

Graph Neural Networks for Parkinson's Disease Detection

Shakeel A. Sheikh
Idiap Research Institute
Martigny, Switzerland
shakeel.sheikh@idiap.ch

Yacouba Kaloga
Idiap Research Institute
Martigny, Switzerland
yacouba.kaloga@idiap.ch

Md. Sahidullah
Institute for Advancing Intelligence
TCG CREST, India
md.sahidullah@tcgcrest.org

Ina Kodrasi
Idiap Research Institute
Martigny, Switzerland
ina.kodrasi@idiap.ch

Abstract—Despite the promising performance of state-of-the-art approaches for Parkinson's Disease (PD) detection, these approaches often analyze individual speech segments in isolation, which can lead to sub-optimal results. Dysarthric cues that characterize speech impairments from PD patients are expected to be related across segments from different speakers. Isolated segment analysis fails to exploit these inter-segment relationships. Additionally, not all speech segments from PD patients exhibit clear dysarthric symptoms, introducing label noise that can negatively affect the performance and generalizability of current approaches. To address these challenges, we propose a novel PD detection framework utilizing Graph Convolutional Networks (GCNs). By representing speech segments as nodes and capturing the similarity between segments through edges, our GCN model facilitates the aggregation of dysarthric cues across the graph, effectively exploiting segment relationships and mitigating the impact of label noise. Experimental results demonstrate the advantages of the proposed GCN model for PD detection and provide insights into its underlying mechanisms.

Index Terms—dysarthria, pathological speech, Parkinson's disease, Graph Convolutional Networks, distance measures

I. INTRODUCTION

Speech production, a highly intricate phenomenon, necessitates a series of complex processes involving the coordinated function of various articulators [1]. Each process is critical and relies on the seamless integration of neural, muscular, and auditory systems to produce spoken language [1]. These processes are influenced by factors such as neurological impairments, emotional state, context, and prior knowledge, highlighting the complexity of speech production [2]. In particular, neurological disorders such as Parkinson's disease (PD) affect the control of the muscles involved in speech production, leading to speech dysarthria characterized by imprecise articulation, abnormal speech rhythm, breathiness, and hypernasality [3]–[5].

To detect speech dysarthria, clinicians commonly use audio-perceptual tests to evaluate and analyze the speech patterns of patients [6]. However, these manual diagnostic tests are often labor-intensive, time-consuming, and subjective [7]. To address these challenges, various automatic methods for PD detection have been proposed, which can be generally grouped into two main categories, i.e., classical machine

learning [8]–[13] and deep learning approaches [14]–[18]. These approaches leverage a variety of speech representations, such as e.g., the short-time Fourier transform (STFT) [19], Mel frequency cepstral coefficients [20], Mel spectrograms [11], or self-supervised embeddings like wav2vec2 (w2v2) [16], [21], integrated with diverse classifiers and architectures to achieve pathological speech detection. Despite the promising performance of state-of-the-art automatic PD detection approaches, they generally focus on analyzing individual speech segments in isolation. Such individual analysis are sub-optimal for two reasons. First, the manifestation of dysarthric cues - such as imprecise articulation, abnormal speech rhythm, breathiness, or hypernasality - is expected to be closely related across representations of different speech segments from various PD speakers. By analyzing speech segments in isolation, state-of-the-art PD detection approaches fail to take advantage of this relationship. Second, not all speech segments from PD patients necessarily contain dysarthric cues, resulting in label noise present in the training data. By analyzing speech segments in isolation, such label noise may considerably affect the accuracy and generalization of state-of-the-art PD detection approaches.

To tackle these challenges, in this paper we propose using graph convolutional networks (GCNs) [22] for PD detection. Specifically, we construct a graph where nodes represent speech segments and edges reflect the relationships between these segments based on their similarity in input representations. Given their proven effectiveness for pathological speech detection in [11], [16], [21], we use w2v2 embeddings of speech segments as input representations. The message-passing mechanism of GCNs aggregates information across connected nodes, allowing to exploit the similarity of pathological cues across multiple segments of speech as well as to potentially overcome label noise issues arising from segments of PD speakers lacking strong pathological cues. Experimental results establish the advantages of the proposed approach for PD detection and provide further insights on the model mechanisms.

II. BACKGROUND AND MOTIVATION

Graphs represent relationships (edges) between objects (nodes) and are extensively used to model data across numerous domains [23]–[25]. To learn from such graph-structured

This work was supported by the Swiss National Science Foundation project CRSII5_202228 on “Characterisation of motor speech disorders and processes”.

data, graph neural networks (GNNs) have been developed. Various types of GNNs have emerged, including GCNs [22], graph attention networks [26], and GraphSAGE (Sample and aggreGatE) [27]. Among these, GCNs are particularly valued for their ability to effectively extract task-relevant information while maintaining relatively low computational complexity, making them a preferred option within the graph community [28].

In recent years, several works in speech processing have exploited the properties of GCNs. Most GCN-based studies, such as those on emotion recognition [29], [30], speaker verification [31], and anti-spoofing [32], focus on graph classification problems where a separate graph is created for speech segments based on time frames. However, in PD detection where pathological speakers share common attributes related to PD, it is more logical to establish inter-speaker connections by building a graph that connects all speech segments rather than creating a separate graph for each speaker. This approach naturally frames the problem as a node classification task, where each node represents a segment of speech. To the best of our knowledge, such an approach is the first in the speech processing area that exploits a graph-structure among the speech segments.

To further highlight our rationale for using GCNs for PD detection, Fig. 1 presents a t-SNE visualization of the w2v2 embeddings (averaged across time) for speech segments from 5 healthy speakers and 5 PD patients from the PC-GITA database [33] (cf. Section IV). In Fig. 1a, the colors represent the different speakers from whom the segments are extracted, while in Fig. 1b, the colors indicate whether the segments are labeled as healthy (blue) or pathological (red). The latent space depicted in Fig. 1b demonstrates that segments sharing the same label tend to cluster closely together, even across different speakers. Furthermore, the overlap between clusters of healthy and PD segments in Fig. 1(b) indicates that not all segments from a pathological speaker exhibit PD biomarkers. By clustering segments with shared characteristics across both speaker and PD groupings, GCNs can effectively diffuse the relevant information.

III. GCN FOR PD DETECTION

We approach PD detection as a graph node classification problem where some speech segments are labeled (i.e., the training data) and the aim is to predict the labels of the unlabeled segments (i.e., the testing data). For each available segment $i \in \{1, 2, \dots, n\}$, with n being the total number of segments in the database, the m -dimensional node representation \mathbf{x}_i is obtained from the w2v2 embeddings computed as outlined in Section IV.

A. Graph Construction

The most straightforward idea for constructing the graph is to calculate the distance between segment representations \mathbf{x}_i and set a threshold below which an edge is established. However, this approach has two key drawbacks. First, it relies on a threshold to be tuned which increases the computational




Fig. 1: t-SNE (perplexity 100) of the w2v2 embeddings (averaged across time) of segments from 5 healthy and 5 pathological speakers from the PC-GITA database: (a) each color represents a speaker and (b) each color represents the label of the segment, with blue denoting healthy segments and red denoting PD segments.

cost of the method, particularly since the threshold is a real number. Second, this method inadequately controls the local connectivity of the graph. While we can adjust this threshold to achieve a certain number of edges, we lack control over their distribution, potentially resulting in some nodes being overly connected while others remaining isolated. To address these issues, we adopt a topological approach as in [34]. To this end, we connect each segment to the k segments with the most similar neighbors, where k is an integer hyper-parameter of the approach (cf. Section IV). This method enables precise local control over graph connectivity, efficiently connecting similar but spatially separated components within the latent space. To this end, we compute the $(n \times n)$ -dimensional kernel \mathbf{K}^d where the entry in the i -th row and j -th column \mathbf{K}_{ij}^d represents the similarity between the representations \mathbf{x}_i and \mathbf{x}_j . This similarity is computed as

$$\mathbf{K}_{ij}^d = \exp(-d(\mathbf{x}_i, \mathbf{x}_j)/h), \quad (1)$$

where d is a user-defined distance measure (cf. Section IV) and h is the average distance across all distinct representations pairs, i.e., $h = \frac{2}{n(n-1)} \sum_{i>j} d(\mathbf{x}_i, \mathbf{x}_j)$. The neighborhood of segment \mathbf{x}_i is represented by the i -th column of the kernel and is denoted by $\mathbf{K}^d(:, i)$. We then compute the Euclidean distance between all column pairs in the kernel \mathbf{K}^d . A segment i is connected to a segment j if $\mathbf{K}^d(:, i)$ is among the top k closest columns to $\mathbf{K}^d(:, j)$, or vice versa, if $\mathbf{K}^d(:, j)$ is among the top k closest columns to $\mathbf{K}^d(:, i)$. This procedure results in a symmetric graph structure described by the $(n \times n)$ -dimensional adjacency matrix \mathbf{A} , where $\mathbf{A}_{i,j} = 1$ if segments i and j are connected in the graph and $\mathbf{A}_{i,j} = 0$ otherwise.

B. PD Detection

Following the method outlined in Section III-A, the entire dataset is transformed into a graph $\mathbb{G} = (\mathbf{A}, \mathbf{X})$, with \mathbf{A} the previously defined adjacency matrix and \mathbf{X} the $n \times m$ -dimensional matrix of node representations. We process this graph using a GCN [22], which updates node representations iteratively by propagating information through neighboring

nodes, exploiting both the initial representations \mathbf{X} and the structure \mathbf{A} . At each layer l of the GCN with L layers, node representations $\mathbf{X}^{(l)}$ are updated as

$$\mathbf{X}^{(l+1)} = \text{ReLU} \left(\tilde{\mathbf{A}} \mathbf{X}^{(l)} \mathbf{W}^{(l)} \right), \quad (2)$$

with $\mathbf{X}^{(0)}$ being the input node representation \mathbf{X} , $\mathbf{W}^{(l)}$ being the trainable weights, $\tilde{\mathbf{A}}$ being a matrix term (derived from the adjacency matrix \mathbf{A}) used to stabilize the gradient, and the ReLU non-linearity applied element-wise.¹ The output representations of the GCN are then processed by a linear layer followed by *softmax* to predict the class label of each segment. While the graph is constructed using all segments in the database (i.e., training and test segments) the weights $\mathbf{W}^{(l)}$ and the linear layer are trained for PD detection using only the segments belonging to the training data.

IV. EXPERIMENTAL SETTINGS

In the following, we present the various experimental settings used for generating the experimental results in Section V.

Dataset. We use the PC-GITA pathological speech dataset, which contains recordings from gender- and age-balanced groups of 50 patients diagnosed with PD and 50 healthy speakers [35]. In line with the state-of-the-art literature, we train PD detection models using only controlled speech recordings of 10 sentences and a phonetically balanced text for each speaker. In addition, [11] has recently shown that w2v2 embeddings can successfully extract pathological cues also from spontaneous speech. Hence, to show the applicability of the proposed model for different speech modes, we also analyze the performance of PD detection models using only spontaneous speech recordings for each speaker.

It should be noted that further analysis on additional databases is required to fully validate the benefits of the proposed model, which remains a topic for future investigation. Currently, the PC-GITA database is the only high-quality open-source database of healthy and PD speech. As demonstrated in [16], other existing databases, such as UA-Speech [36] and TORGO [37], are unsuitable for validating pathological speech detection models, as models trained on them may incorrectly interpret recording artifacts as indicators of pathology.

Input representations. Prior to computing input representations, available utterances are segmented into 500 ms segments with a 50% overlap. For a given speech segment, we extract embeddings from the first layer of the transformer module of w2v2 as in [11], [38]. We use the XLR53 version of w2v2, which has been pre-trained on 56K hours of unlabeled multilingual audio data [39]. The extracted embeddings are averaged across time to obtain the input representation \mathbf{x}_i for the i -th speech segment.

Baselines. To effectively evaluate our proposed GCN-based PD detection method, we consider two baseline models. The first baseline model consists of a fully connected (FC)

layer applied to the input representations. This model represents the state-of-the-art model when using w2v2 embeddings for PD detection [11], [16], [21], [38]. The second baseline model is a k-nearest neighbors (KNN) approach considered here due to its similarity to our proposed approach (in terms of using a distance measure to compare input representations). The comparison between the GCN and KNN approaches allows us to assess whether the improvement in PD detection performance achieved through the GCN approach is merely due to the distance measure or due to the added benefit of using a GCN. To investigate the impact of the used distance measure for the KNN and GCN approaches, we consider three different distance measures, i.e., the Euclidean distance $d(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$, the cosine distance $d(\mathbf{x}_i, \mathbf{x}_j) = 1 - \mathbf{x}_i \cdot \mathbf{x}_j / \|\mathbf{x}_i\|_2 \|\mathbf{x}_j\|_2$, and the Manhattan distance $d(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_1$.

Training and Evaluation. Evaluation is performed using a 10-fold stratified speaker-independent cross-validation framework. In each fold, 80% of the speakers are used for training, 10% for validation, and 10% for testing. Performance is evaluated using the speaker-level accuracy computed through soft voting of the softmax output obtained for all segments belonging to a speaker. To account for the bias introduced by the exact split of speakers into training, validation, and test sets, PD detection models are trained using 5 distinct splits. The reported performance represents the mean and standard deviation of the obtained accuracy across these different splits.

For all considered models, results are obtained through hyperparameter tuning on the validation set, with tailored grid searches conducted for each model. For the FC baseline, the hyperparameter search includes the learning rate $lr \in \{10^{-3}, 10^{-2}\}$. For the KNN baseline, the hyperparameter search includes the number of neighbors $k \in \{1, 2, 3, 5, 7, 10\}$. For the proposed GCN approach, the hyperparameter search includes the learning rate $lr \in \{10^{-3}, 10^{-4}\}$, the number of neighbors $k \in \{1, 2, 3, 5, 7, 10\}$, and the layer depth $L \in \{2, 3, 4, 5\}$. To implement the GCN model, we use the PyTorch Geometric library, which provides an efficient and flexible framework for handling graph data and performing graph-based computations.

V. EXPERIMENTAL RESULTS

In the following, the PD detection accuracy of the proposed GCN-based approach is compared to the FC and KNN baselines for the controlled and spontaneous speech modes. Additional analysis on the impact of the number of neighbors k and the number of layers L on the performance of the GCN-based approach are also provided.

Performance of baselines and proposed approach. Table I presents the performance of the FC, KNN (for various distance measures), and GCN (for various distance measures) approaches for both controlled and spontaneous speech modes. For completeness, the optimal number of neighbors k for the KNN and GCN approaches and the optimal layer depth L for the GCN approach (obtained through hyperparameter tuning on the validation set) are also presented. It can be observed

¹For additional details, the reader is referred to [22].

TABLE I: PD detection accuracy (Acc) using the FC baseline, the KNN baseline, and the proposed GCN approach for both controlled (Ctl) and spontaneous (Spt) speech modes. Several distance measures (Dist) are used for the KNN and GCN approaches, i.e., Euclidean distance (E), cosine distance (C), and Manhattan distance (M). The number of neighbors k and layer depth L (when applicable) are obtained through hyperparameter tuning on the validation set.

Model	Dist	k	L	Acc (Ctl)	k	L	Acc (Spt)
FC	-	-	-	79.0 ± 2.9	-	-	84.2 ± 3.1
KNN	E	5	-	72.4 ± 1.8	5	-	82.6 ± 1.9
KNN	C	5	-	76.8 ± 1.3	5	-	83.8 ± 1.8
KNN	M	1	-	73.4 ± 2.3	1	-	80.8 ± 1.5
GCN	E	3	3	82.0 ± 2.2	5	2	91.8 ± 1.3
GCN	C	5	2	82.2 ± 2.2	3	3	92.8 ± 2.4
GCN	M	3	3	85.8 ± 2.3	7	3	92.0 ± 1.2

that the proposed GCN approach considerably outperforms the FC and KNN baselines for both speech modes, independently of the distance measure used. This confirms the effectiveness of the GCN method in utilizing message passing to integrate information across speech samples from diverse patients. In terms of distance measures, in the spontaneous speech mode we observe minimal variation in performance of the GCN approach, with the cosine distance yielding the highest accuracy of 92.8%. In contrast, the controlled speech mode exhibits more variation, with the Manhattan distance achieving the best performance of 85.8%. Furthermore, Table I shows that the highest performance for the KNN baseline is achieved when using the cosine distance measure. However independently of the distance measure used, the simple FC baseline yields a better performance than KNN, which indicates that using distance measures in their simplest form is insufficient. To fully exploit distance information, incorporating graph-based computations as in the proposed approach is necessary, as illustrated by the higher performance of GCN.

Impact of the number of neighbors k in the GCN approach. To investigate the impact of the number of neighbors k in the GCN approach, in the following we fix the layer depth L and examine the performance when varying k . For each distance measure and speech mode, the layer depth L is fixed to the values presented in Table I. Fig. 2 illustrates the influence of the number of neighbors on the proposed GCN approach for both speech modes. It can be observed that as the number of neighbors for a particular node increases, the detection performance decreases for all considered distance measures. This decline is higher in the controlled speech mode (Fig. 2a) than the spontaneous speech mode (Fig. 2b). These results suggest that sparser graphs created from speech segments of various patients tend to yield better PD detection performance.

Impact of the layer depth L in the GCN approach. To investigate the impact of the layer depth L in the GCN approach, in the following we fix the number of neighbors k and examine the performance when varying L . For each




Fig. 2: Impact of number of neighbors k on PD detection accuracy of the proposed GCN approach using different distance measures for (a) controlled speech and (b) spontaneous speech.




Fig. 3: Impact of layer depth L on PD detection accuracy of the proposed GCN approach using different distance measures for (a) controlled speech and (b) spontaneous speech.

distance measure and speech mode, k is fixed to the values presented in Table I. The results presented in Fig. 3 show that there is typically a decline in PD detection performance of the GCN as the layer depth increases, with the decline being more pronounced for the spontaneous speech mode setting in Fig. 3b. Such an effect is expected and well-documented in the GCN literature and it arises from saturation that occurs after certain depths in the GCN [40].

In summary, the presented results show that a shallow depth of 2 or 3 layers and a small neighborhood size (less than 5) typically ensures an advantageous performance when using a GCN for PD detection. This demonstrates that graph-based approaches are highly effective compared to the state-of-the-art FC layers, with only a modest increase in computational complexity.

VI. CONCLUSION

In this paper we have proposed a novel node classification approach for PD detection using GCNs, constructing a graph from speech segments using all available speakers. Experimental results have shown that GCNs considerably improve the PD detection accuracy compared to baseline methods, regardless of the distance measure used to define connections between segments. Future research will explore incorporating additional meta data into the graph construction process such as gender or age information.

REFERENCES

- [1] Bernd J Kröger and Trevor Bekolay, *Neural Modeling of Speech Processing and Speech Learning*, Springer, 2019.
- [2] Donald B Freed, *Motor Speech Disorders: Diagnosis and Treatment*, Plural Publishing, 2023.
- [3] Celia Stewart et al., “Speech dysfunction in early Parkinson’s disease,” *Movement Disorders: Journal of the Movement Disorder Society*, vol. 10, no. 5, pp. 562–565, Sept. 1995.
- [4] F. L. Darley, Arnold E Aronson, and Joe R Brown, “Differential diagnostic patterns of dysarthria,” *Journal of Speech and Hearing Research*, vol. 12, no. 2, pp. 246–269, Oct. 1969.
- [5] L B-Ravary and Steve W Beet, *Automatic speech signal analysis for clinical diagnosis and assessment of speech disorders*, Springer Science & Business Media, Aug. 2012.
- [6] Parvaneh Janbakhshi, *Automatic pathological speech assessment*, Ph.D. thesis, EPFL, 2022.
- [7] Michaela Pernon et al., “Perceptual classification of motor speech disorders: the role of severity, speech task, and listener’s expertise,” *Journal of Speech, Language, and Hearing Research*, vol. 65, no. 8, pp. 2727–2747, 2022.
- [8] Ina Kodrasi et al., “Automatic discrimination of apraxia of speech and dysarthria using a minimalistic set of handcrafted features,” in *Proc. Annual Conference of the International Speech Communication*, Shanghai, China, Oct. 2020, pp. 4991–4995.
- [9] Ina Kodrasi and Hervé Bourlard, “Spectro-temporal sparsity characterization for dysarthric speech detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1210–1222, April 2020.
- [10] Evaldas Vaiciukynas et al., “Detecting Parkinson’s disease from sustained phonation and speech signals,” vol. 12, no. 10, pp. e0185613.
- [11] Shakeel A Sheikh and Ina Kodrasi, “Impact of speech mode in automatic pathological speech detection,” in *Proc. European Signal Processing Conference*, Lyon, France, Aug. 2024, pp. 3127–3131.
- [12] Jun Wang et al., “Towards automatic detection of amyotrophic lateral sclerosis from speech acoustic and articulatory samples,” pp. 1195–1199.
- [13] Biswajit Karan et al., “An investigation about the relationship between dysarthria level of speech and the neurological state of Parkinson’s patients,” *Biocybernetics and Biomedical Engineering*, vol. 42, no. 2, pp. 710–726, Apr. 2022.
- [14] Parvaneh Janbakhshi and Ina Kodrasi, “Adversarial-free speaker identity-invariant representation learning for automatic dysarthric speech classification,” in *Proc. Annual Conference of the International Speech Communication*, Incheon, Korea, Sept. 2022, pp. 2138–2142.
- [15] Parvaneh Janbakhshi and Ina Kodrasi, “Experimental investigation on stft phase representations for deep learning-based dysarthric speech detection,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Singapore, May 2022, pp. 6477–6481.
- [16] Guilherme Schu et al., “On using the UA-speech and Torgo databases to validate automatic dysarthric speech classification approaches,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Rhodes Island, Greece, June 2023, pp. 1–5.
- [17] Farhad Javanmardi et al., “Pre-trained models for detection and severity level classification of dysarthria from speech,” *Speech Communication*, vol. 158, pp. 103047, Feb. 2024.
- [18] J C Vásquez-Correa et al., “Convolutional neural network to model articulation impairments in patients with Parkinson’s disease,” in *Proc. Annual Conference of the International Speech Communication*, Stockholm, Sweden, Aug. 2017, pp. 314–318.
- [19] J.C. Vásquez-Correa et al., “Convolutional neural network to model articulation impairments in patients with Parkinson’s disease,” in *Proc. Annual Conference of the International Speech Communication*, Stockholm, Sweden, Aug. 2017, pp. 314–318.
- [20] Sara Hawi et al., “Automatic Parkinson’s disease detection based on the combination of long-term acoustic features and Mel frequency cepstral coefficients (MFCC),” *Biomedical Signal Processing and Control*, vol. 78, pp. 104013, 2022.
- [21] Farhad Javanmardi et al., “Wav2vec-based detection and severity level classification of dysarthria from speech,” in *Proc. Annual Conference of the International Speech Communication*, Dublin, Ireland, Aug. 2023, pp. 1–5.
- [22] Thomas N. Kipf and Max Welling, “Semi-supervised classification with graph convolutional networks,” in *Proc. International Conference on Learning Representations*, Toulon, France, April 2017, pp. 2645–2656.
- [23] Kristian Kersting et al., “Benchmark data sets for graph kernels,” 2016.
- [24] Zonghan Wu et al., “A comprehensive survey on graph neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [25] Antonio Ortega et al., “Graph signal processing: Overview, challenges, and applications,” *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [26] Petar Veličković et al., “Graph Attention Networks,” in *Proc. International Conference on Learning Representations*, Vancouver, Canada, May 2018.
- [27] Will Hamilton et al., “Inductive Representation Learning on Large Graphs,” in *Proc. Annual Conference on Neural Information Processing Systems*, California, USA, Dec. 2018, pp. 1025–1035.
- [28] Felix Wu et al., “Simplifying graph convolutional networks,” in *Proc. International conference on machine learning*, California, USA, June, 2019, pp. 6861–6871.
- [29] Amir Shriram and Tanaya Guha, “Compact graph architecture for speech emotion recognition,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Toronto, Canada, June, 2021, pp. 6284–6288.
- [30] Deeksha Chandola et al., “SERC-GCN: Speech emotion recognition in conversation using graph convolutional networks,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Seoul, Korea, April, 2024.
- [31] Liang He et al., “A study on graph embedding for speaker recognition,” in *Proc. International Conference on Acoustics, Speech and Signal Processing*, Seoul, Korea, April, 2024, IEEE, pp. 10741–10745.
- [32] Feng Chen et al., “Graph-based spectro-temporal dependency modeling for anti-spoofing,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Rhodes Island, Greece, June, 2023, pp. 1–5.
- [33] J. R. Orozco-Arroyave et al., “New Spanish speech corpus database for the analysis of people suffering from Parkinson’s disease,” in *Proc. Language Resources and Evaluation Conference*, Reykjavik, Iceland, May 2014, pp. 342–347.
- [34] Jia Chen et al., “Graph multiview canonical correlation analysis,” *IEEE Transactions on Signal Processing*, vol. 67, no. 11, pp. 2826–2838, 2019.
- [35] J. R. Orozco-Arroyave et al., “New Spanish speech corpus database for the analysis of people suffering from Parkinson’s disease,” in *Proc. International Conference on Language Resources and Evaluation*, Reykjavik, Iceland, May 2014, pp. 342–347.
- [36] Heejin Kim et al., “Dysarthric speech database for universal access research,” in *Proc. Annual Conference of the International Speech Communication*, Brisbane, Australia, Sept. 2008, pp. 1741–1744.
- [37] Frank Rudzicz et al., “The TORGO database of acoustic and articulatory speech from speakers with dysarthria,” *Language Resources and Evaluation*, vol. 46, pp. 523–541, Mar. 2012.
- [38] Mahdi Amiri and Ina Kodrasi, “Adversarial robustness analysis in automatic pathological speech detection approaches,” in *Proc. Annual Conference of the International Speech Communication*, Rhodes Islands, Greece, Sept. 2024, pp. 1415–1419.
- [39] Alexis Conneau et al., “Unsupervised cross-lingual representation learning for speech recognition,” *arXiv preprint arXiv:2006.13979*, 2020.
- [40] Sitao Luan et al., “Break the ceiling: Stronger multi-scale deep graph convolutional networks,” in *Proc. International Conference on Neural Information Processing Systems*, Vancouver, Canada, Dec. 2019, pp. 10945–10955.