NU at WASSA 2024 Empathy and Personality Shared Task: Enhancing Personality Predictions with Knowledge Graphs; A Graphical Neural Network and LightGBM Ensemble Approach

Emmanuel Osei-Brefo
University of Newcastle, UK
emmanuel.osei-brefo@newcastle.ac.uk

Huizhi Liang
University of Newcastle, UK
Huizhi.Liang@newcastle.ac.uk

Abstract

This paper proposes a novel ensemble approach that combines Graph Neural Networks (GNNs) and LightGBM to enhance personality prediction based on the personality Big 5 model. By integrating BERT embeddings from user essays with knowledge graph-derived embeddings, our method accurately captures rich semantic and relational information. Additionally, a special loss function that combines Mean Squared Error (MSE), Pearson correlation loss, and contrastive loss to improve model performance is introduced. The proposed ensemble model, made of Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs), and LightGBM, demonstrates superior performance over other models, with significant improvements in prediction accuracy for the Big Five personality traits achieved. Our system officially ranked 2^{nd} at the Track 4: PER track.

1 Introduction

Personality prediction is a complex task that benefits from understanding both the semantic content of text and the relationships between entities. Traditional machine learning models often fail to capture this relational information. To address this, we propose a novel ensemble approach integrating BERT embeddings, knowledge graph features, Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs), and LightGBM. In this paper, we describe our participation in WASSA 2024 Shared Track 4: Personality Prediction (PER). This year's Track 4, as outlined by (Giorgi et al., 2024), is similar to last year's Shared Track 4 in terms of predicting the Big Five personality traits (OCEAN). However, unlike the 2023 session, where each essay writer was asked to complete the Ten Item Personality Inventory (Barriere et al., 2023), this year's session does not require this step. Our method aims to enhance prediction accuracy by leveraging both semantic and relational data. In recent years, psychologists have developed a number of personality-testing questions (Zhang et al., 2022).

The Big Five model, comprising Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism (OCEAN), is widely used for personality assessment (Barriere et al., 2023). Recent studies have shown the effectiveness of incorporating deep learning techniques in personality prediction (Mehta et al., 2020; Digman, 1990). Knowledge graphs (KGs) represent entities and their relationships, providing valuable contextual information (Peng et al., 2023). GNNs, particularly GCNs and GATs, can effectively process graph-structured data by capturing the structural relations between nodes (Zhang Si, 2019).

Graph Neural Networks (GNNs): Graph Neural Networks (GNNs) are employed in various domains such as social analysis, fraud detection (Akoglu et al., 2015), natural language processing, and computer vision due to their ability to capture structural relations between data, providing more insights compared to isolated data analysis (Zhang Si, 2019). Graph Convolutional Networks (GCNs) enhance this by aggregating information from neighboring nodes, enabling comprehensive extraction of interdependent data. BERT has been extensively used in several tasks to generate token or sentence representations enriched with prior knowledge (Osei-Brefo and Liang, 2022). Our main contributions for participating in the WASSA 2024 Shared Track 4: Personality Prediction (PER) are as follows:

- Integrating BERT embeddings with knowledge graph features.
- Development of an ensemble model that combines GCN, GAT, and LightGBM.

Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis, pages 412–419
August 15, 2024 ©2024 Association for Computational Linguistics

- Introduction of a novel loss function that combines Mean square error loss, Pearson correlation loss, and contrastive loss.

2 Related Work

Previous studies have explored the correlation between personality traits and empathy perception, highlighting the importance of agreeableness and conscientiousness in predicting empathy (Omitaomu et al., 2022; Melchers MC, 2016; Giorgi et al., 2024). Techniques such as text generation adversarial networks and multitask detection models have been employed to enhance personality prediction (Sun et al., 2018; Tu et al., 2022). Recent advancements include the use of dynamic deep graph convolutional networks and the integration of psychological language dictionaries with Transformer language models for improved personality detection (Yang et al., 2023; Kerz et al., 2022). Our approach builds on these methods by combining BERT embeddings, knowledge graphs, GCN, GAT, and LightGBM in a novel ensemble model.

3 Methodology

Figure 1 depicts the architecture of our proposed personality traits prediction system. As can be seen in figure 1, our methodology encompasses several key steps, which are:

- **BERT Pre-training**: A BERT pre-training is utilised to accurately represent the personality features extracted from the sentences in an individual's essays.
- **Knowledge Graph Integration**: These sentence representations are then combined with vector representations derived from knowledge graphs, which include demographic features and the Interpersonal Reactivity Index (IRI).
- **Graph Structure Processing**: During the graph structure data processing, we leverage the comprehensive mapping and syntactic analysis capabilities of multi-layer neural networks, specifically Graph Convolutional Networks (GCN) and Graph Attention Networks (GTA), in conjunction with LightGBM.
- **Personality Trait Modelling and Prediction**: This integrated approach allows the joint modelling and prediction of the individual's personality traits with high accuracy.

3.1 Feature Extraction

The feature extraction method employed combines BERT embeddings with knowledge graph embeddings to capture both semantic and relational information. The process involved is captured in Algorithm 1 presented in Appendix B.

3.2 Graphical Neural Network

In the OCEAN prediction task, a Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs) which are two advanced graph neural network models, were employed in conjunction with LightGBM to predict OCEAN traits.
LightGBM, a gradient boosting framework, is used to complement the graph neural networks. It excels in handling large-scale data and provides efficient training with lower memory usage.

3.3 Mathematical Formulation for GCN and GAT

For a graph $G = (V, E)$ with node features X:

GCN Layer:

$$H^{(l+1)} = \sigma\left(\tilde{D}^{-1/2}\tilde{A}\tilde{D}^{-1/2}H^{(l)}W^{(l)}\right)$$

Where $\tilde{A} = A + I$ is the adjacency matrix with added self-loops, $\tilde{D}$ is the degree matrix, $W^{(l)}$ are the trainable weights, and σ is an activation function.
The adjacency matrix (Zheng et al., 2023) represents the connections between nodes in the graph. Each element in the matrix indicates whether a pair of nodes is connected, and the addition of self-loops ensures that each node is connected to itself. This is critical for the GCN since it allows the model to consider each node's own features in addition to its neighbors' features during convolution.
The adjacency matrix produced is shown in Figure 2 in the Appendix A.

GAT Layer:

$$H_i^{(l+1)} = \sigma\left(\sum_{j\in\mathcal{N}(i)} \alpha_{ij} W H_j^{(l)}\right)$$

where α_{ij} are the attention coefficients computed as:

$$\alpha_{ij} = \frac{\exp(\mathrm{LeakyReLU}(\mathbf{a}^T[WH_i^{(l)}||WH_j^{(l)}]))}{\sum_{k\in\mathcal{N}(i)}\exp(\mathrm{LeakyReLU}(\mathbf{a}^T[WH_i^{(l)}||WH_k^{(l)}]))}$$

with $\mathbf{a}$ as the attention mechanism's weight vector.

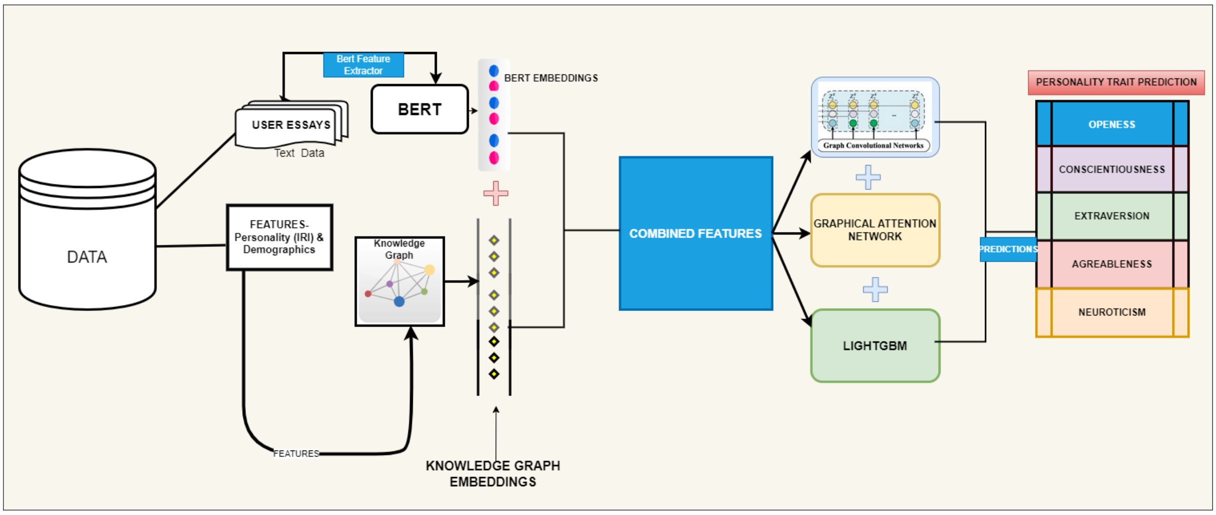

Figure 1: System Architecture of our proposed personality traits prediction system

3.4 Ensemble Prediction

The ensemble model combines predictions from GCN, GAT, and LightGBM using weighted average. The procedure involved is shown in Algorithm 4 in the Appendix B.

3.4.1 Proposed Loss Functions

The loss function used in our methodology is designed to enhance the learning process by incorporating Mean Squared Error (MSE) loss, Pearson correlation loss, and contrastive learning loss. Each component serves a specific purpose:

- **MSE Loss:** It focuses on reducing the prediction error by minimizing the difference between predicted and actual values.
- **Pearson Loss:** Acts as a regularizer to ensure a strong correlation between predictions and targets, enhancing the alignment of predicted and actual OCEAN scores.
- **Contrastive Loss:** It adds an additional layer of learning by emphasizing the relationships between pairs of examples, which is crucial for capturing subtle differences in text and the combined effects of various features.

This multi-faceted approach improves the overall learning process by balancing error minimization, correlation enhancement, and relationship learning. The proposed loss functions are detailed below:

Mean Squared Error (MSE) Loss:

$$\ell_m = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2$$

Contrastive Loss:

$$\ell_c = \frac{1}{N}\sum_{i=1}^{N}(1-y_i)\cdot D_i^2 + y_i \cdot \max(0, m - D_i)^2$$

where D_i is the Euclidean distance between a pair of samples, y_i is the binary label indicating if the samples are similar, and m is the margin.

Pearson Loss: The Pearson correlation coefficient between predictions $\hat{y}$ and targets y is given by:

$$\rho = \frac{\sum_{i=1}^{N}(\hat{y}_i - \bar{\hat{y}})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{N}(\hat{y}_i - \bar{\hat{y}})^2}\sqrt{\sum_{i=1}^{N}(y_i - \bar{y})^2}}$$

The Pearson loss ℓ_ρ is then given by:

$$\ell_\rho = 1 - \rho$$

Total Loss:

$$\ell_T = \frac{1}{2}\ell_m + \lambda_{\text{reg}} \cdot \ell_\rho + \lambda_{\text{c}} \cdot \ell_c$$

where N is the number of pairs, ℓ_m is the Mean Squared Error loss, ℓ_c is the Contrastive loss, and ℓ_ρ is the Pearson loss. The parameters λ_{reg} and λ_{c} are the regularization weights for the Pearson loss and Contrastive loss, respectively.

4 Experiments

4.1 Data Collection and Preprocessing

The dataset used in this study was the Track 4 dataset provided by the organizers, which was subsequently merged with the Task 3 dataset. This merging process involved integrating the essays of each individual from Task 4, resulting in a comprehensive dataset. The final training and development data included features such as:

- **Essays:** Detailed personal essays written by the participants.
- **Demographic Information:** Gender, education, race, age, and income.
- **IRI Features:** Scores from the Interpersonal Reactivity Index.
- **OCEAN Traits:** Target scores for the OCEAN personality traits.

This enriched dataset provided a robust foundation for our feature extraction and model training processes.

4.2 Feature Extraction

- **BERT Embeddings:** We utilized the BERT model to generate embeddings from user essays. These embeddings captured the contextual information and semantic nuances of the text, offering a rich representation of the user's language usage.
- **Knowledge Graph (KG) Construction:** For each user, a knowledge graph was constructed by extracting entities and relationships from their essays using spaCy. These entities and relationships were then represented in a directed graph using NetworkX.
- **KG Embeddings:** The Node2Vec algorithm was employed to generate embeddings for the nodes within the knowledge graph, capturing both structural and relational information.
- **Combined Feature Vector:** The BERT embeddings were combined with the node embeddings from the knowledge graph to form a comprehensive feature vector for each user.

4.3 Models Used

The models used are the Graph Convolutional Network (GCN) to capture the local neighborhood structure within the feature graph, the Graph Attention Network (GAT) that introduced attention mechanisms to allow the model to weigh the importance of different neighboring nodes, and the LightGBM model, which complemented the two graph models by providing robust predictions based on the extracted features. These three models leveraged the strengths of each other to improve overall prediction accuracy as an ensemble strategy. The configuration of hyper-parameters for the proposed model is shown in Table 1.

Hyper-parameters	Description	size
batch size	size of mini-batch used	32
Learning Rate	Used for Adam Optimisation	$1 \times 10^{-}4$
Optimiser used	Type of optimiser used	Adam optimisation
Number of iterations	Number of epochs used	500

Table 1: Hyper-parameters used for the the Ensemble model

4.4 Model Evaluation

The performance of our models was evaluated using the Pearson correlation coefficient for each OCEAN trait and the average Pearson correlation across all traits. This evaluation metric was chosen because it measures the linear correlation between predicted and actual values, providing insight into the model's predictive accuracy.

4.5 Results and Discussion

The ensemble model is evaluated using the Pearson correlation coefficient. A significant improvement in prediction accuracy for the Big Five personality traits compared to baseline models was observed for the test data represented as Ensemble(b) model (unofficial test results) in Table 3. Tables 2 and 3 show the comparative performance of different models on the OCEAN traits for the validation set obtained from using 20% of the training dataset and all the samples of the test data provided by the organisers respectively. The ensemble model outperformed the baseline models, demonstrating the effectiveness of combining GNNs and LightGBM. Figures 3, 4 and 5 in Appendix C show the Loss and Pearson correlations plots per epoch for the GCN, GTA and Ensemble models, respectively. The integration of knowledge graph features and BERT embeddings proves to be particularly beneficial.

Model	$\overline{\rho}$	O	C	E	A	N
LSTM	0.032	0.182	0.085	0.134	-0.258	0.0148
MLP	0.066	0.254	-0.336	0.517	-0.249	0.144
LLM(GPT 3.5)	0.162	0.227	0.185	0.149	0.049	0.200
Ensemble	**0.482**	**0.579**	**0.27**	**0.662**	**0.302**	**0.600**

Table 2: Performance comparison of models on OCEAN traits for the validation Dataset, which is 20% of the training data provide by organisers, where $\overline{\rho}$ represents the average Pearson

Model	$\overline{\rho}$	O	C	E	A	N
LSTM	0.077	0.088	0.296	-0.406	0.206	0.199
MLP	-0.051	-0.178	-0.111	-0.103	0.185	-0.047
LLM(GPT 3.5)	0.095	0.153	-0.069	0.265	0.176	-0.05
Ensemble(a)	0.069	-0.103	0.102	-0.085	0.154	0.279
Ensemble(b)	**0.302**	**0.089**	**0.322**	**0.263**	**0.380**	**0.457**

Table 3: Performance comparison of models on OCEAN traits for the test Dataset: Where Ensemble(a) represents the official pearson results released by the competition organisers and Ensemble(b) is the unofficial results obtained after further improvements to the models during post-competition phase. The LLM represents Open AI's GPT 3.5 turbo model

The main difference in the performance between Ensemble(a) model and Ensemble(b) model is the setting of hyper-parameters. We tuned the hyper-parameters of Ensemble(b) model during the post-submission stage, where different combinations of custom loss functions and optimal weights for the Ensemble models were explored. We selected the best setting of hyper-parameters for Ensemble(b) model.

In the context of personality detection, GCNs effectively captured the relationships and contextual information between nodes by leveraging the graph's topological structure and node characteristics. This capability significantly aided in the prediction of personality traits. The use of a fixed-weight matrix for convolution in GCNs ensured simplicity and scalability.

Additionally, the integration of GCN, GATs and LightGBM enhanced this approach by incorporating attention mechanisms. These mechanisms allowed the model to weigh the importance of different neighbors when aggregating information. The ensemble approach's success underscores the importance of combining different model strengths to achieve better prediction accuracy. The sensitivity analysis (detailed in Appendix D) further demonstrated the impact of different combinations of MSE, Pearson, and contrastive losses on the model's performance, highlighting the optimal weights for each component.

5 Conclusion and Future Work

This paper has proposed an ensemble approach to leverage BERT embeddings and knowledge graph embeddings for GNNs and LightGBM that significantly enhances personality prediction. The introduction of a specialised loss function that combines MSE, Pearson correlation loss, and contrastive losses was crucial for balancing error minimization, correlation enhancement, and the learning of relationships. Future work will involve the refinement of these methods with additional data sources to improve their performance.

Limitations

The process of extracting meaningful features from text using BERT and constructing knowledge graphs requires substantial computational resources. Additionally, the quality of the extracted features can vary depending on the preprocessing and entity extraction methods used, potentially impacting the model's performance.

References

Leman Akoglu, Hanghang Tong, and Danai Koutra. 2015. Graph based anomaly detection and description: a survey. *Data mining and knowledge discovery*, 29:626–688.

Valentin Barriere, Jo ao Sedoc, Shabnam Tafreshi, and Salvatore Giorgi. 2023. Findings of wassa 2023 shared task on empathy, emotion and personality detection in conversation and reactions to news articles. In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, I& Social Media Analysis*, pages 511–525.

John M Digman. 1990. Personality structure: Emergence of the five-factor model. *Annual review of psychology*, 41(1):417–440.

Salvatore Giorgi, Joao Sedoc, Valentin Barriere, and Shabnam Tafreshi. 2024. Findings of wassa 2024 shared task on empathy and personality detection in interactions. In *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, I& Social Media Analysis*.

Elma Kerz, Yu Qiao, Sourabh Zanwar, and Daniel Wiechmann. 2022. Pushing on personality detection from verbal behavior: A transformer meets text contours of psycholinguistic features. *arXiv preprint arXiv:2204.04629*.

Yash Mehta, Samin Fatehi, Amirmohammad Kazameini, Clemens Stachl, Erik Cambria, and Sauleh Eetemadi. 2020. Bottom-up and top-down: Predicting personality with psycholinguistic and language model features. In *2020 IEEE international conference on data mining (ICDM)*, pages 1184–1189. IEEE.

Haas BW Reuter M Bischoff L Montag C Melchers MC, Li M. 2016. Similar personality patterns are associated with empathy in four different countries. *Front Psychol*.

Damilola Omitaomu, Shabnam Tafreshi, Tingting Liu, Sven Buechel, Chris Callison-Burch, Johannes C.

Eichstaedt, Lyle Ungar, and João Sedoc. 2022. Empathic conversations: A multi-level dataset of contextualized conversations. *ArXiv*, abs/2205.12698.

Emmanuel Osei-Brefo and Huizhi Liang. 2022. UoR-NCL at SemEval-2022 task 6: Using ensemble loss with BERT for intended sarcasm detection. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 871–876, Seattle, United States. Association for Computational Linguistics.

Ciyuan Peng, Feng Xia, Mehdi Naseriparsa, and Francesco Osborne. 2023. Knowledge graphs: Opportunities and challenges. *Preprint*, arXiv:2303.13948.

Xiangguo Sun, Bo Liu, Jiuxin Cao, Junzhou Luo, and Xiaojun Shen. 2018. Who am i? personality detection based on deep learning for texts. In *2018 IEEE international conference on communications (ICC)*, pages 1–6. IEEE.

Geng Tu, Jintao Wen, Hao Liu, Sentao Chen, Lin Zheng, and Dazhi Jiang. 2022. Exploration meets exploitation: Multitask learning for emotion recognition based on discrete and dimensional models. *Knowledge-Based Systems*, 235:107598.

Tao Yang, Jinghao Deng, Xiaojun Quan, and Qifan Wang. 2023. Orders are unwanted: dynamic deep graph convolutional network for personality detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13896–13904.

Shunxiang Zhang, Hanqing Xu, Guangli Zhu, Xiang Chen, and KuanChing Li. 2022. A data processing method based on sequence labeling and syntactic analysis for extracting new sentiment words from product reviews. *Soft Computing*, pages 1–14.

Xu Jiejun Maciejewski Ross Zhang Si, Tong Hanghang. 2019. Graph convolutional networks: a comprehensive review. *Computational Social Networks*.

Ruiling Zheng, Peifeng Su, and Xian'an Jin. 2023. Arithmetic-geometric matrix of graphs and its applications. *Applied Mathematics and Computation*, 442:127764.

A The Adjacency Matrix

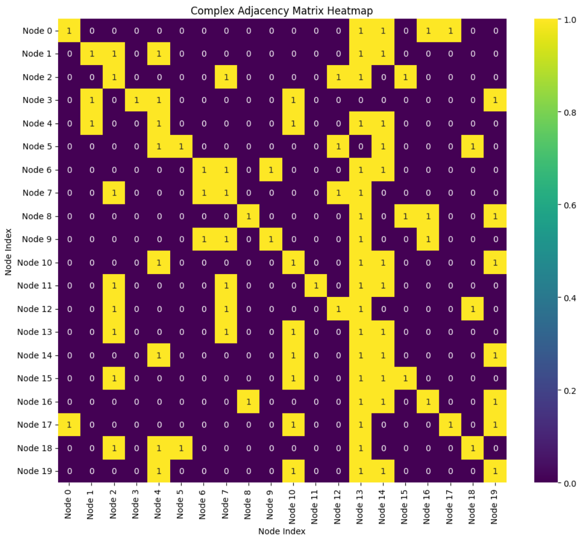

Figure 2: The Adjacency Matrix for the combined features used

B Algorithms

Algorithm 1 BERT Embedding Extraction

Input: Text data T
Initialize BERT tokenizer and model
for each text t in T **do**
    Tokenize t using the BERT tokenizer
    Pass the tokenized text through the BERT model
    Extract the embeddings from the last hidden layer
end for
return embeddings E
Output: BERT embeddings E

Algorithm 2 Knowledge Graph Construction and Embedding

Input: Text data T, additional features F
for each text t in T **do**
    Construct a knowledge graph G from t and F
    Generate node embeddings using Node2Vec
    Aggregate node embeddings to form a fixed-size graph embedding
end for
return embeddings K
Output: Knowledge graph embeddings K

Algorithm 3 Ensemble Learning: Training GCN, GAT, and LightGBM

Input: Training data D, validation data V
Initialize GCN and GAT models
for each epoch **do**
    Train GCN on D and validate on V
    Train GAT on D and validate on V
    Compute training and validation losses
end for
Save the best performing GCN and GAT models
Initialize LightGBM models for each OCEAN trait
for each trait **do**
    Train the LightGBM model on D
    Validate the model on V
    Save the best performing LightGBM model
end for
return trained GCN, GAT, and LightGBM models
Output: Trained GCN, GAT, and LightGBM models

Algorithm 4 Ensemble Prediction

Input: Models M, test data T
for each model m in M **do**
    Generate predictions on T
end for
Average the predictions from all models
return ensemble predictions P
Output: Ensemble predictions P

C Loss and Pearson Plots

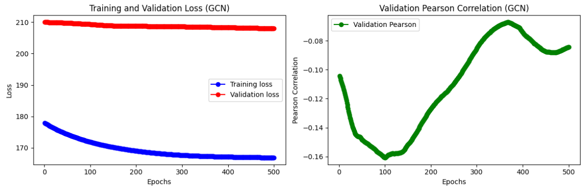

Figure 3: Loss and Pearson plots for GCN per epoch

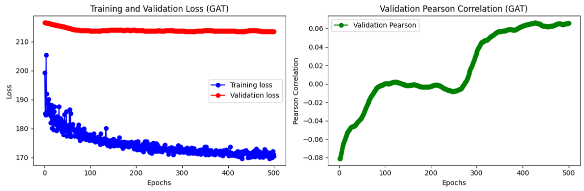

Figure 4: Loss and Pearson plots for GTA per epoch

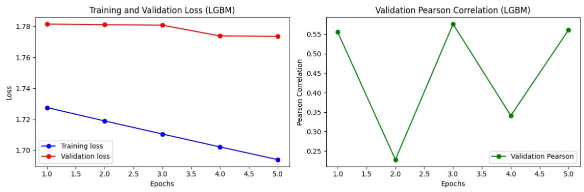

Figure 5: Loss and Pearson plots for the Ensemble GNN+LightGBM per epoch

D Sensitivity Analysis of Loss components

LOSSES	% Weights			Pearson Corr, ρ					
	ℓ_m	ℓ_p	ℓ_c	O	C	E	A	N	$\overline{\rho}$
ℓ_m	1	0	0	0.522	0.431	0.654	0.213	0.475	0.459
ℓ_m + ℓ_p	0.5	0.5	0	0.542	0.237	0.630	0.274	0.558	0.482
ℓ_m+ ℓ_c	0.5	0	0.5	0.548	0.265	0.614	0.314	0.601	0.468
ℓ_m + ℓ_p + ℓ_c	0.5*	0.4*	0.1*	0.579	0.269	0.662	0.302	0.598	**0.482**
	0.5	0.3	0.2	0.567	0.242	0.627	0.273	0.604	0.462
	0.5	0.2	0.3	0.544	0.257	0.605	0.291	0.518	0.443
	0.5	0.1	0.4	0.545	0.254	0.648	0.319	0.540	0.461

Table 4: Sensitivity analysis to find the effect of each loss component: with optimal combination of 50% of MSE, 40% OF Pearson Loss and 10% of Constrastive loss