Monitoring Behavioral Changes Using Spatiotemporal Graphs: A Case Study on the StudentLife Dataset

Anoushka Harit*
Cancer Research UK, Cambridge Institute, University of Cambridge, CB2 0RE, UK
ah2415@cam.ac.uk

Zhongtian Sun, Jongmin Yu
Department of Applied Mathematics and Theoretical Physics, University of Cambridge, CB3 0WA, UK
{zs440, jy522}@cam.ac.uk

Noura Al Moubayed
Department of Computer Science, Durham University, DH1 3LE, UK
noura.al-moubayed@durham.ac.uk

Abstract

This paper introduces a novel method for monitoring behavioral changes in university students by constructing spatiotemporal graphs from smartphone sensor data. Utilizing the StudentLife dataset, which collects multimodal data from smartphone sensors over a 10-week period, we capture detailed aspects of student behavior, including location, physical activity, and self-reported stress. By representing this data as spatiotemporal graphs, we model behavioral evolution across both temporal and spatial dimensions, employing a Spatiotemporal Graph Neural Network (STGNN) to detect patterns associated with stress, sleep quality, and academic performance. This method enables a dynamic, high-resolution analysis of student well-being, offering a more comprehensive understanding of behavior over time.

1 Introduction

University students often experience various stressors due to academic pressures, social interactions, and lifestyle changes. Monitoring these factors can provide insights into their mental health and help predict behavioural trends affecting their well-being. Mobile sensing technologies now enable continuous data collection from students' smartphones, presenting new opportunities to track behaviours in real time. This paper introduces a novel approach to monitor behavioral changes in university students using spatiotemporal graphs based on the StudentLife dataset [Wang et al., 2014], which collects data on physical activity, phone usage, and self-reported stress and sleep quality. By structuring this data as an evolving spatiotemporal graph, we capture both spatial aspects (students' behavior at specific times and locations) and temporal aspects (behavioral changes over time), enabling early diagnostic insights into mental health.

Graph representation learning has shown strong performance across various deep learning tasks and applications [Sun et al., 2021, 2022, 2023a,b, Harit et al., 2024, Wynn et al., 2024]. We employ a Spatiotemporal Graph Neural Network (STGNN) [Yu et al., 2017], integrating Graph Convolutional Networks (GCNs) [Kipf and Welling, 2016] to capture spatial relationships and Recurrent Neural

*I would like to express my sincere gratitude to my mentors, William Prew and Florian Markowetz, at the Markowetz Lab, for their valuable feedback and insightful suggestions, which greatly improved the quality of this paper.

Networks (RNNs) [Elman, 2002] to model temporal changes in behavior over time. This hybrid model uniquely combines spatial dependencies such as behavioral similarities across students with temporal dynamics to provide a comprehensive analysis of behavior trends and predictive insights into future stress levels. Unlike models like MFDS-STGCN [Zhou et al., 2024], which primarily apply spatiotemporal modeling to static datasets, our approach adapts dynamically to evolving student data and incorporates clustering to detect group patterns in stress and activity. This additional layer allows for more responsive monitoring of well being over time. This work contributes to the field in three key ways:

1. **Spatiotemporal graph modelling:** We create dynamic graphs to capture relationships between students' behaviours over time.
2. **Graph-based deep learning:** By applying STGNN, we model complex behavioural patterns and predict mental health trends.
3. **Behavioral clustering for early detection:**We identify groups of students with similar behavioural trends, allowing for early detection of stress and mental health issues.

1.1 Related work

Mobile sensing has revolutionized behavioural analysis, particularly in mental health monitoring. [Wang et al., 2014] introduced the StudentLife dataset, showcasing the potential of smartphone sensors in tracking student behaviour, but early approaches treated this data as an independent time series, missing the critical spatial dynamics and evolving nature of behaviours.

The introduction of Spatiotemporal Graph Neural Networks (STGNNs) by [Yu et al., 2017] bridged this gap by integrating spatial and temporal dimensions, but their application in student behaviour monitoring remains unexplored. Existing models struggle with continuous, high-dimensional smartphone data and often overlook the spatial correlations critical for understanding shared environments and behaviour patterns.

Our approach leverages STGNNs to model the StudentLife dataset [Wang et al., 2014], combining graph convolutional networks (GCNs) for spatial recognition with recurrent neural networks (RNNs) for temporal trends. This enables us to capture complex behaviours across multiple dimensions physical activity, phone usage, sleep quality, and stress levels while predicting future behaviours and identifying early signs of mental health issues. By addressing limitations in existing models, our work provides a scalable, real-time solution for comprehensive mental health monitoring, offering crucial insights into evolving student well-being.

2 Methodology

Our methodology leverages spatiotemporal graph modelling to analyze student behaviour dynamics using the StudentLife dataset. We construct a Spatiotemporal Graph Neural Network (STGNN) to capture both spatial dependencies (behavioural similarities among students at the same time) and temporal dependencies (how behaviours evolve). By integrating graph convolutional networks (GCNs) [Kipf and Welling, 2016] for spatial relationships and recurrent neural networks (RNNs) [Sun et al., 2023a] for temporal changes, we provide a comprehensive model for monitoring and predicting behavioural patterns. The following sections outline data preprocessing, graph construction, and the detailed architecture of the STGNN. For data preprocessing, we drew inspiration from [Bonafonte et al., 2023], adapting their techniques to structure the multimodal sensor data for our analysis.

2.1 Dataset Description

The StudentLife dataset [Wang et al., 2014] comprises sensor-based smartphone data from 48 Dartmouth College students over a 10-week term. It includes GPS-based location data to track movement, activity data from accelerometers and gyroscopes to monitor physical activities, and phone usage records, such as app usage, call logs, and messaging. Additionally, students provided self-reported data through surveys on stress, sleep quality, and mood. The data is discretized into daily time slots, we extract features such as location, activity levels, phone usage, and survey results. PHQ-9 scores [Kroenke et al., 2001] ranging from 0 to 27 are used as the target variable for the depression prediction task.

2.2 Spatiotemporal Graph Construction

We represent the StudentLife data as a spatiotemporal graph $G = (V, E)$, where:

1. Nodes (V) represent students at specific time steps.
2. Edges (E) capture behavioural similarities (spatial edges) or temporal continuity (temporal edges).

Spatial edges connect students at the same time step based on behavioural similarity. The similarity between two students i and j at time t, denoted S_{ij}^t, is computed as:

$$S_{ij}^t = \exp\left(-\frac{\left\|x_i^t - x_j^t\right\|_2^2}{\sigma^2}\right)$$

where x_i^t is the feature vector for student i at time t, and σ controls sensitivity. If S_{ij}^t exceeds a threshold δ, an edge is created between v_i and v_j.

Temporal edges link nodes across consecutive time steps for the same student, modelling the temporal evolution of behaviours. For each student i, an edge connects v_i^t and v_i^{t+1}, indicating how the student's behaviour changes over time.

2.3 Feature Representation

Each node v_i^t is associated with a feature vector $x_i^t = [l_i^t, a_i^t, p_i^t, s_i^t]$, where:

1. l_i^t : Location data (latitude and longitude),
2. a_i^t : Activity data (e.g., steps or sedentary time),
3. p_i^t : Phone usage data (e.g., screen time),
4. s_i^t : Self-reported stress, sleep quality, mood and PHQ-9 score Kroenke et al. [2001](for depression).

These features are normalized across all students to account for variations in behaviour.

2.4 Spatiotemporal Graph Neural Network

We apply a Spatiotemporal Graph Neural Network (STGNN) to learn from the spatiotemporal graph.

1. **Graph Convolution for Spatial Dependencies:** For a graph at time t, we update node embeddings h_i^t using graph convolution:

$$h_i^t = \sigma\left(\sum_{j \in \mathcal{N}(i)} w_{ij}^t \cdot h_j^t + W h_i^t\right)$$

where $\mathcal{N}(i)$ represents neighboring nodes of i, w_{ij}^t is the edge weight, W is a learnable weight matrix, and σ is the activation function (e.g., ReLU).

2. **Recurrent Layers for Temporal Dependencies:** To capture behavioural changes over time, we apply recurrent layers (e.g., LSTM [Hochreiter, 1997] or GRU[Chung et al., 2014]) to the node embeddings h_i^t :

$$h_i^{t+1} = \text{RNN}\left(h_i^t, h_i^{t+1}\right)$$

The RNN captures how each student's behaviour evolves from time t to $t + 1$.

2.5 Clustering and Behavioral Analysis

Once we have learned the node representations, we apply clustering (e.g., K-means or spectral clustering) to group students based on behavioural similarities. The clustering objective is to minimize the within-cluster variance:

$$\min \sum_{k=1}^{K} \sum_{h_i^t \in C_k} \left\|h_i^t - \mu_k^t\right\|_2^2$$

where C_k is the set of nodes in cluster k, and μ_k^t is the cluster centroid at time t. This helps us identify groups of students exhibiting similar stress patterns or activity levels.

3 Results and Discussions

In this study, we evaluate the Spatiotemporal Graph Neural Network (STGNN) against baseline models, including Graph Convolutional Networks (GCN) [Kipf and Welling, 2016], Long Short-Term Memory (LSTM) networks [Hochreiter, 1997], and MFDS-STGCN [Zhou et al., 2024]. The main tasks are stress level classification (evaluated with accuracy and F1 score) and depression score prediction (using RMSE). Each model was trained using all 125 features, as well as individual feature groups (WiFi, GPS, Phone Log, Activity, and Academic), to assess the impact of each data type.

To gain additional insights into behavioral patterns, we applied K-means clustering to the STGNN embeddings, identifying group-level trends: high mobility patterns correlated with lower stress, while sedentary behaviors were associated with higher stress levels. These clustering results highlight the model's ability to uncover meaningful behavioral profiles that could support targeted interventions. For robustness, we conducted 50 rounds of training and testing with different random initializations, reporting the mean and standard deviation for each metric.

3.1 Clustering and Behavioral Analysis Results

To identify distinct behavioral patterns, we applied K-means clustering on STGNN embeddings, which capture spatial and temporal dimensions of student behavior (e.g., mobility, stress, phone usage, activity levels). We determined $k = 2$ clusters based on the optimal silhouette score, maximizing intra-cluster cohesion and inter-cluster separation. Clustering was performed using normalized feature vectors to maintain equal weighting across behavioral metrics.

Table 1: Clustering and Behavioral Analysis Results

Cluster	Students	Mobility Score	Stress Score	Phone Usage (hrs/day)	Activity Level
Cluster 1	25	8.5 $\pm$ 1.2	3.2 $\pm$ 0.8	4.1 $\pm$ 0.9	High
Cluster 2	30	2.4 $\pm$ 0.9	7.1 $\pm$ 1.1	6.5 $\pm$ 1.2	Low

The analysis identified two distinct clusters in Table 1: Cluster $\mathbf{1}$ (high mobility, moderate phone usage, low stress) suggests an active lifestyle correlated with reduced stress, while Cluster 2 (low mobility, high phone usage, high stress) reflects a sedentary profile, potentially linked to isolation or academic pressures. These insights indicate a correlation between behavior patterns and stress levels, suggesting that targeted interventions may benefit high-stress, low-mobility students.

3.2 Stress Level Classification Results

The accuracy and F1 score for stress level classification are summarized in Table 2. Across all feature sets, the STGNN outperforms both GCN [Kipf and Welling, 2016] and LSTM [Hochreiter, 1997]. With all features combined, STGNN achieved the highest performance with an F1 score of 65.4 ± 3.1 and an accuracy of 67.1 ± 2.7. The MFDS-STGCN model also performed well, achieving an F1 score of 63.2 ± 3.2 and an accuracy of 64.0 ± 2.9, though slightly below STGNN. In comparison, the GCN and LSTM models showed lower performance, with F1 scores of 61.5 ± 3.0 and 60.4 ± 3.5, respectively. The superior performance of STGNN highlights the importance of incorporating both spatial and temporal dependencies, as this approach provides a more comprehensive understanding of student behaviors, which is crucial for accurate stress prediction.

Table 2: Stress Level Classification Results (F1 and Accuracy)

Model	Metric	All Features	WiFi	GPS	Phone Log	Activity	Academic
STGNN	F1 Score	65.4 ± 3.1	62.7 ± 3.5	58.3 ± 3.4	59.9 ± 3.8	60.1 ± 3.1	57.2 ± 3.0
	Accuracy	67.1 ± 2.7	64.5 ± 2.9	60.9 ± 3.2	61.3 ± 3.7	61.4 ± 3.0	58.7 ± 2.9
MFDS-STGCN	F1 Score	63.2 ± 3.2	61.0 ± 3.4	57.0 ± 3.1	58.2 ± 3.6	58.5 ± 3.2	56.3 ± 3.0
	Accuracy	64.0 ± 2.9	62.1 ± 3.1	59.0 ± 3.0	59.5 ± 3.5	59.0 ± 3.1	57.0 ± 2.9
GCN	F1 Score	61.5 ± 3.0	60.2 ± 3.2	55.4 ± 3.1	56.5 ± 3.5	56.7 ± 3.4	55.1 ± 3.2
	Accuracy	62.4 ± 2.7	60.4 ± 3.0	58.0 ± 2.9	58.7 ± 3.3	58.8 ± 2.8	56.2 ± 2.7
LSTM	F1 Score	60.4 ± 3.5	57.8 ± 3.0	54.1 ± 3.3	55.7 ± 3.4	55.5 ± 3.1	54.6 ± 3.5
	Accuracy	61.0 ± 3.1	58.3 ± 3.2	56.2 ± 3.1	57.0 ± 3.5	56.8 ± 3.0	55.0 ± 3.2

WiFi and Phone Log features contributed the most to model performance among individual feature groups. The WiFi feature group, which captures student mobility patterns, resulted in an F1 score of 62.7 ± 3.5 for STGNN. This indicates that mobility patterns, particularly time spent at different locations, are closely linked to students' stress levels.

3.3 Depression Score Prediction Results

Table 3 reports the RMSE for the depression score prediction task using PHQ-9 as the target variable. The STGNN model achieved the lowest RMSE with all features (3.5 ± 3.1) and showed strong performance with only WiFi data (3.5 ± 3.4), highlighting mobility as a key predictor of depression. The MFDS-STGCN model followed closely with an RMSE of 3.6 ± 3.2 using all features but did not surpass STGNN. The GCN and LSTM models had higher RMSE values across most feature sets, emphasizing the effectiveness of spatial and temporal modeling in STGNN and MFDS-STGCN.

Table 3: Root Mean Square Error (RMSE) for the regression of PHQ-9 Depression Score

Model	Mean RMSE	All Features	WiFi	GPS	Phone Log	Activity	Academic
STGNN	3.5 ± 3.1	3.5 ± 3.4	3.9 ± 3.5	3.7 ± 3.3	3.6 ± 3.1	3.7 ± 3.2	3.7 ± 3.2
MFDS-STGCN	3.6 ± 3.2	3.6 ± 3.4	4.1 ± 3.3	4.0 ± 3.4	3.9 ± 3.3	3.8 ± 3.4	3.9 ± 3.3
GCN	3.8 ± 3.4	3.7 ± 3.5	4.0 ± 3.4	4.2 ± 3.5	4.1 ± 3.6	4.0 ± 3.5	4.1 ± 3.6
LSTM	4.1 ± 3.2	4.0 ± 3.3	4.3 ± 3.4	4.1 ± 3.2	4.2 ± 3.3	4.3 ± 3.4	4.2 ± 3.5

3.4 Discussions

Our results demonstrate STGNN's superiority over GCN, LSTM, and MFDS-STGCN in both stress detection and depression prediction, achieving an F1 score of 65.4% and an accuracy of 67.1% in stress classification. By integrating spatial relationships and temporal dynamics, STGNN offers a more comprehensive model of student behavior.

WiFi and Phone Log data emerged as the most predictive features, with WiFi mobility patterns particularly effective for identifying mental health trends. Notably, WiFi features alone outperformed the full feature set for depression prediction, suggesting the potential for focused, efficient mental health monitoring. While MFDS-STGCN performed well, it did not reach the predictive accuracy of STGNN, indicating the added value of STGNN's enhanced spatiotemporal integration.

Our approach builds upon previous StudentLife dataset studies by enhancing scalability through passive data collection, which supports real-world applicability. Although GCN, LSTM, and MFDS-STGCN capture certain behavioral patterns, they lack STGNN's effective integration of spatiotemporal information, essential for nuanced mental health predictions. This is ongoing work, and we plan to extend the model to include additional behavioral features and improve predictive accuracy, aiming for a robust, scalable solution for mental health monitoring across diverse student populations.

References

Irene Bonafonte, Cristina Bustos, Abraham Larrazolo, Gilberto Lorenzo Martinez Luna, Adolfo Guzmán Arenas, Xavier Baró, Isaac Tourgeman, Mercedes Balcells, and Agata Lapedriza. Analyzing the contribution of different passively collected data to predict stress and depression. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 1–4. IEEE, 2023.

Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

Jeffrey L Elman. Finding structure in time. *Psycholinguistics: Critical Concepts in Psychology*, 4: 343, 2002.

Anoushka Harit, Zhongtian Sun, Jongmin Yu, and Noura Al Moubayed. Breaking down financial news impact: A novel ai approach with geometric hypergraphs. *arXiv preprint arXiv:2409.00438*, 2024.

S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Kurt Kroenke, Robert L Spitzer, and Janet BW Williams. The phq-9: validity of a brief depression severity measure. *Journal of general internal medicine*, 16(9):606–613, 2001.

Zhongtian Sun, Anoushka Harit, Jialin Yu, Alexandra I Cristea, and Noura Al Moubayed. A generative bayesian graph attention network for semi-supervised classification on scarce data. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2021.

Zhongtian Sun, Anoushka Harit, Alexandra I Cristea, Jialin Yu, Lei Shi, and Noura Al Moubayed. Contrastive learning with heterogeneous graph attention networks on short text classification. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE, 2022.

Zhongtian Sun, Anoushka Harit, Alexandra I Cristea, Jingyun Wang, and Pietro Lio. Money: Ensemble learning for stock price movement prediction via a convolutional network with adversarial hypergraph model. *AI Open*, 4:165–174, 2023a.

Zhongtian Sun, Anoushka Harit, Alexandra I Cristea, Jingyun Wang, and Pietro Lio. A rewiring contrastive patch performermixer framework for graph representation learning. In *2023 IEEE International Conference on Big Data (BigData)*, pages 5930–5939. IEEE Computer Society, 2023b.

Rui Wang, Fanglin Chen, Zhenyu Chen, Tianxing Li, Gabriella Harari, Stefanie Tignor, Xia Zhou, Dror Ben-Zeev, and Andrew T Campbell. Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*, pages 3–14, 2014.

Adam Wynn, Jingyun Wang, Zhongtian Sun, and Atsushi Shimada. Analysing learner behaviour in an ontology-based e-learning system: A graph neural network approach. 2024.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.

Dongbo Zhou, Hongwei Yu, Jie Yu, Shuai Zhao, Wenhui Xu, Qianqian Li, and Fengyin Cai. Mfds-stgcn: Predicting the behaviors of college students with fine-grained spatial-temporal activities data. *IEEE Transactions on Emerging Topics in Computing*, 2024.