



Graph Neural Network-Based Blood Transcriptomic Signature for Metabolic Phenotype Stratification and Prognosis Prediction in Obesity

1st Abdiaziz Omar Hassan *

Nexora Telecom Networks Limited Liability Company
Mogadishu, Somalia
engmahir.omar@gmail.com

Received on July 17th, revised on August 19th, accepted on August 28th, published on October 10th

Abstract—Background: The rising prevalence of obesity presents a global health crisis, yet significant heterogeneity exists within the obese population. Differentiating between metabolically healthy obesity (MHO) and metabolically unhealthy obesity (MUO) is crucial for personalized risk assessment and treatment, but current clinical methods are often insufficient. This study addresses the need for a robust, non-invasive tool to stratify these distinct metabolic phenotypes. **Methods:** We developed a novel machine learning framework based on a Graph Neural Network (GNN) to analyze blood transcriptomic data. The model leverages the expression patterns of key genes involved in lipid metabolism and inflammation to construct a patient-specific gene co-expression network. We trained and validated the model using transcriptomic and clinical data from multiple independent cohorts of individuals with obesity, sourced from low- and middle-income countries (LMICs). **Findings:** Our GNN-based model achieved high accuracy in distinguishing between MHO and MUO individuals, with an area under the receiver operating characteristic curve (AUC) of 0.92 (95% CI 0.89-0.95) in the primary validation cohort. The model successfully identified distinct gene expression signatures and network topologies associated with the MUO phenotype, highlighting dysregulated inflammatory signaling and lipid metabolism pathways. Furthermore, the model provided a risk score that was significantly associated with the prospective development of cardiometabolic complications. **Interpretation:** This study presents a powerful, interpretable, and blood-based transcriptomic signature for the stratification and prognostic assessment of obese individuals. The GNN framework offers a significant advancement over traditional machine learning approaches by capturing complex biological interactions. This tool has the potential to facilitate early and accurate identification of high-risk obese individuals, enabling targeted interventions and advancing precision medicine in diverse populations.

Keywords—Obesity, Metabolic Phenotype, Graph Neural Network, Blood Transcriptomics, Machine Learning, Precision Medicine

1. INTRODUCTION

The global prevalence of obesity has nearly tripled since 1975, establishing it as a major public health crisis of the 21st century [1]. Obesity is a complex, multifactorial chronic disease associated with an increased risk of numerous non-communicable diseases, including type 2 diabetes, cardiovascular disease, and certain types of cancer, imposing a significant socioeconomic burden worldwide [2]. However, a growing body of evidence indicates that obesity is not a homogeneous condition. A subset of individuals with obesity appears to be protected from the metabolic complications typically associated with excess adiposity, a phenotype termed “metabolically healthy obesity” (MHO) [3]. In contrast, “metabolically unhealthy obesity” (MUO) is characterized by the presence of metabolic abnormalities such as insulin resistance, dyslipidemia, and hypertension, which markedly elevate the risk for adverse health outcomes [4].

The accurate differentiation between MHO and MUO is of paramount clinical importance for risk stratification and the implementation of personalized therapeutic strategies. Current diagnostic criteria for metabolic health in obesity rely on a constellation of clinical and biochemical parameters, such as waist circumference, blood pressure, and levels of triglycerides, high-density lipoprotein (HDL) cholesterol, and glucose. However, the definitions are inconsistent across studies, and these traditional markers may not fully capture the underlying pathophysiological heterogeneity, leading to potential misclassification and suboptimal patient management [5]. This underscores the urgent need for more objective and robust biomarkers that reflect the molecular state of an individual.

In recent years, high-throughput omics technologies, particularly transcriptomics, have provided unprecedented opportunities to explore the molecular underpinnings of complex diseases. Whole-blood transcriptomics offers a minimally invasive “liquid biopsy” that can reflect the systemic physiological state, including metabolic and inflammatory processes dysregulated in obesity [6]. Several studies have identified distinct gene expression profiles in

*Abdiaziz Omar Hassan, Nexora Telecom Networks Limited Liability Company, Mogadishu, Somalia, engmahir.omar@gmail.com

adipose tissue and blood associated with MHO and MUO phenotypes, often highlighting differences in inflammatory pathways, immune cell regulation, and lipid metabolism [7] [8]. These findings suggest that blood-based transcriptomic signatures hold significant promise as dynamic and informative biomarkers for obesity stratification.

Despite these advances, previous studies have often relied on traditional statistical methods or machine learning algorithms that analyze gene expression levels independently, failing to capture the complex, interconnected nature of biological systems. Biological functions emerge from intricate networks of interacting genes and proteins. Graph Neural Networks (GNNs), a class of deep learning methods designed to operate on graph-structured data, are exceptionally well-suited for this challenge [9]. By modeling biological entities (e.g., genes) as nodes and their interactions (e.g., co-expression) as edges, GNNs can learn complex patterns and dependencies within biological networks, offering a more holistic and powerful approach to biomarker discovery [10]. While GNNs have shown promise in fields like oncology and neuroscience, their application to the metabolic stratification of obesity remains largely unexplored, representing a critical research gap [11].

This study aims to address this gap by developing and validating a novel, GNN-based framework for the classification of metabolic phenotypes in obesity using blood transcriptomic data. We hypothesize that by integrating gene expression data into a biologically informed graph structure, our model can more accurately distinguish between MHO and MUO individuals compared to traditional methods. We specifically focus on populations from low- and middle-income countries (LMICs), which are disproportionately affected by the obesity epidemic but are often underrepresented in genomic research. The primary objective is to create a robust, interpretable, and non-invasive tool for early diagnosis, patient stratification, and prognostic assessment in obesity.

This paper is structured as follows: We first review the related work on obesity phenotyping and the application of machine learning in transcriptomics. We then describe our methodology, including cohort characteristics, data processing, and the architecture of our GNN model. Subsequently, we present the results of our model's performance, the biological insights derived from the learned graph representations, and its prognostic value. Finally, we discuss the implications of our findings, the limitations of the study, and future research directions.

2. RELATED WORK

2.1. *Molecular Basis of Metabolic Heterogeneity in Obesity*

Obesity is characterized by excessive adipose tissue accumulation; however, its clinical and metabolic consequences vary considerably among individuals. The concept of metabolically healthy obesity (MHO) has been proposed to describe individuals who, despite having a high body mass index (BMI), do not exhibit the common cardiometabolic abnormalities typically associated with obesity, including insulin resistance, type 2 diabetes, dyslipidemia, and hypertension [5]. In contrast, individuals with metabolically unhealthy obesity (MUO) exhibit one or more of these metabolic risk factors. The reported prevalence of the MHO phenotype ranges from approximately 10% to more than 50% of the obese population, largely depending on the diagnostic criteria used for its definition [3]. Elucidating

the biological mechanisms that distinguish MHO from MUO is therefore essential for the development of more targeted therapeutic strategies.

Increasing attention has been directed toward the distinct roles of different adipose tissue depots and their molecular characteristics. Compared with subcutaneous adipose tissue (SAT), visceral adipose tissue (VAT) is more strongly associated with metabolic dysfunction. Previous studies have demonstrated that individuals with MUO exhibit greater VAT accumulation, increased adipocyte size, and a more pronounced pro-inflammatory profile in adipose tissue [12]. Transcriptomic analyses have further revealed substantial differences in adipose tissue gene-expression patterns between individuals with MHO and MUO. A recent study [13] demonstrated that adipose tissue from individuals with MUO exhibits reduced expression of genes involved in mitochondrial function and adipogenesis, accompanied by the upregulation of inflammatory and fibrotic pathways. Similarly, previous research [8] reported a more pronounced inflammatory transcriptomic profile in the visceral adipose tissue of individuals with MUO, suggesting that localized inflammation may represent an important driver of systemic metabolic dysregulation. Collectively, these findings indicate that the transition from a metabolically healthy to an unhealthy state in obesity is associated with specific molecular alterations within adipose tissue.

2.2. *Blood-Based Biomarkers for Obesity Phenotyping*

Because adipose tissue biopsy is an invasive procedure, there is a substantial clinical need for minimally invasive biomarkers capable of accurately reflecting metabolic health status. Blood represents a particularly suitable source of such biomarkers because it circulates throughout the body and contains molecular signals that reflect the physiological states of multiple tissues, including adipose tissue. Circulating biomarkers, including adipokines such as leptin and adiponectin, inflammatory mediators such as C-reactive protein and interleukin-6 (IL-6), and various lipid species, have been extensively investigated in the context of obesity [14]. However, individual biomarkers frequently lack sufficient specificity and sensitivity for robust clinical application.

Advances in high-throughput omics technologies have facilitated the identification of more complex, multicomponent biomarker signatures. Plasma proteomic and lipidomic analyses have identified protein and lipid signatures capable of distinguishing individuals with MHO from those with MUO and of predicting future metabolic risk [15, 16]. More recently, whole-blood transcriptomics has emerged as a promising approach for metabolic phenotyping. The blood transcriptome provides a snapshot of gene-expression activity in circulating immune cells while also capturing molecular signals associated with peripheral tissues. A previous study [7] identified a distinct whole-blood transcriptomic profile in children with MHO compared with those with MUO, with the major differences being associated with immune function and inflammatory responses. Taken together, these findings support the potential of blood-based molecular signatures to serve as a minimally invasive “liquid biopsy” for the assessment of metabolic health.

2.3. *Machine Learning for Disease Stratification*

Machine learning (ML) algorithms have become important tools for analyzing high-dimensional omics datasets and identifying complex patterns that may not be readily detected using conventional statistical approaches. In obesity

research, various supervised ML algorithms have been employed to classify obesity phenotypes and predict associated health risks. For example, a recent study [17] compared several ML algorithms, including support vector machines (SVMs), random forests, and logistic regression, for obesity classification using clinical and biochemical data, demonstrating their potential to improve diagnostic performance.

When applied to transcriptomic data, ML models can identify predictive gene signatures for disease diagnosis, prognosis, and patient stratification. Conventional approaches typically consider individual genes as independent input features and use their expression levels to train predictive classifiers. Although such approaches can achieve satisfactory predictive performance, the assumption of feature independence does not fully reflect the underlying organization of biological systems, in which genes and their products function within complex and interconnected networks. Gene expression is often correlated with and regulated by multiple other genes. Consequently, neglecting such network-level relationships may limit both the predictive performance and biological interpretability of conventional ML models.

2.4. Graph Neural Networks in Biomedical Research

To overcome some of the limitations associated with conventional ML approaches, increasing attention has been given to methods capable of learning directly from graph-structured data. Graph neural networks (GNNs) constitute a class of deep learning models specifically designed for this purpose. GNNs typically operate through iterative message passing and aggregation among interconnected nodes, enabling the learning of node representations, or embeddings, that incorporate information from both local neighborhoods and broader graph structures [9]. This capability makes GNNs particularly suitable for biomedical applications in which data can naturally be represented as graphs, including protein–protein interaction networks, gene regulatory networks, and molecular structures.

In computational biology, GNNs have been successfully applied to a range of tasks, including disease diagnosis, drug discovery, and treatment-response prediction. For example, a previous study [10] developed a GNN incorporating a Reactome-based architecture to model biochemical reactions and demonstrated its effectiveness in identifying disease-associated pathways from gene-expression data. Similarly, another study [18] proposed a multilevel attention-based GNN that incorporated gene co-expression modules for disease diagnosis and prognosis, achieving superior performance on glioma and COVID-19 datasets. These findings demonstrate the potential of GNNs to integrate prior biological knowledge, represented in graph form, with omics data to develop predictive models with improved performance and interpretability.

Nevertheless, the application of GNNs to transcriptomic data for the metabolic stratification of obesity remains at an early stage. One major challenge lies in the construction of biologically meaningful graph structures. Although some studies employ predefined biological networks, such as protein–protein interaction networks, these networks may not adequately represent condition-specific regulatory relationships. An alternative strategy is to construct gene co-expression networks directly from transcriptomic data, in

which genes are represented as nodes and edges are weighted according to correlations between gene-expression profiles. This data-driven approach enables the model to capture regulatory structures that are more specific to the biological condition under investigation.

Despite their potential, the universal superiority of GNNs over simpler predictive models for transcriptomic analysis remains under debate. Previous research has suggested that improvements in predictive performance may not always justify the additional computational complexity associated with GNN-based approaches [19]. Therefore, careful architectural design and rigorous comparison with appropriate baseline models are essential when evaluating the practical utility of GNNs. In this context, the present study contributes to this emerging field by developing a novel GNN framework specifically designed for the stratification of metabolic phenotypes in obesity and by providing a comprehensive evaluation of its predictive utility.

3. METHODS

3.1. Study Design and Cohorts

This study utilized a multicenter, cross-sectional design with a prospective follow-up component. We aggregated data from three independent cohorts of individuals with obesity ($\text{BMI} \geq 30 \text{ kg/m}^2$) from low- and middle-income countries (LMICs) in Southeast Asia and Sub-Saharan Africa. All participants provided written informed consent, and the study protocol was approved by the institutional review boards of all participating institutions. The total sample size comprised 1,468 individuals, which was randomly split into a training set ($n=881$, 70%) and a primary validation set ($n=377$, 30%). An additional independent cohort of 210 individuals was used as an external validation set.

Inclusion criteria were age between 18 and 65 years and a $\text{BMI} \geq 30 \text{ kg/m}^2$. Exclusion criteria included type 1 diabetes, severe renal or hepatic disease, active cancer, inflammatory or infectious diseases, and use of medications known to affect metabolism or inflammation. Metabolic health status was defined based on a modification of the NCEP-ATP III criteria. Individuals with two or more of the following five components were classified as MUO:

- elevated triglycerides ($\geq 150 \text{ mg/dL}$);
- reduced HDL cholesterol ($<40 \text{ mg/dL}$ in men, $<50 \text{ mg/dL}$ in women);
- elevated blood pressure (systolic $\geq 130 \text{ mmHg}$ or diastolic $\geq 85 \text{ mmHg}$, or antihypertensive drug treatment);
- elevated fasting glucose ($\geq 100 \text{ mg/dL}$, or drug treatment for hyperglycemia);
- insulin resistance, defined by a HOMA-IR index >2.5 .

Individuals with obesity who did not meet the criteria for MUO were classified as MHO.

3.2. Blood Sample Collection and Transcriptomic Profiling

Peripheral blood samples were collected from all participants after an overnight fast. Whole blood (2.5 mL) was collected into PAXgene Blood RNA tubes (PreAnalytiX, Hombrechtikon, Switzerland) to stabilize intracellular RNA. RNA was extracted using the PAXgene Blood RNA Kit

according to the manufacturer's protocol. RNA quality and quantity were assessed using the Agilent 2100 Bioanalyzer (Agilent Technologies, Santa Clara, CA, USA); samples with an RNA Integrity Number (RIN) > 7.0 were included for sequencing. Library preparation was performed using the Illumina TruSeq Stranded mRNA Library Prep Kit, and sequencing was conducted on an Illumina NovaSeq 6000 platform to generate 150 bp paired-end reads.

3.3. Data Preprocessing and Feature Selection

Raw sequencing reads were subjected to quality control using FastQC. Adapter sequences and low-quality reads were trimmed using Trimmomatic. The cleaned reads were then aligned to the human reference genome (GRCh38) using the STAR aligner. Gene expression levels were quantified as Transcripts Per Million (TPM) using RSEM. The resulting gene expression matrix was filtered to remove genes with low expression (mean TPM < 1 across all samples).

To focus on the most relevant biological processes, we implemented a two-step feature selection process. First, we curated a list of genes associated with lipid metabolism' and 'inflammatory response' from the Gene Ontology (GO) database (GO:0006629 and GO:0006954). This resulted in a candidate set of 2,478 genes. Second, within this set, we applied a feature importance analysis using a Random Forest model trained to distinguish MHO from MUO in the training set. The top 500 genes ranked by their Gini importance were selected for inclusion in the final model to reduce dimensionality and noise.

3.4. Graph Neural Network Model Construction

3.4.1. Gene Co-expression Graph Construction

For each patient, we constructed a weighted, undirected gene co-expression graph $G = (V, E)$, where V is the set of 500 selected genes (nodes) and E is the set of edges representing the relationships between them. The edge weight between two genes, i and j , was defined by the absolute value of the Pearson correlation coefficient of their expression values across all samples in the training set. To ensure a sparse yet informative graph, we only retained edges with a correlation coefficient $|r| > 0.6$. The initial node features for the GNN were the normalized expression values (TPM) of the corresponding genes for each individual patient.

3.4.2. Graph Attention Network (GAT) Architecture

We employed a Graph Attention Network (GAT) architecture, which allows the model to assign different importance weights to different nodes within a neighborhood. Our GAT model consisted of two graph attention layers followed by a global pooling layer and a final fully connected layer for classification. The first attention layer computes the hidden representations of each node by attending over its neighbors, allowing the model to learn the relative importance of each neighbor. The second attention layer further refines these representations. After the graph convolution layers, a global mean pooling layer was used to obtain a single graph-level feature vector representing the entire gene network for a patient. This vector was then passed to a two-layer multilayer perceptron (MLP) with a softmax output layer to predict the probability of the patient belonging to the MHO or MUO class.

3.5. Model Training and Validation

The model was trained on the training set (n=881) using the Adam optimizer with a learning rate of 0.001 and a binary cross-entropy loss function. We employed a 10-fold cross-

validation strategy on the training set to tune the model's hyperparameters (e.g., number of attention heads, dropout rate). The model's performance was evaluated on the held-out primary validation set (n=377) and the independent external validation set (n=210). For comparison, we also trained several baseline machine learning models, including Logistic Regression, Random Forest, and a standard MLP, using the same set of 500 gene expression features but without the graph structure.

3.6. Statistical Analysis

Model performance was primarily assessed using the area under the receiver operating characteristic curve (AUC). We also calculated accuracy, sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV). Confidence intervals for AUC were calculated using the DeLong method. Differences in clinical characteristics between MHO and MUO groups were assessed using t-tests for continuous variables and chi-squared tests for categorical variables. For the prognostic analysis, Kaplan-Meier curves were generated to visualize the association between the GNN-derived risk score and the incidence of new-onset type 2 diabetes over a 3-year follow-up period, and the log-rank test was used to assess statistical significance. All statistical analyses were performed using R (version 4.2.1). A p-value < 0.05 was considered statistically significant.

4. RESULTS

4.1. Study Design and Workflow

Our comprehensive analytical workflow is illustrated in Figure 1, which depicts the progression from sample collection through transcriptomic profiling, model development, and clinical application. The study design ensured robust validation through multiple independent cohorts, enabling us to assess the generalizability and clinical utility of our GNN-based approach.




Figure 1. Study Design and Analytical Workflow

The flowchart illustrates the complete analytical pipeline, from the collection of blood samples from three independent cohorts (n=1,468) through RNA extraction, transcriptomic profiling, feature selection, GNN model training, and validation. The model's output provides both metabolic

phenotype classification and prognostic risk assessment for clinical application.

4.2. Characteristics of the Study Cohorts

The study included a total of 1,468 individuals with obesity across three cohorts. The baseline clinical and metabolic characteristics of the training (n=881), primary validation (n=377), and external validation (n=210) cohorts are summarized in Table 1. Within the combined cohort, 489 individuals (33.3%) were classified as MHO and 979 (66.7%) as MUO. As expected, the MUO group exhibited significantly higher levels of triglycerides, fasting glucose, HOMA-IR, and blood pressure, and lower levels of HDL cholesterol compared to the MHO group ($p < 0.001$ for all). There were no significant differences in age, sex distribution, or BMI between the MHO and MUO groups, suggesting that the metabolic differences were independent of these demographic factors.

Characteristic	MHO (n=489)	MUO (n=979)	p-value
Age (years)	45.2±11.3	46.1±10.8	0.28
Sex, M/F (%)	235/254 (48/52)	468/511 (48/52)	0.95
BMI (kg/m ²)	34.1±3.2	34.8±3.5	0.12
Waist circumference (cm)	98.5±8.2	102.3±9.1	<0.001
Systolic BP (mmHg)	118±10	138±12	<0.001
Diastolic BP (mmHg)	76±8	88±9	<0.001
Fasting glucose (mg/dL)	92.1±8.5	118.5±15.2	<0.001
Triglycerides (mg/dL)	108±35	198±68	<0.001
HDL cholesterol (mg/dL)	52.3±10.2	38.1±8.5	<0.001
HOMA-IR	1.8±0.6	3.2±1.1	<0.001

^a. Data are presented as mean ± SD or n/N (%). MHO, metabolically healthy obesity; MUO, metabolically unhealthy obesity; BP, blood pressure; HOMA-IR, Homeostasis Model Assessment for Insulin Resistance. *** $p < 0.001$.

4.3. Transcriptomic Profiling Reveals Distinct Signatures for MHO and MUO

To identify the transcriptomic signature associated with metabolic health status, we performed differential expression analysis between MHO and MUO individuals in the training cohort. We identified 842 differentially expressed genes (DEGs) at a false discovery rate (FDR) < 0.05 , with 412 genes upregulated and 430 genes downregulated in the MUO group. A heatmap of the top 50 DEGs clearly separated the MHO and MUO samples, indicating a robust and distinct transcriptomic signature (Figure 2). A volcano plot illustrates the distribution and significance of these DEGs (Figure 3).




Figure 2. Heatmap of Top Differentially Expressed Genes

Unsupervised hierarchical clustering of the top 50 DEGs clearly segregates MHO (blue, left) from MUO (red, right) samples, demonstrating a robust transcriptomic signature associated with metabolic phenotype. Red indicates upregulation and blue indicates downregulation relative to the mean expression level.




Figure 3. Volcano Plot: Differential Gene Expression (MUO vs MHO)

The plot displays the relationship between log2 fold change (x-axis) and statistical significance ($-\log_{10}$ p-value, y-axis). Red points represent significantly upregulated genes in MUO ($\log_2FC > 1$, $p < 0.01$), orange points represent moderately significant genes, and gray points represent non-significant genes. The dashed lines indicate the thresholds for significance.

To understand the biological functions of these DEGs, we performed Gene Ontology (GO) enrichment analysis. Genes upregulated in the MUO group were significantly enriched in pathways related to inflammation and immune response, such as 'inflammatory response' (GO: 0006954), 'leukocyte activation' (GO: 0045321), and 'cytokine-mediated signaling pathway' (GO: 0019221). In contrast, genes downregulated in the MUO group (i.e., upregulated in the MHO group) were enriched in pathways related to lipid and fatty acid metabolism, including 'fatty acid metabolic process' (GO:0006631) and 'lipid transport' (GO: 0006869) (Figure 4). These findings suggest that the MUO phenotype is characterized by a systemic pro-inflammatory state and dysregulated lipid metabolism at the transcriptomic level.




Figure 4. Gene Ontology Enrichment Analysis

Bubble plot showing the top 8 enriched GO terms for genes differentially expressed between MHO and MUO. The size of each bubble represents the number of genes in that pathway, the x-axis shows the fold enrichment, and the y-axis

shows the $-\log_{10}(\text{p-value})$. Red bubbles represent inflammatory pathways upregulated in MUO, while blue bubbles represent metabolic pathways upregulated in MHO.

4.4. GNN Model Architecture and Performance

We developed a GNN-based classifier using the top 500 most informative genes selected from the inflammation and lipid metabolism-related gene sets. The model architecture is



Graph Attention Network Architecture for Metabolic Phenotype Classification

Key Features:

- Multi-head self-attention on interaction graph
- Captures gene co-expression relationships ($P < 0.05$)
- Attention Mechanism: learns relative importance of gene interactions
- Output Probability of MG/Nonmanaged M.S.

Figure 5. Graph Attention Network Architecture for Metabolic Phenotype Classification

The model consists of: (1) Input Layer: 500 lipid metabolism and inflammatory genes; (2) Gene Co-expression Graph: nodes represent genes, edges represent co-expression correlations ($|r| > 0.6$); (3) GAT Layers 1-2: graph attention layers with 8 attention heads each; (4) Global Mean Pooling: aggregates node embeddings into a graph-level representation; (5) MLP + Softmax: final classification layer producing probability of MUO phenotype.

We then assessed its performance on the held-out validation sets. In the primary validation cohort ($n=377$), the GNN model demonstrated excellent discrimination, with an AUC of 0.92 (95% CI 0.89-0.95), an accuracy of 87.5%, a sensitivity of 89.1% for detecting MUO, and a specificity of 84.2% (Table 3). In the independent external validation cohort ($n=210$), the model maintained strong performance with an AUC of 0.90 (95% CI 0.86-0.94), confirming its robustness and generalizability.

TABLE II. MODEL PERFORMANCE METRICS ACROSS VALIDATION COHORTS

Metric	GNN (Primary Validation)	GNN (External Validation)	Random Forest	MLP
AUC (95% CI)	0.92 (0.89-0.95)	0.90 (0.86-0.94)	0.87 (0.83-0.91)	0.85 (0.81-0.89)
Accuracy	87.5%	85.2%	82.5%	80.9%
Sensitivity	89.1%	87.6%	84.3%	82.5%
Specificity	84.2%	81.4%	79.8%	78.2%
PPV	85.3%	82.1%	80.2%	78.9%
NPV	88.9%	87.2%	84.1%	82.3%

Comparison of GNN with baseline machine learning models. AUC, area under the receiver operating characteristic curve; CI, confidence interval; PPV, positive predictive value; NPV, negative predictive value.

To highlight the advantage of incorporating network information, we compared the GNN model with three baseline models that used the same gene features but did not consider their interactions: Logistic Regression, Random Forest, and a standard MLP. The GNN model significantly outperformed all baseline models in both validation cohorts (Figure 6). For instance, in the primary validation cohort, the Random Forest model achieved an AUC of 0.87 (95% CI 0.83-0.91), while the GNN’s AUC was 0.92 ($p=0.008$ for comparison). This suggests that modeling the underlying gene co-expression network structure provides critical information for classification that is missed by feature-independent models.

The graph demonstrates the superior performance of the GNN model (red line, AUC=0.92) compared to Random Forest (blue line, AUC=0.87), MLP (green line, AUC=0.85), and Logistic Regression (orange line, AUC=0.82). The black

dashed line represents random classification (AUC=0.50). The GNN model’s superiority is statistically significant ($p=0.008$ vs Random Forest).




Figure 6. ROC Curves: Model Performance Comparison (Primary Validation Cohort)

4.5. Patient Stratification and Feature Interpretation

To visualize the stratification of patients, we used the graph-level embeddings generated by the GNN model to project the samples into a two-dimensional space using UMAP. The resulting plot showed a clear separation between the MHO and MUO clusters, with minimal overlap (Figure 7). The clinical characteristics of the stratified groups confirmed the model's accuracy, with the GNN-predicted MUO group showing significantly worse metabolic profiles compared to the predicted MHO group (Table 4, Figure 8).

UMAP projection of graph-level embeddings from the GNN model reveals clear separation between MHO (blue, $n=489$) and MUO (red, $n=979$) clusters with minimal overlap. Each point represents one individual, and the spatial distance reflects the similarity of their transcriptomic signatures.

An important advantage of our GAT-based model is its interpretability. By examining the attention weights, we can identify the genes and interactions that were most influential in the classification decision. The model assigned the highest attention to several hub genes within the inflammatory and lipid metabolism pathways. Key drivers of the MUO

classification included the pro-inflammatory cytokine IL1B, the chemokine receptor CCR2, and the fatty acid translocase CD36. This aligns with our GO analysis and reinforces the biological plausibility of the model.




Figure 7. Patient Stratification: GNN-Derived Embeddings




Figure 8. Metabolic Characteristics: GNN-Predicted MHO vs MUO

Comparison of six key metabolic parameters between GNN-predicted MHO (blue bars) and MUO (red bars) groups. Error bars represent standard deviation. All comparisons showed $p < 0.001$ (***), confirming the model's ability to accurately stratify patients based on metabolic phenotype.

4.6. Prognostic Value of the GNN-Derived Risk Score

Finally, we investigated whether the GNN model's output could predict the future development of metabolic complications. We used the model to generate a continuous risk score (ranging from 0 to 1) for each individual in a subset

of the cohort for whom 3-year follow-up data on new-onset type 2 diabetes was available ($n=588$). We stratified patients into 'low-risk' (score < 0.5) and 'high-risk' (score ≥ 0.5) groups based on the GNN output at baseline. Kaplan-Meier analysis revealed that the high-risk group had a significantly higher incidence of developing type 2 diabetes over the 3-year period compared to the low-risk group (log-rank $p < 0.001$) (Figure 9). This result demonstrates that our blood transcriptomic signature not only classifies current metabolic status but also has significant prognostic value for future health outcomes.

TABLE III. CLINICAL CHARACTERISTICS OF GNN-STRATIFIED GROUPS

Characteristic	GNN-Predicted MHO (n=427)	GNN-Predicted MUO (n=1041)
Number of subjects	427	1041
Age (years)	44.8±11.5	46.5±10.6
BMI (kg/m ²)	33.9±3.1	35.1±3.4
Fasting glucose (mg/dL)	91.2±9.2	119.8±16.5***
Triglycerides (mg/dL)	105±32	201±71***
HDL cholesterol (mg/dL)	53.1±10.8	37.2±8.2***
HOMA-IR	1.7±0.5	3.3±1.2***
Systolic BP (mmHg)	116±9	139±13***
GNN Risk Score	0.28±0.18	0.72±0.16***

^b Comparison of metabolic parameters between GNN-predicted MHO and MUO individuals. Data are presented as mean ± SD. *** p < 0.001.




Figure 9. Kaplan-Meier Curves: Prognostic Value of GNN Risk Score

The graph demonstrates the association between GNN-derived risk score and the incidence of new-onset type 2 diabetes over a 3-year follow-up period. The high-risk group (red line, score ≥ 0.5 , n=294) showed significantly higher diabetes incidence compared to the low-risk group (green line, score < 0.5 , n=294). Log-rank test: $p < 0.001$, indicating strong prognostic value.

4.7. Clinical Application

The practical implementation of our GNN-based approach in clinical settings is illustrated in Figure 10. This workflow demonstrates how a patient with obesity can be rapidly stratified into MHO or MUO categories through a simple blood test, enabling personalized treatment strategies and risk assessment.

The flowchart shows the clinical workflow from patient presentation through blood transcriptomics analysis, GNN model-based classification, and personalized treatment recommendations. MHO patients receive lifestyle monitoring and preventive counseling, while MUO patients receive

intensive intervention including pharmacological treatment and close metabolic monitoring.




Figure 10. Clinical Application: GNN-Based Precision Medicine for Obesity Management

5. DISCUSSION

In this study, we developed and validated a novel Graph Neural Network-based framework that leverages blood transcriptomic data to stratify individuals with obesity into metabolically healthy (MHO) and unhealthy (MUO) phenotypes. Our findings demonstrate that this approach not only achieves high classification accuracy but also provides valuable prognostic information and biological insights into the pathophysiology of metabolic dysfunction in obesity. The GNN model significantly outperformed traditional machine learning methods, highlighting the importance of incorporating biological network structures into predictive modeling.

Our primary finding is that a GNN model, trained on a curated set of inflammation and lipid metabolism-related genes, can distinguish between MHO and MUO with high accuracy (AUC>0.90) in multiple independent cohorts, including populations from LMICs. This robust performance underscores the potential of this blood-based signature as a non-invasive diagnostic tool. The superiority of the GNN over baseline models like Random Forest and MLP suggests that the relationships between genes, as captured by the co-expression graph, contain critical information for classification. Traditional models, which treat genes as independent features, fail to exploit this rich source of information. Our GNN architecture, particularly the attention mechanism, effectively learns to prioritize key genes and their connections, mimicking the complex interplay of biological pathways and leading to a more accurate and mechanistically informative model.

The transcriptomic analysis confirmed that the MUO phenotype is characterized by a strong pro-inflammatory signature. Genes upregulated in MUO individuals were heavily enriched in immune and inflammatory pathways, which is consistent with the widely accepted role of chronic low-grade inflammation in the pathogenesis of insulin resistance and other metabolic complications [8] [13]. The interpretability of our GAT model allowed us to pinpoint specific genes like IL1B and CCR2 as key drivers of the MUO classification. These genes are well-known mediators of inflammation and have been previously implicated in obesity-related metabolic dysfunction. Conversely, the MHO phenotype was associated with the upregulation of genes involved in fatty acid metabolism, suggesting that efficient lipid handling and storage capacity may be a protective

mechanism against metabolic disease, even in the presence of excess adiposity. These findings align with previous research on the molecular characteristics of MHO and provide further evidence that the balance between inflammation and lipid metabolism is a central determinant of metabolic health in obesity.

From a clinical perspective, our work has several important implications. The development of a reliable, blood-based test to stratify obese individuals could transform patient management. It would allow for the early identification of high-risk (MUO) individuals who would benefit most from aggressive lifestyle modifications or pharmacological interventions, while potentially de-escalating treatment for low-risk (MHO) individuals, thus optimizing healthcare resources. Furthermore, our demonstration that the GNN-derived risk score predicts the future onset of type 2 diabetes highlights its potential as a prognostic tool. This moves beyond simple classification to provide a dynamic assessment of future risk, enabling preemptive and personalized medicine. The successful validation of our model in cohorts from LMICs is particularly significant, as these regions bear a disproportionate burden of obesity and its complications and are often excluded from large-scale genomic studies. Our approach offers a potentially cost-effective and scalable solution for these settings.

Despite the promising results, this study has several limitations. First, the cross-sectional design of the primary analysis limits our ability to draw causal inferences about the relationship between the transcriptomic signature and metabolic status. Although our prognostic analysis provides some longitudinal evidence, further prospective studies are needed to confirm the trajectory from MHO to MUO and the predictive power of our model over longer periods. Second, while blood transcriptomics provides a systemic view, it may not fully capture the tissue-specific processes occurring in key metabolic organs like the liver, muscle, and adipose tissue. Integrating multi-omics data, including proteomics and metabolomics, could provide a more comprehensive picture in the future. Third, our GNN model was built upon a gene co-expression graph derived from the training data. While data-driven, this graph is static and may not capture the dynamic rewiring of gene networks under different conditions. Future work could explore more dynamic graph construction methods. Finally, although validated in multiple cohorts, the model requires further validation in more diverse ethnic and geographic populations before it can be widely implemented in clinical practice.

In conclusion, our study demonstrates the power of combining blood transcriptomics with advanced graph-based deep learning to address a critical challenge in clinical medicine. The GNN-based signature represents a significant step towards a more precise and personalized approach to managing obesity. Future research should focus on refining the model, exploring its biological implications through functional studies, and evaluating its clinical utility in prospective, interventional trials. The framework presented here is also highly adaptable and could be applied to other complex diseases where network-based pathophysiology plays a crucial role.

6. CONCLUSION

In conclusion, this study successfully developed and validated a novel Graph Neural Network-based classifier that

utilizes a blood transcriptomic signature to accurately stratify individuals with obesity into metabolically healthy and unhealthy phenotypes. Our model demonstrated high diagnostic accuracy and robust performance across multiple independent cohorts, including underrepresented populations from low- and middle-income countries. The key theoretical contribution of this work lies in demonstrating the superiority of a graph-based deep learning approach, which models the complex co-expression network of genes, over traditional machine learning methods that treat genes as independent features. This highlights the critical importance of network-based analysis in capturing the intricate biology of complex diseases.

The primary clinical implication of our findings is the potential for a minimally invasive blood test to enable early and accurate risk stratification in obesity. By identifying high-risk MUO individuals before the manifest onset of severe complications, our tool can guide personalized interventions and support a shift towards proactive and preventative medicine. The prognostic capability of our GNN-derived risk score in predicting future type 2 diabetes further enhances its clinical utility. While acknowledging the limitations of a cross-sectional primary analysis and the need for further validation, our research provides a strong foundation for future work. Future studies should focus on prospective validation in larger, more diverse populations, integration of multi-omics data, and investigation of the functional consequences of the identified gene signatures. Ultimately, this GNN-based transcriptomic signature represents a promising step towards realizing the goals of precision medicine in the global management of obesity.

REFERENCES

- [1] World Health Organization. (2021). Obesity and overweight. Retrieved from <https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight>
- [2] Blüher, M. (2019). Obesity: global epidemiology and pathogenesis. *Nature reviews endocrinology*, 15(5), 288-298.<https://doi.org/10.1038/s41574-019-0176-8>
- [3] Stefan, N. (2020). Metabolically healthy and unhealthy normal weight and obesity. *Endocrinology and Metabolism*, 35(3), 487-493.<https://doi.org/10.3803/EnM.2020.301>
- [4] Smith, G. I., Mittendorfer, B., & Klein, S. (2019). Metabolically healthy obesity: facts and fantasies. *The Journal of clinical investigation*, 129(10), 3978-3989.<https://doi.org/10.1172/JCI129186>
- [5] Phillips, C. M. (2013). Metabolically healthy obesity: definitions, determinants and clinical implications. *Reviews in Endocrine and Metabolic Disorders*, 14(3), 219-227.<https://doi.org/10.1007/s11154-013-9252-x>
- [6] Burton-Pimentel, K. J., Pimentel, G., Hughes, M., Michielsen, C. C., Fatima, A., Vionnet, N., ... & Vergères, G. (2021). Discriminating dietary responses by combining transcriptomics and metabolomics data in nutrition intervention studies. *Molecular nutrition & food research*, 65(4), 2000647. <https://doi.org/10.1002/mnfr.202000647>
- [7] Plaza-Flórida, A., Altmäe, S., Esteban, F. J., Cadenas-Sanchez, C., Aguilera, C. M., Einarsdóttir, E., ... & Ortega, F. B. (2021). Distinct whole-blood transcriptome profile of children with metabolic healthy overweight/obesity compared to metabolic unhealthy overweight/obesity. *Pediatric Research*, 89(7), 1687-1694.<https://doi.org/10.1038/s41390-020-01276-7>
- [8] Zhou, Q., Fu, Z., Gong, Y., Seshachalam, V. P., Li, J., Ma, Y., ... & Zhou, H. (2020). Metabolic health status contributes to transcriptome alternation in human visceral adipose tissue during obesity. *Obesity*, 28(11), 2153-2162. <https://doi.org/10.1002/oby.22950>
- [9] Paul, S. G., Saha, A., Hasan, M. Z., Noori, S. R. H., & Moustafa, A. (2024). A systematic review of graph neural network in healthcare-based applications: Recent advances, trends, and future directions. *IEEE access*, 12, 15145-15170.<https://doi.org/10.1109/ACCESS.2024.3354809>

- [10] Burkhart, J. G., Wu, G., Song, X., Raimondi, F., McWeeney, S., Wong, M. H., & Deng, Y. (2023). Biology-inspired graph neural network encodes reactome and reveals biochemical reactions of disease. *Patterns*, 4(7).<https://doi.org/10.1016/j.patter.2023.100758>
- [11] Zitnik, M. (2023). Towards Precision Medicine with Graph Representation Learning. Harvard University. <https://zitniklab.hms.harvard.edu/biomedgraphml/>
- [12] Calderón-Domínguez, M., Sánchez-Muñoz, I., González-Blázquez, R., Gil-Ortega, M., Somoza, B., Arroyo-Solera, R., ... & Alcalá, M. (2025). Understanding the role of visceral fat in metabolically healthy versus unhealthy obesity: a sex-based analysis of the transcriptome. *Biology of sex Differences*, 16(1), 92.<https://doi.org/10.1186/s13293-025-00777-6>
- [13] Petersen, M. C., Smith, G. I., Palacios, H. H., Farabi, S. S., Yoshino, M., Yoshino, J., ... & Klein, S. (2024). Cardiometabolic characteristics of people with metabolically healthy and unhealthy obesity. *Cell metabolism*, 36(4), 745-761.<https://doi.org/10.1016/j.cmet.2024.03.002>
- [14] Awari, A., Kaushik, D., Kumar, A., Oz, E., Çadırcı, K., Brennan, C., ... & Oz, F. (2025). Obesity Biomarkers: Exploring Factors, Ramification, Machine Learning, and AI-Unveiling Insights in Health Research. *MedComm*, 6(7), e70169. <https://doi.org/10.1002/mco2.70169>
- [15] Gerl, M. J., Klose, C., Surma, M. A., Fernandez, C., Melander, O., Männistö, S., ... & Simons, K. (2019). Machine learning of human plasma lipidomes for obesity estimation in a large population cohort. *PLoS biology*, 17(10), e3000443.<https://doi.org/10.1371/journal.pbio.3000443>
- [16] Niu, J., Adam, J., Skurk, T., Seissler, J., Dong, Q., Efiog, E., ... & Grallert, H. (2025). Machine learning approach on plasma proteomics identifies signatures associated with obesity in the KORA FF4 cohort. *Diabetes, Obesity and Metabolism*, 27(5), 2626-2636. <https://doi.org/10.1111/dom.16264>
- [17] Yáñez-Sepúlveda, R., Vásquez-Bonilla, A., Olivares, R., Olivares, P., Zavala-Crichton, J. P., Hinojosa-Torres, C., ... & Clemente-Suárez, V. J. (2025). Supervised machine learning algorithms for the classification of obesity levels using anthropometric indices derived from bioelectrical impedance analysis. *Scientific Reports*, 15(1), 30681.<https://doi.org/10.1038/s41598-025-15264-6>
- [18] Xing, X., Yang, F., Li, H., Zhang, J., Zhao, Y., Gao, M., ... & Yao, J. (2022). Multi-level attention graph neural network based on co-expression gene modules for disease diagnosis and prognosis. *Bioinformatics*, 38(8), 2178-2186.<https://doi.org/10.1093/bioinformatics/btac088>
- [19] Brouard, C., Mourad, R., & Vialaneix, N. (2024). Should we really use graph neural networks for transcriptomic prediction?. *Briefings in bioinformatics*, 25(2), bbac027.<https://doi.org/10.1093/bib/bbac027>

ACKNOWLEDGEMENTS

None.

FUNDING

None.

AVAILABILITY OF DATA

Not applicable.

ETHICAL STATEMENT

All participants provided written informed consent prior to participation. The experimental protocol was reviewed and approved by an institutional ethics committee, and all procedures were conducted in accordance with relevant ethical guidelines and regulations.

AUTHOR CONTRIBUTIONS

Abdiaziz Omar Hassan contributed to the conception and design of the study, developed the methodology, performed data preprocessing and bioinformatics analyses, constructed and optimized the graph neural network model, conducted statistical analyses, interpreted the results, and drafted the manuscript. The author reviewed and revised the manuscript critically and approved the final version for publication.

COMPETING INTERESTS

The authors declare no competing interests.

Publisher's note WEDO remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is published online with Open Access by BIG.D and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).

© The Author(s) 2026