Node-Based Gaussian Graphical Model for Identifying Discriminative Brain Regions from Connectivity Graphs

Bernard Ng[1,2(✉)], Anna-Clare Milazzo[1], and Andre Altmann[1]

[1] FIND Lab, Stanford University, Stanford, CA, USA
bernardyng@gmail.com
[2] Parietal Team, Neurospin, INRIA Saclay, Paris, France

Abstract. Despite that the bulk of our knowledge on brain function is established around brain regions, current methods for comparing connectivity graphs largely take an edge-based approach with the aim of identifying discriminative connections. In this paper, we explore a node-based Gaussian Graphical Model (NBGGM) that facilitates identification of brain regions attributing to connectivity differences seen between a pair of graphs. To enable group analysis, we propose an extension of NBGGM via incorporation of stability selection. We evaluate NBGGM on two functional magnetic resonance imaging (fMRI) datasets pertaining to within and between-group studies. We show that NBGGM more consistently selects the same brain regions over random data splits than using node-based graph measures. Importantly, the regions found by NBGGM correspond well to those known to be involved for the investigated conditions.

Keywords: Brain · Connectivity · fMRI · Node-Based Gaussian Graphical Model

1 Introduction

Brain function is known to be largely mediated via the interactions between brain regions. Disruptions of these brain connections can result in severe consequences. Functional magnetic resonance imaging (fMRI) provides a non-invasive means for studying functional connectivity. The standard approach for analyzing brain connectivity estimated from fMRI data is to perform univariate test on each connection [1]. However, the bulk of our knowledge on brain function originates from lesion studies and task-based functional imaging studies [2], and hence, is organized around brain regions, e.g. hippocampus is involved in memory processes. Further, how to design treatments to target a specific brain connection is unclear. Devising methods for identifying brain regions that give rise to altered connection patterns is thus beneficial.

A modeling perspective that aids identification of discriminative brain regions is to treat the brain as a graph and characterize each brain region with a set of graph measures [3]. Under this perspective, brain regions are abstracted as graph nodes with connectivity modeled by graph edges. Graph measures, such as node degree (ND), clustering coefficient (CC), and betweenness centrality (BC), are widely employed, but only capture very specific attributes of the graphs, and thus, might at times fail to

© Springer International Publishing Switzerland 2015
L. Zhou et al. (Eds.): MLMI 2015, LNCS 9352, pp. 1–8, 2015.
DOI: 10.1007/978-3-319-24888-2_6

isolate the discriminative nodes. A simple example illustrating this shortcoming is shown in Fig. 1. Despite that the connections to node 1 are different between the two graphs, ND, CC, and BC of node 1 are exactly the same. Recently, a probabilistic graphical model has been put forth that more fully exploits the connectivity information in the graphs for identifying discriminative brain regions [4]. Promising results have been shown in a schizophrenia study.

In this paper, we explore a node-based Gaussian graphical model (NBGGM) [5] for identifying brain nodes that drive differences between connectivity graphs. Compared to the model in [4], NBGGM has much fewer modeling assumptions. The classical GGM (Section 2.1) provides sparse inverse covariance estimates with zero elements reflecting conditional independence between node pairs. The output of GGM is thus a sparse set of edges. In contrast to GGM, NBGGM incorporates a row-column overlap norm (Section 2.2) that promotes selection of entire rows and columns, resulting in sparsity patterns similar to that in Fig. 1(d). Relevant nodes can hence be identified from the sparsity patterns. The original NBGGM (Section 2.3) is designed for finding relevant nodes from a pair of graphs (or multiple graphs), but not from two sets of graphs, which often arises in fMRI studies. For example, one might be interested in studying the differences between two experimental conditions within the same group of subjects, i.e. a within-group study. Alternatively, one might be interested in comparing two groups of subjects, e.g. Alzheimer's disease (AD) patients vs. healthy controls, i.e. a between-group study. Both types of studies entail comparing two sets of graphs in drawing group inference. To facilitate group analysis, we propose an extension of NBGGM (Section 2.4) via incorporation of stability selection [6]. The underlying idea is that if we subsample the subjects many times and apply NBGGM on the "average" graphs of each subsample, nodes that are truly discriminative are likely to be selected over a large fraction of subsamples, whereas false nodes are unlikely to be persistently selected. Also, stability selection has the property of being insensitive to the choice of regularization parameters [6]. We evaluate NBGGM on two fMRI datasets pertaining to within and between-group studies, and compare its performance against using weighted ND, CC, and BC in identifying discriminative brain regions from functional connectivity graphs.

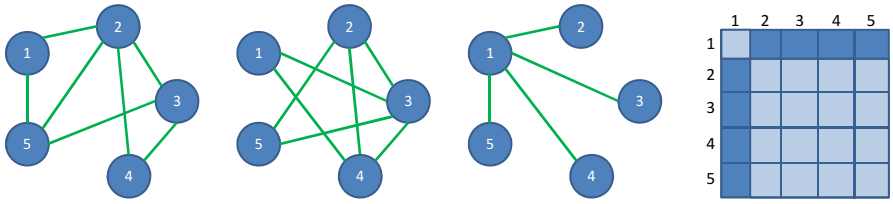

(a) Graph A (b) Graph B (c) Edge Differences (d) Adjacency Matrix

Fig. 1. Motivation example. (a) Edges between nodes 1 and 2 as well as between nodes 1 and 5 in Graph A are absent in Graph B. (b) Edges between nodes 1 and 3 as well as between nodes 1 and 4 in Graph B are absent in Graph A. (c) All edge differences are associated with node 1, but ND = 2, CC = 1, and BC = 0 for node 1 of both graphs. (d) Adjacency matrix encoding edge differences with non-zeros in dark blue. NBGGM promotes selection of entire rows and columns, which produces sparsity patterns similar to (d), hence enables node identification.

2 Methods

2.1 Gaussian Graphical Model

Let $\mathbf{X}$ be a $n \times d$ time series matrix with n being the number of time samples and d being the number of nodes. Assuming $\mathbf{X}$ follows a centered multivariate Gaussian distribution, $N(\mathbf{0}, \boldsymbol{\Sigma})$, a sparse estimate of $\boldsymbol{\Sigma}^{-1}$ can be obtained by minimizing the penalized negative log data likelihood over the space of positive definite matrices [5]:

$$\min_{\boldsymbol{\Theta}>0} tr(\mathbf{C}\boldsymbol{\Theta}) - logdet(\boldsymbol{\Theta}) + \lambda_1 \|\boldsymbol{\Theta}\|_1 , \tag{1}$$

where $\mathbf{C}$ is the sample covariance of $\mathbf{X}$, $\|\Theta\|_1 = \sum_{i,j}|\Theta_{ij}|$ is the l_1 norm of Θ, and λ_1 controls the level of sparsity. The key property of (1) is that 0's in its solution indicate the corresponding node pairs are estimated to be conditionally independent given the other nodes. $\|\Theta\|_1$ promotes a sparse estimate of $\boldsymbol{\Sigma}^{-1}$, but does not impose any structure on the sparsity pattern. Constraining the sparsity pattern to be similar to Fig. 1(d) is useful for identifying relevant nodes. How to exert such a constraint is discussed next.

2.2 Row-Column Overlap Norm

A widely-used approach to impose structure on the sparsity patterns is to employ the group least absolute shrinkage and selection operator (LASSO) penalty [7]:

$$\|\mathbf{W}\|_{2,1} = \sum_{g=1}^{G} \|\mathbf{W}_g\|_2 , \tag{2}$$

where $\mathbf{W}_g$ is a vector with elements predefined to be in group g and G is the number of non-overlapping groups. Minimizing (2) promotes all elements within each group g to be jointly selected (or jointly set to 0). To select entire rows and columns, i.e. impose a sparsity pattern similar to Fig. 1(d), one might be tempted to predefine each column of $\boldsymbol{\Theta}$ as a group and apply (2). However, due to the symmetric constraint on $\boldsymbol{\Theta}$, selecting the entire column j of $\boldsymbol{\Theta}$, i.e. $\boldsymbol{\Theta}_{ij}$ for all i**,** requires $\boldsymbol{\Theta}_{ji}$ for all i to also be selected. Since all elements of each column are enforced to be jointly selected, the entire $\boldsymbol{\Theta}$ would be selected in theory. This problem is due to overlaps between groups. In practice, assuming nodes a and b are relevant, (2) would only select $\boldsymbol{\Theta}_{aa}$, $\boldsymbol{\Theta}_{ab}$, $\boldsymbol{\Theta}_{ba}$, and $\boldsymbol{\Theta}_{bb}$ [5], which does not help node identification. One way to deal with group overlaps arising from the symmetry of $\boldsymbol{\Theta}$ is to use a row-column overlap norm [5]:

$$\Omega(\boldsymbol{\Theta}) = \min_{\mathbf{V}} \sum_{j=1}^{d} \|\mathbf{V}_j\|_2 \quad s.t. \quad \boldsymbol{\Theta} = \mathbf{V} + \mathbf{V}^T , \tag{3}$$

where $\mathbf{V}$ is a $d \times d$ matrix and $\mathbf{V}_j$ is its j^{th} column. Defining $\mathbf{V}_j$ as a group enforces entire columns to be selected, while imposing $\boldsymbol{\Theta} = \mathbf{V} + \mathbf{V}^{\mathrm{T}}$ ensures $\boldsymbol{\Theta}$ is symmetric. Hence, adding (3) to (1) would produce the desired effect of generating sparse symmetric estimates of $\boldsymbol{\Sigma}^{-1}$ with entire rows and columns selected. How to adapt (1) and (3) for identifying discriminative nodes from a pair of graphs is discussed next.

2.3 Node-Based Gaussian Graphical Model

To identify nodes that give rise to differences between a pair of weighted graphs, one can combine (1) and (3) as follows [5], which we refer to as NBGGM:

$$\min_{\boldsymbol{\Theta}^1>0,\boldsymbol{\Theta}^2>0} \sum_{k=1}^{2}\left(tr(\mathbf{C}^k\boldsymbol{\Theta}^k)-logdet(\boldsymbol{\Theta}^k)\right)+\lambda_1\sum_{k=1}^{2}\left\|\boldsymbol{\Theta}^k\right\|_1+\lambda_2\Omega\left(\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2\right). \quad (4)$$

Penalizing $\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2$ encourages $\boldsymbol{\Theta}^1$ and $\boldsymbol{\Theta}^2$ to be similar, while enforcing sparsity highlights their key distinctions. Importantly, $\Omega(\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2)$ imposes structured sparsity such that entire rows and columns of $\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2$ are jointly selected, which aids node identification. Given weighted graphs, all edges associated with any node will always display some differences between graphs (let it be small). Thus, in effect, $\Omega(\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2)$ helps find nodes with edges displaying larger differences in edge weights. Jointly minimizing all terms in (4) is non-trivial. A widely-used strategy is to introduce auxiliary variables to decouple the objective into sub-problems that are easily solvable. This strategy is the core of alternating direction method of multipliers [5], which can efficiently solve (4). With the implementation in [5], we observed that the selected nodes are sometimes not visually apparent from $\boldsymbol{\Theta}^1-\boldsymbol{\Theta}^2$. We instead recommend determining the selected nodes based on non-zero columns in $\mathbf{V}$, which we found to provide an unambiguous answer. For the choice of λ_1 and λ_2, we bypass selecting a specific combination by using stability selection, as discussed next. Note that NBGGM can also be used for finding co-hubs between a pair of graphs by applying (3) on $[\boldsymbol{\Theta}^1;\boldsymbol{\Theta}^2]$, i.e. $\boldsymbol{\Theta}^1$ and $\boldsymbol{\Theta}^2$ concatenated column-wise [5]. Also, (4) can be extended for multi-graphs [5].

2.4 Group Analysis with NBGGM

In fMRI connectivity studies, the task of comparing two sets of graphs, $\{\mathbf{C}^1(p)\}$ and $\{\mathbf{C}^2(q)\}$, is often encountered. Here, p and q denote subject indices. For within-group studies, $p = q$ and the analysis of interest is comparing conditions 1 and 2. In the case of between-group studies, $p \neq q$ and the analysis of interest is comparing two groups of subjects. A natural way to employ (4) in finding discriminative nodes is to apply it on the subject averages of $\{\mathbf{C}^1(p)\}$ and $\{\mathbf{C}^2(q)\}$. However, declaring the selected nodes as significant can be dangerous, since sparse methods tend to be unstable, i.e. perturbations to the data can easily result in a different set of nodes being selected [6]. To obtain a stable set of nodes, one strategy is to employ stability selection [6]:

1. Randomly subsample $\{\mathbf{C}^1(p)\}$ and $\{\mathbf{C}^2(q)\}$ by half, and compute their respective averages. For within-group analysis, choose the same set of subjects.
2. For each (λ_1, λ_2) combination in $[\lambda_1^{\max}, \lambda_1^{\min}] \times [\lambda_2^{\max}, \lambda_2^{\min}]$, apply (4) to the two average sample covariance matrices. Let $\mathbf{Z}^m(\lambda_1, \lambda_2)$ be a $d\times1$ vector with elements corresponding to the nodes selected for subsample m set to 1.
3. Repeat steps 1 and 2 for $M = 1000$ times.
4. Compute the proportion of subsamples, $\boldsymbol{\pi}_i(\lambda_1, \lambda_2)$, that node i is selected for each (λ_1, λ_2) combination.
5. Declare node i as significant if $max_{(\lambda 1,\, \lambda 2)}\boldsymbol{\pi}_i(\lambda_1, \lambda_2) \geq \pi_{\text{th}}$.

A π_{th} that controls the expected number of false positives, $E(F)$, is given by [6]:

$$E(F) \leq \frac{1}{2\pi_{th} - 1} \frac{\gamma^2}{d}, \quad (5)$$

where F is the number of false positives and γ is the expected number of selected nodes, which can be approximated by: $1/M \cdot \sum_m \sum_i (\mathrm{U}_{(\lambda 1, \lambda 2)} \mathbf{Z}_i^m(\lambda_1, \lambda_2))$. $\mathrm{U}_{(\lambda 1, \lambda 2)}$ denotes the union over all (λ_1, λ_2) combinations. We highlight two insights on (5) that have major implications on applying stability selection. First, (5) is a conservative bound on the family-wise error rate (FWER) = $P(F \geq 1)$, since $E(F) = \sum_{f=1}^{\infty} P(F \geq f) > P(F \geq 1)$. To control FWER at $\alpha = 0.05$ with multiple comparison correction, i.e. $P(F \geq 1) \leq \alpha/d$, even for $\gamma = 1$, π_{th} based on (5) is >1. In [6], π_{th} is recommended to be set between 0.6 and 0.9. In this work, we set π_{th} to 0.8. Second, although stability selection does not require choosing a specific (λ_1, λ_2) pair, for $n/2 > d$, a "small enough" $(\lambda_1^{min}, \lambda_2^{min})$ pair could lead to all nodes being selected for all subsamples, thus $max_{(\lambda 1, \lambda 2)} \boldsymbol{\pi}_i(\lambda_1, \lambda_2) = 1$. Hence, all nodes would be declared as significant. Since averaging the sample covariance matrices is equivalent to computing the sample covariance of their associated time series concatenated together, even with e.g. 20 subjects, $n/2$ would be in the thousands range, whereas d is typically in the hundreds range. Thus, to mitigate declaring all nodes as relevant, we set λ_1^{min} and λ_2^{min} such that <50% of the nodes would be selected by (4). λ_1^{max} and λ_2^{max} are set such that only ~1% of the nodes are selected.

3 Materials

Two fMRI datasets from a within-group and a between-group study were used for method evaluation. For the within-group case, fMRI data were collected from 19 healthy subjects as they think about joyful events (happy state) and sad events (ruminative state) in a self-driven manner without external stimulus. The data of the two mental states were acquired at two separate scan sessions, each lasting 8 min. Data acquisition was performed on a 3T GE scanner with TR = 2 s, TE = 30 ms, and flip angle = 77°. For each mental state, the fMRI data of each subject were motion corrected, normalized to MNI space, and spatially smoothed with at 6 mm FWHM Gaussian kernel using FSL. Motion artifacts, white matter and cerebrospinal fluid confounds, and average global signals were regressed out from the voxel time series. A highpass filter at 0.01 Hz was subsequently applied to remove scanner drifts. To define brain nodes, we employed the atlas in [8], which comprises 90 functionally-defined regions that span 14 widely-observed networks. Regions in the cerebellum were excluded due to incomplete coverage. Voxel time series within each region were averaged to generate brain region time series. Pearson's correlation matrices were then computed from these regional time series, which served as input to NBGGM and the contrasted graph measures. For the between-group case, resting state fMRI data of 6 min duration were collected from 20 AD subjects and 20 matched healthy controls (HC) with a similar acquisition protocol. The same preprocessing steps were performed except a bandpass filter at 0.01 to 0.1 Hz was employed. Also, all 90 regions in [8] were used to generate brain region time series.

4 Results and Discussion

In this work, we focused on method assessments with real fMRI data since substantive synthetic experiments had been performed in [5]. To quantitatively evaluate NBGGM on real data, we randomly split the subjects into two halves (S_1 and S_2) in 20 different ways for each dataset, and examined the consistency of the identified brain regions between S_1 and S_2. Consistency was estimated using the Dice coefficient (DC) = $2|\mathbf{z}_1 \cap \mathbf{z}_2|/(|\mathbf{z}_1|+|\mathbf{z}_2|)$, where $\mathbf{z}_1$ and $\mathbf{z}_2$ are $d\times 1$ vectors with 1 indicating the corresponding brain regions were selected for S_1 and S_2, respectively, and 0 otherwise. Here, $\cap$ is the intersection operator and $|\cdot|$ denotes cardinality. For comparisons, we computed weighted versions of ND, CC, and BC [3] with negative Pearson's correlations set to 0, as required for computing these graph measures. These measures are widely used for node characterization. Specifically, ND is the number of edges connected to a node, CC is the fraction of neighbors of a node being neighbors of each other, and BC is the fraction of all shortest paths that contain a given node [3]. To identify discriminative brain regions with each graph measure, we applied a paired t-test and a two-sample test, respectively, in contrasting happy vs. ruminative and AD vs. HC. No brain regions were found to be significant at $p < 0.05$ with false discovery rate (FDR) correction [9] for S_1 and S_2 of all 20 random splits. To permit comparisons with NBGGM, we took the top regions for each graph measure based on the magnitude of the t-values, with the number of regions matching that found with NBGGM on S_1 and S_2 of each split, and computed the DC. We also tried applying stability selection to the graph measures, but the selection frequency was <0.8 (threshold used for NBGGM) for almost all regions, resulting in DC being close to 0, hence not presented here.

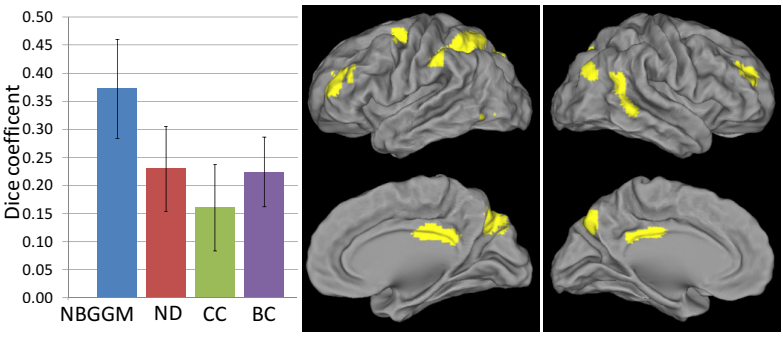

Fig. 2. Happy vs ruminative. NBGGM provided significantly higher consistency in the identified brain regions than the contrasted methods. Regions found by NBGGM shown in yellow.

Within-Group Analysis. The average DC over the 20 splits for happy vs. ruminative are shown in Fig. 2. Comparing NBGGM against each graph measure in a pairwise fashion, NBGGM's average DC is significantly higher based on a Wilcoxon signed rank test at $p < 0.05$ with Bonferroni correction. Applying NBGGM on all subjects detected brain regions associated with emotion processing (Fig. 2), whereas no brain regions were significant using the contrasted graph measures at $p < 0.05$ with FDR correction.

In particular, the left frontal pole is associated with happiness and sadness, and the right frontal pole tends to react to sadness, anger, and fear [10]. The cingulate cortex is involved with processing of anger, fear, sadness, and happiness [11], with the posterior division mainly involved with happy emotions. NBGGM also detected the precuneus, which is responsible for emotional judgment of self and others [12].

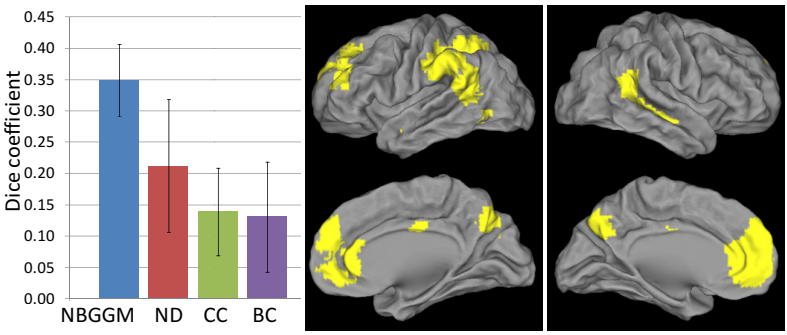

Fig. 3. AD vs HC. NBGGM provided significantly higher consistency in the identified brain regions than the contrasted methods. Regions found by NBGGM shown in yellow.

Between-Group Analysis. The average DC over the 20 splits for AD vs. HC are shown in Fig. 3. NBGGM's average DC is again significantly higher than the contrasted graph measures based on a Wilcoxon signed rank test at $p < 0.05$ with Bonferroni correction. Applying NBGGM on all subjects detected brain regions associated with AD (Fig. 3), whereas none of the graph measures found any significant brain regions at $p < 0.05$ with FDR correction. Brain regions found by NBGGM included the precuneus, the posterior cingulate cortex, and the medial prefrontal cortex, which are known to be affected in AD [13], with memory decline being the primary symptom. Other brain regions included the bilateral angular gyri, which are involved with memory retrieval [14]. The bilateral angular gyri, in addition, belong to the language network, and their detection along with the lateral occipital cortex as well as the supramarginal gyrus, which are part of the visuospatial network, might be linked to specific cognitive impairments in AD beyond the predominant amnestic variant, namely primary progressive aphasia and posterior cortical atrophy, respectively [15].

5 Conclusions

We proposed incorporating stability selection into NBGGM to extend its application to the identification of discriminative nodes from two sets of weighted graphs. A key advantage of NBGGM is that it does not require thresholding the weighted graphs and more holistically uses the connectivity information than popular node-based graph measures, such as weighted ND, CC, and BC that capture only local graph information. We showed on real fMRI data that NBGGM provides significantly higher consistency in discriminative node identification than ND, CC, and BC. Further, the regions identified by NBGGM conformed well to prior neuroscience knowledge.

Acknowledgements. Bernard Ng is supported by the Lucile Packard Foundation for Children's Health, Stanford NIH-NCATS-CTSA UL1 TR001085 and Child Health Research Institute of Stanford University.

References

1. Li, K., Guo, L., Nie, J., Li, G., Liu, T.: Review of methods for functional brain connectivity detection using fMRI. Comput. Med. Imaging Graph. **33**, 131–139 (2009)
2. Friston, K., Holmes, A., Worsley, K., Poline, J.B., Frith, C., Frackowiak, R.: Statistical parametric maps in functional imaging: A general linear approach. Human Brain Mapp. **2**, 189–210 (1995)
3. Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: uses and interpretations. NeuroImage **52**, 1059–1069 (2010)
4. Venkataraman, A., Kubicki, M., Golland, P.: From connectivity models to region labels: identifying foci of a neurological disorder. IEEE Trans. Med. Imaging **32**, 2078–2098 (2013)
5. Mohan, K., London, P., Fazel, M., Witten, D., Lee, S.: Node-based learning of multiple Gaussian graphical models. J. Mach. Learn. Research **15**, 445–488 (2014)
6. Meinshausen, N., Bühlmann, P.: Stability Selection. J. Roy. Statist. Soc. Ser. B **72**, 417–473 (2010)
7. Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. J. Royal Stat. Soc. Series B **68**, 49–67 (2006)
8. Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D.: Decoding subject-driven cognitive states with whole-brain connectivity patterns. Cereb. Cortex **22**, 158–165 (2012)
9. Nichols, T., Hayasaka, S.: Controlling the familywise error rate in functional neuroimaging: A comparative review. Stat. Methods Med. Research **12**, 419–446 (2003)
10. Damasio, A.R., Grabowski, T.J., Bechara, A., Damasio, H., Ponto, L.L., Parvizi, J., Hichwa, R.D.: Subcortical and cortical brain activity during the feeling of self-generated emotions. Nat. Neurosci. **3**, 1049–1056 (2000)
11. Vogt, B.A.: Pain and emotion interactions in subregions of the cingulate gyrus. Nat. Rev. Neurosci. **6**, 533–544 (2005)
12. Ochsner, K.N., Knierim, K., Ludlow, D.H., Hanelin, J., Ramachandran, T., Glover, G., Mackey, S.C.: Reflecting upon feelings: an fMRI study of neural systems supporting the attribution of emotion to self and other. J. Cogn. Neurosci. **16**, 1746–1772 (2004)
13. Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V.: Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. Proc. Natl. Acad. Sci. USA **101**, 4637–4642 (2004)
14. Seghier, M.L.: The angular gyrus: Multiple functions and multiple subdivisions. Neuroscientist **19**, 43–61 (2013)
15. Ranasinghe, K.G., Hinkley, L.B., Beagle, A.J., Mizuiri, D., Dowling, A.F., Honma, S.M., Finucane, M.M., Scherling, C., Miller, B.L., Nagarajan, S.S., Vossel, K.A.: Regional functional connectivity predicts distinct cognitive impairments in Alzheimer's disease spectrum. Neuroimage Clin. **23**, 385–395 (2014)