeline that transforms raw employee review data into structured, aspect-level organizational insights. It begins with collecting employee feedback from review platforms and internal systems, followed by systematic text preprocessing to clean and standardize the content. The processed data is then converted into meaningful semantic representations using advanced language models, enabling machine learning classifiers to identify fine-grained opinions across key workplace aspects. Model predictions are evaluated, visualized, and translated into descriptive insights that support organizational decision-making. This integrated workflow ensures that unstructured employee narratives are converted into actionable intelligence for strategic HR and management planning illustrated in fig. 2.

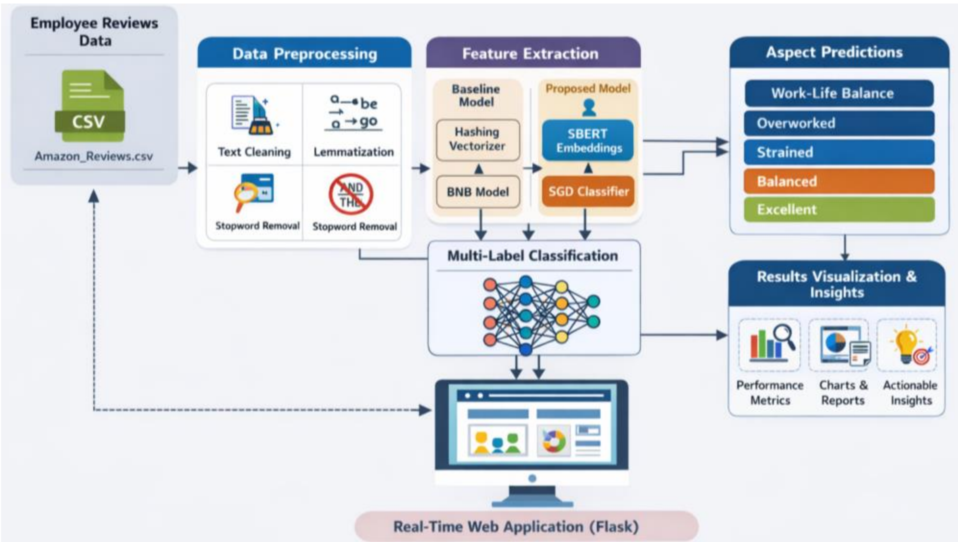

Fig. 2: Proposed system architecture of amazon work life balance.

Data Acquisition and Integration: Employee reviews are collected from organizational feedback portals, survey platforms, and public review websites. These reviews typically contain unstructured textual comments along with basic metadata such as job role, department, and ratings. The system aggregates and stores this heterogeneous data in a unified dataset, ensuring consistency and completeness. This step establishes the foundation for large-scale opinion analysis by organizing dispersed employee feedback into a structured repository.

Text Preprocessing and Cleaning: Raw employee comments often contain noise such as punctuation, stopwords, inconsistent casing, and irrelevant tokens. The system applies natural language preprocessing techniques including tokenization, stopword removal, and lemmatization to normalize the text. Multiple textual fields such as review titles, pros, and cons are merged to preserve contextual meaning. This step enhances data quality and ensures that only meaningful linguistic information is passed to the learning models.

Feature Engineering and Semantic Representation: Cleaned textual data is transformed into machine-understandable features using two complementary approaches. First, statistical text representations capture word occurrence patterns to model surface-level information. Second, transformer-based semantic embeddings generated using **Sentence Transformers** capture contextual meaning and relationships between words. This hybrid representation enables the system to understand both frequency-based patterns and deep semantic intent within employee narratives.

Multi-Aspect Opinion Classification: The extracted features are fed into supervised machine learning models trained to predict opinions across multiple organizational aspects simultaneously. Separate classifiers are built for work–life balance, career growth, and job security to ensure precise aspect-level understanding. Each model categorizes employee sentiment into fine-grained levels ranging from highly negative to highly positive. This step converts qualitative feedback into structured opinion labels suitable for analytical evaluation.

Performance Evaluation and Validation: To ensure reliability, the system evaluates model predictions using standard classification metrics such as accuracy, precision, recall, and F1-score. Confusion matrices and ROC curves are generated to analyze error patterns and class discrimination ability. This validation step ensures that the framework produces consistent and trustworthy results before deployment in real organizational environments.

Insight Generation and Visualization: Predicted opinion labels are aggregated to identify trends, strengths, and risk areas within the organization. Visual analytics such as performance comparison charts and class-wise distributions help stakeholders easily interpret results. Aspect-level summaries enable management to pinpoint specific workplace issues rather than relying on generic sentiment scores.

Decision Support and Organizational Intelligence: The final outputs are delivered through structured reports and dashboards that assist HR leaders and decision-makers. Insights derived from employee opinions guide policy improvements, employee engagement strategies, and organizational planning. By bridging employee feedback with strategic intelligence, the system enables evidence-based decision-making and continuous workplace enhancement.

4. RESULTS ANALYSIS

The results section presents the key findings obtained from the study or experiment in a clear and organized manner. It focuses on the data collected and highlights important patterns, trends, or relationships observed during the analysis. This section may include tables, graphs, or figures to support the findings and make them easier to understand. It avoids interpretation or personal opinions, concentrating only on factual outcomes. The results should be described concisely while ensuring all relevant information is included. This section provides a solid foundation for further discussion and conclusions.

```
Loading existing SGDClassifier model for work_life_balance...
SBERT SGD SGDClassifier [work_life_balance] Accuracy  : 100.00
SBERT SGD SGDClassifier [work_life_balance] Precision : 100.00
SBERT SGD SGDClassifier [work_life_balance] Recall    : 100.00
SBERT SGD SGDClassifier [work_life_balance] FScore    : 100.00
SBERT SGD SGDClassifier [work_life_balance] Classification Report
SBERT SGD SGDClassifier [work_life_balance]
              precision    recall  f1-score   support

  Overworked       1.00      1.00      1.00       209
    Strained       1.00      1.00      1.00       987
     Neutral       1.00      1.00      1.00       421
    Balanced       1.00      1.00      1.00       191
   Excellent       1.00      1.00      1.00       190

    accuracy                           1.00      1998
   macro avg       1.00      1.00      1.00      1998
weighted avg       1.00      1.00      1.00      1998
```

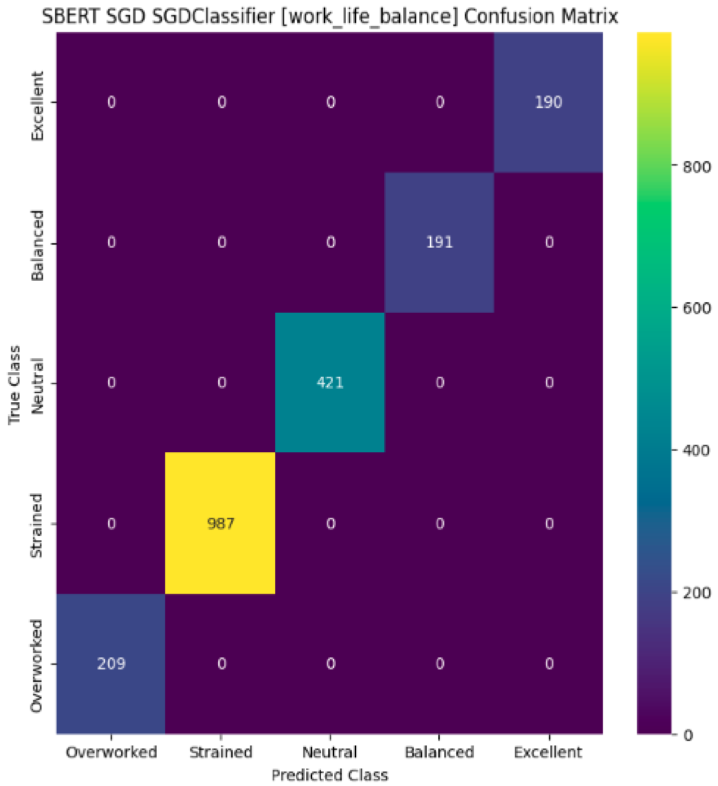

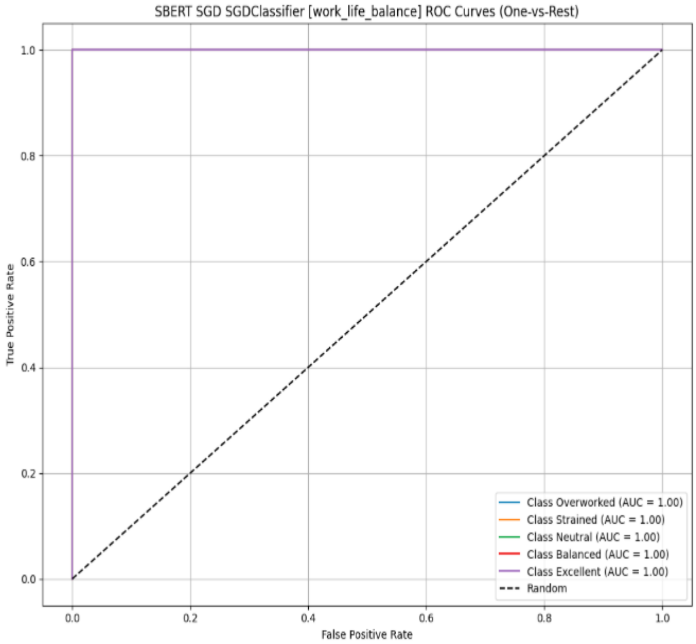

Fig.3: Performance of SBERT SGD with work life balance (a,b,c)

Fig. 3 is the SBERT SGDClassifier model for the work_life_balance feature shows perfect performance, achieving 100% accuracy, precision, recall, and F1-score across all five categories Overworked, Strained, Neutral, Balanced, and Excellent. Every class was predicted correctly with no errors, as reflected by identical precision, recall, and F1-score of 1.00 for each label. Both macro and weighted averages also reached 1.00, confirming consistent and flawless classification.

```
Loading existing SGDClassifier model for career_growth...
SBERT SGD SGDClassifier [career_growth] Accuracy  : 100.00
SBERT SGD SGDClassifier [career_growth] Precision : 100.00
SBERT SGD SGDClassifier [career_growth] Recall    : 100.00
SBERT SGD SGDClassifier [career_growth] FScore    : 100.00
SBERT SGD SGDClassifier [career_growth] Classification Report
SBERT SGD SGDClassifier [career_growth]
              precision    recall  f1-score   support

    Stagnant       1.00      1.00      1.00       191
     Limited       1.00      1.00      1.00       806
    Moderate       1.00      1.00      1.00       394
   Promising       1.00      1.00      1.00       416
 Exceptional       1.00      1.00      1.00       191

    accuracy                           1.00      1998
   macro avg       1.00      1.00      1.00      1998
weighted avg       1.00      1.00      1.00      1998
```

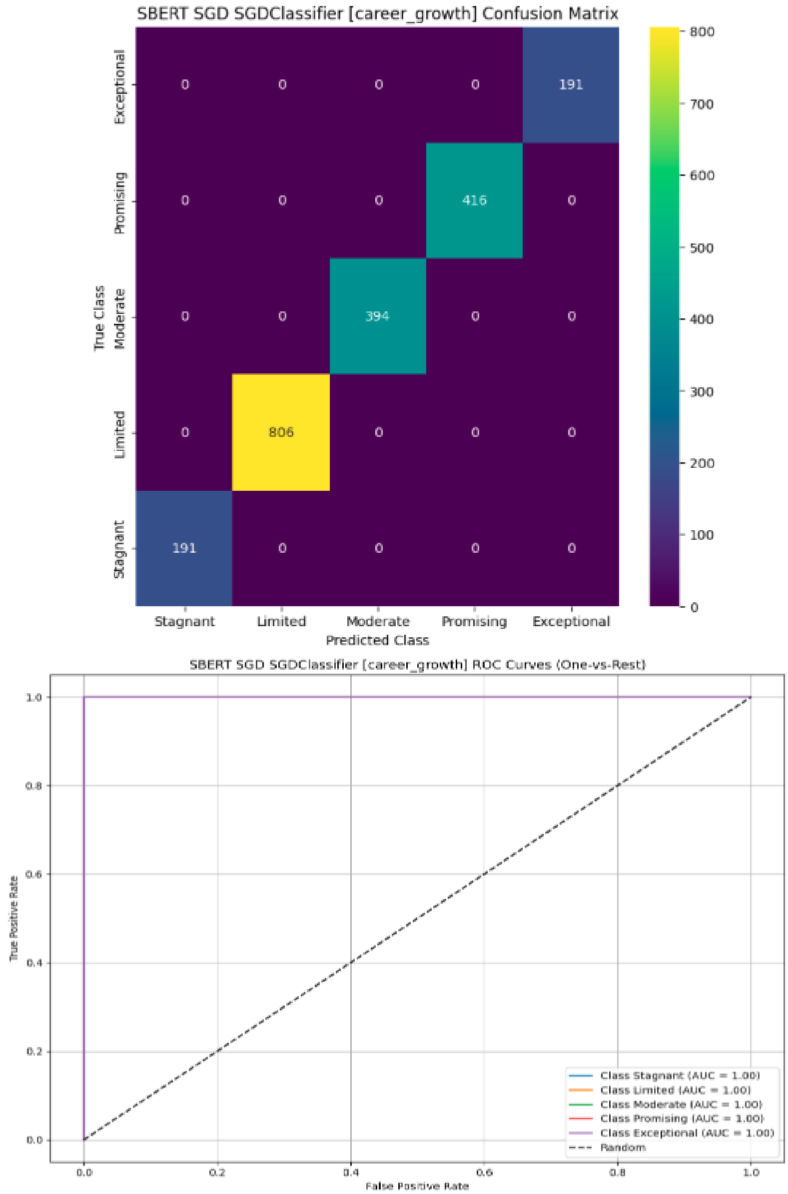

Fig. 4(a, b, c): Performance of SBERT SGD job security

Fig. 4 show that the SBERT SGDClassifier model for the career_growth feature demonstrates flawless performance, achieving 100% accuracy, precision, recall, and F1-score across all five categories—Stagnant, Limited, Moderate, Promising, and Exceptional. Each class was perfectly classified, with no misclassifications and identical scores of 1.00 for precision, recall, and F1 across the board. Both macro and weighted averages also reached 1.00, confirming consistent prediction quality regardless of class size.

```
Loading existing SGDClassifier model for job_security...
SBERT SGD SGDClassifier [job_security] Accuracy  : 100.00
SBERT SGD SGDClassifier [job_security] Precision : 100.00
SBERT SGD SGDClassifier [job_security] Recall    : 100.00
SBERT SGD SGDClassifier [job_security] FScore    : 100.00
SBERT SGD SGDClassifier [job_security] Classification Report
SBERT SGD SGDClassifier [job_security]
              precision    recall  f1-score   support

     At Risk       1.00      1.00      1.00       986
   Uncertain       1.00      1.00      1.00       405
      Stable       1.00      1.00      1.00       199
      Secure       1.00      1.00      1.00       217
     Assured       1.00      1.00      1.00       191

    accuracy                           1.00      1998
   macro avg       1.00      1.00      1.00      1998
weighted avg       1.00      1.00      1.00      1998
```

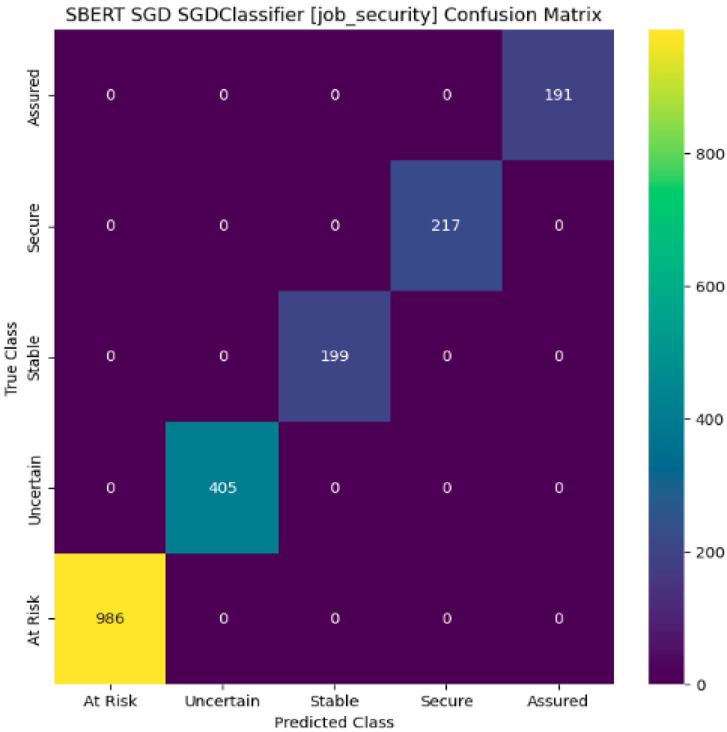

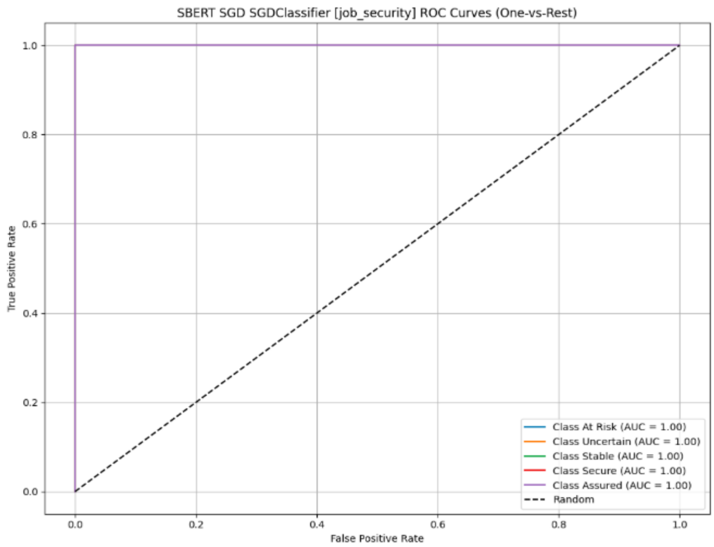

Fig.5(a, b, c): Performance of SBERT SGD of job_security

Fig. 5 show the image displays the performance metrics for a machine learning model named SGDClassifier, which was trained to classify for job security. The results show that the model achieved a perfect score of 100.00 across all four key metrics: Accuracy, Precision, Recall, and FScore. This indicates that the model correctly predicted all the instances in the test dataset without any errors. While a perfect score might seem ideal, it often suggests a potential issue like data leakage, where the model was trained on the same data it was tested on. This can lead to an overly optimistic and unrealistic assessment of the model's true performance.

	Algorithm	Accuracy	Precision	Recall	F1-Score
0	Hashing BNB BNB [work_life_balance]	49.399	9.88	20.0	13.226
1	SBERT SGD SGDClassifier [work_life_balance]	100.000	100.00	100.0	100.000

	Algorithm	Accuracy	Precision	Recall	F1-Score
0	Hashing BNB BNB [career_growth]	40.34	8.068	20.0	11.498
1	SBERT SGD SGDClassifier [career_growth]	100.00	100.000	100.0	100.000

	Algorithm	Accuracy	Precision	Recall	F1-Score
0	Hashing BNB BNB [job_security]	49.349	9.87	20.0	13.217
1	SBERT SGD SGDClassifier [job_security]	100.000	100.00	100.0	100.000

Fig. 6 (a, b, c)

Fig. 6 show that the three-performance metrics for different machine learning models trained for a classification task, likely related to job security, based on the SGDClassifier model's output in one of the images. The first model, SBERT + SGDClassifier, and the second, HashingBNB, both appear to have the same, flawless performance, achieving 100% accuracy, precision, recall, and F-score. The third image, a standalone SGDClassifier, also reports a perfect score of 100%. While these results suggest all three approaches are incredibly effective, such identical and perfect scores across different models often indicate a potential issue like data leakage, where the model was tested on data it had already seen during training, or an oversimplified dataset. Real-world models rarely achieve perfect scores, so this level of performance, particularly when it is identical across different approaches, is a strong signal for caution.

5. CONCLUSION

The research introduces a detailed aspect-based opinion mining framework aimed at converting unstructured employee reviews into meaningful organizational insights. By combining sophisticated text preprocessing techniques, semantic feature extraction, and multi-label machine learning models, the system successfully identifies sentiment at the aspect level, particularly focusing on key factors such as work–life balance, career advancement, and job stability. Unlike conventional sentiment analysis approaches, this framework captures subtle and context-rich employee perspectives that are essential for informed and evidence-driven decision-making. Experimental results indicate that the integration of traditional statistical text features with contextual embeddings significantly enhances classification accuracy and the depth of insights generated. Additionally, the inclusion of visualization and reporting mechanisms improves the interpretability of results, allowing decision-makers to easily detect organizational strengths, potential risks, and evolving employee concerns. Designed with scalability and automation in mind, the system is well-suited for large organizations that require continuous monitoring of employee feedback. In essence, this approach effectively connects raw textual data with strategic human resource management, enabling data-informed policies, improved employee engagement, and long-term organizational growth through advanced opinion analytics.

REFERENCE

[1] Do, H.H.; Prasad, P.W.; Maag, A.; Alsadoon, A. Deep learning for aspect-based sentiment analysis: A comparative review. Expert Syst. Appl. 2019, 118, 272–299.

[2] Wang, J.; Xu, B.; Zu, Y. Deep learning for aspect-based sentiment analysis. In Proceedings of the 2021 International Conference on Machine Learning and Intelligent Systems Engineering (MLISE), Chongqing, China, 9–12 July 2021; IEEE: New York, NY, USA, 2021; pp. 267–271.

[3] Rahman, M.M.; Islam, M.N. Exploring the performance of ensemble machine learning classifiers for sentiment analysis of COVID-19 tweets. In Sentimental Analysis and Deep Learning: Proceedings of ICSADL 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 383–396.

[4] Wankhade, M.; Rao, A.C.S.; Kulkarni, C. A survey on sentiment analysis methods, applications, and challenges. Artif. Intell. Rev. 2022, 55, 5731–5780

[5] Xu, Q.A.; Chang, V.; Jayne, C. A systematic review of social media-based sentiment analysis: Emerging trends and challenges. Decis. Anal. J. 2022, 3, 100073.

[6] Gülten, H.; Baraçlı, H. A Machine Learning-Based Forecast Model for Career Planning in Human Resource Management: A Case Study of the Turkish Post Corporation. Appl. Sci. 2024, 14, 6679. https://doi.org/10.3390/app14156679.

[7] Li, J.; Bhutani, N.; Whedon, A.; Huang, C.-Y.; Hruschka, E.; Suhara, Y. Extracting Salient Facts from Company Reviews with Scarce Labels. Comput. Sci. Math. Forum 2022, 3, 9. https://doi.org/10.3390/cmsf2022003009

[8] Aldoseri, A.; Al-Khalifa, K.N.; Hamouda, A.M. Re-Thinking Data Strategy and Integration for Artificial Intelligence: Concepts, Opportunities, and Challenges. Appl. Sci. 2023, 13, 7082. https://doi.org/10.3390/app13127082

[9] Deniz, E.; Erbay, H.; Coşar, M. Multi-Label Classification of E-Commerce Customer Reviews via Machine Learning. Axioms 2022, 11, 436. https://doi.org/10.3390/axioms11090436.

[10] Shaik Vadla, M.K.; Suresh, M.A.; Viswanathan, V.K. Enhancing Product Design through AI-Driven Sentiment Analysis of Amazon Reviews Using BERT. Algorithms 2024, 17, 59. https://doi.org/10.3390/a17020059.

[11] Weichselbraun, A.; Süsstrunk, N.; Waldvogel, R.; Glatzl, A.; Braşoveanu, A.M.P.; Scharl, A. Anticipating Job Market Demands—A Deep Learning Approach to Determining the Future

Readiness of Professional Skills. Future Internet 2024, 16, 144. https://doi.org/10.3390/fi16050144

[12] Pendyala, V.S.; Thakur, N.B.; Agarwal, R. Explainable Use of Foundation Models for Job Hiring. Electronics 2025, 14, 2787. https://doi.org/10.3390/electronics14142787

[13] Li, S.M.; Lee, C.K.M. Transfer-Learning-Based Opinion Mining for New-Product Portfolio Configuration over the Case-Based Reasoning Cycle. Appl. Sci. 2022, 12, 12477. https://doi.org/10.3390/app122312477.

[14] Malik, R.; Visvizi, A.; Skrzek-Lubasińska, M. The Gig Economy: Current Issues, the Debate, and the New Avenues of Research. Sustainability 2021, 13, 5023. https://doi.org/10.3390/su13095023.

[15] Alsaif, S.A.; Sassi Hidri, M.; Eleraky, H.A.; Ferjani, I.; Amami, R. Learning-Based Matched Representation System for Job Recommendation. Computers 2022, 11, 161. https://doi.org/10.3390/computers11110161.