

Defining Patient Specific Functional Parcellations in Lesional Cohorts via Markov Random Fields

Naresh Nandakumar^{1(⊠)}, Niharika S. D'Souza¹, Jeff Craley¹, Komal Manzoor², Jay J. Pillai², Sachin K. Gujar², Haris I. Sair², and Archana Venkataraman¹

¹ Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, USA nnandak1@jhu.edu artment of Neuroradiology Johns Hopkins School of Medi

² Department of Neuroradiology, Johns Hopkins School of Medicine, Baltimore, USA

Abstract. We propose a hierarchical Bayesian model that refines a population-based atlas using resting-state fMRI (rs-fMRI) coherence. Our method starts from an initial parcellation and then iteratively reassigns the voxel memberships at the subject level. Our algorithm uses a maximum a *posteriori* inference strategy based on the neighboring voxel assignments and the Pearson correlation coefficients between the voxel time series and the parcel reference signals. Our method is generalizable to different initial atlases, ensures spatial and temporal contiguity in the final network organization, and can handle subjects with brain lesions, whose rs-fMRI data varies tremendously from that of a healthy cohort. We validate our method by comparing the intra-network cohesion and the motor network identification against two baselines: a standard functional parcellation with no reassignment and a recently published method with a purely data-driven reassignment procedure. Our method outperforms the original functional parcellation in intra-network cohesion and both methods in motor network identification.

Keywords: Rs-fMRI \cdot Markov Random Field Patient-specific networks

1 Introduction

Resting-state fMRI (rs-fMRI) captures the intrinsic communication patterns in the brain. Neural activity at rest is organized into Resting State Networks (RSNs), which are determined by both spatial coherence and temporal synchrony at the voxel level [1]. Non-invasive methods for identifying RSNs are of particular interest when planning neurosurgery, where the goal is to localize (and subsequently avoid) cruicial motor and language functionality, also known as the eloquent cortex. Accurately identifying these key functional networks will determine safe resection margins and may also help us to predict the functional

[©] Springer Nature Switzerland AG 2018

G. Wu et al. (Eds.): CNI 2018, LNCS 11083, pp. 88–98, 2018. https://doi.org/10.1007/978-3-030-00755-3_10

outcomes of surgery [2]. Task-fMRI, where the subject is performing an explict language or motor paradigm, is currently the most popular technique for mapping eloquent brain areas. However, recent interest has moved towards rs-fMRI to overcome both the unreliability of certain task activations and the inability of critically ill patients to complete a cognitively demanding protocol [2].

While a standard functional parcellation is often sufficient to identify RSNs in healthy individuals, a subject specific approach is necessary for brain lesion patients due to a compensatory mechanism known as neural plasticity. For example, it has been shown that motor and language functionality in patients with low-grade gliomas can migrate to other parts of the brain [3]. Therefore, data from glioma patients may not conform to a standard functional atlas.

Previous work has focused on deriving subject-specific functional atlases from the original rs-fMRI data. These techniques include, but are not limited to, spatially constrained hierarchical parcellations [4] and data-driven clustering techniques [5]. In contrast, we approach this problem as one of *atlas refinement* rather than atlas construction. This strategy allows us to work with *any* previously validated anatomical or functional parcellation. Liu et. al have introduced a method to obtain subject-specific RSNs by iteratively reassigning the voxel memberships based on the Pearson correlation coefficients between the voxel time series and a collection of network reference signals. This method was originally derived to determine functional network differences in healthy subjects and showed promise for clinical populations [6]. These reference signals are calculated as a weighted average of the previous iteration's reference signal and the current average time series for the RSN. Their method does not explicitly consider spatial contiguity in RSNs, and requires the user to specify various parameters.

We develop a Bayesian model that uses both spatial and temporal information to iteratively refine an initial functional parcellation on a patient-specific basis. Our model uses a Markov Random Field (MRF) prior to encourage spatial contiguity within the functional parcels. We employ a maximum *a posteriori* (MAP) inference strategy for voxel-wise network assignment until a predefined convergence criteria is met. Our method builds on prior work in Bayesian network modeling [7] and MRF priors for rs-fMRI data [8]. We validate our method on rs-fMRI data from 67 glioma patients. Our initial atlas is the Yeo 17 network functional parcellation [9], which is one of the most widely cited functional atlases in the literature. We compare the performance of our method with the original parcellation (no reassignment) and with reassignment according to Liu, whose paper references the same parcellation. Our validation metrics include the intra-network cohesion amongst the final RSNs and motor network identification via task based fMRI concordance using three distinct task paradigms.

2 A Bayesian Model for Voxel Reassignment

Our model infers an underlying (i.e. latent) network architecture that integrates both spatial contiguity and temporal synchrony across the brain. At each voxel, we leverage the time series data, the neighborhood network membership, and a



Fig. 1. A graphical model of our framework where shaded nodes represent observed random variables.

binary tumor label, indicating if the voxel lies within the lesion or not. Let X_v be the network assignment for voxel v. In our framework, X_v gets assigned to one of K+1 values, where K is the number of networks (or region parcels) defined in the initial atlas. An assignment of $X_v = 0$ indicates no network membership for voxel v if it belongs to the glioma. Mathematically, let \mathbf{X}_{-v} be the current network assignments of all other voxels in the brain. Likewise, \mathbf{y}_v is the time series data at voxel v, $\boldsymbol{\mu}_k$ is the current reference signal for network $k \in \{1...,K\}$, and $b_v \in \{0,1\}$ is the binary tumor label such that $b_v = 0$ implies that voxel vis tumorous. Our setup is illustrated in Fig. 1. As seen, the assignment for X_v depends on its immediate spatial neighbors. The relationship between X_v and \mathbf{X}_{-v} is captured by an MRF prior while the relationship between X_v and \mathbf{y}_v, b_v is captured by the data likelihood. For visualization purposes the 2D representation in Fig. 1 shows four neighbors per pixel. However, we have implemented a 3D model, which has six neighbors per voxel.

We encourage spatial contiguity in our latent network assignments by stipulating that voxel v will be more likely to assume the state of its neighboring voxels. We model the MRF prior after the Potts model [10]:

$$P(X_v = k | \mathbf{X}_{-v}) = \frac{1}{Z_x} \Psi(X_v, \mathbf{X}_{-v}, k) \propto \left\{ 1 + \exp\left[-\left(\beta + \sum_{i \in ne(v)} \mathbb{1}_{X_i = k}\right) \right] \right\}^{-1}$$
(1)

where β controls the influence of the neighbor voxel network memberships on voxel v. Here, ne(v) denotes the neighbors of voxel v, and the sum $\sum_{i \in ne(v)} \mathbb{1}_{X_i=k}$ captures how often these neighbors are assigned to network k. Notice that this sum will be zero for network k when $X_i \neq k$ for all $i \in ne(v)$.

The likelihood $P(\mathbf{y}_v|X_v = k; \boldsymbol{\mu}_k, b_v)$ is modeled after a rescaled version of the Pearson correlation coefficient between the reference $\boldsymbol{\mu}_k$ and data \mathbf{y}_v :

$$\rho = \frac{\operatorname{cov}(\boldsymbol{y}_v, \boldsymbol{\mu}_k)}{\sigma_{\boldsymbol{y}_v} \sigma_{\boldsymbol{\mu}_k}} \tag{2}$$

where ρ is subsequently shifted and scaled to be between [0, 1] to allow for a normalizable density. The final likelihood model is given by

$$P(\boldsymbol{y}_{v}|X_{v}=k;\boldsymbol{\mu}_{k},b_{v}) = \frac{1}{Z_{y}}\Psi(\boldsymbol{y}_{v},\boldsymbol{\mu}_{k},b_{v}) = \frac{1}{Z_{y}}\left\{\frac{(\rho_{\boldsymbol{y}_{v},\boldsymbol{\mu}_{k}}+1)}{2}\times b_{v}\right\}.$$
 (3)

The rescaled Pearson correlation coefficient goes to one for strong positive correlations and zero for strong negative correlations. The label b_v sets the likelihood to zero for tumorous voxels, which corresponds to no network membership.

2.1 Approximate Posterior Inference

The observed rs-fMRI time series y_v are conditionally independent given $\{X_v\}$. Based on the model factorization, our posterior distribution can be written as

$$P(X_v = k | \boldsymbol{X}_{-v}, \boldsymbol{y}_v; \theta) = \frac{1}{Z} \Psi(X_v, \boldsymbol{X}_{-v}, k) \Psi(\boldsymbol{y}_v, \boldsymbol{\mu}_k, b_v)$$
(4)

where $\Psi(X_v, \mathbf{X}_{-v}, k)$ models the prior, $\Psi(\mathbf{y}_v, \boldsymbol{\mu}_k, b_v)$ models the likelihood under the belief that $X_v = k$, and Z is a normalization constant that combines both Z_x and Z_y . We have derived an update procedure based on maximizing the following log-posterior over all possible network assignments:

$$X_{v}^{*} = \operatorname*{argmax}_{k} \Big\{ -\log Z + \log \Psi(X_{v}, \boldsymbol{X}_{-v}, k) + \log \Psi(\boldsymbol{y}_{v}, \boldsymbol{\mu}_{k}, b_{v}) \Big\}.$$
(5)

2.2 Implementation Details

We have derived an algorithm based on max product message passing to ensure atlas stability [11]. Our algorithm iterates between two main steps: updating the network assignments $\{\boldsymbol{X}_v\}$ and updating the reference signals $\{\boldsymbol{\mu}_k\}$. Let $\boldsymbol{Y} \in \mathbb{R}^{x \times y \times z \times T}$ be the aggregated rs-fMRI data across all (x, y, z) spatial coordinates and let $\boldsymbol{X}^{(t)} \in \mathbb{R}^{x \times y \times z}$ be the assignment information at iteration t. Let $\boldsymbol{B} \in \mathbb{R}^{x \times y \times z}$ be the binary tumor matrix. The values stored in \boldsymbol{B} are 0 at tumorous voxels and 1 elsewhere. We initialize our algorithm with the Yeo atlas and then the Hadamard product $\boldsymbol{X}^{(0)} = \boldsymbol{X} \odot \boldsymbol{B}$, which defaults all tumorous voxel assignments to 0 due to unreliable signal at these locations. We then parcellate \boldsymbol{Y} by the assignments in $\boldsymbol{X}^{(0)}$ and calculate the initial reference signals $\boldsymbol{\mu}_k^{(0)}$ for $k \in \{1 \dots K\}$ with

$$\boldsymbol{\mu}_{k}^{(t)} = \frac{\sum_{v=1}^{V} \boldsymbol{Y}_{v} \cdot \mathbb{1}(\boldsymbol{X}_{v}^{(t)} = k)}{\sum_{v=1}^{V} \mathbb{1}(\boldsymbol{X}_{v}^{(t)} = k)}.$$
(6)

At each main iteration t, we determine the voxel assignments I times according to our MAP rule in Eq. (5) initializing with $\mathbf{X}^{(t-1)}$. Using the assignments in i-1, we employ a flooding schedule to simultaneously determine the network values $\widehat{\mathbf{X}}^{(i)}$ at iteration i. The updated assignments $\mathbf{X}^{(t)}$ is given by majority

Algorithm 1. Max product message passing procedure for atlas refinement. Here, c is a prespecified threshold on the network consistency between iterations.

1:	procedure MRFREFINEMENT(X, I	$(oldsymbol{B},oldsymbol{Y})$
2:	$\boldsymbol{X}^{(0)} \gets \boldsymbol{X} \odot \boldsymbol{B}$	
3:	$ig\{oldsymbol{\mu}_1^{(0)}\cdotsoldsymbol{\mu}_K^{(0)}ig\}\leftarrowoldsymbol{Y},oldsymbol{X}^{(0)}$	\triangleright Eq. (6)
4:	$t \leftarrow 1$	
5:	$M \leftarrow 0$	$\triangleright \text{ Membership retention } M \in [0, 1]$
6:	while $M < c$ do	$\triangleright \text{ Convergence threshold } c \in [0, 1]$
7:	$\widehat{\boldsymbol{X}}^{(1)} \gets \boldsymbol{X}^{(t-1)}$	
8:	for $i = 2: I$ do	
9:	for $v \in V$ do	`
10:	$\widehat{oldsymbol{X}}_v^{(i)} \gets \operatorname{argmax}_k \Big\{ \log$	$\Psi(X_v, \widehat{\boldsymbol{X}}_{-v}^{(i-1)}, k) + \log \Psi(\boldsymbol{y}_v, \boldsymbol{\mu}_k^{(t-1)}, b_v) \bigg\}$
11:	$\boldsymbol{X}^{(t)} \gets mode(\{\widehat{\boldsymbol{X}}^{(i)}\}_{i=1}^{I})$	
12:	$ig\{oldsymbol{\mu}_1^{(t)}\cdotsoldsymbol{\mu}_K^{(t)}ig\}\leftarrowoldsymbol{Y},oldsymbol{X}^{(t)}$	\triangleright Eq. (6)
13:	$M \leftarrow \frac{\sum_{v} \mathbb{1}(X_{v}^{(t-1)} = X_{v}^{(t)}) \cdot B_{v}}{\sum_{v} B_{v}}$	\triangleright Fraction of retained voxel memberships
14:	$t \leftarrow t+1$	
15:	$\mathbf{return}\; \boldsymbol{X}^t$	

vote over the determined network values $\{\widehat{X}^{(i)}\}_{i=1}^{I}$. Given the new assignments, we update the network signals by Eq. (6) for $k \in \{1 \dots K\}$ and check if the convergence criteria has been met by calculating the fraction of non-zero assigned voxels retaining the same network membership between iterations t-1 and t. If the membership consistency between iterations is less than a specified stopping criteria, we repeat the procedure. Each voxel of interest has six neighbors as determined by adjacency in each of the three coordinate directions. Algorithm 1 presents our pseudo-code where the subject-specific inputs are B and Y.

2.3 Baseline Comparisons

We compare our Bayesian approach with the original parcellation and with the voxel reassignment method described by Liu [6]. We use the Yeo 17 network atlas due to its strong reproducibility and the large sample size used for construction. We confine our experiments to the more conservative cortical ribbon version of the Yeo atlas to get a more detailed parcellation. The method of Liu initializes the reference signals to the average time series defined by the original parcellation. From here, Liu reassigns voxel v by considering the maximum correlation between its time series and all K reference signals. A confidence value for each voxel is also computed as the ratio of the maximum correlation over the second highest correlation. The reference signal updates are only taken from voxels that have confidence values which exceed a predetermined threshold. They are computed as weighted combinations of the previous iteration's reference signals with the updated reference signals. The corresponding weights are nonlinear functions of the signal-to-noise ratio, the inter-subject variability, and the iteration number. We applied the Liu baseline with the parameter suggestions provided in [6],

which were optimized for the 17 network Yeo atlas. We initialize each method, original, Liu, and proposed as described in Sect. 2.2.

3 Experimental Results

Dataset and Preprocessing: Our dataset includes task and rs-fMRI for 67 glioma patients who underwent preoperative mapping as part of their clinical workup. The fMRI were acquired using a 3.0 T Siemens Trio Tim (TR = 2000 ms, TE = 30 ms, flip angle = 90°, field of view = 24 cm, acquisition matrix = $64 \times 64 \times 33$, slice thickness = 4 mm, and slice gap of 1 mm). We manually segmented each patient's tumor using MIPAV. Figure 2 illustrates tumor segmentations for three different patients to motivate the heterogeneity of our cohort. The fMRI was processed using SPM8. Both rs-fMRI and task-fMRI underwent slice timing correction, motion correction and registration to the MNI-152 template. The rs-fMRI was scrubbed using the ArtRepair toolbox in SPM, linearly detrended, underwent nuisance regression utilizing CompCor [12], bandpass filtered from 0.01 to 0.1 Hz, and spatially smoothed with a 6 mm FWHM Gaussian kernel.

General Linear Model (GLM) in SPM8 was used to derive activation maps for the three motor tasks [2]. Our dataset includes three different motor paradigms that were designed to target distinct parts of the motor homonculus [13]: finger tapping, tongue moving, and foot tapping. Since the task-fMRI data was acquired for clinical purposes, only 42 patients performed the finger task, 35 patients performed the tongue task, and 20 patients performed the foot task.

The population-based atlas contains 17 distinct functional networks confined to the cortical ribbon [9]. For both methods, a network retention convergence criteria of 0.98 was used. We chose $\beta = -0.5$ and I = 100 iterations for our model and a confidence value of 1.5 for the Liu baseline. Different combinations of reference signal calculation between updates for both our method and the Liu baseline were explored; we have reported the optimal results in each case.

3.1 Evaluating Resting State Network Cohesion

Our intra-network cohesion metric quantifies the temporal synchrony between voxels that belong to the same network [1]. Let V_k be the voxels assigned to



Fig. 2. The tumor segmentations (yellow) for three different patients are shown. (Color figure online)

network k, we define the Network Cohesion (NC) as the average correlation between voxels assigned to network k with the network signal μ_k .

$$NC_k = \frac{\sum_{j \in V_k} \rho_{\boldsymbol{y}_j, \boldsymbol{\mu}_k}}{|V_k|} \tag{7}$$

Figure 3 illustrates the difference in NC between our proposed method and both the original parcellation (left) and the Liu baseline (right). A value greater than zero is considered to be more temporally synchronous while a value less than zero is considered to be less temporally synchronous. In all 17 networks, our method outperforms the original atlas with significance p < 0.005. This highlights the importance of our subject-specific approach for glioma patients, whose functional networks are substantially reorganized due to tumor presence.



Fig. 3. Difference in intra-network cohesion between our method and the original parcellation (left) and the Liu baseline (right).

Naturally, the Liu baseline achieves higher NC due to its correlation-based voxel reassignment procedure. Figure 4 shows the original parcellation, and the final network assignment using our method and Liu's method in a single patient. Each distinct color represents one of the 17 networks. We observe an overall lack of spatial contiguity in the Liu baseline, as highlighted in the white circle. This might be due to spurious noise within rs-fMRI signal at the voxel level, resulting in some spatially discontiguous reassignment. The large grey area in the right hemisphere is the excluded tumorous region for this subject.

Figure 5 shows the proportion of voxels retained in the original network membership between our method (left) and the Liu baseline (right). We observe substantial reorganization in the networks defined from our method. Along with higher NC, this further motivates our approach, showing that many voxels in the original parcellation may not belong to the proper RSN for this cohort. We observe an even larger reorganization in the Liu networks. In the following section we conjecture that the displacement in the Liu networks may be too large, because while the Liu baseline provides more temporally cohesive RSNs, it fails to identify functionally consistent motor networks.



Fig. 4. Left: Original network assignment. Middle: Our final network assignment. Right: Liu's final network assignment. For visualization, we have dilated the networks according to the liberal Yeo mask.



Fig. 5. Network retention for our method (left) and Liu's method (right).

3.2 Motor Network Concordance, as Validated by Task-fMRI

Our second experiment quantifies the rs-fMRI concordance between the pseudoground truth motor network in each patient and the motor RSN identified by each of the methods. Specifically, we will use the GLM activation map across three distinct motor tasks to define seed locations for motor functionality. The seed is defined as a group of highly activated voxels within the activation map. The Yeo atlas separates the motor network into two different parcels [9]. Our measure of task concordance will be the maximum correlation between the reference signals of these two RSNs and the average time series associated with the GLM activation seed. We determine that a method is better at motor network identification by having a higher positive correlation with significance p < 0.05.

Figure 6 illustrates the performance gain of our method. The pink boxplots show the difference in task concordance between our method and the original atlas, while the blue shows the difference in task concordance between our method and the Liu baseline. The tasks are ordered as finger, tongue, and foot from left to right. Table 1 summarizes the results and corresponding p-values for this experiment. The values in bold show when our method outperforms other methods with a student t-test with significance threshold $\alpha = 0.05$.



Fig. 6. Difference in task concordance between our method and both the original atlas (pink) and the Liu baseline (blue). Our method achieves significantly better performance in five out of the six comparisons. (Color figure online)

Task	Sample size	Ours vs. Original	Ours vs. Liu
Finger	42	$3.6\mathrm{e}{-3}$	$1.2\mathrm{e}{-5}$
Tongue	35	$7.0\mathrm{e}{-3}$	$3.7\mathrm{e}{-2}$
Foot	20	0.45	$9.8\mathrm{e}{-5}$

Table 1. P-values for our method vs. the original atlas and the Liu baseline.

Our method outperforms the Liu method in each of the three tasks. In addition, our method performs better than the original atlas in the finger and tongue task, but not the foot task. This latter result can be due to the local area of the motor homonculus that foot activaton lies in [13] or the smaller sample size. By observing p-values reported for the finger and foot task, we conclude that no reassignment would be preferrable to the Liu baseline in this experiment. However, the Liu method RSNs were the most temporally cohesive. Though network cohesion is a desirable property for RSNs [1], we have demonstrated that higher cohesion does not always lead to a functionally consistent motor network. We conjecture that (1) Liu is too liberal in the voxel reassignment, and (2) both spatial and temporal consistency are required for RSN identification.

In summary, our method balances both spatial contiguity with temporal synchrony to help describe functional networks in patients who have undergone localized neural plasticity. We observe that our method shows more cohesive RSNs for tumor patients than a population-based functional atlas. We also determine that the motor network refined by our method is a closer representation to the actual motor network in these patients. This combination of results give us confidence in our method for characterizing RSNs in a lesional population.

4 Conclusion

We have formulated a Bayesian model that can refine a population atlas on a patient-specific basis. Our model considers both spatial contiguity as well as temporal synchrony between voxels, all while handling large and variable brain lesions. Our method outperforms established baselines for identifying a functionally consistent motor network. The use of the MRF prior along with iterative voxel reassignment shows a viable balance between properties of interest in resulting RSNs. These methodological improvements broaden the applications in which one can use rs-fMRI for analysis. We have generated a method that can be translated to other patient cohorts with anatomical brain lesions, like stroke or traumatic brain injury. Our performance in assessing RSN cohesion shows that our method captures subject-specific functional organization well, even in a pathological population. Our method outperforms both baselines in terms of motor network identification, which is an important step for preoperative planning for neurosurgical resections to avoid permanent motor network damage.

Future work with our method will involve different initial atlases. Specifically, we aim to observe how our method performs with atlases of different network numbers, and different initial size (voxel membership) of networks. Methodologically, we aim to vary the number of neighbors considered in our prior model, assigning varying weights to neighbors of different geodisic distances from the center voxel. Clinically, we aim to study reorganization of sites near the glimoa, which is known to show the most neural plasticity in these patients [3].

Acknowledgements. This work was supported by the RSNA Research & Education Foundation Carestream Health RSNA Research Scholar Grant; Contract grant number: RSCH1420.

References

- 1. Lee, W.H., Frangou, S.: Linking functional connectivity and dynamic properties of resting-state networks. Sci. Rep. 7(1), 16610 (2017)
- Sair, H.I., et al.: Presurgical brain mapping of the language network in patients with brain tumors using resting-state fMRI: comparison with task fMRI. Hum. Brain Mapp. 37(3), 913–923 (2016)
- Duffau, H.: Lessons from brain mapping in surgery for low-grade glioma: insights into associations between tumour and brain plasticity. Lancet Neurol. 4(8), 476– 486 (2005)
- Blumensath, T., Jbabdi, S., et al.: Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. Neuroimage 76, 313–324 (2013)
- Craddock, R.C., et al.: A whole brain fMRI atlas generated via spatially constrained spectral clustering. Hum. Brain Mapp. 33(8), 1914–1928 (2012)
- Wang, D., et al.: Parcellating cortical functional networks in individuals. Nat. Neurosci. 18(12), 1853 (2015)
- Venkataraman, A., et al.: Bayesian community detection in the space of group-level functional differences. IEEE Trans. Med. Imaging 35(8), 1866–1882 (2016)
- Ryali, S.: A parcellation scheme based on von Mises-fisher distributions and Markov random fields for segmenting brain regions using resting-state fMRI. NeuroImage 65, 83–96 (2013)

- Thomas, B., Yeo, B.T., et al.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 106(3), 1125–1165 (2011)
- Wainwright, M.J., Jordan, M.I., et al.: Graphical models, exponential families, and variational inference. Found. Trends® Mach. Learn. 1(1-2), 1-305 (2008)
- 11. Kschischang, F.R., Frey, B.J., Loeliger, H.-A.: Factor graphs and the sum-product algorithm. IEEE Trans. Inf. Theory 47(2), 498–519 (2001)
- Behzadi, Y., et al.: A component based noise correction method (compcor) for bold and perfusion based fMRI. Neuroimage 37(1), 90–101 (2007)
- Jack Jr., C.R., et al.: Sensory motor cortex: correlation of presurgical mapping with functional MR imaging and invasive cortical mapping. Radiology 190(1), 85–92 (1994)