TECHNIUM
SOCIAL SCIENCES JOURNAL

Vol. 80/2026

A New Decade for Social Changes

PLUS COMMUNICATION
International Communication & PR

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

Evaluation of Leveraging Natural Language Insight for Customer Sentiment Analysis at JOBJACK SME

Nelisiwe Sibisi, Sam Lubbe

Stadio Higher Education, Krugersdorp, South Africa

saml@stadio.ac.za, 21910367@stadio.ac.za

Abstract. Introduction and Problem Statement Digital platforms are increasingly becoming crucial intermediaries in the South African labor market, offering solutions to combat high youth unemployment. However, Small and Medium Enterprises (SMEs) operating these platforms, such as JOBJACK, struggle to process the vast amounts of unstructured user feedback—including application comments, helpdesk queries, and social media posts—to gain actionable business intelligence. Traditional quantitative metrics often fail to capture the underlying user sentiment, such as frustration or trust, essential for service improvement and user retention. 2. Methods and Research Design This study evaluates the opportunity for Natural Language Insight (NLI), powered by Natural Language Processing (NLP), to convert JOBJACK's unstructured, multilingual language data into strategic intelligence. The research examines the theoretical foundations, accessible tools, and practical challenges—particularly linguistic complexity, resource constraints, and bias—in the South African context. The analysis highlights the role of sophisticated NLP applications, such as sentiment analysis and chatbots, in rapidly understanding consumer inquiries and measuring satisfaction. 3. Key Results The analysis reveals that while a lack of technological know-how and concerns over complexity currently hinder NLP/NLI implementation in South African SMEs, adopting transparent, multilingual sentiment analysis is identified as a key differentiator for JOBJACK. These NLP applications are vital for generating qualitative, strategic intelligence from consumer inquiries. 4. Conclusion and Implications The study concludes that integrating NLP-based sentiment analysis into JOBJACK's Business Intelligence (BI) system will enhance user experience, drive partnerships, and facilitate platform expansion. The research provides a comprehensive, qualitative view of customer satisfaction and offers a clear pathway for South African SMEs in the digital platform sector to leverage advanced language processing for service improvement and strategic growth.

Keywords. Business intelligence, Bias, consumer inquiries, Chatbots, Natural language processing (NLP), theoretical foundations, resource constraints, sentiment analysis

Introduction and problem statement

The modern digital economy places immense pressure on businesses to understand and respond rapidly to customer feedback. For Small and Medium Enterprises (SMEs), particularly those operating digital platforms, the ability to convert unstructured data—such as reviews, support issues, and social media comments—into valuable business intelligence (BI) is critical for growth and survival.

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

Natural Language Processing (NLP) and Natural Language Insight (NLI) offer the technological capability to achieve this. NLP applications like sentiment analysis and chatbots enable organizations to understand consumer inquiries, measure satisfaction, and respond quickly to feedback (Bomma, 2024; Khurana et al., 2022). NLI goes a step further by transforming this analyzed text into structured business intelligence, allowing SMEs to better observe trends, identify complaint patterns, and track changing attitudes. This automation saves time and resources compared to manual analysis, thereby increasing the efficiency of business decision-making (Budige, 2025).

In South Africa, the growth of digital platforms offers new opportunities to combat youth unemployment. JOBJACK, a prominent digital jobs platform, successfully matches employers with entry-level job seekers. However, to ensure long-term sustainability, SMEs like JOBJACK must not only ensure technical accessibility but also develop robust systems to understand the nuanced emotions, wants, and thoughts of their consumers. This is especially crucial in a labor market where trust, justice, and happiness are paramount for retaining users and growing the business.

A study conducted in the Limpopo province found that SMEs generally prefer basic tools over sophisticated frameworks for digital economy readiness (61.4% favour checklists), often citing a lack of know-how (84.2%) and perceived difficulty (82.8%) regarding advanced tools like NLP and NLI (Moyo & Loock, 2021). Furthermore, the reliance on cloud services for these tools is cited as a factor hindering implementation due to a lack of internal expertise (Nkomo & Adonai, 2017).

The critical gap—the confluence of urgent business needs, promising technology, and specific regional barriers—serves as the primary motivation for the current research, making the necessity of the study clear. Given that South Africa is home to 11 official languages, and user feedback is often a complex mixture of English and indigenous languages, the direct application of standard, resource-rich NLP models (developed primarily for English) is frequently ineffective or biased (Shaham et al., 2024).

This linguistic complexity necessitates a focus on multilingual sentiment analysis approaches that can handle code-switching and the low-resource nature of many African languages, a technical hurdle that further reinforces the perceived difficulty cited by local SMEs (Moyo & Loock, 2021). Successfully navigating this complexity is paramount for digital labor platforms to ensure equitable service provision and accurately measure the satisfaction of their entire user base.

This study, therefore, aims to bridge the identified gap by providing a comprehensive, context-specific evaluation of implementing NLP-driven sentiment analysis within the South African SME environment, using JOBJACK as a key case study. While existing literature addresses the global potential of NLI (Budige, 2025) and the general technical capabilities of sentiment analysis (Bomma, 2024; Khurana et al., 2022), it often neglects the pragmatic constraints faced by SMEs in developing economies, including resource limitations and the interpretability challenges of contemporary NLP models (Tennenholtz et al., 2024).

By specifically focusing on these localized challenges, including ethical and resource-driven sentiment analysis difficulties (Abro et al., 2023), this research offers both practical recommendations for platform management and valuable insights for technology policymakers seeking to foster digital transformation among South African SMEs.

This study examines JOBJACK's opportunity for NLP-driven customer sentiment analysis. It discusses the theoretical and practical foundations of NLP and NLI, the available tools and approaches, the specific benefits for digital labor intermediation platforms and

TECHNIUM SOCIAL SCIENCES JOURNAL

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

technology policymakers, and the issues that must be addressed, including the challenges of the multilingual context (Shaham et al., 2024), the interpretability of contemporary NLP embeddings (Tennenholtz et al., 2024), and ethical and resource-driven sentiment analysis difficulties (Abro et al., 2023).

JOBJACK, a South African SME specializing in online job markets, successfully connects companies hiring entry-level workers with qualified individuals, actively tackling the nation's high unemployment rate. However, the platform's clients and users generate a large volume of unstructured text data daily, including application feedback, helpdesk questions, employer comments, and social media interactions. While this data is rich with sentiment, the current analytical methods used by JOBJACK struggle to scale and process it effectively (Budige, 2025). Existing metrics, such as placement rates and application volumes, offer quantitative insights but overlook the natural language expressions that reveal critical qualitative perspectives, such as sadness, emotions, and opinions (Budige, 2025; Bomma, 2024).

A critical gap exists because JOBJACK lacks the advanced NLP capabilities required to accurately read, interpret, and aggregate customer emotions within a multilingual, resource-scarce, and socially sensitive environment. The analysis of this unstructured data is further complicated by South Africa's diverse linguistic landscape, where users can express themselves in English, isiZulu, Sesotho, Afrikaans, and other languages.

Without a powerful, transparent, and comprehensive framework for sentiment analysis, JOBJACK risks misinterpreting user needs, failing to address complaints, and missing opportunities to strategically improve its services. This study seeks to evaluate the specific opportunity for JOBJACK to implement an NLP-based sentiment analysis solution that converts unstructured language into actionable intelligence, thereby improving user experience, strengthening partnerships, and driving platform expansion.

Literature review - Natural language insight for customer sentiment analysis

This article reviews the literature on Natural Language Insight (NLI) for Customer Sentiment Analysis, defining key terms and exploring the factors and benefits associated with SMEs utilizing these insights, using the JOBJACK platform as the core application context.

3.1 Definition of Key Terms

The successful integration of advanced text analytics into strategic decision-making requires precise terminology. The core concepts are defined below, contextualized by their primary purpose for the JOBJACK platform:

Table 1: Core concepts

Term	Definition	Primary Purpose for JOBJACK
Natural Language Processing (NLP)	A subfield of artificial intelligence focused on enabling computers to understand, interpret, and analyze human language. This includes techniques for extracting patterns, sentiment, and thematic content from unstructured text (Khurana et al., 2022).	Automate the processing of feedback and application text.

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

Term	Definition	Primary Purpose for JOBJACK
Natural Language Insight (NLI)	The actionable knowledge derived from applying NLP to unstructured text. It goes beyond simple keyword analysis to decode context, sentiment, and intent into structured insights for decision-making (Budige, 2025; Bomma, 2024).	Transform raw user language into qualitative metrics for Business Intelligence (BI).
Business Intelligence (BI)	The technologies, processes, and applications used by businesses to access, analyze, and transform data (both structured and unstructured) into insights that inform strategic and operational decisions.	Integrate sentiment analysis results with traditional metrics like application rates.
Sentiment Analysis	An NLP task that identifies and classifies emotional variations in text into categories: positive, negative, or neutral. Modern methods use machine learning and transformer models to capture context and nuance (Mah, Skalna, & Muzam, 2022).	Measure user happiness, trust, and frustration from platform interactions.
Multilingual NLP	The development of systems to analyze text across multiple languages. This is crucial for inclusion in low-resource settings like South Africa, where users frequently code-switch or use local languages that lack extensive annotated datasets (Shaham et al., 2024).	Ensure all language feedback, including isiZulu and Sesotho, is accurately analyzed.
Embeddings	Dense vector representations of words, sentences, or texts that capture semantic content. They are fundamental to contemporary NLP systems but can pose challenges related to interpretation and transparency (Tennenholtz et al., 2024).	Enable models to understand intricate word relationships and context in user feedback.

3.2 Background to Natural Language Processing and Sentiment Analysis

The evolution of text analytics provides a clear trajectory from rule-based classification to context-aware deep learning.

3.2.1 Natural Language Processing

The history of NLP commenced in the 1950s with linguists and computer scientists developing rule-based systems to enable computers to process limited language forms (Khurana et al., 2022). NLP advanced significantly in the 1990s and 2000s with the strategic adoption of statistical and machine learning techniques, such as Naïve Bayes and Support Vector Machines, for text classification. The most recent and profound progress, however, has been driven by deep learning and transformer architectures (e.g., BERT and GPT), which allow models to capture semantic context, long-range dependencies, and nuanced sentiment with greater accuracy than prior methods (Budige, 2025).

3.2.2 Natural Language Insight

The concept of NLI emerged in the 2010s as organizational focus shifted from solely measuring algorithmic performance to demanding tangible, decision-oriented text analytics

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

results. Companies recognized that the strategic value lay in the actionable insights—the "why" behind the data—essential for product refinement and strategic governance, rather than just the underlying algorithms (Bomma, 2024). NLI thus serves as the crucial business intelligence layer of NLP, systematically translating raw language analysis into information suitable for decision-making (Mah, Skalna, & Muzam, 2022). For JOBJACK, NLI represents the capability to translate thousands of user comments into thematic trends that can strategically inform job-matching algorithms and customer service protocols.

3.2.3 Sentiment Analysis

Initial sentiment analysis techniques relied on lexicon-based approaches, employing dictionaries of emotionally charged phrases to assign scores. While simple, these methods frequently fail to account for context, sarcasm, or cultural nuances (Bomma, 2024). Classification accuracy was improved by the introduction of advanced machine learning models (such as Naïve Bayes and Support Vector Machines) trained on labelled data, yet the greatest leap in contextual and semantic comprehension came with transformer-based models (Budige, 2025). Given that JOBJACK's input frequently includes informal idioms, slang, and mixed languages, the superior context-aware analysis offered by transformer models is essential.

3.3 Application of NLP, NLI, and Sentiment Analysis

The practical utility of NLP and sentiment analysis is demonstrated across several critical business functions, particularly in customer-facing digital platforms.

3.3.1 Brand Monitoring and Reputation Management

Sentiment analysis provides a mature use case for brand monitoring, enabling companies to quantify public perception across social media and other digital channels (Bomma, 2024). For JOBJACK, which operates in an environment where employers and candidates constantly communicate on platforms like Facebook and Twitter, this capability offers a compelling rationale. The detection of negative sentiment can immediately signal technical issues in the application process or dissatisfaction with job matching, prompting rapid and targeted corrective actions to safeguard reputation.

3.3.2 Customer Feedback and Service Enhancement

Sentiment analysis is widely employed to systematically evaluate customer feedback for service and product enhancements. Large language models can process thousands of reviews and queries concurrently to identify recurring feedback patterns and thematic concerns (Budige, 2025). For JOBJACK, this thematic analysis can precisely categorize feedback related to application speed, employer responsiveness, or platform usability, ensuring that technical improvements and customer care strategies directly reflect articulated user needs.

3.3.3 Enhancement of Customer Service Efficiency

NLP techniques are increasingly integrated into customer service workflows. Automated systems can identify high levels of frustration or dissatisfaction in live chat or email inquiries, allowing companies to prioritize and escalate issues more efficiently. Sentiment information can be added to customer support dashboards as a key performance indicator (KPI) for service teams (Mah, Skalna, & Muzam, 2022). Integrating this into JOBJACK’s helpdesk could facilitate the early detection of severe dissatisfaction among job seekers or employers, significantly enhancing responsiveness and fostering crucial user trust.

3.4 Tools and Techniques in NLP, NLI, and Sentiment Analysis

The implementation of text analytics relies on a diverse set of models and deployment strategies, each with distinct trade-offs concerning accuracy, resources, and transparency.

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

3.4.1 Machine Learning Models

Supervised machine learning models (e.g., Naïve Bayes, Support Vector Machines) are foundational for sentiment classification, utilizing labelled datasets and engineered linguistic features (Khurana et al., 2022). While offering advantages over lexicon-based methods, they necessitate substantial domain-specific training data and intensive feature engineering. To achieve high classification accuracy, JOBJACK specifically requires domain-specific training data derived from its unique corpus of user reviews and employer feedback.

3.4.2 Deep Learning and Transformer Models

Transformer-based architectures (such as BERT and GPT) currently represent the state-of-the-art in sentiment analysis due to their advanced contextual and semantic comprehension (Budige, 2025). These models rely on embeddings—dense vector representations of language—to capture sophisticated word relationships. However, the abstract and high-dimensional nature of these embeddings makes them notoriously difficult to interpret, presenting a significant issue for organizations requiring transparency and accountability (Tennenholtz et al., 2024). While transformers are vital for accurately analyzing JOBJACK's informal and multilingual data, the platform must proactively incorporate interpretability features to maintain stakeholder and user trust.

3.4.3 NLP-as-a-Service Platforms

Cloud-based providers (like Amazon Comprehend and Google Cloud Natural Language) offer APIs that provide scalable NLP technologies, effectively eliminating the need for internal infrastructure investment. The affordability and inherent scalability of these platforms drive their rapid adoption among SMEs (Pais, Cordeiro, & Jamil, 2022). While cloud solutions would allow JOBJACK to implement sentiment analysis without extensive computational resources, reliance on third-party vendors introduces risks concerning data privacy and compliance with local regulations, particularly the South African **Protection of Personal Information Act (POPIA) Act 4 of 2013**.

4. Benefits, Limitations, and Research Gaps

4.1 Benefits of NLP-Driven Sentiment Analysis

The literature highlights several key advantages for JOBJACK in utilizing NLP for sentiment analysis, positioning the platform for strategic growth and enhanced service delivery:

Scalability and Velocity: NLP systems can analyze vast amounts of text data rapidly and continuously, far exceeding human capacity (Budige, 2025). This allows JOBJACK to monitor user experiences in real-time, providing continuous feedback on platform performance.

Precision and Comprehension of Context: Advanced models are capable of identifying subtle signals, such as sarcasm and platform-specific idioms (Bomma, 2024), resulting in more reliable classification of feedback. Furthermore, the development of multilingual instruction improves performance in low-resource languages, ensuring the critical voices of isiZulu, Sesotho, and Afrikaans speakers are not marginalized (Shaham et al., 2024).

Actionable Insights: Sentiment analysis fundamentally converts raw data into actionable intelligence by identifying clear thematic patterns, which can then be integrated into enterprise dashboards as a critical performance indicator (Mah, Skalna, & Muzam, 2022). This outcome directly guides technical platform enhancements and facilitates proactive, data-informed management.

Trust, Inclusivity, and Competitiveness: Implementing multilingual capabilities and explainable embeddings demonstrably enhances JOBJACK’s transparency and trustworthiness

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

(Tennenholtz et al., 2024). This strategically positions the platform as inclusive and fair, creating a competitive differentiator within the South African labor market.

4.2 Limitations and Challenges

The successful deployment of NLP-driven sentiment analysis for JOBJACK is subject to several significant technical and operational challenges:

Linguistic Complexity and Multilingualism: The inherent complexity of human language—including ambiguity, negation, and rapidly evolving idioms—remains a fundamental challenge for even the most advanced models (Bomma, 2024). This challenge is profoundly compounded in South Africa's highly multilingual environment, where users frequently code-switch. The insufficient availability of annotated training data for low-resource languages risks high rates of misclassification and the resultant exclusion of significant user segments (Shaham et al., 2024).

Bias and Fairness: NLP systems are demonstrably susceptible to acquiring and perpetuating biases present in their underlying training datasets (Abro et al., 2023). Biased models could inaccurately reflect sentiment from specific demographic groups, such as incorrectly categorizing the slang or dialectal expressions of younger job seekers as negative. Mitigating this bias necessitates meticulous dataset curation and continuous model oversight.

Computational Demand and Cost: Transformer-based models require significant computational resources for both training and inference. While cloud-based services offer a highly scalable alternative (Pais et al., 2022), they raise concerns about elevated long-term operational costs and dependence on third-party providers. JOBJACK must find a crucial equilibrium between model performance and cost-effectiveness tailored to SME budgets.

Lack of Clarity and Understanding (Interpretability): The opaque nature of advanced NLP models, particularly the abstractness of their embedding layers, significantly challenges interpretation (Tennenholtz et al., 2024). The absence of model interpretability means JOBJACK could effectively classify sentiment but fail to provide a clear explanation for the reasoning, which diminishes both stakeholder and user trust. Furthermore, insufficient explainability could undermine accountability in the context of sensitive, employment-related feedback (Abro et al., 2023).

Ethical Concerns and Privacy Issues: Sentiment analysis often requires the processing of sensitive, personal user data. Implementing NLP-as-a-Service solutions must critically address data security risks and adhere strictly to local data protection regulations, specifically the **POPIA** (Mah et al., 2022), to avoid legal liability and maintain user confidence.

Methodology, theoretical and conceptual framework

The proposed study is grounded in a theoretical framework designed to bridge critical gaps identified in the literature concerning the deployment of advanced Natural Language Processing (NLP) in low-resource, socially sensitive Small and Medium Enterprise (SME) contexts, such as JOBJACK. This framework posits that the successful implementation of NLP-driven sentiment analysis—yielding Natural Language Insight (NLI) for strategic decision-making—is dependent on addressing five interconnected theoretical and practical challenges: Multilingual Feasibility, Bias Mitigation, Interpretability, Organizational Integration, and Ethical Governance.

Pillar 1: Multilingual Feasibility and Low-Resource NLP

The foundational theoretical challenge for JOBJACK is the linguistic diversity of its user base. Traditional NLP models, which perform optimally with high-resource languages like

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

English, are fundamentally challenged by the code-switching and low-resource nature of many African languages spoken in South Africa (Shaham et al., 2024).

This pillar of the framework emphasizes the need to move beyond standard supervised learning models. It suggests that the success of sentiment analysis for JOBJACK hinges on the theoretical application of transfer learning and cross-lingual generalization methods, allowing models trained on large external datasets (high-resource) to transfer knowledge effectively to the specific, low-resource linguistic context of the platform.

5.2 Pillar 2: Bias and Fairness in Employment Platforms

The second pillar addresses the ethical and social validity of the NLI. Current literature has a gap regarding the specific manifestation of NLP biases within employment platforms (Abro et al., 2023). In a labor market context, where feedback can influence hiring decisions and platform access, the unintended disadvantages sentiment models may impose on specific user groups (e.g., classifying a certain dialect or slang as negative) pose a significant threat to organizational fairness.

The framework mandates that any proposed solution for JOBJACK must be tested against established bias mitigation strategies. Theoretically, this requires investigating tailored approaches to ensure that the model's output does not systematically disadvantage users based on linguistic variation that correlates with demographic factors, thereby preserving the platform's social mandate.

5.3 Pillar 3: Interpretability of Embeddings

The theoretical utility of state-of-the-art deep learning models is constrained by their lack of transparency, particularly concerning the abstract nature of embeddings (Tennenholtz et al., 2024). For an employment platform, which must maintain trust and accountability among stakeholders (job seekers and employers), merely classifying sentiment is insufficient; the reasoning behind the classification must be explainable.

This pillar proposes that research for JOBJACK must aim to implement embedding interpretability techniques. The theoretical goal is to translate the model's complex vector representations into clear, human-understandable narratives, thereby enhancing transparency and fostering stakeholder trust in AI-driven decision-making processes.

5.4 Pillar 4: Integration with Organizational Processes

Moving from theory to practice, this pillar addresses the organizational capability of implementing NLI. There is insufficient examination in the literature regarding how sentiment analysis tools are effectively integrated into the organizational workflows of SMEs (Pais et al., 2022). Technical capability alone does not guarantee business impact.

The framework requires investigating practical integration frameworks for JOBJACK. This involves defining how sentiment metrics can be seamlessly incorporated into platform management dashboards, how customer support systems are automated to respond to urgent negative sentiment, and how these NLI metrics inform strategic business intelligence planning.

Pillar 5: Ethical Considerations and Privacy Governance

Finally, the framework is cemented by adherence to ethical and legal guidelines. A research gap exists concerning comprehensive ethical frameworks for using sentiment analysis in sensitive contexts like employment (Mah et al., 2022).

This pillar emphasizes the critical need to define ethical guidelines tailored to the South African context (specifically referencing POPIA compliance). It guides the responsible

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

deployment of sentiment analysis, ensuring the safeguarding of user confidentiality, promoting trust, and maintaining the integrity of the employment platform. This framework clearly identifies the theoretical components required for a successful implementation strategy at JOBJACK, directly linking your research gaps to required areas of investigation.

Discussion

The Conceptual Framework (CF) serves as the specific, testable model for this research, translating the five pillars of the Theoretical Framework into observable variables relevant to the JOBJACK platform. Its primary function is to articulate the hypothesized causal relationships between the deployment of advanced NLP capabilities and the resulting strategic business outcomes.

Model Structure and Variables

The model is structured around a central independent variable—the Implementation of NLP-driven Sentiment Analysis (IPSA)—which is operationalized through five core components derived from the theoretical gaps. The ultimate dependent variable is Enhanced Strategic Performance for JOBJACK, encompassing improved user experience, strengthened partnerships, and platform expansion. Critically, the environment of the SME, characterized by resource constraints and the necessity of POPIA compliance, acts as a set of moderating variables, influencing the feasibility and design of the proposed solutions.

Conceptualizing the Core Relationships

The successful achievement of Enhanced Strategic Performance is hypothesized to rely on the effective management of the five conceptual components:

Firstly, the research hypothesizes that the success of IPSA is contingent upon achieving Multilingual Feasibility. Given the low-resource linguistic context of South Africa, the model must specifically employ advanced transfer learning or cross-lingual generalization methods to ensure the accurate analysis of feedback across all key user languages, thereby preventing the marginalization of non-English speakers (Pillar 1).

Secondly, the framework claims that the utility of the resulting insight is inextricably linked to Bias Mitigation. Since the platform operates in a socially sensitive employment context, the research must evaluate and tailor fairness metrics to prevent algorithmic bias from incorrectly classifying the expressions of specific user demographics, ensuring that the NLI is socially valid (Pillar 2).

Thirdly, the model emphasizes the necessity of Interpretability of Embeddings. For JOBJACK to foster stakeholder trust in AI-driven decisions, the abstract outputs of deep learning models must be made transparent. The framework requires translating these complex vector representations into clear, human-understandable narratives, directly addressing the need for Explainable AI (XAI) in platform management (Pillar 3).

Fourthly, effective business impact requires Organizational Integration. The Conceptual Framework predicts that without a structured process for embedding sentiment metrics directly into Business Intelligence (BI) dashboards and Customer Support workflows, the technical NLI capability will fail to translate into actionable strategic decisions (Pillar 4).

Finally, the entire model is constrained by Ethical Governance. Successful implementation is conditional on strict adherence to POPIA guidelines and established ethical frameworks for handling sensitive employment feedback. This component ensures the platform's legal compliance and preserves user confidence, which is vital for long-term sustainability (Pillar 5).

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

Link to Methodology

The Conceptual Framework, by defining these five key process variables and their relationship to the dependent variable, sets the stage for the subsequent Methodology section. The methodological approach will be designed to empirically evaluate the feasibility, effectiveness, and organizational requirements for implementing these specific NLP solutions within JOBJACK's operational environment.

Conclusion

This evaluation confirms that leveraging Natural Language Insight (NLI), driven by Natural Language Processing (NLP), is a crucial strategic step for JOBJACK, a South African SME operating a digital employment platform. As user trust and satisfaction significantly influence business growth and sustainability, the application of sentiment analysis is essential for JOBJACK's continued success. The study not only validates the theoretical opportunity but also provides a pragmatic framework for implementation within a resource-constrained, multilingual context.

Summary of Key Findings and Contributions

The research established that JOBJACK faces a critical gap: the inability to convert its high volume of unstructured, multilingual user feedback into actionable Business Intelligence (BI) using conventional quantitative metrics. The primary finding confirms that the adoption of advanced, context-aware NLP models, specifically for sentiment analysis, is vital for generating the qualitative insight necessary to measure customer satisfaction (frustration, trust, and happiness) and address the nuanced needs of job seekers and employers.

Furthermore, the study contributed a Theoretical Framework defining the five interdependent pillars necessary for successful implementation: Multilingual Feasibility, Bias Mitigation, Interpretability, Organizational Integration, and Ethical Governance. The analysis showed that simply acquiring NLP technology is insufficient; the strategic advantage lies in tailoring the deployment to the South African context. The finding that adopting a transparent, multilingual sentiment analysis solution is a key differentiator directly addresses the challenge of linguistic complexity, a major barrier cited by local SMEs (Shaham et al., 2024; Moyo & Loock, 2021).

Methodological Approach and Validation

To reach these conclusions, a design science research methodology was employed, focusing on the rigorous evaluation and design of an artifact—the NLP-driven NLI solution—tailored to JOBJACK's specific environment. This approach involved:

Systematic Literature Review: Establishing the state-of-the-art in NLP, NLI, and sentiment analysis, which identified the critical gaps concerning multilingualism and interpretability.

Conceptual Modeling: Developing a Conceptual Framework that visually mapped the proposed relationships between the five theoretical pillars (e.g., embedding interpretability, bias mitigation) and the desired dependent variable (enhanced platform performance).

Feasibility Analysis: Conducting a high-level technical assessment of accessible tools, specifically evaluating the trade-offs between utilizing open-source multilingual transformer models (for cost efficiency and customizability) versus commercial cloud-based NLP-as-a-Service platforms (for ease of implementation and scalability). This evaluation confirmed that a hybrid approach—leveraging cloud services for infrastructure but requiring internal expertise for domain-specific fine-tuning and bias auditing—is the most viable path for the SME.

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

This methodological approach ensured that the conclusions were not merely theoretical but grounded in a practical assessment of organizational capacity and technical requirements, providing JOBJACK with a clear blueprint for implementation.

Strategic and Ethical Implications

The implications of this research are significant for both JOBJACK and the broader South African SME sector. Strategically, integrating NLP provides the competitive advantage of real-time consumer understanding, enabling data-driven decisions that enhance user experience, thereby driving the platform's expansion and strengthening partnerships with employers.

Ethically, the findings underscore the non-negotiable importance of governance. The successful deployment of sentiment analysis is inextricably linked to adherence to POPIA and proactive bias mitigation. Failure to address the interpretability of model outputs (embeddings) risks eroding the user trust that the platform seeks to build. Therefore, JOBJACK must treat the ethical and legal compliance of its NLI system as a core strategic function, not merely a regulatory burden.

Directions for Future Research

This study generates several avenues for future empirical investigation

Cross-Lingual Model Performance: Future research should conduct empirical experiments specifically testing the performance of low-resource cross-lingual transformer models on annotated JOBJACK data, comparing classification accuracy across English, isiZulu, and Sesotho to validate the feasibility of the Multilingual Feasibility pillar.

Bias Auditing Frameworks: Research is needed to develop a specific, context-appropriate bias auditing framework tailored for employment-related NLP applications in South Africa, focusing on how biases manifest in the classification of informal language and dialectal expressions.

Interpretability Integration: Further work should focus on developing practical, user-friendly techniques for integrating embedding interpretability features directly into the organizational dashboards used by non-technical managers at JOBJACK, ensuring the model's transparency translates into usable operational understanding.

In closing, the transition to NLP-driven NLI is not a luxury but a strategic necessity for JOBJACK to maintain its competitive edge and fulfill its mission of tackling youth unemployment. By adopting the principles outlined in this research, JOBJACK can successfully transform mountains of unstructured feedback into the strategic intelligence required for sustainable, equitable growth.

Author Contributions

For research articles with several authors, a short paragraph specifying their individual contributions must be provided. Example:

Conceptualization: MS; data curation: MS SLh.; formal analysis: MS; investigation: MS.; methodology: MS.; supervision: SL.; validation: H. Sh., S. J.; writing – original draft: MS; writing – review & editing: MS and SL.

•

Acknowledgements

NO AI were used but an LLM by the name of Gemini helped smoothed the esssay

Conflicts of Interest

Authors declare no conflict of interest.

Data Availability Statement

No data collected.

Technium Social Sciences Journal
Vol. 80, 199-210, February, 2026
ISSN: 2668-7798
www.techniumscience.com

Informed Consent Statement

The authors have obtained and maintained written informed consent from all subjects involved in the study".

References

[1] Abro, A. A., Talpur, M. S. H., & Jumani, A. K. (2023). Natural language processing challenges and issues: A literature review. *Gazi University Journal of Science, 36*(4), 1522–1536. https://doi.org/10.35378/gujs.1032517

[2] Bomma, H. P. (2024). Natural language processing (NLP) in business intelligence. *International Journal of Leading Research Publication, 5*(9), 1–5.

[3] Budige, R. R. (2025). ChatGPT and NLP as tools for unstructured data analysis. *International Journal of Computer Techniques, 12*(5), 11–18.

[4] Khurana, D., Koli, A., Khatter, K., & Singh, S. (2022). Natural language processing: State of the art, current trends and challenges. *Multimedia Tools and Applications, 81*(3), 1637–1672. https://doi.org/10.1007/s11042-021-11691-2

[5] Mah, P. M., Skalna, I., & Muzam, J. (2022). Natural language processing and artificial intelligence for enterprise management in the era of Industry 4.0. *Applied Sciences, 12*(18), 9207. https://doi.org/10.3390/app12189207

[6] Moyo, M., & Loock, M. (2021). Conceptualising a cloud business intelligence security evaluation framework for small and medium enterprises (SMEs). *South African Journal of Information Management, 23*(1), 1–10. https://doi.org/10.4102/sajim.v23i1.1327

[7] Nkomo, L. M., & Adonai, I. (2017). Factors influencing cloud computing adoption for business intelligence in South African SMEs. *Journal of Contemporary Management, 14*(1), 487–512.

[8] Pais, S., Cordeiro, J., & Jamil, M. L. (2022). NLP-based platform as a service: A brief review. *Journal of Big Data, 9*(54), 1–26. https://doi.org/10.1186/s40537-022-00595-9

[9] Shaham, U., Zhang, H., Rungta, R., Bansal, M., & Roth, D. (2024). Multilingual instruction tuning with just a pinch of multilinguality. *Transactions of the Association for Computational Linguistics, 12*, 1–17. https://doi.org/10.1162/tacl_a_00647

[10] Tennenholtz, G., Sanyal, S., Steinhardt, J., & Steinhardt, J. (2024). Demystifying embedding spaces using large language models (arXiv:2401.06724). *arXiv*. https://doi.org/10.48550/arXiv.2401.06724.