Graph Emotion Decoding from Visually Evoked Neural Responses

Zhongyu Huang[1,2], Changde Du[1,2], Yingheng Wang[3,4], and Huiguang He[1,2,5]

[1] Research Center for Brain-Inspired Intelligence, National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China
[2] School of Artificial Intelligence, University of Chinese Academy of Sciences, China
[3] Department of Electronic Engineering, Tsinghua University, Beijing, China
[4] Department of Biomedical Engineering, Johns Hopkins University, Baltimore, USA
[5] Center for Excellence in Brain Science and Intelligence Technology, Chinese Academy of Sciences, Beijing, China
{huangzhongyu2020, huiguang.he}@ia.ac.cn

Abstract. Brain signal-based affective computing has recently drawn considerable attention due to its potential widespread applications. Most existing efforts exploit emotion similarities or brain region similarities to learn emotion representations. However, the relationships between emotions and brain regions are not explicitly incorporated into the representation learning process. Consequently, the learned representations may not be informative enough to benefit downstream tasks, e.g., emotion decoding. In this work, we propose a novel neural decoding framework, Graph Emotion Decoding (GED), which integrates the relationships between emotions and brain regions via a bipartite graph structure into the neural decoding process. Further analysis shows that exploiting such relationships helps learn better representations, verifying the rationality and effectiveness of GED. Comprehensive experiments on visually evoked emotion datasets demonstrate the superiority of our model[1].

Keywords: Neural decoding · Graph neural networks · Brain · Emotion.

1 Introduction

Human emotions are complex mental states closely linked to the brain's responses to our diverse subjective experiences. Generally, emotions can be perceived in various ways, such as visual signals [5], audio signals [4], physiological signals [10], or functional neuroimaging techniques [16]. Since functional magnetic resonance imaging (fMRI) has a high spatial resolution and allows a direct, comprehensive assessment of the functions of individual brain regions, it is widely used to explore the relationships between emotions and brain regions.

Recently, Horikawa et al. [8] exploited linear regression to design a neural encoding and decoding model (Fig. 1), building a bridge between emotional experiences and the corresponding brain fMRI responses. The emotional experiences

[1] The code is publicly available at https://github.com/zhongyu1998/GED.

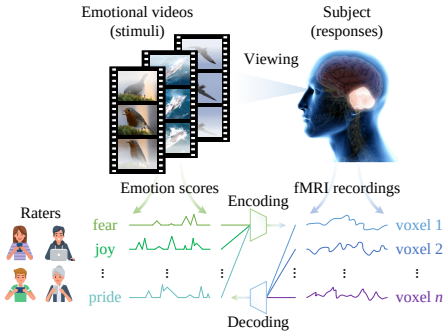

Fig. 1. Encoding and decoding visually evoked emotional responses in fMRI.

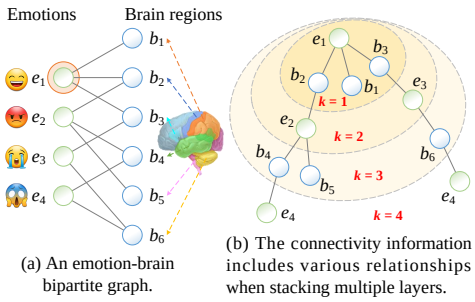

Fig. 2. A toy case of an emotion-brain bipartite graph and the connectivity information captured by embedding propagation layers.

are evoked by videos, and each video was previously annotated by multiple raters using 34 emotion categories (e.g., joy, fear, and sadness; *multi-label*). As shown in Fig. 1, the decoding process aims to predict scores of individual emotions for presented videos from fMRI recordings, and the encoding process aims to predict responses of individual voxels for presented videos from emotion scores.

Although Horikawa et al. [8] make great progress in emotion understanding, their linear regression model processes emotions independently and thus ignores their interconnections. A psychoevolutionary theory proposed by Plutchik [17] shows that different emotions are interrelated. For example, "amusement" is closer to "joy" and "satisfaction" than "horror" and "sadness". Thus, in multi-label emotion classification tasks, emotions with high correlations tend to be labeled together. A pioneering work [21] exploits Graph Neural Networks (GNNs) to model the dependencies between different emotions, and similarly, we can improve the emotion decoding tasks by incorporating such emotion correlations.

In addition, since we use the brain responses to decode emotions, it is intuitive to exploit the associations between different brain regions to promote our emotion decoding tasks. For example, previous studies [12, 16] demonstrate that the perception of aversive emotions (e.g., fear, anxiety) is processed in similar brain regions, such as the insula, amygdala, and medial thalamus. Accordingly, in fMRI-based emotion classification tasks, we can leverage the functional connectivity to analyze the associations between brain regions and construct brain networks [1]. In this way, different emotion categories correspond to different brain networks with more distinguishable graph topological properties and graph structural features. Several latest advances [3, 13, 15, 22] apply similar ideas and achieve remarkable results in analyzing brain diseases and neurological disorders.

Despite their effectiveness, these approaches are insufficient to learn satisfactory emotion representations (or embeddings) for downstream emotion decoding tasks. The primary reason is that only the emotion similarities (i.e., correlations between different emotions) or brain region similarities (i.e., associations between different brain regions) are considered in the learning process. However,

the embedding function lacks the explicit modeling of the relationships between emotions and brain regions. As a result, the learned embeddings may not be informative enough to benefit emotion decoding.

In this work, we propose to integrate the relationships between emotions and brain regions into the representation learning process. Inspired by Wang et al. [18], we construct an emotion-brain bipartite graph (as shown in Fig. 2(a)) to model such relationships and design a GNN to propagate embeddings recursively on this graph. In this way, each embedding propagation layer in the GNN can refine the embedding of an emotion (or a brain region) by aggregating the embeddings of its connected brain regions (or emotions). After stacking multiple (k) embedding propagation layers, the high-order (kth-order) connectivity information (as shown in Fig. 2(b)) can be captured by the learning process and incorporated into the embeddings.

Next, we explain how decoding tasks benefit from high-order connectivity information with a toy case illustrated in Fig. 2(a), where the green and blue nodes denote emotions and brain regions, respectively. Suppose we are interested in an emotion e_1 and aim to decode for e_1. Stacking one embedding propagation layer helps e_1 learn from its 1-hop neighbors b_1, b_2, and b_3 (immediately connected brain regions). And stacking two layers helps e_1 capture the information from its 2-hop neighbors e_2 and e_3 (potentially related emotions, since they have common neighbors b_2 and b_3, respectively). Now the emotion correlations between e_1 and e_2 (e_3) are captured, and the associations between different brain regions can be captured in a similar way. After stacking multiple layers, e_1 can capture most of the high-order connectivity information and perceive most nodes in the graph, thereby aggregating their embeddings and integrating the relationships between them into the neural decoding process. For example, after stacking four layers, all the emotions and brain regions have been perceived by e_1, and their information has been passed along the connected path, as shown in Fig. 2(b). Although we mainly take this idea for decoding tasks in this work, it can be similarly applied to encoding tasks, which is left for future work.

Our contributions can be summarized into three folds:

- We highlight the significance of exploiting the relationships between emotions and brain regions in neural decoding.
- We propose a novel neural decoding framework, Graph Emotion Decoding (GED), which builds a bridge between emotions and brain regions, and captures their relationships by performing embedding propagation.
- We verify the rationality and effectiveness of our approach on visually evoked emotion datasets. Experimental results demonstrate the superiority of GED.

2 Methodology

In this section, we present our proposed approach in detail. The pipeline of our approach is illustrated in Fig. 3. We first construct an emotion-brain bipartite graph with the fMRI recordings and emotion scores on the training set and initialize embeddings for all nodes. Then we stack multiple embedding propagation

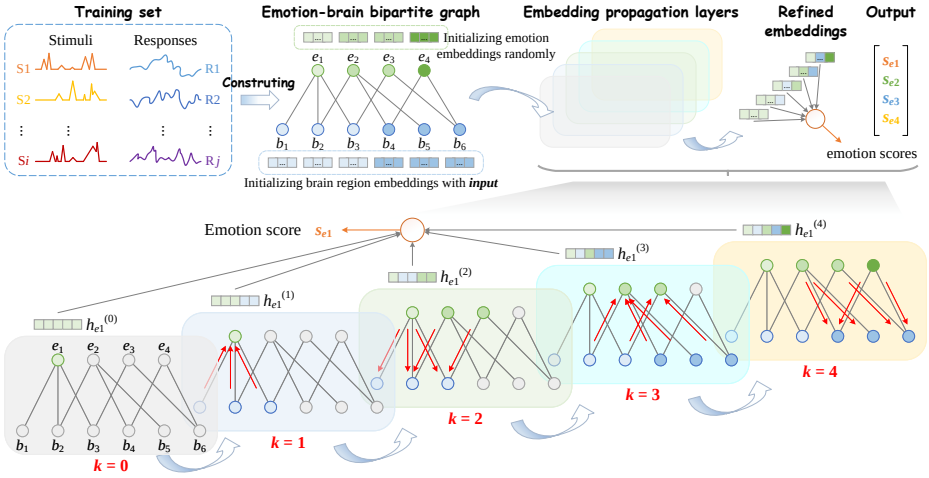

Fig. 3. The pipeline of Graph Emotion Decoding (GED). The component below illustrates how to propagate the information and predict an accurate emotion score s_{e_1} by stacking multiple embedding propagation layers. Each colored panel shows the results of k iterations of propagation starting from node e_1, where the colored nodes indicate that their information has been received by e_1, and the red arrows show the direction of information flow when aggregating neighborhoods from the previous layer.

layers and refine the embeddings by integrating the relationships between nodes on the graph. At last, the refined emotion embeddings are fed into downstream emotion decoding tasks to predict the decoded emotion scores.

2.1 Constructing an Emotion-Brain Bipartite Graph

Let $G = (U, V, E)$ denote an emotion-brain bipartite graph. The vertex set of G consists of two disjoint subsets U and V, where U denotes the emotion set, and V denotes the brain region set. We use each vertex in U and V to represent an emotion category and a brain region, respectively. There are 34 emotion categories in our decoding task, and a total of 370 brain regions consisting of 360 cortical areas (180 cortical areas per hemisphere; defined by the Human Connectome Project (HCP) [6]) and 10 subcortical areas. The complete list of 34 emotion categories and 10 subcortical areas is presented in Table 2 of Appendix.

Lindquist et al. [14] find that when participants are experiencing or perceiving an emotion category, the activation in corresponding brain regions is greater than in a neutral control state. This evidence guides us to construct the edge set E in G as follows. First, on the training set, we classify each stimulus into the corresponding emotion category set according to its highest emotion score. If a stimulus gets the highest score in multiple emotion categories, it will be assigned to these sets simultaneously. Then, for each stimulus in each emotion category set, we sort brain regions in descending order by the average response

of each brain region and select l brain regions corresponding to the top l average responses. Next, the potentially active brain regions for each emotion category are selected by voting. More specifically, the brain regions that got the top m votes among all stimuli in an emotion category set are taken as the potentially active ones for this emotion category. Thus we construct m edges for each emotion vertex, and each edge connects this emotion category (a vertex in U) to one of the potentially active brain regions (vertices in V). In our experiments, we treat l and m as hyper-parameters and select them based on the performance in one random training fold, see Fig. 4 in Appendix for details.

2.2 Initializing Embeddings for Emotions and Brain Regions

Our emotion decoding task takes the voxels in brain regions of interest from fMRI recordings as the initial input. Since each brain region contains a different number of voxels, we use the average pooling to reduce the feature size of each brain region to the same dimension. Specifically, each brain region is surrounded by a cuboid and divided into n equal parts along each axis of the cuboid. The responses of all voxels in each sub-cuboid are then averaged, and the average result is taken as the value of this sub-cuboid. If there are no voxels in a sub-cuboid, we simply take 0 as the value of this sub-cuboid. As a result, we initialize an embedding vector $\boldsymbol{h}_b^{(0)} \in \mathbb{R}^d$ for each brain region b, where $d = n \times n \times n$ denotes the embedding size. For each emotion e, we use Glorot initialization [7] to initialize its embedding vector $\boldsymbol{h}_e^{(0)} \in \mathbb{R}^d$, which can be refined by propagating embeddings on graph G and optimized in an end-to-end manner.

2.3 Embedding Propagation Layers

In general, GNNs follow a neighborhood aggregation strategy. They iteratively update the embedding of a node by aggregating embeddings of its neighbors, and each layer increases the size of the influence distribution by aggregating neighborhoods from the previous layer [20]. After k iterations of aggregation (stacking k layers), a node's embedding captures the structural information within its k-hop network neighborhood [19]. Hence, we can leverage GNNs to perform embedding propagation between the interrelated emotions and brain regions, thus refining their embeddings by integrating the relationships into the learning process.

For an emotion e and each of its connected brain regions $u \in \mathcal{N}_e$, where $\mathcal{N}_e$ denotes the set of 1-hop neighbors of emotion e in graph G, we take their embeddings $\boldsymbol{h}_e$ and $\boldsymbol{h}_u$ as the information to be propagated. Specifically, we consider e's own information $\boldsymbol{h}_e$ to retain its original features, and aggregate the information $\boldsymbol{h}_u$ along each edge (e, u) from e's neighborhood to refine e's embedding. We simply take the element-wise mean of the vectors in $\{\boldsymbol{h}_e\} \cup \{\boldsymbol{h}_u, u \in \mathcal{N}_e\}$ as our aggregation mechanism, commonly known as the "MEAN aggregator" [19]. Formally, the kth embedding propagation layer is:

$$\boldsymbol{a}_e^{(k)} = \text{MEAN}\left(\left\{\boldsymbol{h}_e^{(k-1)}\right\} \cup \left\{\boldsymbol{h}_u^{(k-1)}, u \in \mathcal{N}_e\right\}\right), \ \boldsymbol{h}_e^{(k)} = \text{ReLU}\left(\boldsymbol{W}^{(k)} \cdot \boldsymbol{a}_e^{(k)}\right) \quad (1)$$

where $\boldsymbol{h}_e^{(k)} \in \mathbb{R}^{d_k}$ is the hidden representation (embedding vector) of emotion e at the kth iteration/layer, and d_k is the embedding size of layer k with $d_0 = d$; $\boldsymbol{h}_e^{(k-1)}, \boldsymbol{h}_u^{(k-1)} \in \mathbb{R}^{d_{k-1}}$ are embedding vectors of emotion e and brain region u generated from the previous layer, respectively; $\boldsymbol{W}^{(k)} \in \mathbb{R}^{d_k \times d_{k-1}}$ is a learnable weight matrix. And for a brain region b, we can similarly obtain the embedding by aggregating information from its connected emotions. In this way, each embedding propagation layer can explicitly exploit the first-order connectivity information between emotions and brain regions. After stacking multiple embedding propagation layers, the high-order connectivity information can be captured by the representation learning process and incorporated into the embeddings.

After propagating with K layers, we obtain a set of intermediate embeddings $\{\boldsymbol{h}_e^{(0)}, \boldsymbol{h}_e^{(1)}, \cdots, \boldsymbol{h}_e^{(K)}\}$ from different layers for emotion e. Since the embeddings obtained in different layers reflect different connectivity information, they can flexibly leverage different neighborhood ranges to learn better structure-aware representations [20]. Thus, we combine these intermediate embeddings together to get the decoded emotion score s_e for emotion e, formulated as follows:

$$s_e = \sigma\left(\sum_{k=0}^{K} \boldsymbol{v}_{(k)}^T \boldsymbol{h}_e^{(k)}\right) \tag{2}$$

where $\boldsymbol{v}_{(k)} \in \mathbb{R}^{d_k}$ denotes a learnable scoring vector for embeddings obtained in the kth iteration/layer, shared by all emotions; σ is the logistic sigmoid nonlinearity, used to convert the result $\sum_k \boldsymbol{v}_{(k)}^T \boldsymbol{h}_e^{(k)}$ into an emotion score $s_e \in [0, 1]$ (also can be interpreted as the probability of the existence of emotion e).

3 Experiments

3.1 Dataset Description

We conduct experiments on visually evoked emotion datasets [8], which provide the measured fMRI responses to 2,196 visual stimuli (2,181 unique emotionally evocative short videos + 15 duplicates). For visual stimuli, the videos were originally collected by Cowen and Keltner [2]. Each video was previously annotated by a wide range of raters using self-report scales of 34 emotion categories (e.g., joy, fear, and sadness; *multi-label*). The human raters used binarized scores to report whether these emotions exist in each video, and the final emotion category scores were averaged among these raters, ranging between 0 and 1. The value of each emotion category score can be viewed as the intensity of this emotion. For fMRI responses, we use the fMRI data preprocessed by Horikawa et al. [8], where the value of each voxel represents the activity information of this voxel.

3.2 Experimental Setup

In general, there are two manners for emotion decoding, i.e., predict the emotion categories in a multi-label manner or predict the emotion scores in a regressive

manner. Since human emotions are subjective mental states, different people usually have different feelings for the same stimulus. Even the same person may have different feelings for the same stimulus at different time periods. In view of these facts, we believe that it is more appropriate to predict the emotion score than the existence of a specific emotion. Accordingly, our experiments take the emotion scores obtained in Section 3.1 as the prediction targets and transform the multi-label classification task into a more rigorous regression task.

Evaluation Protocol. We perform decoding analyses between the fMRI responses to each visual stimulus and its emotion scores in a cross-validated manner. Specifically, we perform the 10-fold cross-validation on 2,196 pairs of {fMRI recordings, emotion scores} for each subject (a *within-subject* design). In each trial of 10-fold cross-validation, we use 1,976 random pairs (9 folds) as the training set to train a model and then take the remaining 220 pairs (1 fold) as the validation set (a.k.a. testing set) to evaluate the model. We repeat this process 10 times and use each fold as the testing set exactly once.

Evaluation Metric. We take the Mean Absolute Error (MAE) as our evaluation metric, which is widely used in regression tasks, defined as follows. Let $\boldsymbol{y}_i \in \mathbb{R}^c$ and $\hat{\boldsymbol{y}}_i \in \mathbb{R}^c$ denote the true and predicted (or decoded) emotion category scores for stimulus i, respectively, where c is the number of emotion categories. Let y_{ij} and $\hat{y}_{ij}$ denote the true and predicted score of a specific emotion e_j in stimulus i, respectively, where $\hat{y}_{ij} = s_{e_j}$ is calculated by Eq. (2). The MAE value (a.k.a. L1-norm) α is calculated as:

$$\alpha = \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{c} |y_{ij} - \hat{y}_{ij}| \tag{3}$$

where N denotes the number of stimuli.

Baselines. We compare our proposed model GED with several state-of-the-art baselines: Fully-connected Neural Network (FNN), Graph Convolutional Network (GCN) [11], BrainNetCNN [9], and BrainGNN [13]. FNN directly takes fMRI recordings as the input, mainly used to verify the effectiveness of other graph-based models. GCN takes a random graph as the input graph to verify the rationality of the graph construction in Section 2.1. The last two are state-of-the-art brain-network-based models, which are treated as our main baselines.

3.3 Performance Comparison

Table 1 presents a summary of the emotion decoding results. All models are evaluated using the standard 10-fold cross-validation procedure. We report the average and standard deviation of testing MAEs across the 10 folds (or 10 trials). The results show that our model achieves outstanding performance and outperforms baselines by a considerable margin in most subjects. Furthermore, we have the following observations and conclusions:

Table 1. Emotion decoding results (MAE, lower is better) in five distinct subjects.

Model	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5
FNN	2.053 ± 0.059	2.071 ± 0.079	2.078 ± 0.086	2.039 ± 0.065	2.065 ± 0.078
GCN	1.817 ± 0.036	1.822 ± 0.028	1.819 ± 0.032	1.821 ± 0.028	1.818 ± 0.034
BrainNetCNN	1.723 ± 0.031	1.719 ± 0.034	1.736 ± 0.032	1.740 ± 0.031	1.741 ± 0.029
BrainGNN	1.678 ± 0.027	1.675 ± 0.026	1.684 ± 0.029	1.693 ± 0.026	**1.681 ± 0.024**
GED-1	2.010 ± 0.087	1.992 ± 0.076	2.014 ± 0.073	1.950 ± 0.092	2.010 ± 0.099
GED-2	1.755 ± 0.024	1.753 ± 0.028	1.769 ± 0.022	1.775 ± 0.022	1.778 ± 0.023
GED-3	1.675 ± 0.028	1.673 ± 0.029	1.684 ± 0.024	1.694 ± 0.027	1.698 ± 0.024
GED-4	**1.643 ± 0.028**	**1.647 ± 0.029**	**1.659 ± 0.026**	**1.668 ± 0.028**	**1.674 ± 0.027**

- FNN shows poor performance in all subjects, which indicates that traditional neural networks are insufficient to capture the complex relationships between emotions and brain regions. Other graph-based models consistently outperform FNN in all cases, demonstrating the importance of incorporating the relationships into the learning process.
- The brain-network-based models and GED all surpass GCN, which takes a random graph as the input graph. This fact highlights the importance of graph construction and verifies the rationality of our construction in Sec. 2.1.
- GED-4 significantly outperforms brain-network-based models in most cases (except Subject 5). This justifies the rationality of the analysis in Section 1 and the effectiveness of our model. However, GED performs on par with BrainGNN in Subject 5. The reason might be that Subject 5 was reluctant to use the custom-molded bite bar to fix his head, thus causing some potential head movements during fMRI scanning (see [8] for details). As a result, the head movements introduce noise to the collected fMRI responses and further affect the graph construction, leading to the sub-optimal performance.

3.4 Ablation Study

As the embedding propagation layer plays a pivotal role in GED, we perform ablation studies on GED to investigate how the layer size affects the performance. Specifically, we search the layer size in {1, 2, 3, 4} and report the corresponding results in Table 1, where the suffix "-k" indicates the number of layers is k. As shown in Table 1, GED-1 exhibits better performance than FNN in all cases. Such improvement verifies that exploiting the relationships between emotions and brain regions can help learn better representations, thereby enhancing the decoding performance. After stacking two layers, GED can implicitly exploit both the emotion correlations and the associations between brain regions (as explained in Section 1), and thus GED-2 achieves substantial improvement over GED-1. Moreover, GED-2, GED-3, and GED-4 are consistently superior to GED-1 in all cases, and their performance gradually improves with the increase of layers. These facts justify the effectiveness of stacking multiple embedding propagation layers and empirically show that the high-order connectivity information can significantly facilitate the emotion decoding tasks.

4 Conclusion and Future Work

In this work, we propose a novel neural decoding framework, Graph Emotion Decoding (GED), which integrates the relationships between emotions and brain regions via a bipartite graph structure into the neural decoding process. Further analysis verifies the rationality and effectiveness of GED. In conclusion, we take an important step forward to better understand human emotions by using graph-based neural decoding models. It would be interesting to apply the proposed idea to the neural encoding process, and we leave the explorations of emotion encoding tasks for future work.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China (61976209, 62020106015), CAS International Collaboration Key Project (173211KYSB20190024), and Strategic Priority Research Program of CAS (XDB32040000).

References

1. Bullmore, E., Sporns, O.: Complex brain networks: Graph theoretical analysis of structural and functional systems. Nature Reviews Neuroscience **10**(3), 186–198 (2009)
2. Cowen, A.S., Keltner, D.: Self-report captures 27 distinct categories of emotion bridged by continuous gradients. Proceedings of the National Academy of Sciences **114**(38), E7900–E7909 (2017)
3. Cui, H., Dai, W., Zhu, Y., Li, X., He, L., Yang, C.: BrainNNExplainer: An interpretable graph neural network framework for brain network based disease analysis. In: ICML 2021 Workshop on Interpretable Machine Learning in Healthcare (2021)
4. Dellaert, F., Polzin, T., Waibel, A.: Recognizing emotion in speech. In: Proceeding of the 4th International Conference on Spoken Language Processing. vol. 3, pp. 1970–1973 (1996)
5. Ekman, P.: Facial expression and emotion. American Psychologist **48**(4), 384 (1993)
6. Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., et al.: A multi-modal parcellation of human cerebral cortex. Nature **536**(7615), 171–178 (2016)
7. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the 13th International Conference on Artificial Intelligence and Statistics. pp. 249–256 (2010)
8. Horikawa, T., Cowen, A.S., Keltner, D., Kamitani, Y.: The neural representation of visually evoked emotion is high-dimensional, categorical, and distributed across transmodal brain regions. iScience **23**(5), 101060 (2020)
9. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. NeuroImage **146**, 1038–1049 (2017)
10. Kim, J., André, E.: Emotion recognition based on physiological changes in music listening. IEEE Transactions on Pattern Analysis and Machine Intelligence **30**(12), 2067–2083 (2008)

11. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: Proceedings of the 5th International Conference on Learning Representations (2017)
12. von Leupoldt, A., Sommer, T., Kegat, S., Baumann, H.J., Klose, H., Dahme, B., Büchel, C.: Dyspnea and pain share emotion-related brain network. NeuroImage **48**(1), 200–206 (2009)
13. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: BrainGNN: Interpretable brain graph neural network for fMRI analysis. Medical Image Analysis **74**, 102233 (2021)
14. Lindquist, K.A., Wager, T.D., Kober, H., Bliss-Moreau, E., Barrett, L.F.: The brain basis of emotion: A meta-analytic review. The Behavioral and Brain Sciences **35**, 121–202 (2012)
15. Ma, J., Zhu, X., Yang, D., Chen, J., Wu, G.: Attention-guided deep graph neural network for longitudinal alzheimer's disease analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 387–396 (2020)
16. Phan, K.L., Wager, T., Taylor, S.F., Liberzon, I.: Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. NeuroImage **16**(2), 331–348 (2002)
17. Plutchik, R.: A general psychoevolutionary theory of emotion. In: Theories of Emotion, pp. 3–33. Elsevier (1980)
18. Wang, X., He, X., Wang, M., Feng, F., Chua, T.S.: Neural graph collaborative filtering. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 165–174 (2019)
19. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: Proceedings of the 7th International Conference on Learning Representations (2019)
20. Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.i., Jegelka, S.: Representation learning on graphs with jumping knowledge networks. In: Proceedings of the 35th International Conference on Machine Learning. pp. 5453–5462 (2018)
21. Xu, P., Liu, Z., Winata, G.I., Lin, Z., Fung, P.: EmoGraph: Capturing emotion correlations using graph networks. arXiv preprint arXiv:2008.09378 (2020)
22. Zhang, Y., Tetrel, L., Thirion, B., Bellec, P.: Functional annotation of human cognitive states using deep graph convolution. NeuroImage **231**, 117847 (2021)