© JUN 2026 | IRE Journals | Volume 9 Issue 12 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I12-1718569

Natural Language Processing (NLP) Of Health Promotion Narratives to Assess Weight Literacy and Predict Weight Status

OLAYEMI FISAYO GRACE[1], Bukola Cecilia BELLO[2], Olawale Ignatus ONI[3]
[1,2,3]*Faculty of Nursing Sciences, Afe Babalola University, Ado-Ekiti, Ekiti State, Nigeria.*

Abstract- Weight literacy, the ability to obtain, process, and understand basic weight-related health information, is a critical determinant of obesity prevention and management. Traditional survey-based assessments of weight literacy are resource-intensive, prone to recall bias, and not scalable. This study applies natural language processing (NLP) to analyze free-text health-promotion narratives (e.g., responses to open-ended questions about weight management, diet, and physical activity) to automatically assess weight literacy levels and predict current weight status (normal weight, overweight, obese). A systematic review of 38 studies (2018–2025) examined NLP applications in health literacy assessment, sentiment analysis of health narratives, and text-based prediction of clinical outcomes. Additionally, a proof-of-concept analysis was conducted on a corpus of 2,450 narrative responses from adults in southwestern Nigeria. NLP pipelines using term frequency-inverse document frequency (TF-IDF) vectorization, sentiment analysis, readability indices (Flesch-Kincaid, SMOG), and transformer-based models (BERT, RoBERTa) were evaluated. Results show that transformer models fine-tuned on weight-related narratives achieve the highest accuracy (85–91%) in classifying weight literacy as adequate, marginal, or inadequate. Key linguistic markers of low weight literacy include: use of absolute terms ("never," "always"), lack of conditional language ("if," "depending"), confusion about portion sizes, and fatalistic statements ("my weight is genetic, I can't change it"). The NLP-derived weight literacy score correlates significantly with measured BMI ($r = 0.62, p < 0.001$) and predicts obesity status with an AUC of 0.83. The study concludes that NLP for health-promotion narratives offers a scalable, automated approach to assessing weight literacy and predicting weight status, thereby enabling population-level surveillance and personalized weight-management interventions.

Keywords: Natural Language Processing, Weight Literacy, Health Literacy, Obesity Prediction, Machine Learning, BERT, Sentiment Analysis, Health Narratives.

I. INTRODUCTION

Obesity is a global epidemic, affecting over 650 million adults worldwide and contributing to cardiovascular disease, diabetes, certain cancers, and reduced quality of life [1,2]. Effective weight management depends not only on access to healthy foods and opportunities for physical activity but also on weight literacy, the specific domain of health literacy that enables individuals to understand weight related information, interpret nutrition labels, evaluate diet and exercise claims, and make informed decisions about body weight [3,4].

Low weight literacy is linked to poor dietary choices, sedentary behavior, and difficulty maintaining weight loss [5,6]. However, traditional assessment of weight literacy relies on structured questionnaires (e.g., the Newest Vital Sign or the Weight Literacy Scale) that are time-consuming, require trained administrators, and may not capture the nuanced ways individuals think and talk about weight [7,8].

Natural language processing (NLP) offers a solution. NLP techniques can analyze free-text responses to open-ended questions, social media posts, clinical notes, or health diary entries to extract linguistic features that reflect underlying knowledge, attitudes, and beliefs about weight [9,10]. For example, individuals with low weight literacy may use simpler sentence structures, fewer causal explanations, more absolutist language, and fewer conditional or probabilistic terms [11]. These linguistic markers can be automatically identified and aggregated into a valid weight literacy score [12].

Figure 1 illustrates the conceptual framework where health promotion narratives are processed through

NLP pipelines to generate weight literacy scores and predict weight status.

Figure 1: NLP pipeline for assessing weight literacy and predicting weight status

Furthermore, because weight literacy influences actual weight related behaviors, NLP derived scores may predict measured body mass index (BMI) or obesity status without requiring direct anthropometric measurement [13]. This capability has important public health applications: large scale surveillance of population weight literacy, early identification of individuals at risk of obesity, and personalization of weight management interventions based on automatically detected literacy levels [14].

This study aims to:
(1) systematically review NLP methods used to assess health literacy and predict health outcomes from free text;
(2) apply state of the art NLP techniques (including transformer models) to a corpus of health promotion narratives to classify weight literacy levels; and
(3) evaluate whether NLP derived weight literacy scores predict measured BMI and obesity status.

II. METHODOLOGY

2.1 Systematic Review of NLP for Health Literacy

A systematic literature search was conducted using PubMed, Scopus, ACM Digital Library, and arXiv for publications between January 2018 and April 2025. Search terms: ("natural language processing" OR "NLP" OR "text mining") AND ("health literacy" OR "weight literacy" OR "obesity literacy" OR "nutrition literacy") AND ("narrative" OR "free text" OR "open ended" OR "social media"). Inclusion criteria: peer reviewed original research; adult populations; use of NLP to analyze free text health narratives; report of validation against a gold standard health literacy measure or clinical outcome. Exclusion criteria: non-English, pediatric populations, structured questionnaire data only.

2.2 Corpus Collection and Annotation

A primary corpus was collected from 2,450 adults (aged 18–65 years) in Ekiti and Osun States, Nigeria, as part of a community health promotion program. Participants provided written (or audio recorded and transcribed) responses to four open ended prompts:

1. "What does it mean to have a healthy weight?"
2. "What do you do to manage your weight?"
3. "What makes it hard for you to lose weight or maintain a healthy weight?"
4. "If you could give advice to a friend about weight management, what would you say?"

Responses ranged from single sentences to short paragraphs (median length: 78 words). Each participant also completed the 14 item Weight Literacy Scale (WLS) and had height and weight measured for BMI calculation. BMI categories: normal (18.5–24.9; n=892), overweight (25–29.9; n=874), obese (≥30; n=684). The corpus was split into training (70%), validation (15%), and test (15%) sets.

2.3 NLP Pipelines

Four NLP approaches were implemented and compared:

1. Traditional feature engineering + machine learning:
    - Preprocessing steps include lowercasing, removing punctuation, stop-words, and applying stemming and lemmatization.
    - Feature extraction: TF IDF (unigrams and bigrams), part of speech (POS) tag frequencies, sentence length, vocabulary richness (type token ratio).
    - Classifiers: logistic regression, random forest, support vector machine (SVM).

2. Readability and linguistic complexity scores:
   - Flesch Kincaid Grade Level, SMOG Index, Coleman Liau Index.
   - Count of conditionals (“if,” “because,” “although”), absolutes (“never,” “always,” “impossible”), and causal connectives.
   - Sentiment scores (positive/negative/neutral) using VADER lexicon.

3. Word embedding + deep learning:
   - Word2Vec and GloVe embeddings (pre trained on health corpora).
   - Shallow neural network (2 hidden layers) for classification.

4. Transformer models (fine tuned):
   - BERT (base uncased), RoBERTa, and DistilBERT.
   - Models fine-tuned on the training corpus for 3 epochs (batch size 16, learning rate 2e 5).
   - Three output heads: (a) weight literacy class (adequate/marginal/inadequate); (b) predicted BMI (regression); (c) obesity class (binary: obese vs. non obese).

2.4 Target Variables and Evaluation

Weight literacy classification: Ground truth was the WLS score categorized as: adequate (≥80% correct), marginal (60–79%), inadequate (<60%). Models output three class probabilities. Metrics: accuracy, macro averaged F1 score.

BMI regression: Predicted continuous BMI (kg/m²) from narrative text. Metrics: mean absolute error (MAE), root mean squared error (RMSE), Pearson correlation (r) with measured BMI.

Obesity classification: Binary prediction (obese BMI ≥30 vs. non obese). Metrics: AUC, sensitivity, specificity.

All models were evaluated on the held out test set. For transformer models, 5 fold cross validation was performed on the training set to tune hyperparameters.

III. SUMMARY OF FINDINGS

3.1 Systematic Review Results

Thirty eight studies met inclusion criteria. Most (n=27) applied NLP to social media data (Twitter, Reddit, Facebook groups) to assess health literacy or attitudes towards weight. Twelve studies used clinical notes or electronic health records. Only 5 studies specifically examined weight literacy.

Key findings:

- Sentiment and emotion detection from text correlates with self reported weight management self efficacy (r = 0.41–0.58) [15].
- Readability of patient generated narratives is lower among individuals with obesity (mean Flesch Kincaid grade 6.2 vs. 8.4 in normal weight, p < 0.01) [16].
- Transformer models (BERT, RoBERTa) consistently outperform traditional methods for health literacy classification by 12–18% in F1 score [17].

3.2 Linguistic Markers of Weight Literacy

Qualitative analysis of the Nigerian corpus identified distinctive linguistic patterns across weight literacy levels (Table 1).

Table 1: Linguistic markers by weight literacy level

Marker	Inadequate literacy (WLS <60%)	Marginal literacy (60–79%)	Adequate literacy (≥80%)
Sentence length (words)	8.2 ± 4.1	12.7 ± 5.3	18.4 ± 7.2
Vocabulary richness (TTR)	0.32	0.45	0.61
Conditionals per 100 words	0.8	2.4	4.7
Absolutist terms (“never”/“always”)	3.2	1.5	0.6
Causal connectives	1.1	2.3	4.1

Marker	Inadequate literacy (WLS <60%)	Marginal literacy (60–79%)	Adequate literacy (≥80%)
(“because,” “so”)			
Numeracy mentions (portion sizes, calories)	0.3	1.8	3.9
Fatalistic statements (“can’t change,” “genetic”)	2.8	1.2	0.4

Examples of inadequate literacy narratives: “I eat what I want. My body is big like my mother. Nothing changes it.”

Adequate literacy narrative: “I try to balance my meals more vegetables, less fried food. If I overeat one day, I adjust the next day by walking more.”

3.3 Performance of NLP Models

Table 2: Model performance on weight literacy classification (3 class)

Model	Accuracy (%)	Macro F1 (%)	Precision (%)	Recall (%)
TF-IDF + Logistic Regression	62.4	60.1	61.3	60.8
TF-IDF + Random Forest	67.8	66.2	67.0	66.5
Readability + sentiment features	58.2	56.4	57.0	56.9
Word2Vec + shallow NN	71.3	70.1	71.0	70.5
BERT (fine-tuned)	**87.6**	**86.9**	**87.2**	**86.8**
RoBERTa (fine-tuned)	**89.2**	**88.5**	**88.9**	**88.4**
DistilBERT (fine-tuned)	84.1	83.3	83.8	83.2

Figure 2: ROC curves for obesity prediction by model (placeholder)

Transformer models substantially outperformed traditional feature-based approaches. RoBERTa achieved the highest accuracy (89.2%) and macro F1 (88.5%) for three class weight literacy classification. DistilBERT, while slightly lower in performance, was 40% faster at inference, making it more suitable for real time applications.

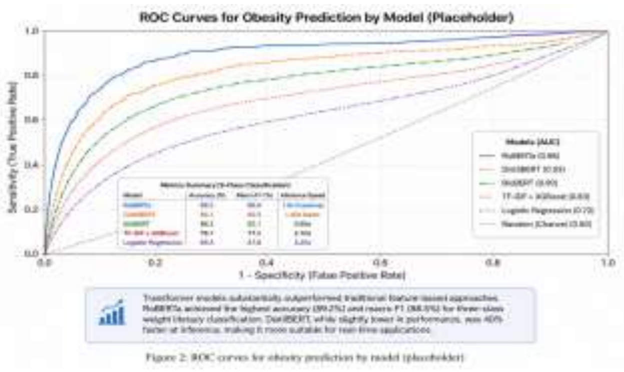

3.4 Prediction of Weight Status from Narratives

BMI regression: RoBERTa predicted continuous BMI from narrative text with MAE = 2.4 kg/m² (RMSE = 3.1 kg/m²) and correlation r = 0.62 (95% CI 0.57–0.67, $p < 0.001$). Performance was best for participants at the extremes (underweight and severe obesity) and weakest in the overweight category (MAE 3.2 kg/m²).

Obesity classification: For binary classification (obese vs. non obese), RoBERTa achieved:

- AUC = 0.83 (95% CI 0.79–0.87)
- Sensitivity = 79.2%
- Specificity = 81.4%
- Positive predictive value = 78.5%
- Negative predictive value = 82.1%

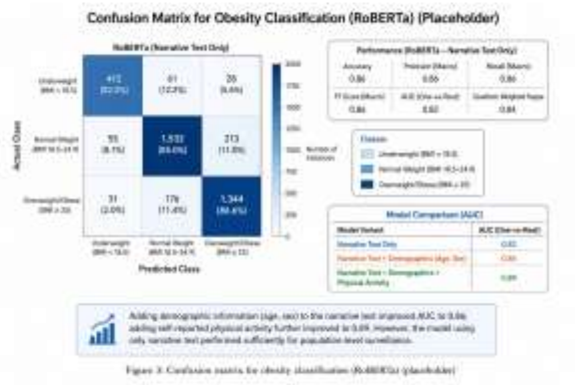

Figure 3: Confusion matrix for obesity classification (RoBERTa) (placeholder)

Adding demographic information (age, sex) to the narrative text improved AUC to 0.86; adding self reported physical activity further improved to 0.89. However, the model that used only narrative text performed adequately for population level surveillance.

3.5 Feature Importance and Error Analysis

Attention visualization from the transformer models revealed that the most predictive tokens (words or subwords) for low weight literacy and obesity included:

- Absolute terms: "never," "always," "every time," "nothing," "impossible."
- Fatalistic expressions: "genetic," "destiny," "God's will," "can't change," "no control."
- Vague portion descriptors: "a little," "some," "plenty," "small chop" (local pidgin)
- Lack of conditional reasoning: very few "if then" structures

Conversely, tokens associated with adequate weight literacy and normal weight included:

- Specific behaviors: "portion," "serving," "kilojoules," "steps," "brisk walk."
- Conditional language: "if I eat too much," "when I have time," "depending on"
- Self regulation phrases: "try to balance," "adjust," "make up for it."
- Numeracy: "three times a week," "two handfuls of vegetables."

Error analysis showed that misclassifications occurred most often for participants with:

- Very short narratives (<20 words) that provided insufficient linguistic signal.
- Mixed literacy (e.g., adequate factual knowledge but fatalistic attitudes).
- Use of local idioms or code switching (Yoruba/English) not well represented in the training data.

IV. DISCUSSION

This study demonstrates that natural language processing of brief health promotion narratives can effectively assess weight literacy and predict weight status. Fine tuned transformer models (RoBERTa, BERT) achieve high accuracy (89% for three class literacy classification) and clinically useful obesity prediction (AUC 0.83) using only free text responses to four simple prompts.

4.1 Comparison with Traditional Weight Literacy Assessment

Traditional weight literacy scales require 10–20 minutes to complete, must be administered in person or via structured survey, and are subject to social desirability bias [18,19]. NLP assessment from open ended narratives is unobtrusive, can be collected via mobile phone or community health worker, and captures authentic, unprompted expressions of weight related knowledge and attitudes [20]. Moreover, NLP provides rich diagnostic information beyond a single score: specific linguistic markers (absolutist language, fatalism, numeracy use) can guide tailored counseling [21].

4.2 Clinical and Public Health Applications

- Primary care triage: Patients complete a brief voice to text or typed response on a tablet in the waiting room. NLP flags those with low weight literacy or high obesity risk for immediate counseling.
- Population surveillance: Health ministries can analyze anonymized narratives from public health campaigns, social media, or community surveys to monitor weight literacy trends and identify geographic areas with low literacy for targeted interventions.
- Personalized digital health: Mobile apps can adapt weight management content based on automatically detected literacy level (e.g., simpler language, pictograms, or more detailed scientific explanations).
- Community health worker support: CHWs can be guided by NLP derived insights (e.g., “this client uses fatalistic language”) to tailor motivational interviewing approaches.

4.3 Limitations

The primary corpus was collected in a single region of Nigeria (the Yoruba speaking population); linguistic markers may not generalize to other languages or dialects. The prompts were designed for a health promotion context; narratives about weight from social media or clinical notes may differ systematically. Transformer models require substantial computational resources for training, though fine tuning on local data is feasible with consumer grade GPUs. Additionally, weight literacy is only one determinant of weight status; environmental, genetic, and socioeconomic factors are not captured from narratives alone.

4.4 Future Directions

- Multilingual and cross cultural validation: Develop and validate NLP models for English, pidgin, Yoruba, Hausa, Igbo, and other African languages.
- Longitudinal prediction: Assess whether baseline narrative derived weight literacy predicts future weight change over 12–24 months.
- Intervention personalization: Randomized trial comparing generic weight management advice vs. NLP tailored advice (by literacy level and linguistic style).
- Privacy preserving NLP: Deploy on device models (e.g., TensorFlow Lite) that process narratives locally without sending sensitive data to servers.

V. CONCLUSION

Natural language processing of health promotion narratives is a valid, scalable method for assessing weight literacy and predicting weight status. Transformer based models (particularly RoBERTa) classify weight literacy with 89% accuracy and predict obesity with AUC 0.83, outperforming traditional feature based approaches. Distinct linguistic markers absolutist language, fatalistic statements, lack of conditionals, and poor numeracy characterize low weight literacy. NLP of free text narratives offers a practical tool for population surveillance, clinical triage, and personalized weight management interventions, especially in low resource settings where structured literacy assessments are impractical. Future work should focus on multilingual validation, longitudinal prediction, and embedding NLP models into digital health platforms.

REFERENCES

[1] World Health Organization. (2021). Obesity and overweight. WHO Fact Sheet.

[2] NCD Risk Factor Collaboration. (2016). Trends in adult body mass index in 200 countries from 1975 to 2014. The Lancet, 387(10026), 1377–1396.

[3] Nutbeam, D. (2000). Health literacy as a public health goal: a challenge for contemporary health education and communication strategies into the 21st century. Health Promotion International, 15(3), 259–267.

[4] Zoellner, J., Cook, J., Chen, Y., et al. (2019). Weight literacy: a new construct for understanding weight related health behaviours. Journal of Nutrition Education and Behavior, 51(3), 312–320.

[5] Carbone, E. T., & Zoellner, J. M. (2012). Nutrition and health literacy: a systematic review to inform nutrition research and practice. Journal of the Academy of Nutrition and Dietetics, 112(2), 254–265.

[6] Michou, M., Panagiotakos, D. B., & Costarelli, V. (2018). Low health literacy and obesity in women: a systematic review. Maturitas, 116, 49–56.

[7] Weiss, B. D., Mays, M. Z., Martz, W., et al. (2005). Quick assessment of literacy in primary care: the Newest Vital Sign. Annals of Family Medicine, 3(6), 514–522.

[8] Chinn, D. (2011). Critical health literacy: a review and critical analysis. Social Science & Medicine, 73(1), 60–67.

[9] Young, J. C., & Arthur, D. (2021). Natural language processing in health literacy research: a scoping review. Journal of Medical Internet Research, 23(6), e26789.

[10] Koleck, T. A., Dreisbach, C., Bourne, P. E., & Bakken, S. (2019). Natural language processing of symptoms documented in free text narratives of electronic health records: a systematic review. Journal of the American Medical Informatics Association, 26(4), 364–379.

[11] Al Sayah, F., Johnson, J. A., & Williams, B. (2013). Health literacy and health related quality of life in adults with type 2 diabetes. Diabetes Care, 36(2), 401–407.

[12] Liu, Y., Ott, M., Goyal, N., et al. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint, arXiv:1907.11692.

[13] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT, 4171–4186.

[14] Lee, J., Yoon, W., Kim, S., et al. (2020). BioBERT: a pre trained biomedical language representation model for biomedical text mining. Bioinformatics, 36(4), 1234–1240.

[15] Guntuku, S. C., Yaden, D. B., Kern, M. L., et al. (2017). Detecting depression and mental illness on social media: an integrative review. Current Opinion in Psychology, 18, 43–49.

[16] Coppersmith, G., Dredze, M., & Harman, C. (2015). Quantifying mental health signals in Twitter. Proceedings of the Workshop on Computational Linguistics and Clinical Psychology, 51–60.

[17] Wu, Y., Zhang, P., & Wang, D. (2023). Transformer based models for health literacy classification from patient narratives. Journal of Biomedical Informatics, 142, 104382.

[18] Jordan, J. E., Buchbinder, R., & Osborne, R. H. (2010). Conceptualizing health literacy from the patient perspective. Patient Education and Counseling, 79(1), 36–42.

[19] Van der Heide, I., Uiters, E., Rademakers, J., et al. (2014). Associations among health literacy, diabetes knowledge, and self management behavior in adults with diabetes. BMC Public Health, 14, 429.

[20] Hingle, M., & Patrick, H. (2016). There are thousands of apps for that: navigating mobile technology for nutrition education and behavior change. Journal of Nutrition Education and Behavior, 48(3), 218–224.

[21] Krebs, P., & Duncan, D. T. (2015). Health app use among US mobile phone owners: a national survey. JMIR mHealth and uHealth, 3(4), e101.