Using Spatial Characteristics to Aid Automation of SOM Segmentation of Functional Image Data

Patrick O'Driscoll
Applied Physics Program
Department of Statistics
Rice University, Houston, Texas 77005
Email: patrick.odriscoll@rice.edu

Erzsébet Merényi
Departments of Statistics, and
Electrical and Computer Engineering
Rice University, Houston, Texas 77005
Email: erzsebet@rice.edu

Robert Grossman
Department of Neurosurgery
Houston Methodist Neurological Institute
Houston, Texas 77030
Email: RGrossman@houstonmethodist.org

***Abstract*—We propose a new similarity measure, Combined Connectivity and Spatial Adjacency ($CCSA$), to be used in hierarchical agglomerative clustering (HAC) for automated segmentation of Self-Organizing Maps (SOMs, Kohonen [1]). The $CCSA$ measure is specifically designed to assist segmentation of large, complex, functional image data by exploiting general spatial characteristics of such data. The proposed $CCSA$ measure is constructed from two strong indicators of cluster structure: the degree of localization of data points in physical space and the degree of connectivity of SOM prototypes (as defined by Taşdemir and Merényi [2]). The new measure is expected to enhance cluster capture in large functional image data cubes such as hyperspectral imagery or fMRI brain images, where many relevant clusters exist with widely varying statistical properties and in complex relationships both in feature space and in physical (image) space. We demonstrate the effectiveness of our approach using the $CCSA$ measure on progressively complex synthetic spatial data and on real fMRI brain data.**

I. Background and Motivation

We propose a new automated SOM clustering algorithm which is expected to benefit spatial image data in general and large functional spatial image data (such as spectral image cubes or fMRI data cubes) in particular. Spatial image cubes record physical measurements that are made on regular grids (in 2D pixel footprints for spectral imagery, or in 3D voxels for fMRI). The (typically) many-dimensional response signals — such as spectral signatures of surface materials in remote sensing imagery, or time courses of neuron responses in fMRI — in each grid cell are the high-dimensional (high-D) feature vectors input to clustering. The high spatial, and high spectral or temporal resolution admits the delineation of many relevant clusters, and the cluster structure in these data is usually very complex, posing a large challenge in general. In addition, the measured signals (the feature vectors) are almost always mixtures of responses from different species at cluster boundaries (spectral signatures of different materials in a pixel, or time courses of neurons contained in a voxel but belonging to different functional brain areas) causing overlaps. The $CCSA$ measure is designed to address this difficulty in the context of automated SOM segmentation, by incorporating cluster localization in the measurement grid.

While the severity of the cluster overlap, in general, depends on the pixel or voxel size over which multiple signals are integrated, there are particular differences in the mixing characteristics of different functional image data. In particular, fMRI data that we use for a demonstration in this paper, has mixed time courses over relatively large voxels (voxel size is a trade-off for high temporal resolution), where the similarity of the individual time courses — on average — is higher than, for example, the average similarity of signatures in a hyperspectral image cube. This results in relatively high density at cluster boundaries, and small distances of SOM prototypes at these boundaries. This challenge cannot be overcome simply by using larger SOMs because the inherent cohesion in the data only causes overfitting by too many prototypes as shown in O'Driscoll et al. [3]. The $CCSA$ measure offers the possibility of tuning for the general characteristics of particular types of functional image data.

A. Related Works on Automated SOM Clustering

Interactive SOM clustering uses a variety of visualizations to aid the user in extracting clusters. While the results are generally superior to results of automated clusterings, it takes considerable time and it is often non-repeatable due to user subjectivity, the level of expertise, and other biases. This is further exacerbated by the growing complexity of data such as in large functional image data sets. The SOM, however, is outstanding at learning the manifold structure of these data, therefore automation of the SOM segmentation that matches the success of interactive SOM clustering is highly desirable for speed and repeatability of high-quality clustering. Following we review automated SOM clustering methods applied to spatial image data.

Vesanto and Alhoniemi [4] used a hierarchical agglomerative clustering (HAC) with single, average, and complete linkages combined with SOM based distance during the construction phase, and used a gap criterion for pruning the dendrogram. This method produced clusters of prototypes that were contiguous in SOM space and similar in feature space. They produced clusters that represented the majority of the true clusters well, but often missed outliers and had poor representation of the edge of the cluster boundaries.

Chuang et al. [5] and Chiu et al. [6] used fuzzy C-means to group SOM prototypes of fMRI brain images. Both analyzed only a few brain slices and were able to discover a handful of clusters. Although the fuzzy C-means method was successful

978-1-5090-6638-4/17/$31.00

for their purpose, the approach assumes hyper-elliptical data structure therefore dramatically limiting the discovery of the natural data structure.

Peltier et al. [7] relied on a HAC method using the minimum Euclidean distance between SOM prototypes for identifying clusters of whole-brain fMRI data. Despite the increased complexity of the data (compared to using a few brain slices), this produced more interpretable clusters and a better alignment with expected functional brain areas than the fuzzy C-means approaches. However, this method still only yielded a handful of highly localized clusters.

Liao et al. [8] proposed a spatio-temporal metric to segment SOMs of whole-brain fMRI data with a HAC method. This metric folded in information about the SOM lattice distance between the prototypes, imposing a proximity constraint on the clusters in SOM space, while using the correlation between prototypes in feature space to group prototypes. The use of this similarity measure resulted in at most a minimal improvement over Peltier et al. [7]. This was most likely a result of the small (10 x 10) SOM employed, where the contiguity constraint had a negligible impact on the overall cluster structure.

A HAC method using the nearest neighbor rule for points in feature space was successfully applied to multiple types of data, ranging from hyperspectral satellite images by Gonçalves et al. [9] to EEG data by Sommer and Golz [10]. Gonçalves et al. [9] is of particular interest because it incorporated information from feature space to remove prototypes with large feature space dispersion from clustering at the initialization of the HAC method. These prototypes represented areas of transition between two or more clusters, thus, by extension, implicitly included some information from physical space.

Taşdemir et al. [11] used, in an agglomerative HAC method, a connectivity measure between and within clusters, which was computed from connectivities of prototypes. Connectivity between two prototypes (CONN) was defined in Taşdemir and Merényi [2] as the number of data points which select that pair of prototypes as their best matching unit (BMU) and second BMU. This approach was successful on both synthetic data and real hyperspectral imagery, albeit with relatively low cluster complexity. The use of connectivity as a similarity measure in automated clustering appears a natural extension of its successful application in interactive SOM segmentation through the CONNvis visualization (proposed by Taşdemir and Merényi [2]) in previous work, including interactive clustering of fMRI data (O'Driscoll et al. [3]).

The similarity measures in the above works used information from the feature space, from the SOM lattice, or combinations thereof, to extract clusters. We propose to utilize, in addition to the above, an as yet untapped resource, the exploitation of cluster relations in the physical space of spatial image data. Our novel contribution is the quantification of cluster localization in physical space. This is something that the interactive analyst (consciously or otherwise) also applies to spatial data. We show that automated SOM segmentation with our proposed measure, combining cluster contiguity constraints in physical, feature, and SOM space, is as good — or better — than interactive clustering, for large, complex, functional image data.

II. Proposed Clustering Method for Spatial Data

To cluster data with a SOM, we use the usual two-step approach. The first step is to obtain a learned SOM, the second is to group the SOM prototypes into clusters, using our proposed Connectivity and Spatial Adjacency ($CCSA$) similarity measure in hierarchical agglomerative clustering (HAC).

A. Construction of the $CCSA$ Similarity Measure

The $CCSA$ uses information from three different spaces, illustrated in Fig. 1. Physical space is the grid of pixels or voxels where the feature vectors are measured and where they are indexed by the locations of the measurements. This grid is 2D for hyperspectral and other common image data, and 3D in the case of fMRI data. Feature space is the space where the feature vectors (spectral signatures at the grid cells or time courses in voxels), as well as the SOM prototypes, exist. SOM space is the SOM lattice of neurons (typically 1D or 2D). $CCSA$ uses physical and SOM space explicitly in its formulation, whereas it uses feature space implicitly through a prototype connectivity measure described below.

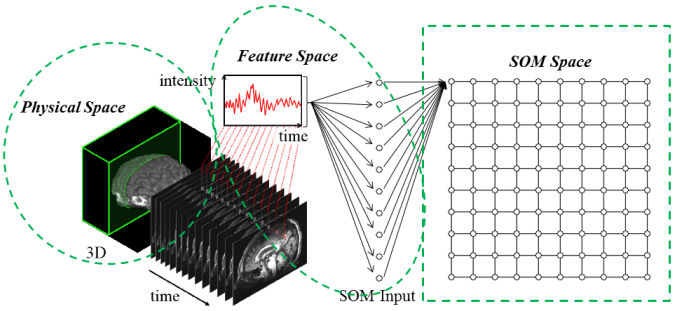

Fig. 1. Illustration of physical space, feature space, and SOM space.

$CCSA$ is constructed from two terms, as defined below. The variables used in this formulation are listed in Table I.

The first term of $CCSA$ determines the degree of localization (the degree of co-occurrence) between two clusters I and J in physical (image) space. In spatial images, tight physical proximity of two clusters often indicates a high probability that they are similar and can be merged. If instances (patches) of the same pair of clusters occur in similar proximity to one another within multiple disconnected spatial areas, this probability increases. Examples are the occurrence of the same pair of land-cover clusters in each others' vicinity (*e.g.*, patches of the same crop developed slightly differently) in different spatial regions in a remote sensing image; or a series of similarly activated spatially adjacent regions at symmetric physical locations in a fMRI brain image (*e.g.*, subregions of the visual cortex). Therefore it is important to appropriately reward the merging of clusters that show the same close configuration within multiple disconnected spatial regions as in the above example. At the same time, we need to appropriately penalize the potential merge of clusters that appear in multiple regions but which are already well-distinguished from others (not systematically co-occurring with another cluster), and should not be merged.

TABLE I. DEFINITION OF VARIABLES IN CONSTRUCTION OF $CCSA$

I, J	Cluster labels for clusters I and J
$\mathbf{X_I}$, $\mathbf{X_J}$	Sets of data points in physical space (*i.e.,* sets of image locations $[x, y]$ or $[x, y, z]$) belonging to I and J
ix, jx	Index of data points of cluster I and J, *i.e.*, $ix \in \mathbf{X_I}$ and $jx \in \mathbf{X_J}$
$\|.\|$	Cardinality of a given set
d_{Ph}	Distance in physical space, city-block if physical space is discrete and Euclidean if physical space is continuous
d_{SOM}	Distance in SOM lattice, city-block
$A(.)$	Surface area in physical space at given radius For discrete physical space the following are used: 2D $A(d_{Ph}(ix, jx)) = 4 \times d_{Ph}(ix, jx)$ 3D $A(d_{Ph}(ix, jx)) = 4 \times d_{Ph}(ix, jx)^2 + 2$ For continuous physical space the following are used: 2D $A(d_{Ph}(ix, jx)) = 2\pi d_{Ph}(ix, jx)$ 3D $A(d_{Ph}(ix, jx)) = 4\pi d_{Ph}(ix, jx)$
$CADJ$	Cumulative adjacency as defined by Taşdemir and Merényi [2]
$\mathbf{C_I}$, $\mathbf{C_J}$	Sets of prototypes (*i.e.,* sets of SOM grid locations $[x, y]$) belonging to clusters I and J
ip, jp	Index of prototypes of clusters I and J, *i.e.*, $ip \in \mathbf{C_I}$ and $jp \in \mathbf{C_J}$
RF_{ip}, RF_{jp}	Receptive field of prototypes ip and jp
t_{SOM}	User specified threshold to limit contribution to $CCSA$ from SOM neighborhood of radius t_{SOM}

We measure this localization by the adaptation of a radial distribution function (RDF). RDF, or pair correlation function, which is traditionally used in statistical mechanics and material science to describe how the density of particles varies as a function of distance from a reference particle. The RDF is the probability of finding a particle at a given distance from a reference particle and is calculated by computing the density of particles on a surface area at a given distance. To convert this function to a single numerical value, we modify the RDF by integrating the density of consecutive spherical shells around a given point, giving us a total sum of the surface densities. Then by averaging over all points in the cluster and scaling this term we obtain eq. (1) under the hypothesis that two clusters are merged. Since this is an average over all points in a given cluster, it allows us to reward multiple localized occurrences of the same cluster in physical space.

$$RDF(I,J) = \frac{1}{|\mathbf{X_I}| + |\mathbf{X_J}|} \sum_{ix \in \mathbf{X_I}} \sum_{\substack{jx \in \mathbf{X_J} \\ ix \neq jx}} \frac{1}{A(d_{Ph}(ix, jx))} \tag{1}$$

where d_{Ph} is the city-block distance in case of a discrete physical space (like the pixel or voxel grid of images), in contrast to Euclidean distance in the continuous spaces of physics problems from where we borrow the RDF idea. (We choose city-block distance in discrete space to help reduce rounding errors.) In the rare event of $|\mathbf{X_I}| = 1$ $RDF(I, I)$ is undefined, and we set it to a value of $1/A(1)$, which is the situation of two neighboring pixels belonging to the same cluster.

To quantify how the proposed merger of clusters I and J improves the cluster localization in physical space, we use an average improvement rate, known as scaled RDF ($sRDF$), as in eq. (2). $sRDF$ is the mean of the normalized improvement from merging cluster I with cluster J compared to cluster I, and vice versa. The normalization is performed by taking the ratios of an inter-cluster RDF term and an intra-cluster RDF term. The intra-cluster term, $RDF(I, I)$, is the localization of the cluster with itself. The inter-cluster term, $RDF(I, J)$, is the localization from one cluster to another in physical space. Therefore the ratio of $RDF(I, J)$ to $RDF(I, I)$ gives the improvement by merging clusters I and J when compared to the RDF of cluster I. By taking the mean of these ratios, we determine the favorability of the merger, compared to the unmerged clusters, in terms of their mutual localization in physical space.

$$sRDF(I,J) = \frac{1}{2}\left(\frac{RDF(I,J)}{RDF(I,I)} + \frac{RDF(J,I)}{RDF(J,J)}\right) \tag{2}$$

The Second term of $CCSA$ uses the cumulative adjacency of the SOM prototypes in feature space to determine the connectivity between clusters I and J. Cumulative adjacency between pairs of prototypes, $CADJ(ip, jp)$, is defined by Taşdemir and Merényi [2] as the number of data points for which prototype ip is the BMU and prototype jp is the second BMU. The ratio of the sum of the $CADJ(I, J)$ values from a given prototype ip to any prototype in $\mathbf{C_J}$ over the receptive field of prototype ip gives the proportion of connections of prototype ip to prototypes in $\mathbf{C_J}$. Taking the mean of this ratio yields $RCADJ$, as in eq. (3). $RCADJ$ quantifies how connected cluster I is to cluster J in feature space. To avoid contributions from (possibly) many weakly connected prototypes (each of which may be connected by only a few noisy points, but together can unduly distort $RCADJ(I, J)$) one can exploit the topology-preserving property of the SOM and constrain the neighborhood radius, t_{SOM}, from within which prototypes are included in the sum in eq. (3). The recommended value for t_{SOM} is the SOM radius within which topology violations are considered local, thus not harmful for cluster capture. (See Taşdemir and Merényi [12] for more information.) This is an extension of SOM-contiguity constraints imposed in Vesanto and Alhoniemi [4], Liao et al. [8], and Gonçalves et al. [9]. $RCADJ$ combines information contained in both feature and SOM space.

$$RCADJ(I,J) = \frac{1}{|\mathbf{C_I}| + |\mathbf{C_J}|} \sum_{jp \in \mathbf{C_J}} \sum_{\substack{ip \in \mathbf{C_I} \\ ip \neq jp}} \frac{dCADJ(ip, jp)}{|RF_{ip}|} \tag{3}$$

$$dCADJ(ip, jp) = \begin{cases} CADJ(ip, jp) & d_{SOM}(ip, jp) \leq t_{SOM} \\ 0 & otherwise \end{cases} \tag{4}$$

This definition of $RCADJ$ breaks down when I and J are the same, and cluster I consists of only one prototype. In this case, we use eq. (5) instead of eq. (3). This is the ratio of the largest $CADJ$ connection to any SOM prototype, rp, to the size of the receptive field of prototype ip. In either case, $RCADJ$ will have a value of 0 only when there are no connections from any prototype in cluster I to cluster J.

$$RCADJ(I,I) = \frac{argmax_{rp}|CADJ(ip,rp)|}{|RF_{ip}|} \quad (5)$$

Since we want to merge clusters with a high degree of connections, similarly to $sRDF$, we use the average normalized improvement of intra- and inter-cluster $RCADJ$ terms to evaluate the merge, as per eq. (6).

$$sRCADJ(I,J) = \frac{1}{2}\left(\frac{RCADJ(I,J)}{RCADJ(I,I)} + \frac{RCADJ(J,I)}{RCADJ(J,J)}\right) \quad (6)$$

Similarly to $sRDF$, a larger value of $sRCAD(I,J)$ indicates stronger connectivity between clusters I and J when compared to the internal connectivity of either cluster.

Since we only wish to merge clusters that express both a high degree of localization in physical space and a high degree of connectivity in feature and SOM space, we use the product of the $sRDF(I,J)$ and $sRCADJ(I,J)$ terms yielding our similarity measure $CCSA(I,J)$, as in eq. (7). Because of its construction $CCSA(I,J)$ is non-negative, and it reports a value of 0 when no connections exist between the prototypes of clusters I and J, *i.e.,* the clusters are cleanly separated and should not be merged. Thus, the larger the value of $CCSA(I,J)$ the more advantageous it is to merge clusters I and J.

$$CCSA(I,J) = sRDF(I,J) \times sRCADJ(I,J) \quad (7)$$

B. Overview of the Clustering Algorithm

We use the standard hierarchical agglomerative clustering (HAC) framework, with our $CCSA$ similarity measure (eq. (7)). We initialize the algorithm by excluding empty prototypes and assuming each prototype is its own cluster. In every step of the hierarchical procedure, we merge the two clusters for which $CCSA$ is largest. We continue merging clusters until there are only a pre-defined number of clusters remaining, or the maximum value of $CCSA$ is 0. $CCSA = 0$ indicates a complete segmentation of the prototypes, and $CCSA$ can no longer advise further merges of clusters.

C. CONNvis Visualization

CONNvis is a visualization of the $CONN$ similarity measure and is described in detail in Taşdemir and Merényi [2]. Since we use CONNvis in all of our SOM overlays, we review it here briefly to help the reader see the connectivity between prototypes and better understand the $CCSA$ similarity measure. The connectivity between two prototypes ip and jp is the symmetric version of $CADJ(ip,jp)$, *i.e.,* $CONN(ip,jp) = CONN(jp,ip) = CADJ(ip,jp) + CADJ(jp,ip)$. The CONNvis visualization shows the connection strength, the $CONN(ip,jp)$ value, for every pair of prototypes (represented by black dots in Fig. 2 (b), and other figures of SOM visualizations) by the width of a line segment connecting prototypes ip and jp. The line width is proportional to the binned normalized $CONN(ip,jp)$ value. (Binning is done to aid human visual inspection, but not for computation of the $CCSA$ measure. For a full description of the non-linear binning see Taşdemir and Merényi [2].) An additional, local information is also expressed by CONNvis: colors of the connections indicate the ranking of a prototype's connections by strength, where strength is the $CONN(ip,jp)$ value. A prototype's highest-ranking connection indicates its most important neighbor, thereby revealing the relative local density of the manifold at prototype ip. Red, blue, green and yellow code the first-, second-, third-, and fourth-ranking neighbors in this order. Beyond the fourth rank, gray shades proportional to the connection strengths are used. This is an arbitrary coloring scheme (and number of bins) that works well in practice, but other colors or numbers of bins could be used to best suit a given visualization purpose.

III. Data and Results

We first show how the clustering works on increasingly complicated synthetic spatial data sets (Lsun, Clown, and 6D 20-class data sets). We also compare the Lsun and Clown results to previously published clusterings by other authors.

A. 2D Lsun Data

The Lsun data from Ultsch [13], shown in Fig. 2, is composed of three clusters in 2D feature / physical space: two rectangular clusters and one spherical cluster, with reasonably clear separation. To demonstrate how our clustering algorithm progresses, we show three different stages of clustering of a 10 x 10 SOM in Fig. 2 (a), (c) and (e). At each of the stages, the segmentation provides a consistent set of clusters that are highly connected in SOM space and highly localized in physical space. Although there is a noticeable misclassification in the spherical cluster in Fig. 2 (e), this is a direct result of the lack of connection between the true cluster and the small subgroup of three light blue prototypes in Fig. 2 (f). (Notice that the connection from a light blue prototype that seemingly links it to the yellow cluster is through a dead prototype indicated by the absence of a black dot at its grid cell.) In this case, our similarity measure is expected to fail due to multiplication between the two terms in eq. (7). The HAC terminates with a maximum $CCSA = 0$ value after finding 4 clusters. One can confirm this evaluation by examining the CONNvis representation in Fig. 2 (f).

By comparison, experiments by Taşdemir et al. [11] echo these results. They cluster the exact same SOM as we did, using HAC with $CONN$ linkage, average linkage, centroid linkage, Ward's measure, and k-means. Ward's method yields correct classification; $CONN$ and average linkage perform the same way as in our case (described above); and all of these outperform centroid linkage and k-means.

B. 2D Clown Data

The Clown data set and a learned SOM has generously been provided by Drs. Vesanto and Alhoniemi [4]. This is a 2D synthetically generated clown face and body composed of nine true clusters of varying shapes, sizes, densities and proximities thus representing non-trivial clustering challenges (Fig. 3 (a)). The left eye (on the right of the figure) comprises three discrete subunits, the right eye and the nose are large, dense and regularly shaped, the mouth is thin and U-shaped,

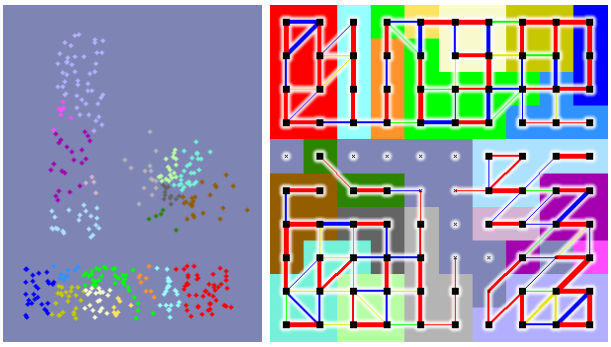

(a) Feature space, 20 clusters (b) SOM space, 20 clusters

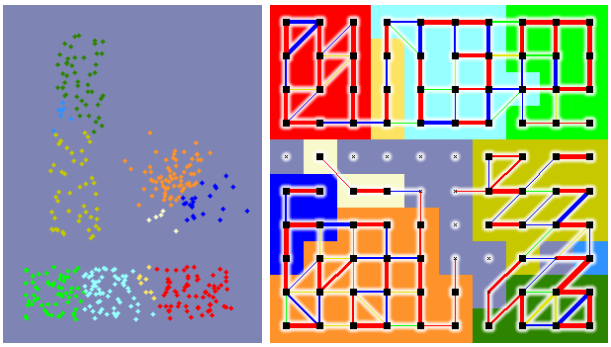

(c) Feature space, 10 clusters (d) SOM space, 10 clusters

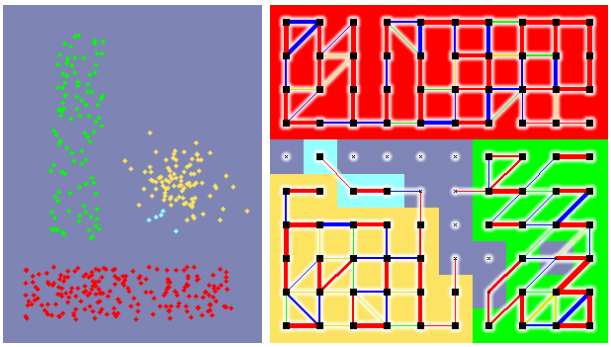

(e) Feature space, 4 clusters (f) SOM space, 4 clusters

Fig. 2. Lsun data from Ultsch [13] partitioned into 20, 10, and 4 clusters at three different stages of our automatic HAC / $CCSA$ segmentation of a 10 x 10 SOM. Colors of data points in a), c) and e), and the same colors of blocks of SOM cells in b), d) and f) indicate cluster labels. The colors of $CONN$ connections (red, blue, green and yellow) in the CONNvis overlay may unavoidably coincide with cluster colors, therefore we stress that the coloring of connecting line segments between prototypes in the SOM lattice has no relation to cluster assignments. It is also important to note that the color assignment to clusters is arbitrary at each of the clustering stages, thus color consistency across different segmentations is at most coincidental. At each stage, the clustering produces clusters strongly connected in SOM space and highly localized in physical space.

and the body has a larger volume and much lower density than the other parts. In addition, there are three outliers sets. The learned SOM from Vesanto and Alhoniemi [4] is a 19 x 17 SOM with a hexagonal lattice structure. We visualize it on a regular square grid. We cluster this SOM with HAC / $CCSA$ until we reach 16 clusters. This stopping criterion is experimentally determined and it is larger than the true number of clusters (nine), but produces a segmentation that well represents the true clusters shown in Fig. 3 (a). Two of the three clusters of the left eye are merged, and one of the outliers is combined with the left eye. The extra clusters (beyond the true ones) consist primarily of prototypes that are not connected to their true clusters or they are weakly connected, as outlined in Fig. 3 (b), and at the same time are highly localized in physical space as seen in Fig. 3 (a). We expect our $CCSA$ similarity measure to preserve these types of clusters, as they are more connected and localized to themselves than to the true cluster.

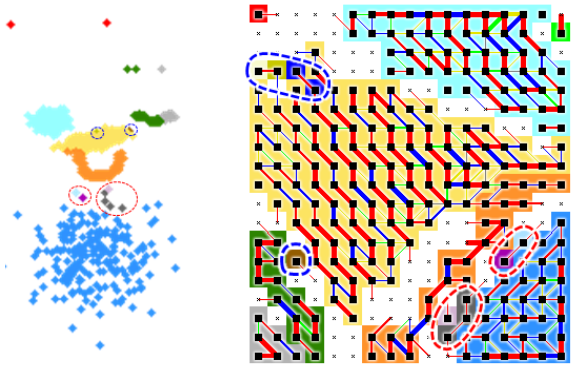

(a) Feature space, 16 clusters (b) SOM space, 16 clusters

Fig. 3. Clown data from Vesanto and Alhoniemi [4] segmented into 16 clusters using our HAC / $CCSA$ automated clustering. Colors indicate cluster labels as described in the caption of Fig. (2). (a) Clusters in feature / physical space; (b) Clusters in the segmented SOM with CONNvis overlay. The right eye (light blue), and most of the body (medium blue) is perfectly identified. There are a few clusters weakly connected to the body (circled in red). The nose (yellow) and mouth (orange) show acceptable separation. The nose has a few small, poorly connected and highly localized clusters (circled in blue). The left eye is grouped into 2 subclusters instead of the three true ones. Two of the three outliers are identified well, the third is grouped with the left eye.

In their original clustering of this same SOM [4], Vesanto and Alhoniemi present very similar results using HAC with single, average, and complete linkage, identifying eight clusters. The slight differences are: four subclusters in the left eye (vs. two in our clustering); one of the three outliers identified correctly (vs. two in ours); and misclassification between the body, mouth, and left eye (vs. oversegmentation at the boundaries of these in ours).

C. 6D 20-class Data Cube

This 6D synthetically generated "spectral" data cube from Merényi et al. [14] has 20 different spatial areas with different 6-element feature vectors attached to the (x,y) locations within each of the 20 areas. The spatial layout of the classes in the physical space can be seen in Fig. 4. (Although this is not the exact truth label map, which we omit for space considerations, it is almost a perfect match with only a few stray points in the fields of a few classes.) Here, some of the class regions are embedded in others: these are the smaller squares of light and medium green and lilac colors. A 1-pixel class of hot pink color is also embedded but hard to see in this image. In feature space, the 6D signatures have varying degrees of correlation (thus separation). We cluster the 20 x 20 SOM of the 6D 20-class data produced in Merényi et al. [14]. We obtain best results when stopping at 21 clusters, as seen in Fig. 4. There are only a few areas of confusion resulting from

prototypes on the fringes of a few clusters in SOM space where the neighboring clusters are also neighbors in physical space and are more connected to the incorrect class than to the correct one in SOM space. This is caused by noise present in this data set. In most cases, we get clean-cut clusters despite the appearance of weak connections between clusters adjacent in both SOM and physical space. This symmetric and binned representation of the connectivity hides (visually) the more nuanced relations which are exploited by the HAC / $CCSA$ algorithm. Examples of the misclassified prototypes are circled in Fig. 4 (b). The 21st cluster (black, containing just a single point) is the result of a minor topology violation in the SOM.

In summary, the clustering of the 6D 20-class data cube is very successful and more successful than the 2D data set examples. The success can be attributed to the larger number of connections that exist within the clusters, as each prototype is connected to a larger number of other prototypes in the 6D 20-class data than in the 2D data sets. Furthermore, areas of misclassification in the 6D 20-class data sets are attributed to either topology violations or stronger connections to the incorrect cluster than to the true one; both cases are rare and directly caused by noise.

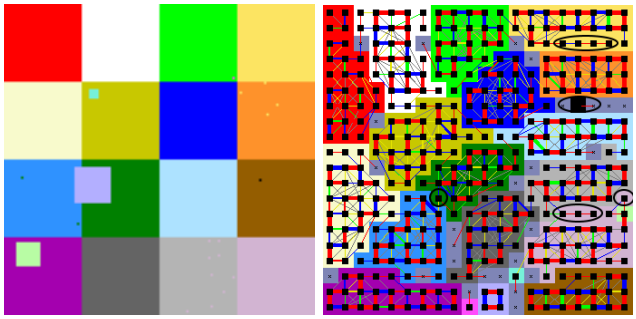

(a) Physical space, 21 clusters (b) SOM space, 21 clusters

Fig. 4. Automated HAC / $CCSA$ clustering of the 6D 20-class data cube. a) Physical space representation of clusters; b) SOM space representation of the clustered 20 x 20 SOM with CONNvis overlay. Colors and representations are the same as described in the caption of Fig. 2. The 21 clusters produced match the true classes almost perfectly, with minimal cluster confusion. Misclassified prototypes are circled in black. They all contain only a few data points.

D. fMRI Brain Image

We apply our HAC / $CCSA$ clustering to a whole-brain fMRI data set collected at the Methodist Hospital (using a Siemens Vario 3 Tesla scanner with a temporal resolution of 1.3 s and voxel size of 3.3 x 3.3 x 5 mm^3). We use the learned 40 x 40 SOM of this data set from our previous work, O'Driscoll et al. [3], where we presented interactive clustering by segmenting the same SOM. Here we compare the results of the automated clustering to the previous interactive results.

The fMRI data comprises (as feature vectors) the time courses (of neural activity encoded by blood oxygen level dependence signals) at each voxel, which is generated in response to a subject's brain activity during the genesis of willed movement. In this case, the movement is prompted by a visual stimulus of the perception of an unpleasant human face. The subject squeezes a ball (for approximately 10 seconds) placed into their right hand if he/she wishes the visual stimulus (the unpleasant face) to be removed. This produces clear activation of a number of brain areas (including the visual cortex, emotional and cognitive areas, the motor cortex, and others) involved in the decision of making a movement. The objective of clustering the time courses is to capture these areas and further investigate the temporal relations of their activation. The data processing and cluster extraction for the interactive process is described in O'Driscoll et al. [3].

We extracted 100 clusters using our proposed automated clustering algorithm. Since our main interest is to trace the activations from the visual cortex to the motor cortex, we filter, in a post-processing step, the clusters whose mean time courses have the highest correlation to the time course of the visual cortex. We perform this filtering for both the interactively and automatically identified clusters, using a correlation threshold of 0.6. (See O'Driscoll et al. [3] for details of the filtering.) The filtering process passes 26 of the automatically identified clusters, and 10 of the interactively identified clusters. The clusters highly correlated to the visual cortex are shown in SOM space in Fig. 5 (a). The majority of prototypes that were assigned to the interactively identified clusters Fig. 5 (b), are also contained in the automatically identified clusters. From the larger number of prototypes labeled (assigned to clusters) in Fig. 5 (a) it is apparent that the cluster of the visual cortex is more correlated with a larger portion of the automatically segmented SOM. This is due to the ability of the $CCSA$ to use the asymmetric $CADJ$ connections in contrast to the symmetric $CONN$ values, used in the visualization for interactive cluster capture.

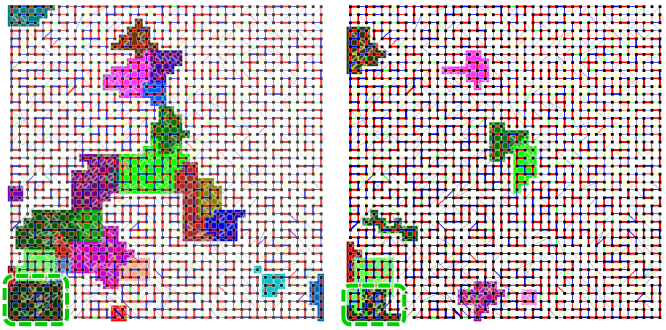

(a) HAC / $CCSA$, 26 clusters (b) Interactive, 10 clusters

Fig. 5. SOM space representation of (a) automatically and (b) interactively identified clusters in the fMRI brain data. The cluster colors are not consistent between the two maps (reconciliation is extremely difficult). The automated HAC / $CCSA$ clustering identified more clusters in contiguous regions from the visual cortex (circled at the lower left hand side). There exists a large degree of overlap between the two clusterings, and the HAC / $CCSA$ clustering is able to label more of the SOM clusters that are highly correlated with the visual cortex.

While the interactively identified clusters are co-located with many known functional brain regions, the automated clustering discovers more of the relevant clusters and with a more complete coverage in physical (brain) space, including those highlighted by the interactive clustering. Part of the reason for this is that we took a conservative approach with the interactive clustering in O'Driscoll et al. [3] and avoided labeling of highly overlapping cluster boundaries which are hard to discern by visual inspection of the CONNvis overlay. The automated method is not constrained by such considerations.

The interactively identified clusterings produced a good

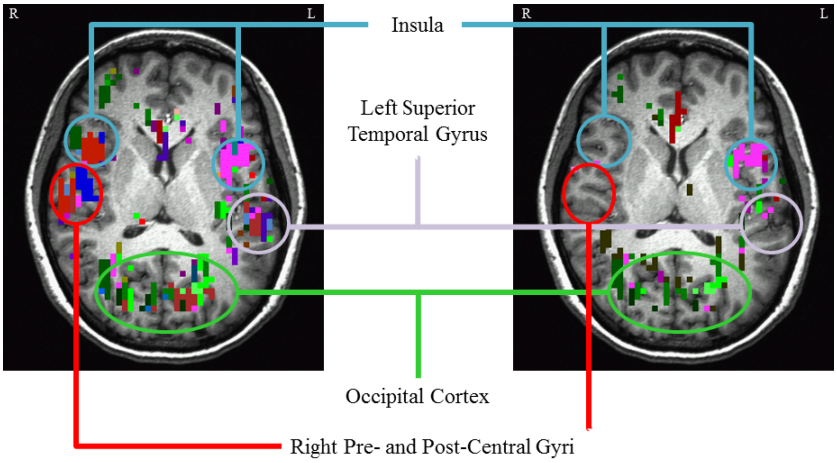

(a) Axial fMRI slice showing clusters obtained by our proposed method (left) and interactive method (right).

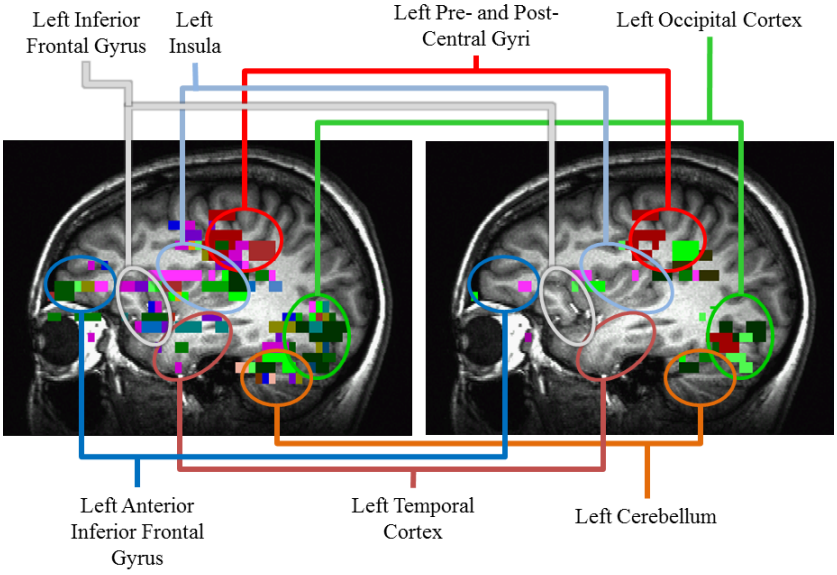

(b) Sagittal fMRI slice showing clusters obtained by our proposed method (left) and interactive method (right).

Fig. 6. Comparison of fMRI clusters obtained with our automatic HAC / $CCSA$ method to clusters extracted with interactive method. The color labels between the two different clusterings are not consistent (irreconcilable).

representation of areas attributed to visual processing (left/right occipital cortex), to motor and sensory function of the right hand (left pre- and post-central gyri, and the supplementary motor area) and limited representation of brain areas connected to the following: the bilateral activity of the subject in the motor and sensory areas (right pre- and post-central gyri), areas thought to be used in processing of human faces (left/right temporal gyri), decision making (left/right frontal gyri), memory (left/right cingulate), emotional response (left/right insula), general motor function (left/right cerebellum), components of the basal ganglia (left/right caudate nucleus, thalamus, lentiform nucleus), and the left and right precuneus.

The automatically identified clusters include all of these regions, more clearly and more completely, as well as additional real clusters. Fig. 6 shows selected fMRI slices placed over an anatomical underlay, to demonstrate the typical improvement in the discovery of functional regions. Fig. 6 (a) compares, in an axial slice, activation regions which were captured by the two methods. Differences are pointed out by ovals of different colors over the corresponding brain areas. The improvement is most notable in the right insula and the pre- and post-central gyri. Both methods perform equally well on the left insula, cuneus, and middle occipital gyrus. Fig. 6 (b) shows, in a sagittal slice, the same improvement compared to interactively identified clusters. The cerebellum, superior temporal gyrus, inferior and middle frontal gyrus, and the insula are much better captured by the automatic method. Both clusterings represent the pre- and post-central gyrus, the cuneus, and middle occipital gyrus equally well.

IV. Conclusion, Discussion, Future Work

We propose the $CCSA$ similarity measure, which advances the state-of-the-art by combining information from three distinct spaces: physical, feature, and SOM space. We use

$CCSA$ in a hierarchical clustering algorithm to automatically cluster spatial image data. The $CCSA$-based HAC clustering performs acceptably (but not excellently) on synthetic 2D data sets, and its performance increases with increasing data complexity. In particular, $CCSA$ performs significantly better on (both synthetic and real) data with more complex cluster relations and higher-D feature space.

For the Lsun and Clown data sets, the feature space comprises the 2D coordinates of the image points, therefore the feature and physical spaces coincide in these cases, and are reflected accordingly in the calculation of the $CCSA$ similarity measure. This coincidence may result in a certain degree of degeneracy of the $CCSA$ measure which, in turn, can have the counterintuitive effect of $CCSA$ being less discriminative for lower-dimensional data than for more complex, high-D data. As the data becomes more complex, each prototype, on average, is connected to more prototypes in the SOM, allowing the $CCSA$ to inform the HAC about a larger number of potential cluster merges (as $CSSA = 0$ only when two clusters are completely separated), thus providing a more significant advantage to the HAC than in the case of simple data.

We assess the $CCSA$-based clustering performance on a single-subject fMRI data set and compare it to interactively extracted clusters from the same data in O'Driscoll et al. [3]. This demonstrates increased number and quality of medically interesting functional regions discovered by our HAC / $CCSA$ method.

We observe that the vast majority of missclassifications that occur in the 2D and 6D synthetic data sets are a result of prototypes with a small receptive field. This can (and does) occur in more complex, large data sets. Therefore, one future step is to extend our $CCSA$ measure to better handle such prototypes. In addition, we can refine the initialization of the HAC, for example by removing prototypes with small receptive fields.

In follow-up work, we will examine performance on various other large and complex functional image data, such as hyperspectral remote sensing images in terrestrial and astronomical applications, as these have significantly different characteristics in both feature and physical space. Another important aspect is to examine if this method produces as consistent clustering of fMRI data across multiple subjects as the interactive approach in O'Driscoll et al. [3].

A future extension to the $CCSA$ for HAC will be to devise a stopping criterion for the $CCSA$ to "sense" the correct number of clusters. We surmise that we can achieve this by infusing information about general data characteristics into the $CCSA$.

ACKNOWLEDGMENTS

We would like to thank J. Vesanto and E. Alhoniemi for the Clown data set, and the learned SOM. We would also like to thank C. Karmonik, Houston Methodist Hospital for collecting the fMRI data set.

REFERENCES

[1] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biol. Cybern.*, vol. 43, pp. 59–69, January 1982.

[2] K. Taşdemir and E. Merényi, "Exploiting data topology in visualization and clustering of Self-Organizing Maps," *IEEE Trans. on Neural Networks*, vol. 20, no. 4, pp. 549–562, 2009.

[3] P. O'Driscoll, E. Merényi, C. Karmonik, and R. Grossman, "The effect of SOM size and similarity measure on identification of functional and anatomical regions in fMRI data," in *Advances in Self-Organizing Maps and Learning Vector Quantization*. Springer, 2016, pp. 251–263.

[4] J. Vesanto and E. Alhoniemi, "Clustering of the Self-Organizing Map," *IEEE transactions on neural networks*, vol. 11, no. 3, pp. 586–600, 2000.

[5] K.-H. Chuang, M.-J. Chiu, C.-C. Lin, and J.-H. Chen, "Model-free functional MRI analysis using Kohonen clustering neural network and fuzzy C-means," *IEEE Transactions on Medical Imaging*, vol. 18, no. 12, pp. 89–94, 1999.

[6] M.-J. Chiu, C.-C. Lin, K.-H. Chuang, J.-H. Chen, and K.-M. Huang, "Tissue segmentation-assisted analysis of fMRI for human motor response: An approach combining artificial neural network and fuzzy C means," *Journal of Digital Imaging*, vol. 14, no. 1, pp. 38–47, 2001.

[7] S. J. Peltier, T. A. Polk, and D. C. Noll, "Detecting low-frequency functional connectivity in fMRI using a Self-Organizing Map (SOM) algorithm," *Human Brain Mapping*, vol. 20, pp. 220–226, 2003.

[8] W. Liao, H. Chen, Q. Yang, and X. Lei, "Analysis of fMRI data using improved Self-Organizing Mapping and spatio-temporal metric hierarchical clustering," *IEEE Transactions on Medical Imaging*, vol. 27, no. 10, pp. 1472–1482, 2008.

[9] M. L. Gonçalves, M. L. A. Netto, J. A. F. Costa, and J. Z. Júnior, "An unsupervised method of classifying remotely sensed images using Kohonen Self-Organizing Maps and agglomerative hierarchical clustering methods," *International Journal of Remote Sensing*, vol. 29, no. 11, pp. 3171–3207, 2008.

[10] D. Sommer and M. Golz, "Clustering of EEG-segments using hierarchical agglomerative methods and Self-Organizing Maps," in *International Conference on Artificial Neural Networks*. Springer, 2001, pp. 642–649.

[11] K. Taşdemir, P. Milenov, and B. Tapsall, "Topology-based hierarchical clustering of Self-Organizing Maps," *IEEE Transactions on Neural Networks*, vol. 22, no. 3, pp. 474–485, 2011.

[12] E. Merényi, K. Taşdemir, and L. Zhang, "Learning highly structured manifolds: Harnessing the power of SOMs," in *Similarity-Based Clustering*, ser. Lecture Notes in Computer Science, M. Biehl, B. Hammer, M. Verleysen, and T. Villmann, Eds. Berlin Heidelberg: Springer Verlag, 2009, vol. 5400, pp. 138–168.

[13] A. Ultsch, "Maps for the visualization of high-dimensional data spaces," in *Proc. Workshop on Self Organizing Maps*, 2003, pp. 225–230.

[14] E. Merényi, A. Jain, and T. Villmann, "Explicit magnification control of Self-Organizing Maps for "forbidden" data," *IEEE Transactions Neural Networks*, vol. 18, no. 3, pp. 786–797, 2007.