

Multi-object Spring Level Sets (MUSCLE)

Blake C. Lucas^{1,2}, Michael Kazhdan², and Russell H. Taylor²

¹ Johns Hopkins Applied Physics Laboratory, Laurel, MD, USA

² Johns Hopkins University, Baltimore, MD, USA

{blake,misha}@cs.jhu.edu, rht@jhu.edu

Abstract. A new data structure is presented for geometrically modeling multi-objects. The model can exhibit elastic and fluid-like behavior to enable interpretability between tasks that require both deformable registration and active contour segmentation. The data structure consists of a label mask, distance field, and springls (a constellation of disconnected triangles). The representation has sub-voxel precision, is parametric, re-meshes, tracks point correspondences, and guarantees no self-intersections, air-gaps, or overlaps between adjacent structures. In this work, we show how to apply existing registration algorithms and active contour segmentation to the data structure; and as a demonstration, the data structure is used to segment cortical and subcortical structures (74 total) in the human brain.

Keywords: registration, segmentation, tracking, active contour, level set, mesh.

1 Introduction

Deformable registration [1, 2] has become a popular technique for reliably and automatically segmenting multiple objects with little prior knowledge of anatomy or imaging technology. The output of the algorithm is a displacement field describing where voxel locations in the target image map to in the source image. This 3D-to-3D mapping allows any geometric structures identified in the source image to be warped into the target image to create a segmentation of the target [3]. To overcome the registration algorithm's lack of prior knowledge, there is a strong assumption that the target's anatomy is a smooth elastic deformation of the source's anatomy. This is rarely the case, leading to what is known as "atlas bias" where the warped source image still resembles the source image, possibly more than the target. One way to reduce bias is to repeat the registration with different source images and combine segmentations afterwards to create an "average" [4]. Methods that combine multiple segmentations produce "average" segmentations that are smoother than any individual segmentation [5]. Alternatively, one can reduce bias and increase the image segmentation's fidelity by following elastic registration with "fluid-like" segmentation [6, 7]. Segmentation reduces atlas bias from the registration phase because the deformation is driven by information in the target image alone. Traditional geometric data structures used to cascade registration and segmentation either lose information and / or do not preserve important geometric properties. We now review those data structures and the challenges each one poses.

2 Previous Work

Meshes. Triangle meshes have the ability to encode surface labels and track point correspondences. Displacement fields produced by registration can be directly applied to triangle meshes. Even though displacement fields represent smooth elastic deformations, the deformed triangle mesh can have sharp edges and cusps that were not present in the original mesh. To reduce artifacts, the mesh must be smoothed and re-sampled (re-meshed) to produce a higher quality mesh. Re-meshing is challenging and interferes with point correspondences in non-intuitive ways [8]. Re-meshing becomes more difficult when deformations are fluid-like, objects share boundaries or slide against each other, and practical geometric constraints are enforced. These constraints include: the mesh should not intersect adjacent structures or itself, and it should not create air-gaps between adjacent structures.

Level Sets. Objects can be represented with sub-voxel precision as level sets [9]. Level sets are functions of 3D space that are stored as images. A common choice for level set function is the signed distance field; in which, image intensities measure the minimum distance to the object's boundary. Level sets naturally enforce that structures do not self-intersect, and it is simple to enforce that they do not overlap or have air-gaps [10]. The drawback to level sets is that they do not maintain point correspondences or surface labels, and application of a displacement field requires re-sampling the level set image, which acts as a low-pass filter by smoothing underlying geometric structures.

Label Masks. It is common for objects to be manually segmented with a painting tool to produce binary masks. These masks are then merged together to create a label mask image. Displacement fields can be applied to labels by warping and re-sampling the label mask image. Re-sampling acts as a low-pass filter that smoothes underlying structures; and because label image values are region indicators, only nearest-neighbor interpolation can be used. Label masks have only voxel precision and have no ability to maintain point correspondences or surface labels.

Overview. Recently, a geometric data structure was presented that has properties of both meshes and level sets [11]. Spring Level Sets (SpringLS) couple a constellation of disconnected triangle surface elements (springls) with a level set. This work extends SpringLS to the multi-object case and shows how to use it for registration and segmentation. The new data structure is a combination of three existing data structures: a constellation of springls, a label mask, and a distance field. The new Multi-Object Level Set (MUSCLE) data structure is simultaneously all previously mentioned representations and addresses the drawbacks of each. To demonstrate, we cascade diffeomorphic Demons [12] registration and active contour segmentation [9] to segment and label 74 structures in MR images of the human brain.

3 Method

Representation. The Springl Constellation + Label Mask + Distance Field (MUSCLE) data structure is depicted in Fig. 1. A springl \mathcal{S}_n is a triangular surface

element consisting of a particle p_n and three springs connecting the particle to each of the triangle's vertices $q_{n,m}$ (Fig. 1a). Each spring maintains an object label $l_n \in \mathcal{L} = \{1, \dots, L\}$ and a correspondence point a_n that maps the particle p_n to a point on the original model. Each object has its own spring constellation that encloses the object's interior. Constellations for all objects (Fig. 1b) are merged into one constellation for storage and manipulation. The label mask $\chi: \Omega \mapsto \mathcal{L}$ (Fig. 1c) maps each voxel in the image domain $\Omega \subset \mathbb{R}^3$ to an object label l . The distance field $\psi: \Omega \mapsto \mathbb{R}$ (Fig. 1d) measures the distance of each voxel to the nearest object represented by level sets $\varphi_l: \Omega \mapsto \mathbb{R}$ (i.e. $\psi(\mathbf{x}) = \min_l |\varphi_l(\mathbf{x})|$). The signed distance φ_l for each object can be recovered at the boundary

$$\Lambda_l = \{\mathbf{x} \mid \exists \mathbf{y} \in \mathcal{N}(\mathbf{x}) \text{ s.t. } \sigma_l(\mathbf{x}) \neq \sigma_l(\mathbf{y})\} \quad (1)$$

of each region l , where $\mathcal{N}(\mathbf{x})$ is the 6-connected neighborhood of voxel \mathbf{x} and the sign $\sigma_l(\mathbf{x})$ is indicated by,

$$\sigma_l(\mathbf{x}) = \begin{cases} -1 & \chi(\mathbf{x}) = l \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

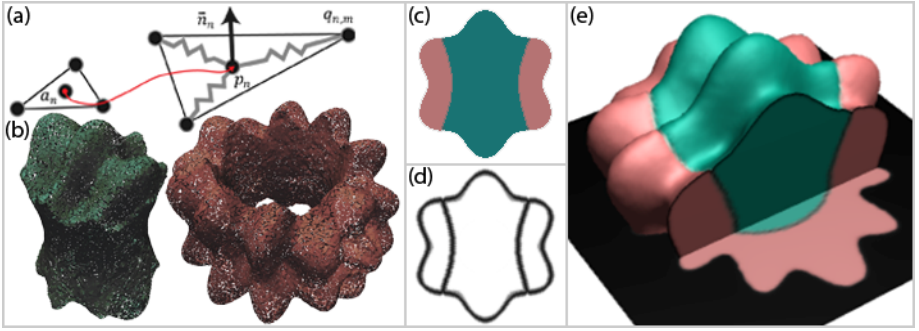


Fig. 1. (a) Diagram of springl. (b) Springl constellation for two objects. (c) label mask. (d) clamped distance field. (e) Raycast rendering of MUSCLE data structure.

We use the convention that distance measurements at locations inside the object have negative values and outside positive values. The partially reconstructed level set $\tilde{\varphi}_l: \Lambda_l \mapsto \mathbb{R}$ is given by $\tilde{\varphi}_l(\mathbf{x}) = \sigma_l(\mathbf{x}) \psi(\mathbf{x})$. $\tilde{\varphi}_l$ measures the signed distance at the boundary (Λ_l) between objects and provides enough information to extract an iso-surface with marching cubes [13] or recover the entire signed distance field for each object with fast-marching [9]. More importantly, the "label mask + distance field" data structure avoids having to store and manipulate independent level set images for each object. For an $M \times M \times M$ image containing L objects, the amount of memory needed to store the MUSCLE data structure is $O(M^3 + LM^2)$. This is compared to $O(LM^2)$ for meshes, $O(LM^3)$ for independent level sets, and $O(M^3)$ for label masks.

Level Set Evolution. MUSCLE, like SpringLS, maintains a mesh and level set representation of the same geometry. To keep them consistent during deformation,

level sets are evolved to minimize their distance to the constellation. This is accomplished by first constructing the clamped distance field for all springls:

$$\omega(\mathbf{x}) = \min\{2d_{max}, d_1(\mathbf{x}) \dots d_N(\mathbf{x})\} \quad (3)$$

where $d_n(\mathbf{x})$ is the distance from location \mathbf{x} to springl n , and $d_{max} = 0.5$ is the clamped distance. Level sets $\tilde{\varphi}_l$ are evolved to minimize the following:

$$E = \sum_l \int \left(\frac{1}{2} (\omega(\mathbf{x}))^2 + \lambda |\nabla \tilde{\varphi}_l(\mathbf{x})| \right) \delta(\tilde{\varphi}_l(\mathbf{x})) d\mathbf{x} \quad (4)$$

where λ is a regularization weight that controls the model's smoothness. Solving the Euler-Lagrange equations, eq. 4 can be minimized with the following scheme:

$$\tilde{\varphi}_l^{z+1}(\mathbf{x}) = \tilde{\varphi}_l^z(\mathbf{x}) - \Delta t \delta_\varepsilon(\tilde{\varphi}_l^z(\mathbf{x})) \left(\omega(\mathbf{x}) \nabla \omega(\mathbf{x}) \cdot \frac{\nabla \tilde{\varphi}_l^z(\mathbf{x})}{|\nabla \tilde{\varphi}_l^z(\mathbf{x})|} + \lambda \nabla \cdot \frac{\nabla \tilde{\varphi}_l^z(\mathbf{x})}{|\nabla \tilde{\varphi}_l^z(\mathbf{x})|} \right) \quad (5)$$

where $\delta_\varepsilon(d)$ is a compactly supported approximation to the dirac delta. The iterative scheme is implemented with Multi-Object Geodesic Active Contours (MOGAC) [14] because it uses the "label mask + distance field" data structure and does not create air-gaps or overlaps between adjacent structures. Equipped with the MUSCLE data structure and evolution scheme, we now discuss how to use them for registration.

Global Registration. Deformable registration is sensitive to the initial alignment of source and target images. To obtain a good initialization, an image-based global registration algorithm such as FLIRT [15] can be used to estimate an affine transformation between source and template images. The 4×4 transformation matrix \mathbf{A} is then applied to the springls constellation (Fig. 2a), labels mask, and distance field. Nearest-neighbor interpolation is used, and must be used, to transform the label mask, and trilinear interpolation is used for the distance field. The interpolator acts as a filter on both the labels and distance field, producing iso-surfaces not quite as smooth as the originals (Fig. 2b). However, transformations can be applied to the constellation without interpolation. The label mask and distance field are then evolved to minimize the distance between their iso-surfaces and the constellation (eq. 5). Fig. 2c illustrates that this method is effective at boosting the fidelity of level sets that undergo global registration.




Fig. 2. Cortical surface after applying Affine transformation to (a) triangle mesh, (b) label mask and distance field (c) MUSCLE after 20 iterations of eq. 5 with $\lambda = 1$

Deformable Registration. Once source and target images have been roughly aligned with global registration, they are more precisely aligned with deformable registration.

Image based registration algorithms output a displacement field $\vec{v}_{ts}: \Omega \mapsto \mathcal{R}^3$ describing the offset of each location in the target image $I_t: \Omega \mapsto \mathcal{R}$ maps in the source image $I_s: \Omega \mapsto \mathcal{R}$. The source image, label mask, or distance field, can be transformed into the target via $I_t(\mathbf{x}) \doteq I_s(\mathbf{x} - \vec{v}_{ts}(\mathbf{x}))$. Displacements fields representing the forward mapping $\vec{v}_{st}: \Omega \mapsto \mathcal{R}^3$ can be applied to mesh vertices $q_{n,m}$ in the springls constellation to obtain new positions $\hat{q}_{n,m}$ via $\hat{q}_{n,m} = q_{n,m} + \vec{v}_{st}(q_{n,m})$. Unfortunately, most registration algorithms produce displacement fields that are not isomorphic (i.e. $\vec{v}_{st}(\mathbf{x} - \vec{v}_{ts}(\mathbf{x})) \neq \vec{v}_{ts}(\mathbf{x})$). A different procedure is needed for transforming label masks and distance fields so that their iso-surfaces are well aligned with springl constellations. To do so, springls are incrementally displaced along linear trajectories from source to target:

$$q_{n,m}(t) = q_{n,m}(0) + t\vec{v}_{st}(q_{n,m}(0)) \quad (6)$$

where $t \in [0,1]$. After each displacement step k s.t. $t = k\Delta t$ for $k = 0, 1, \dots, [1/\Delta t]$, the label mask and distance field are evolved to track the moving mesh via eq. 5. The step size is chosen to be $\Delta t \leq d_{max}/\max_{n,m} \|\vec{v}_{st}(q_{n,m})\|$. The iterative scheme in eq. 5 is repeated 4 times per iteration of eq. 6. Fig. 3a illustrates that applying a displacement field to a mesh can create thin cusp structures that self-intersect. In the MUSCLE framework, the level set representation that is evolved with the mesh cannot develop self-intersections, but it can change topology (Fig. 3b). The simple-point test can be added to the level set method [16] to preserve the object's topology (Fig. 3c).




Fig. 3. Cortical surface after applying displacement field to (a) triangle mesh, (b) MUSCLE (c) MUSCLE with topology-preservation constraint. Back-face of the surface is shown in red, indicating self-intersection. Red circle indicates region with topology change.

Active Contour Segmentation. After registration, active contour methods are used to better align the model with anatomical boundaries visible in the target image. An active contour segmentation framework has been developed for SpringLS [11]. It consists of five phases: advection, relaxation, re-sampling, level set evolution, and hole filling. We extend the advection phase to multiple objects by considering advection equations of the following form,

$$\frac{dp_n}{dt} = \lambda_p(\rho_{l_n}(p_n) - \max_{i \in \mathcal{N}(p_n) \text{ and } i \neq l_n} \rho_i(p_n))\vec{n}_n + \lambda_\sigma \vec{\sigma}(p_n), \quad (7)$$

where $\rho_l: \Omega \mapsto \mathcal{R}$ is the pressure for object l , $\vec{\sigma}: \Omega \mapsto \mathcal{R}^3$ is an external velocity field, and \vec{n}_n is the outward pointing normal for springl n . The first term moves the boundary to classify location p_n as either inside or outside object l_n via the pressure difference between l_n and the maximum pressure from objects in the neighborhood $\mathcal{N}(p_n)$. The second term moves boundaries for all objects towards edges in the image, and $\lambda_p \setminus \lambda_\sigma$ control the relative contributions of each force to the movement of a springl. After all springls have been moved a small distance (less than d_{max}) via eq. 7, the relaxation phase uses inter-springl attraction forces to smooth and regularize the constellation (see Lucas et al. [11] for details). In MUSCLE, springls are only permitted to interact with springls that have the same label l during relaxation. The re-sampling phase removes springls with label l that are too far from $\hat{\phi}_l$'s zero iso-level or are too small. This includes removal of sharp edges and cusp structures (Fig. 3a) that lie far from any level set's zero iso-level. Level sets are then evolved to track the constellation. MUSCLE uses the multi-object level set method outlined in eqs. 3-5. In the final phase, large holes in the constellation are filled by adding springls, and point correspondences for new springls are interpolated based on point correspondences for neighboring springls. MUSCLE fills holes in the constellation with respect to each object individually as opposed to the union of all springls. The five phases are repeated in order until the iso-surfaces stop moving or the maximum number of iterations is reached.

4 Results and Discussion

A MUSCLE was constructed for the brain by extending gyral labels produced by FreeSurfer [17] with fast-marching [9] and combining them with sub-cortical labels and the level set for the central surface produced by CRUISE [18] (Fig. 4). It may seem unusual to parcellate the brain in this way because gyral structures are only present in the GM. However, an emerging application for brain parcellations is to use DTI fiber tracks to assess connectivity of gyral regions [19]. Fiber tracks can only be reliably found in the WM; so to establish connectivity, gyral regions must be extended into the WM. The central surface lies in the "middle" of the GM. We chose to use the central surface in our parcellation because it can be found more reliably [18], and the WM/GM and Pial surfaces have sharp cusps and deep folds that violate the uniform smoothness assumption that is inherent in the regularization of eq. 4.




Fig. 4. (a) gyral labels (top), sub-cortical labels (bottom). (b) Sub-cortical + extended gyral labels (top), level set for central surface (bottom). (c) MUSCLE label mask for 74 structures (top), MUSCLE distance field (bottom). (d) Raycast rendering of MUSCLE parcellation.

MUSCLE brain parcellations were constructed for 10 subjects from the OASIS cross-sectional database [20]. Each parcellation was used to segment the other 9 subjects (90 experiments total) via the following pipeline: 1) each skull-stripped MRI was affine registered to the target; 2) The affine registered image was deformably registered to the target with diffeomorphic Demons [12]; 3) the displacement field and affine transformation were applied to the MUSCLE parcellation; 4) the MUSCLE parcellation was evolved with active contour methods to find the central surface produced by CRUISE. To find the central surface represented by level set $\psi_c : \Omega \mapsto \mathcal{R}$, we use $\rho_l(\mathbf{x}) = H_\varepsilon(\psi_c(\mathbf{x}))$ where $H_\varepsilon(x) = \text{atan}(x/\varepsilon)$ with $\varepsilon = 0.25$, $\lambda_p = 1$, and $\lambda_\sigma = 0$. Fig. 5. shows one example of the registration + segmentation (reg+seg) pipeline. Table 1 compares the accuracy of several approaches. Labeling accuracy is measured with the extended Jaccard metric [7].

MUSCLE reg+seg and Hybrid Warp [6] have similar pipelines. MUSCLE augments Hybrid Warp by providing a data structure with which to represent full brain parcellations that can re-sample / re-mesh under geometric constraints (i.e. no self-intersections, air-gaps, or overlaps). CVS [7] uses a tetrahedral mesh (tet-mesh). Tet-meshes are difficult to use for registration because diffeomorphic deformations require safeguards to avoid tetrahedral inversion. MUSCLE places no restrictions on diffeomorphic deformations and has other useful geometric properties. For example, MUSCLE iso-surfaces are smooth, which is not always the case for meshes (Fig. 3) and rarely the case for iso-surfaces extracted from label masks. It is important to have smooth surfaces (along with point correspondences) because cortical surface analysis is often sensitive to surface curvature [17]. MUSCLE associates gyral labels with springls on the central surface instead of voxels in the GM, which partially explains why cortical labeling accuracy is higher for MUSCLE. In HAMMER and CVS, the sampling density of voxel-based gyral labels is affected by GM thickness. Regions where the GM is thin often report lower labeling accuracy because the Jaccard metric is sensitive to a labeled region's volume size. MUSCLE evenly samples the cortical surface with springls so that gyral labeling is unaffected by GM thickness. MUSCLE reg+seg produces less accurate subcortical segmentations than other works because it uses Demons instead of HAMMER. Subcortical segmentation with Demons and a label mask representation has an accuracy of $54.3 \pm 13.6\%$. MUSCLE's slightly lower accuracy ($53.1 \pm 13.8\%$) is due to smoothing (eq. 4). Subcortical structures have sharper features than the central surface, so they should have a smaller smoothing weight (λ). To improve subcortical accuracy, one could use HAMMER for registration and disable / reduce smoothing in subcortical regions.




Fig. 5. MUSCLE parcellations for (a) source, (b) source after applying displacement field (18 sec) and affine transform (7 sec), (c) target, and (d) source after reg+seg (21 min total, 1.1M - 1.4M springls, 256^3 voxels, 100 active contour iterations). Computation times do not include FLIRT (10 min), Diffeomorphic Demons (1 hr), or CRUISE (2.5 hrs). Dual Intel E5630 (8 cores). Source code available at <http://code.google.com/p/imagesci/>.

Table 1. Performance summary. Distance measurements are for cortical surface only. Accuracies for referenced works do not reflect the same parcellations, datasets, or sample sizes.

Pipeline	Geometric Representation	Template to Subject	Subject to Template	Subcortical Label Acc.	Cortical Label Acc.
MUSCLE reg only	MUSCLE	1.29±0.45 mm	1.05±0.36 mm	53.1±13.8%	80.1±2.6%
MUSCLE reg + seg	MUSCLE	0.40±0.05 mm	0.25±0.01 mm	52.8±13.0%	81.2±2.7%
HAMMER [2, 7]	Label mask	--	--	66.8±6.7%	36.6±6.0%
Hybrid Warp [6]	Triangle mesh	0.45±0.05 mm	0.37±0.05 mm	--	--
CVS [7]	Tet-mesh	1.5 to 2.5 mm	--	70.5±4.6%	54.4±10%

Acknowledgments. This research was supported in part by a fellowship from the Johns Hopkins Applied Physics Laboratory.

References

1. Thirion, J.P.: Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis* 2, 243–260 (1998)
2. Shen, D., Davatzikos, C.: HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. on Medical Imaging* 21, 1421–1439 (2002)
3. Dawant, B.M., Hartmann, S., Thirion, J.P., Maes, F., Vandermeulen, D., Demaerel, P.: Automatic 3-D segmentation of internal structures of the head in MR images using a combination of similarity and free-form transformations. I. Methodology and validation on normal subjects. *IEEE Transactions on Medical Imaging* 18, 909–916 (1999)
4. Aljabar, P., Heckemann, R., Hammers, A., Hajnal, J., Rueckert, D.: Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy. *Neuroimage* 46, 726–738 (2009)
5. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans. on Medical Imaging* 23, 903–921 (2004)
6. Liu, T., Shen, D., Davatzikos, C.: Deformable registration of cortical structures via hybrid volumetric and surface warping. *Neuroimage* 22, 1790–1801 (2004)
7. Postelnicu, G., Zollei, L., Fischl, B.: Combined volumetric and surface registration. *IEEE Transactions on Medical Imaging* 28, 508–522 (2009)
8. Alliez, P., Ucelli, G., Gotsman, C., Attene, M.: Recent advances in remeshing of surfaces. *Shape Analysis and Structuring*, 53–82 (2008)
9. Sethian, J.: *Level set methods and fast marching methods*. Cambridge Univ. Pr. (1999)
10. Losasso, F., Shinar, T., Selle, A., Fedkiw, R.: Multiple interacting liquids. *ACM Transactions on Graphics (TOG)* 25, 812–819 (2006)
11. Lucas, B.C., Kazhdan, M., Taylor, R.H.: SpringLS: A Deformable Model Representation to Provide Interoperability between Meshes and Level Sets. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011, Part II*. LNCS, vol. 6892, pp. 442–450. Springer, Heidelberg (2011)
12. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *Neuroimage* 45, S61–S72 (2009)
13. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *ACM Siggraph Computer Graphics* 21, 163–169 (1987)

14. Lucas, B.C., Kazhdan, M., Taylor, R.H.: Multi-Object Geodesic Active Contours (MOGAC). In: MICCAI, Nice (2012)
15. Jenkinson, M., Smith, S.: A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* 5, 143–156 (2001)
16. Han, X., Xu, C., Prince, J.: A Topology Preserving Level Set Method for Geometric Deformable Models. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 755–768 (2003)
17. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* 31, 968–980 (2006)
18. Han, X., Pham, D.L., Tosun, D., Rettmann, M.E., Xu, C., Prince, J.L.: CRUISE: cortical reconstruction using implicit surface evolution. *Neuroimage* 23, 997–1012 (2004)
19. Zalesky, A., Fornito, A., Harding, I.H., Cocchi, L., Yücel, M., Pantelis, C., Bullmore, E.T.: Whole-brain anatomical networks: Does the choice of nodes matter? *Neuroimage* 50, 970–983 (2010)
20. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open Access Series of Imaging Studies (OASIS). *J. of Cognitive Neuroscience* 19, 1498–1507 (2007)