

Interpreting PET scans by structured patient data: a data mining case study in dementia research

Jana Schmidt · Andreas Hapfelmeier ·
Marianne Mueller · Robert Perneczky ·
Alexander Kurz · Alexander Drzezga · Stefan Kramer

Received: 31 December 2008 / Revised: 23 April 2009 / Accepted: 9 May 2009 /
Published online: 23 July 2009
© Springer-Verlag London Limited 2009

Abstract One of the goals of medical research in the area of dementia is to correlate images of the brain with clinical tests. Our approach is to start with the images and explain the differences and commonalities in terms of the other variables. First, we cluster Positron emission tomography (PET) scans of patients to form groups sharing similar features in brain metabolism. To the best of our knowledge, it is the first time ever that clustering is applied to whole PET scans. Second, we explain the clusters by relating them to non-image variables. To do so, we employ RSD, an algorithm for relational subgroup discovery, with the cluster membership of patients as target variable. Our results enable interesting interpretations of differences in brain metabolism in terms of demographic and clinical variables. The approach was implemented and tested on an exceptionally large data collection of patients with different types of dementia. It comprises 10GB of image data from 454 PET scans, and 42 variables from psychological and demographical data organized in 11 relations of a relational database. We believe that explaining medical images in terms of other variables (patient records, demographic information, etc.) is a challenging new and rewarding area for data mining research.

J. Schmidt · A. Hapfelmeier · M. Mueller · S. Kramer (✉)
Institut für Informatik/I12, TU München, Garching b. München, Germany
e-mail: stefan.kramer@in.tum.de

M. Mueller
e-mail: marianne.mueller@in.tum.de

J. Schmidt
e-mail: jana.schmidt@in.tum.de

R. Perneczky · A. Kurz
Klinik u. Poliklinik für Psychiatrie, u. Psychotherapie, TU München,
München, Germany

A. Drzezga
Nuklearmedizinische Klinik, TU München, München, Germany

Keywords PET · Clustering · Subgroup discovery · Alzheimer’s disease · Dementia · Brain · Neuro imaging · CERAD · CDR

1 Introduction

Every year, 200,000 people are diagnosed with some type of dementia in Germany and 40% of them with Alzheimer’s disease [1]. Due to the “aging society”, this figure is expected to increase continuously over the next decades. Although some forms of dementia like Alzheimer’s disease are well-known and characterized for one 100 years now, the underlying mechanisms are still not sufficiently understood. Therefore, there is great interest in expanding our knowledge of various forms of dementia, including Alzheimer’s disease.

Generally speaking, the symptoms of Alzheimer’s disease are caused by the deposition of pathological proteins in the form of intracellular tangles and extracellular plaques. This deposition is followed by neuron death and deficits in neurotransmitter systems. Furthermore, deficits in glucose metabolism occur, which can be assessed in-vivo by [18] F-FDG Positron-Emission-Tomography (FDG-PET) to detect regional functional pathology.

Given the results from such neuroimaging studies, one of the major goals of medical research is to correlate them with other non-image based variables (e.g., demographic information or clinical data). The usual approach is to select a subset of patients fulfilling specific predefined criteria (e.g., the level of cognitive impairment) and to compare the images associated with those patients to data from healthy controls in a group analysis. However, it is clear that such an approach can never be guaranteed to be complete: If the first step misses an important subset, it is not possible to recover from this omission in subsequent steps.

Therefore, we propose to apply data mining techniques to take the opposite approach: to start with the images and explain the differences and commonalities in terms of non-image variables. In this way, the results of the analysis are less dependent on the choices made in the selection of patients. In fact, the goal is to obtain a complete list of descriptions of subgroups of patients, which are unusual with respect to the Positron emission tomography (PET) images. The approach was implemented and tested on data derived from an exceptionally large pre-existing data collection of patients with different types of dementia, collected at our university hospital. In the first step, we clustered FDG-PET scans of patients to form groups sharing similar features in brain metabolism. In the second step, we explained the clusters by relating them to clinical and other non-image variables. To do so, we employed RSD [10, 17], an algorithm for relational subgroup discovery, with the cluster membership of patients as the target variable. After extracting relevant information from 200GB of data (removing duplicates, intermediate results, and incompletely processed images), we obtained a dataset comprising 10GB of image data from 454 PETs, and 42 variables from clinical and demographical data organized in 11 relations of a relational database. Large image clusters identified metabolic patterns corresponding well to typical findings in major types of dementia. Furthermore, the approach allowed the detection of differences in cognitive performance in presence of comparable brain pathology, thus potentially helping to identify factors supporting compensation (e.g., age, gender, education).

In summary, the contributions of the paper are as follows: First, we present a new application area and task for data mining in a highly relevant area of medical research. Second, we present the first clustering of whole PET scans. Third, we propose, motivated by medical considerations, a new type of correlation analysis based on a loose coupling of clustering and subgroup discovery. Fourth, the procedure itself is novel in the medical area, as the approach

is diametral to current practice in the analysis of PET images and deemed highly relevant by medical experts in the field.

This paper is organized as follows: We start with a description of the data set including the preprocessing steps (Sect. 2). In Sect. 3, we subsequently explain our workflow and how we applied the clustering and subgroup discovery algorithms to the data set. Section 4 presents our results and their interpretation by medical experts. In Sect. 5, we compare our approach to itemset constrained clustering. Section 6 discusses the results from a higher perspective. The paper closes with a review of related work (Sect. 7) and the overall conclusion.

2 Dementia data

The data were provided by the psychiatry and nuclear medicine departments of Klinikum rechts der Isar of Technische Universität München. It consists of demographic information, clinical data, including neuropsychological test results, and PET scans showing the patient's cerebral glucose metabolism. We had access to clinical and demographic data of 4,037 patient visits and 454 PET scans that have been collected between 1995 and 2006.

To increase the quality of the data, we revised the existing psychological data of the 1,100 visits of patients having a corresponding PET or cerebrospinal fluid examination (test for certain protein levels). Our revision included the correction of typing errors and the completion of electronically available test results. For some patients with PET, a revision was not possible due to missing patient records. 257 PETs belong to patients with revised psychological and demographic values. The overall effort for revising the data was approximately four person months.

2.1 Positron emission tomography

Positron emission tomography is a non-invasive medical imaging procedure that has been used for diagnostics of dementia since the early eighties. It displays a three-dimensional map of the glucose metabolisms of the body and is based on the decay of radioactive markers, which are injected into the patient. A scanner records the cell activation, and a computer calculates the three-dimensional image of the metabolism. The recorded PETs in this study indicate the metabolism of the brain, i.e., the transformation of glucose. This reflects the activity of neural cells. The brain of patients suffering from dementia contains regions where the metabolism is clearly lowered. In Alzheimer's disease, the pattern of hypometabolism starts at the hippocampus and spreads over the entire cortex as the disease progresses, sparing few areas such as the motor and primary visual cortices.

Before physicians can actually use the PET scans, the images have to be processed by a sequence of transformation steps. The image preprocessing pipeline of our study is illustrated in the upper left-hand part of Fig. 1. Due to measurement irregularities and the motion of patients during the recording of the images, each image has to be rotated and translated, such that they all fit into the same template. This is achieved by SPM5.¹ Subsequently, the images are forwarded to (X)MedCon,² which transforms the data into raw ASCII files representing the intensity of each voxel as a real value. The last step is the normalization by dividing each voxel value by the mean voxel value of the image. At the end of the preprocessing, each file consists of 69 matrices of 79 rows and 96 columns, summing up to 523,296 voxels. Each

¹ SPM5 release 1 December 2005, based on Matlab 7.3, see <http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>.

² XMedcon 0.9.9.3, <http://xmedcon.sourceforge.net/>.

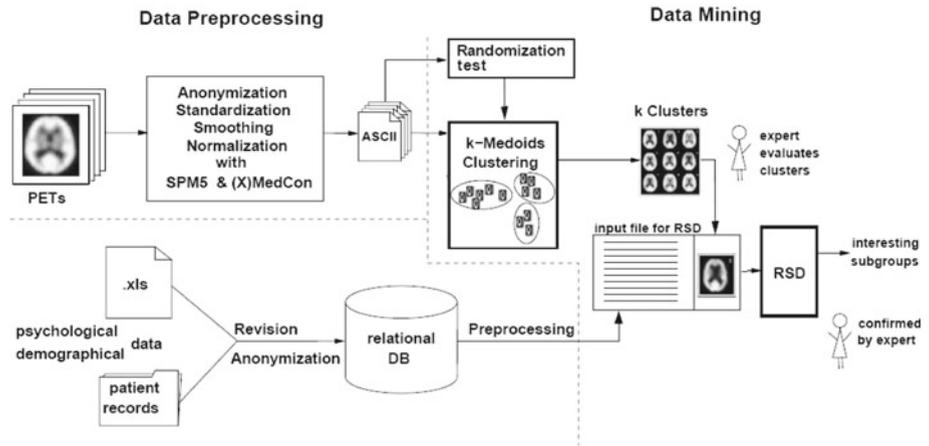


Fig. 1 Workflow of our approach. The upper part shows the processing of image data, the lower part the processing of the structured non-image data. Preprocessing steps are on the left side, data mining and interpretation of results on the right

Fig. 2 Mean image of 20 healthy controls



of the 69 layers reflects a horizontal cut through the brain, presenting a two-dimensional image of the layer with voxels displaying the intensity of metabolism of the corresponding region.

In order to visualize a group of scans, we computed a *mean image* using SPM5. In a mean image, each voxel contains the mean voxel value at that position. The mean image has the advantage of showing the overall properties of a group. Note that each mean image shown in this paper displays only one layer of the three-dimensional mean image. We always choose the same layer (layer 32) to facilitate a comparison of the images. However, for the evaluation of clusters, we take into account all 69 layers.

Figure 2 shows the mean image of 20 healthy controls. The “north” displays the frontal region of the brain from an top-down view. The lighter the spots, the more active is the underlying tissue. Both sides of the brain are symmetric in their metabolism. The lateral regions have an increased metabolism, while the center has lower metabolism. This is due to the fact that the brain cells are residing in the outer areas of the brain, whereas the brain center mostly consists of dendrites.

2.2 Demographic data

We had access to demographic data of 4,037 patient visits (751 with revised data). The data covers gender, age, years of education, type of graduation and profession.

2.3 Neuropsychological data

The clinical data for each patient consist of both psychological test results and information about which of the other tests (e.g., MRI, CT, SPECT and EEG) were completed during a hospital visit. Additionally, the diagnosis for each visit is provided in the form of ICD-10 codes [19].

The assessment instruments are the standard tests to diagnose dementia: Consortium to Establish a Registry for Alzheimer's Disease Neuropsychological Assessment Battery (*CERAD*), Clinical Dementia Rating (*CDR*), and the clock-drawing test *CDT*. These tests are optionally done at each patient visit at the physician's discretion.

The CERAD test consists of eight subtests which evaluate the patient's cognitive abilities in the areas of semantic memory, word finding, visual cognition, orientation, concentration, direct retentiveness, visuo-construction, and delayed retentiveness. These eight subtests are called verbal fluency (achievable score: $0-\infty$), Boston Naming Test (*BNT*) (0–15), Mini-Mental-State Examination (*MMSE*) (0–30), word list learning (0–30), constructional praxis (0–11), word list recall (0–10), word list recognition (0–20) and constructional praxis recall (0–11). The healthier a person, the more points are expected. It is possible to calculate a score (0–100) over all CERAD subtests [3]. This score was also integrated in our study.

The CDR consists of seven scores, which describe the capability of patients to handle their daily life. It delivers scores for memory, orientation, judgment, community, activities, personal care, and a global score. Possible values for the scores are 0, 0.5, 1, 2 or 3, where 0 is the best result.

For the CDT, the patient has to draw a clock which shows ten past eleven. The patient is graded in a range from one to six, depending on how well the clock is drawn, and if it shows the correct time. A score of one reflects the perfect clock.

3 Method

The data mining part of our approach is illustrated on the right-hand side of Fig. 1. In the first step, we apply *k*-Medoids clustering to the image data. The clusters of the best clusterings according to clustering quality and expert evaluation are further interpreted by the corresponding non-image data. For the interpretation of image clusters, we employ relational subgroup discovery (RSD) [10, 17], an algorithm for finding interesting subgroups in data. In subgroup discovery, the goal is to find subgroup descriptions (typically conjunctions of attribute values as in rule learning) for which the distribution of examples with respect to a specified target variable is "unusual" compared to the overall target distribution. In our case, clinical and demographic variables form the subgroup descriptions, and the cluster membership is chosen as the target variable. Before applying RSD, we still remove images with incomplete corresponding non-image data as well as all clusters below a certain size.

In the following, we give a short description of the clustering and subgroup discovery algorithms and explain how we applied them to our data.

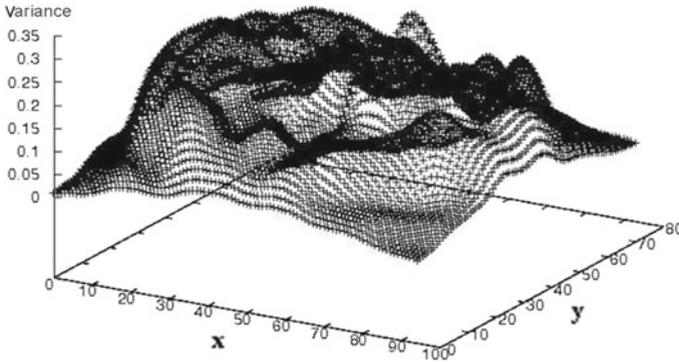


Fig. 3 Distribution of variance (over all patients) of voxels from the 32nd layer of the PET scans

3.1 Clustering

We chose the k -Medoids algorithm [9], a derivate of k -Means [20], with variance-weighted features to cluster PET scans. Alternative clustering approaches and their results will be discussed in Sect. 6. As mentioned earlier, the images have to be normalized and standardized before clustering. The resulting ASCII files (with over 500,000 real-valued entries) were taken as input for k -Medoids. To obtain meaningful and significant clusters using k -Medoids, it is necessary to weight the features according to their variance. This is feasible due to the huge differences in the variance of the intensity in different brain regions. These differences are caused, among others, by activation patterns specific for certain types of dementia and for healthy controls. Moreover, it is clear that not all parts of a PET scan reflect the state of brain tissue: For instance, in the PET image, the brain is surrounded by a dark area which is the same for all scans (cf. Fig. 2). Figure 3 shows the variance distribution of the voxels in the 32nd layer over all PET scans. As regions of low variance are not very informative for the clustering, we chose to weight each voxel at position (i, j, k) by the standard deviation over all PET scans s_{ijk} . Therefore, we obtain the following distance measure based on the Euclidian distance measure with the variables weighted by their standard deviation:

$$\text{dist}(A, B) = \sqrt{\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^l (s_{ijk} \cdot a_{ijk} - s_{ijk} \cdot b_{ijk})^2} \tag{1}$$

Since the optimal number k of clusters is not known in advance, we need a quality measure to compare clusterings with different k . One of the few quality measures for clusterings independent of k is the silhouette coefficient (SC) [9]. For a single cluster C of a clustering \mathcal{C} it is defined as

$$SC_C(\mathcal{C}) = \frac{1}{|C|} \sum_{x \in C} \frac{b(x, C) - a(x, C)}{\max\{a(x, C), b(x, C)\}} \tag{2}$$

where $a(x, C)$ is the distance of object x to the medoid of the cluster C it belongs to and $b(x, C)$ the distance of object x to second nearest medoid. For an entire clustering \mathcal{C}_k we used the mean of $SC_C(\mathcal{C}_k)$ over all C which is referred to as $SC(\mathcal{C}_k)$ in the following.

It always holds $-1 \leq SC \leq 1$. A high SC does not necessarily reflect the best clustering, since $SC(\mathcal{C}_k) = 1$ for all \mathcal{C}_k with $k = 1$ or $k = |X|$. Generally, $SC(\mathcal{C}_k)$ increases with $k \rightarrow |X|$, because $SC_C(\mathcal{C}_k) = 1$ for all clusters C that consist of one example only.

On our data, clustering with more than 30 clusters results in too small clusters despite their high SC s. Clustering with $k < 10$ is not informative as well, because it tends to find results consisting of two or three very large clusters and single outliers forming their own clusters. In this case, the resulting $SC(C)$ is high because outliers forming their own clusters have an SC of 1, which increases the overall SC . Therefore, the number of appropriate k was expected to be between 10 and 30.

As the results of k -Medoids depend both on the initial choice of medoids and the input order of objects, it was tested 5,000 times for each $k \in 2, \dots, 100$, and the clustering C_k^* with the maximum SC was chosen. The initial computation of the distance matrix needed 2 hours, and the 5,000 runs took 2 hours for all k on a 1 GB RAM (1.6 GHz) machine.

For the obtained SC distribution, local maxima exist at $k = 10, 12$, and 16. We hence, chose to further examine these three clusterings using subgroup discovery on the corresponding clinical and demographic data. Since our analysis showed very similar results for clustering with $k = 12$ and $k = 10$, we omit the presentation of the results for $k = 12$, and focus on only those for $k = 10$ and $k = 16$ in Sect. 4.

To determine the significance of a clustering with parameter k , an additional randomization test was performed. All original scans were shuffled to obtain 454 new images in a way that each voxel occurred again in one of the new scans, while keeping its location. So, the variance for each voxel stays the same. Subsequently, a clustering for parameter k was performed 100 times similarly to the original clustering. The resulting SC s were used to create a sorted list. For a given clustering C_k with original data, the p -value can easily be determined by the rank of the corresponding silhouette coefficient $SC(C_k)$ in this list. The final validation of the clustering is done by presenting the mean images to an expert, who interpreted them and explained details.

3.2 Subgroup discovery

Subgroup Discovery (SD) is a method for finding subgroups in a dataset that are sufficiently large and statistically unusual given their distribution on an attribute of interest. In this work, we used the subgroup discovery algorithm RSD and its publicly available Prolog implementation³ [16].

The RSD algorithm is a modification of the CN2 rule learner [5] to find subgroups in a relational dataset. The difficulty of applying a common rule learner to this task is that its covering algorithm is not designed for finding subgroups. It does not take into account examples as soon as they are covered by a rule. This implies that rules discovered in a later iteration are built on a smaller and thus biased subset of examples. Therefore, only the first few rules found by a rule learner are appropriate for subgroup discovery, i.e., they have a sufficiently large coverage. To overcome this problem, RSD assigns the weight $\frac{1}{i+1}$ to each example, where i is the number of rules (subgroup descriptions) covering the example. Initially, the weight of each example is set to one. Whenever a rule is found, the weight is decreased for each covered example. To find large and interesting rules even in later iterations, RSD uses the *modified weighted relative accuracy* ($mWRAcc$) heuristic. For a subgroup with description $Cond$ and target variable $Class$, it is defined as

$$mWRAcc(Class \leftarrow Cond) = \frac{n'(Cond)}{N'} \cdot \left(\frac{n'(Class.Cond)}{n'(Cond)} - \frac{n(Class)}{N} \right) \quad (3)$$

³ <http://labe.felk.cvut.cz/~zelezny/rsd/>.

where N is the number of examples, $n(Class)$ is the number of examples in $Class$, N' is the sum of weights of all examples, $n'(Cond)$ is the sum of weights of examples covered by the subgroup, and $n'(Class.Cond)$ is the sum of weights of examples that are covered by the subgroup and actually fall into the class. $mWRACC$ tends to find rules for examples that are least frequently covered by previously discovered rules. Furthermore, $mWRACC$ ensures a balance between generality and relative accuracy. This results in shorter rules which are thus easier to comprehend compared to the outcome of a rule induction algorithm.

The parameters of RSD were set as follows: For getting the maximal number of interesting subgroups for each class, the output was increased to 20 subgroups for each class. The beam width was set to 15, and the maximal number of literals in each rule was set to 4. For the clustering results with $k = 16$, the running time of RSD on a Pentium 4 (2,8 GHz) with 1 GB RAM was approximately 8 hours.

To determine the significance of a subgroup ($Class \leftarrow Cond_i$), the following likelihood ratio score [8] is used:

$$Sig(Class \leftarrow Cond_i) = 2 \cdot \sum_{j=1}^k n(Class_j \cdot Cond_i) \cdot \log \frac{n(Class_j \cdot Cond_i)}{n(Class_j) \cdot p(Cond_i)} \quad (4)$$

It shows how significantly different the class distribution in a subgroup is from the prior class distribution. The significance can be used to estimate the p -value, which indicates the statistical significance of the rule. It is calculated from the χ^2 -distribution with $k - 1$ degrees of freedom (k =number of clusters). As usual, a subgroup is considered significant, if its p -value is below 0.05. Interesting rules were identified by checking their p -value and by expert validation.

4 Results

This section focuses on the clusterings C_{10}^* and C_{16}^* . First, we present the medical interpretation of the mean images of the obtained clusters in Sect. 4.1. Subsequently, we discuss in Sects. 4.2–4.4 the characteristics of clinical values and subgroup results from RSD with the clusters from the clusterings with $k = 10$ and $k = 16$ as the target variable. Moreover, we relate the subgroup descriptions to medical expert knowledge.

4.1 Clustering of PET scans

Both clusterings (with $k = 10$ and $k = 16$) were significant with a p -value less than 0.01 as determined by the randomization test. Generally, the clusters vary widely in their size and also encompass singleton clusters, i.e., outliers. More specifically, the distribution of cluster sizes for clustering C_{10}^* is (187, 2, 5, 2, 4, 1, 5, 1, 207, 40), meaning that the first cluster consists of 187 PETs, the second of 2, and so forth. In the following, we refer to the first cluster in this list as Cluster 0, and the i th as Cluster $i - 1$. Analogously, the distribution of cluster sizes for clustering C_{16}^* is (2, 42, 1, 8, 105, 104, 28, 40, 8, 1, 3, 61, 1, 1, 48, 1). Thus, both clusterings found outliers and other very small clusters. In fact, one image that was sorted into its own cluster was rotated upside down. It was therefore most different from all other scans. The other singletons scans were strongly deformed, for which SPM5 is unable to compensate. For instance, some patients kept their chin too close to the chest during the recording.

For all clusters containing at least five PETs, we calculated their mean images (five for clustering with $k = 10$, and nine for clustering with $k = 16$). In both clusterings, we found

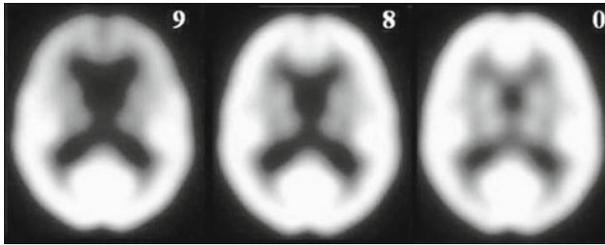


Fig. 4 Mean images of the three largest clusters of the clustering with $k = 10$

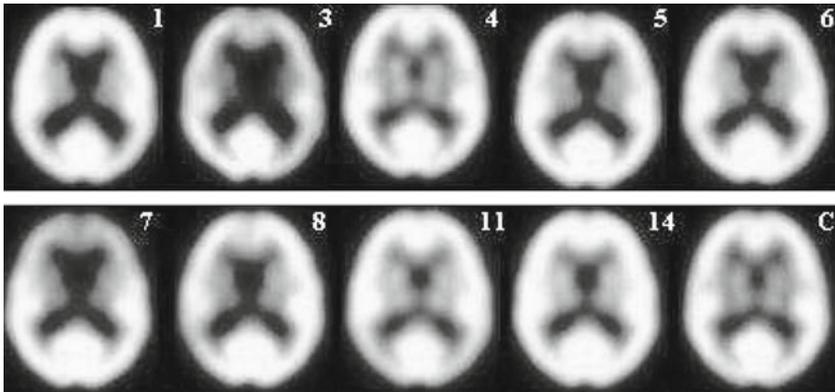


Fig. 5 Mean images of the nine largest clusters of the clustering with $k = 16$. In the second row on the right is the mean image of the healthy control group C

clusters grouping patients with frontotemporal dementia (Cluster 9 in Fig. 4 and Cluster 7 in Fig. 5), nearly healthy patients (Cluster 0 in Fig. 4 and Cluster 4 in Fig. 5) and global hypometabolism (Cluster 3 in Fig. 5 and Cluster 2 of C_{10}^* not illustrated). Clustering with $k = 16$ performed better, because it managed to differentiate more precisely the group of left lateral deficiency (Cluster 8 in Fig. 5) and a typical Alzheimer’s cluster (Cluster 11 of size 61 in Fig. 5). Both were not separated by the clustering with $k = 10$. Even though the remaining clusters cannot be interpreted clearly from a medical point of view, they can be distinguished visually. This can be seen by the differences in metabolism in the occipital and centering regions, which are lighter in Cluster 5 than in Cluster 6 (Fig. 5). Cluster 1 shows highly affected patients.

In summary, the clusterings were judged as meaningful by domain experts (R. Perneczky and A. Drzezga). To further explain the differences in the cognitive areas, we combined the images with clinical data. In the next section, we present a simple correlation analysis and relate the clusters to single non-image variables. This serves as a baseline for the more complex approach based on subgroup discovery.

4.2 Simple correlation analysis of C_{16}^*

In medicine, next to the diagnosis and test results, age and gender are important variables to describe status and progression of a disease. Therefore, we first relate the large clusters of C_{16}^* to those attributes. Looking at the distributions of diagnoses of C_{16}^* , some clusters have a high proportion of patients with Alzheimer’s disease, while others contain only patients with

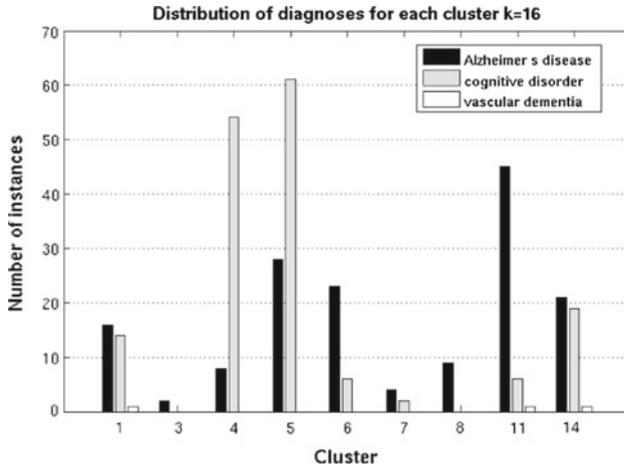


Fig. 6 Distribution of ICD10 codes for clusters of C_{16}^*

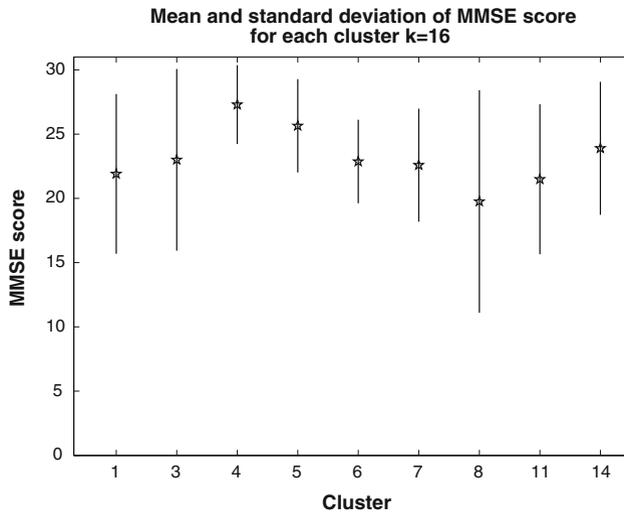


Fig. 7 Distribution of the MMSE score for clusters of C_{16}^*

a cognitive disorder (Fig. 6 clusters 4 and 11). This indicates that k -Medoids clustering is capable of grouping similar images together. This is also supported by the MMSE values of the clusters (Fig. 7). Cluster 4 has the highest MMSE score and a low variance, indicating that its patients are almost healthy. In contrast, the MMSE of Cluster 14 is not bad (around 25), but the high variance indicates that there exist some patients that suffer from a more severe disease. The same is true for clusters 8, 11, and 1. This is an interesting finding, because it states that patients with a similar brain metabolism may have different cognitive abilities. Cluster 8 comprises patients who have a left-lateral metabolic deficiency, so the low MMSE score can be explained well. Again, the high-variation states that there are some patients in Cluster 8 that can compensate the hypermetabolism and still have a good cognitive ability yielding good test results.

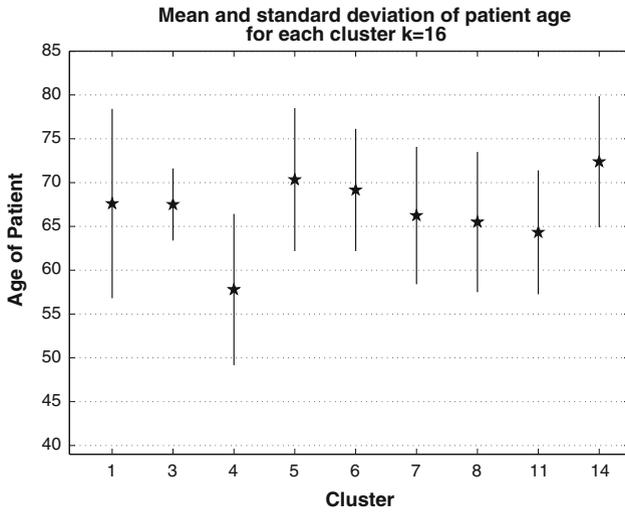


Fig. 8 Distribution of age for clusters of C_{16}^*

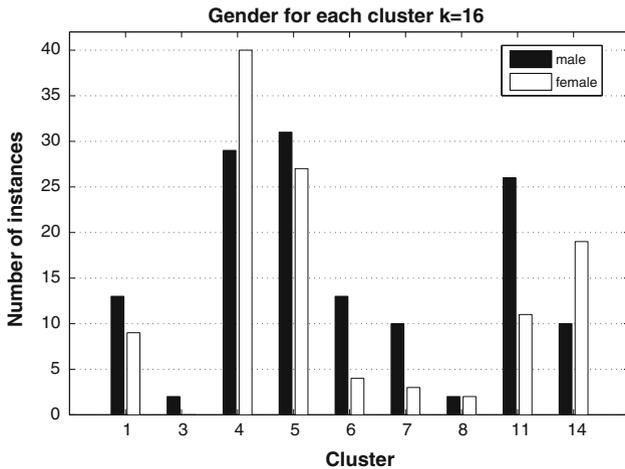


Fig. 9 Distribution of males/females for clusters of C_{16}^*

As age is correlated with the progress of dementia, we also investigated the distribution of age among the clusters (Fig. 8). The difference between Cluster 14 and Cluster 4 is approximately 20 years. The average age of Cluster 14 is around 75 while Cluster 4 has younger people of age 55, which goes hand in hand with the distribution of ICD10 codes that define Cluster 4 as an almost healthy cluster. Cluster 3 also has a very high average age and a low variance characterizing old morbid people with an advanced hypometabolism.

Concerning the gender distribution of each cluster, Clusters 4 and 14 have a higher fraction of women, while Cluster 11 holds more men (Fig. 9). This is quite interesting, because it suggests that there might be some differences in the development of dementia between the genders. Another explanation may be a different behavior as to when a physician is consulted.

Table 1 Interesting subgroups in Cluster 0 of C_{10}^*

Id.	Description	<i>p</i> -value	$\frac{ sg \cap cl }{ cl }$	$\frac{ sg \cap cl }{ sg }$	Class distr.
A1	Age ≤ 54	$<10^{-4}$	0.25	0.85	(29,5,0)
A2	Verbal fluency ≥ 20	7×10^{-3}	0.22	0.64	(25,11,0)
A4	ICD = Alzheimer's & Age: 65–69	5×10^{-2}	0.13	0.7	(14,6,0)
A9	MMST: 26–29 & BNT: 14–15 & ICD = F06.7 & Age: 55–59	2×10^{-2}	0.05	1	(6,0,0)
A11	Activities = 1 & Constructional praxis: 3–5	5×10^{-2}	0.04	1	(4,0,0)
Overall distribution over clusters 0, 8, 9:					(115,118,14)

$\frac{|sg \cap cl|}{|sg|}$ is the proportion of patients in the subgroup *sg* having a PET in cluster *cl*. In contrast, $\frac{|sg \cap cl|}{|cl|}$ expresses the proportion of the patients with a PET in cluster *cl* that are also covered by the subgroup description of *sg*. Bold numbers refer to the cluster described by the subgroups. Each number states the number of patients in the cluster covered by the particular subgroup. Numbers in the last row present the total number of patients in the clusters

This first overview of attributes in the clusters shows their general characteristics. However, simple correlation analysis is not able to detect dependencies among the attributes. Clearly, the correlation of multiple attributes can only be analyzed by more complex methods like subgroup discovery. It allows us to explore the interaction of attributes within a cluster. For instance, Cluster 14 has a high variance in MMSE scores and an unbalanced gender distribution, but so far we cannot see if men have a better MMSE score than women, or if there is no significant difference between the genders. Identifying subgroups helps to combine interesting characteristics and allows to take into account several variables at once.

4.3 Subgroup discovery on C_{10}^*

Our application of subgroup mining requires reliable psychological data and sufficiently large clusters. Thus, we first discard those examples without revised psychological data. Next, we eliminate all clusters with less than five examples after the first filtering step. Thus, we keep Cluster 0, Cluster 8, and Cluster 9, with the sample distribution (115, 118, 14) [initially, it was (187, 207, 40)]. Figure 4 shows the mean images for each of these clusters.

Table 1 displays the interesting subgroups discovered for Cluster 0. Each row in the table presents a subgroup description along with quality measures and the distribution of cases over the clusters. For instance, the first subgroup A1 covers 29 images from Cluster 0, 5 images from Cluster 8, and none from Cluster 9.

The mean image of Cluster 0 resembles the 20 healthy controls (Fig. 2), which leads to the assumption that this cluster contains mainly almost healthy patients. Subgroup A1 shows that 85% of the patients, younger than 55, fall into Cluster 0. Experts confirm that the younger a patient, the better the activity of metabolism and the lower the probability of a dementia diagnosis. Furthermore, subgroup A2 and A9 reveal that the distribution of patients with good test results is especially high in Cluster 0, which also explains the similarity of

Table 2 Interesting subgroups in Cluster 9 of C_{10}^*

Id.	Description	p -value	$\frac{ sg \cap cl }{ cl }$	$\frac{ sg \cap cl }{ sg }$	Class distr.
A41	Constructional praxis: 10–11 & Gender: m & ICD: other diag.	4×10^{-3}	0.57	0.22	(14,14, 8)
A42	Memory: 1 & Activities: 1 & Judgment: 0–0.5 & Community: 1	2×10^{-3}	0.29	0.5	(3,1, 4)
A49	Community: 2–3	2×10^{-3}	0.29	0.4	(1,5, 4)
A52	CDT: ≤ 2 & word list recall: 10 & verbal fluency: 9–12 & other tests: CT	2×10^{-4}	0.21	1	(0,0, 3)
A53	CDT: 3–6 & global: 2–3 & community: 2–3	2×10^{-3}	0.21	0.6	(0,2, 3)
A55	gender: m & verbal fluency: 0–8 & community: 2–3	2×10^{-4}	0.21	1	(0,0, 3)
Overall distribution over clusters 0, 8, 9 :					(115,118, 14)

Bold numbers refer to the cluster described by the subgroups. Each number states the number of patients in the cluster covered by the particular subgroup. Numbers in the last row present the total number of patients in the clusters

the mean image to the healthy controls. Hence, these subgroups confirm the assumption that Cluster 0 contains the healthier patients.

However, there is also one subgroup of patients with Alzheimer's diagnosis (A4) and one describing patients with a very low score for *CDR activities* and *CERAD constructional praxis* (A11). Further research revealed that these people suffer from Alzheimer's disease in an advanced state, which should not fall into this cluster. The corresponding mean images of the outliers in this cluster had the most affinity to Cluster 0 in high-variance areas. With the choice of a larger k , the algorithm can sort them out into different clusters. We can conclude that the subgroup mining step identifies four outliers within Cluster 0.

Cluster 9 (in Fig. 4) describes patients that have a huge deficit in the frontotemporal metabolism. This leads to the conclusion that they are not affected by Alzheimer's disease, but by some other form of dementia. Furthermore, experts assumed no significant reduction of *CERAD constructional praxis* (A41), while *CERAD verbal fluency* is highly impaired (A55) as well as scores in *CDR* (A42). These conclusions could be confirmed by the subgroups in Table 2.

The discovered subgroups for Cluster 8 showed that it consists mainly of patients above the age of 74. However, for a medical interpretation of this cluster, the medical experts could not obtain additional information from the mean image or the other subgroup descriptions. Thus, we assume that this cluster is too heterogeneous and that we need a clustering with a larger k to split this cluster into smaller, more homogeneous clusters.

Table 3 Interesting subgroups in Cluster 4 of C_{16}^*

Id.	Description	<i>p</i> -value	$\frac{ sg \cap cl }{ cl }$	$\frac{ sg \cap cl }{ sg }$	Class distr.
A21	BNT: 14–15 & CDT: 1–2	2×10^{-4}	0.6	0.56	(4,40 ,13,1,3,6,5)
A22	Age: 0–54	$<10^{-4}$	0.39	0.79	(3,27 ,1,0,0,2,1)
A23	Word list recall no: 10 & CERAD-sum: 76–100	$<10^{-4}$	0.42	0.7	(1,29 ,5,0,0,0,2)
A24	Constructional praxis: 10–11 & Constructional praxis recall:10–11	$<10^{-4}$	0.42	0.78	(1,29 ,7,0,1,0,3)
A26	BNT: 14–15 & MMSE: 26–29 & Verbal fluency: ≥ 20	$<10^{-4}$	0.33	0.89	(0,23 ,3,0,0,0,0)
Overall distr. over clusters 1, 4, 5, 6, 7, 11, 14:					(22, 69 ,58,17,13,37,29)

Bold numbers refer to the cluster described by the subgroups. Each number states the number of patients in the cluster covered by the particular subgroup. Numbers in the last row present the total number of patients in the clusters

4.4 Subgroup discovery on C_{16}^*

Analogously to the filtering procedure for C_{10}^* , we only kept seven clusters for the further analysis of C_{16}^* : Clusters 1, 4, 5, 6, 7, 11, and 14, containing (22, 69, 58, 17, 13, 37, 29) images [initially (42, 105, 104, 28, 40, 61, 48)]. Figure 5 shows those clusters and those having more than five images before the filtering steps (Cluster 3 and Cluster 8).

In this clustering, we find a cluster (Cluster 4) that resembles the mean image of the healthy controls. Since no hypometabolism is visible, medical experts interpreted it as representing a group of almost healthy patients. Table 3 shows a subset of the interesting subgroups discovered for Cluster 4. Groups of young patients (A22) and groups of patients with very good test results (A21, A23) were discovered. So, the clustering identified a relatively healthy group that was supported by both subgroup discovery and experts.

Cluster 7 describes patients that have a huge deficit in the frontotemporal metabolism. This leads to the conclusion that they are not affected by Alzheimer's disease, but by some other form of dementia. Furthermore, experts assumed no significant reduction of *CERAD constructional praxis*, while *CERAD verbal fluency* is highly impaired as well as scores in *CDR*. Significant subgroups that show the impairment were found for clustering with $k = 16$ and $k = 10$.

The mean image of Cluster 11 is the prototype of a patient with Alzheimer's disease, which is confirmed by the subgroups displayed in Table 4. 78.4% of the patients in this cluster have the disease, which is visible in the mean image through the reduction of metabolism in the temporoparietal cortex. In fact, this is not the only cluster with a high ratio of patients with Alzheimer's disease. In Cluster 6, 70% of the patients suffer from the disease. This was confirmed by subgroups containing an Alzheimer's diagnosis. Contrary to Cluster 11, in Cluster 6 also the frontal metabolism is reduced. This is reflected by worse (higher) results of *CDR*.

Regarding the subgroups found for Cluster 14 (Table 5), it seems that this cluster describes elderly women (A124, A129, A133) with low test results and therefore a similar state of

Table 4 Interesting subgroups in Cluster 11 of C_{16}^*

Id.	Description	<i>p</i> -value	$\frac{ sg \cap cl }{ cl }$	$\frac{ sg \cap cl }{ sg }$	Class distr.
A101	ICD: Alzheimer's	$< 10^{-4}$	0.78	0.33	(11,5,16,12,3, 29 ,13)
A102	BNT: 14–15 & gender: m & ICD: Alzheimer's	$< 10^{-4}$	0.32	0.71	(0,0,3,1,1, 12 ,0)
A106	age: 55–59 & ICD: Alzheimer's	5×10^{-4}	0.19	0.58	(1,1,0,3,0, 7 ,0)
A107	age: 65–69 & ICD: Alzheimer's	10^{-4}	0.3	0.58	(3,1,3,1,0, 11 ,0)
Overall distr. over clusters 1, 4, 5, 6, 7, 11 , 14:					(22,69,58,17,13, 37 ,29)

Bold numbers refer to the cluster described by the subgroups. Each number states the number of patients in the cluster covered by the particular subgroup. Numbers in the last row present the total number of patients in the clusters

Table 5 Interesting subgroups in Cluster 14 of C_{16}^*

Id.	Description	<i>p</i> -value	$\frac{ sg \cap cl }{ cl }$	$\frac{ sg \cap cl }{ sg }$	Class distr.
A121	Graduation < high school & Gender: f	2×10^{-1}	0.52	0.24	(4,17,15,3,1,8, 15)
A124	ICD: Alzheimer's & Personal: 0–0.5 & Gender: f & Age: 70–73	2×10^{-3}	0.17	0.83	(0,0,0,1,0,0, 5)
A126	wordlist recall: 7–10 & Gender: m & Age: 65–69 & MMSE: 30	2×10^{-1}	0.07	1	(0,0,0,0,0,0, 2)
A127	other tests: MR & Graduation < high school & Age: 74–77	7×10^{-3}	0.24	0.54	(1,1,3,1,0,0, 7)
A129	Gender: f & MMSE: 11–20 & Age: 70–73	10^{-2}	0.17	0.71	(0,0,1,0,0,1, 5)
A133	Gender: f & CERAD-sum: 47–57 & Age: ≥ 78	10^{-2}	0.14	1	(0,0,0,0,0,0, 4)
Overall distr. over clusters 1, 4, 5, 6, 7, 11, 14 :					(22,69,58,17,13,37, 29)

Bold numbers refer to the cluster described by the subgroups. Each number states the number of patients in the cluster covered by the particular subgroup. Numbers in the last row present the total number of patients in the clusters

Table 6 Comparison of the gender distribution in Cluster 14 of C_{16}^*

	Women (<i>n</i>)	Men (<i>n</i>)
Age	75 (19)	70 (10)
MMSE	23.26 (19)	25.22 (9)
global	0.97 (15)	0.8 (5)
CERAD-Sum	53.64 (14)	62.5 (8)
CDT	3.93 (15)	3.66 (9)

dementia. Surprisingly, there is a group of men with high-MMSE scores (A126). Although this subgroup is not significant, it is highly interesting. It indicates that men with the same metabolic patterns as women have a less impaired cognitive ability. Further investigations (Table 6) showed that men and women are in the same age group, but men do have slightly better overall results in the most important psychological tests. Although female and male patients fall into the same cluster (based on brain metabolism), they apparently differ in their cognitive abilities. This finding may possibly be explained by the hypothesis of cognitive reserve, which postulates that some individuals can somehow offset the symptoms of neurodegeneration. Although the neurobiological substrate is still unknown, the higher neuron count in men might be associated with higher reserve. To show the same symptoms of dementia as women, men have to suffer from a larger loss of cells. Another factor discussed in dementia research [14] is the education level, which is also higher among the men in this cluster, compared to the women. Here we now see an example of the power of subgroups. The method allows not only to describe sets of instances, but also to bring up groups with unusual and therefore interesting features, which cannot be found by simple correlation studies. Altogether, we can conclude that the clustering with $k = 16$ produces more meaningful clusters than the clustering with $k = 10$. Even though all subgroups displayed are statistically significant (see the p -values), C_{16}^* differentiates more accurately. For example, it detects an “Alzheimer cluster” (Cluster 11), whereas the distribution of patients with Alzheimer’s disease in C_{10}^* is (39, 48, 5). Therefore, none of the clusters in C_{10}^* shows a preference of being a definite “Alzheimer cluster”.

5 Comparison with constrained clustering

Another possible approach is to apply methods from constraint-based clustering (e.g., by Sese et al. [15]), such that only clusters constrained by descriptions of non-image variables are considered. This is similar to the usual approach in medicine: select a subset of patients fulfilling specific predefined criteria and compare the images associated with those patients. Automating this process results in determining frequent itemsets based on the structured non-image data. This can easily be achieved with the Apriori algorithm [20]. First, we select a subset of patients covered by one frequent itemset. Then we evaluate the similarity of their PET scans by the mean of the pairwise weighted Euclidean distance. In this way, we obtain a mean distance for each itemset.

The gray diamonds in Fig. 10 represent the discovered itemsets, where the x -axis corresponds to their mean distance and the y -axis to the size of the subset covered by the itemset. We found 5,858 frequent itemsets for a *minsupport* of 0.1. For comparison, the white squares represent the large clusters from clustering C_{16}^* . We can see that Cluster 4 (similar to healthy controls) has a lower mean distance (higher similarity of brain activity) than any of the item-

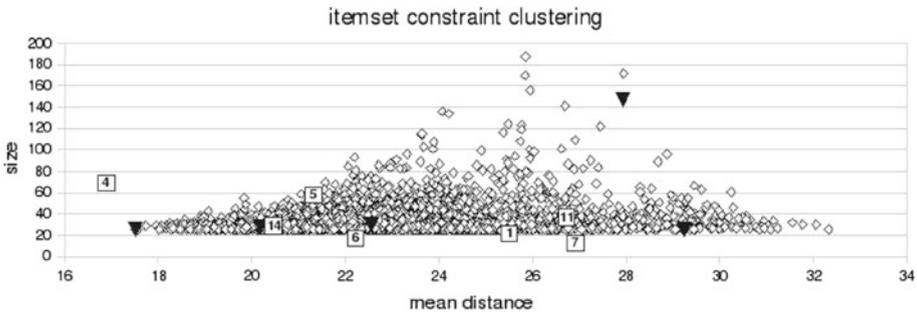


Fig. 10 Mean within-cluster distance (x-axis) and support (y-axis) of resulting itemsets (clusters) with a minimum relative support of 0.1

set constraint groups. However, there is a large number of itemsets with more homogeneous PET scans than Cluster 11 and Cluster 7. This is a result of the clustering approach, where the task is to find a set of disjoint clusters whose union covers the entire example set. Therefore, it is necessary to also put patients with dissimilar PET scans into a cluster.

The clustering task constrained by itemsets means finding a set of disjoint frequent itemsets that together cover all examples. We do that in a greedy way: we determine the itemset with the most homogeneous PET scans. Then we determine the next most homogeneous itemset which does not overlap with the previously determined itemsets and so on. When no more disjoint itemset can be found, we combine the yet uncovered, remaining examples into a final set. The black triangles in Fig. 10 represent the set of disjoint itemsets. It shows that two of the disjoint itemsets have higher mean distance than the clusters in clustering C_{16}^* .

In summary, we can say that itemset-constrained clustering provides physicians with groups of patients that share similar psychological and similar metabolic patterns (itemsets with low mean distance). Additionally, it finds groups with similar psychological patterns and high variation in brain activity (itemsets with high mean distance). However, this approach is not able to detect a complete set of patients with similar brain activity. On the other hand, clustering of PET scans finds groups of patients with similar brain activity. That means our approach can combine all patients with a particular metabolic pattern, even patients that differ in their psychological features. Moreover, we automatically create an overview of the status of disease and provide hypotheses to be further validated by medical experts.

6 Discussion

As they reflect, in some sense, the “ground truth” of the state of a brain, we chose PET images as the starting point of our analysis. The differences between different states are then explained by non-image variables. The presented approach uses k -Medoids and RSD to achieve those tasks, but we also considered and tested other methods for clustering and correlation analysis. Due to the high dimensionality of the data, we also tested the sub-space clustering method PreDeCon [2], but could not identify appropriate parameter settings to obtain significant results. As a further test, we clustered the images with hierarchical agglomerative clustering (complete linkage). This method also found the outliers mentioned above and sorted the relatively healthy patients into one cluster. However, although the two clustering methods produced largely similar results, they differed in the partition of unspecific clusters, which may be explained by the “myopia” of the hierarchical clustering scheme. To

establish a baseline for the subgroup discovery, we performed a simple correlation analysis based on individual variables. Our experiments confirmed that this is clearly not sufficient, as many of the more complex subgroups cannot be detected in this way.

From a medical point of view, we showed that it is possible to obtain meaningful clusters of PET scan images. Medical experts could identify different patterns of disease, and confirmed that the resulting findings (clusters, outliers and subgroups) have medical novelty and significance. As mentioned above, the new approach of reversing the direction of the analysis, i.e., starting with the PET scan images, is considered new in this area.

7 Related work

Data mining in the context of dementia and Alzheimer's disease can be roughly categorized into three categories: First, machine learning approaches for the improvement of differential diagnosis. Second, data mining techniques applied to brain imaging data. Third, transcriptomics and proteomics analysis of dementia and Alzheimer's disease.

One of the earliest work from the first category deals with the prediction of the type of dementia (Alzheimer's disease, vascular dementia, or other) from a small set of demographic variables and the total scores from clinical tests [11]. The latter variables included measures of category fluency, letter fluency, delayed free recall and recognition, simple and complex attention span, visual-constructional abilities, and object naming. Similarly, Corani et al. [6] predicted different types of dementia (Alzheimer's disease, dementia with Lewy bodies, Parkinson's disease with dementia and vascular dementia) from cognitive profiles based on the Cognitive Drug Research (CoDR) system. CoDR consists of a series of computerized tests (tasks), which assess some cognitive faculties of the patient, such as memory, attention, and reaction times. The results from those tests together constitute the cognitive profile of a patient.

Work from the second category includes the one by Fung and Stoeckel [7] and Megalooikonomou et al. [12]. Fung and Stoeckel [7] classify Alzheimer's patients based on their SPECT images. Megalooikonomou et al. [12] give a survey of data mining techniques applied to brain imaging data. Clustering is applied to find groups of inter-related voxels, not to find groups of images. To the best of our knowledge, clustering methods have not been applied to whole PET scans before.

The third category includes a study of gene expression changes in patients with Alzheimer's disease [18] and a study aiming for the discovery of Alzheimer-relevant proteins [4].

The work presented in this paper differs from previous work in its combined analysis of PET images and clinical variables. Similar to work in the first category, the results are mainly useful for diagnostic purposes. Although biological information in the form of transcriptomics or proteomics data is not yet used in our approach, it is easy to incorporate it in the subgroup descriptions or even as the target for subgroup discovery.

8 Conclusion

In the paper, we introduced a new and challenging problem for data mining research: correlating large databases of PET scans with structured patient data. The goal of the work was not to develop completely new methods, but to show that current data mining methods like RSD are able to solve this large (initially 200GB) and complex (image data and 11

relations) problem and produce valid and relevant results. The task itself is critical to gain a better understanding of various forms of dementia. The presented approach aims for more completeness than previous methods. To do so, we first identified clusters of PET scans sharing similar features in brain metabolism. In the second step, we explained the differences and commonalities among those clusters in terms of clinical and demographic variables. To validate the results, we computed p -values of the clusterings and interpreted the clusters and subgroup descriptions in the light of domain knowledge. To the best of knowledge, this type of analysis has not been done before. One of the subproblems, the clustering of whole PET scans (not voxels), also has not been addressed in the literature before.

In future work, we are planning to further improve the quality of the clusterings by taking into account more advanced features (e.g., brain regions) and by developing more advanced methods specifically for high-dimensional PET images. Moreover, standard algorithms for subgroup discovery suffer from similar problems as algorithms for pattern mining and association rule mining [13]. For instance, it is necessary to filter interdependent results, as the refinement (specialization) of an “interesting” subgroup is likely to produce another “interesting” subgroup. Another limiting factor of the subgroup discovery approach is the incompleteness of the psychological data. Finally, the integration of gene expression and proteomic data could aid in the formation of mechanistic hypotheses.

In summary, we believe that explaining medical images in terms of other variables (patient records, demographic information, etc.) is a challenging new and rewarding task for data mining research.

Acknowledgments This work has been supported in part by DFG-grants (Deutsche Forschungsgemeinschaft) Project Numbers: DR 445/3-1, DR 445/4-1 (Drzezga).

References

1. Bär H, Sauer J (2006) Diagnostik und Therapie häufiger Demenzen. In: *Ärzteblatt Thüringen*. Jena, Germany, pp 413–414
2. Böhm C, Kailing K, Kriegel H-P, Kröger P (2004) Density connected clustering with local subspace preferences. In: *ICDM '04: Proceedings of the fourth IEEE international conference on data mining (ICDM'04)*. IEEE Computer Society, Washington, DC, pp 27–34
3. Chandler MJ, Lacritz LH, Hynan LS, Barnard HD, Allen G, Deschner M, Weiner MF, Cullum CM (2005) A total score for the CERAD neuropsychological battery. *Neurology* 65(1):102–106
4. Chen JY, Shen CY, Sivachenko AY (2006) Mining Alzheimer disease relevant proteins from integrated protein interactome data. *Pac Symp Biocomput* 11:367–378
5. Clark P, Niblett T (1989) The CN2 induction algorithm. *Mach Learn* 3(4):261–283
6. Corani G, Edgar C, Marshall I, Wesnes K, Zaffalon M (2006) Classification of dementia types from cognitive profiles data. In: *Proceedings of the 10th european conference on principle and practice of knowledge discovery in databases (PKDD 2006)*. Springer, Heidelberg, pp 470–477
7. Fung G, Stoeckel J (2006) SVM feature selection for classification of SPECT images of Alzheimer's disease using spatial information. *Knowl Inf Syst* 11(2):243–258
8. Kalbfleisch J (1985) *Probability and statistical inference: statistical inference*, vol 2. Springer, Heidelberg
9. Kaufman L, Rousseeuw PJ (1990) *Finding groups in data: an introduction to cluster analysis*. Wiley, New York
10. Lavrač N, Železný F, Flach PA (2003) RSD: Relational subgroup discovery through first-order feature construction. In: Matwin S, Sammut C (eds) *Proceedings of the 12th international conference on inductive logic programming*. Lecture Notes in Artificial Intelligence, vol 2583. Springer, Heidelberg, pp 149–165
11. Mani S, Shankle W, Pazzani MJ, Smyth P, Dick MB (1997) Differential diagnosis of dementia: a knowledge discovery and data mining (KDD) approach. *American Medical Informatics Association (AMIA) Annual Fall Symposium*
12. Megalooikonomou V, Ford J, Shen L, Makedon F, Saykin A (2000) Data mining in brain imaging. *Stat Methods Med Res* 9:359–394

13. Ordóñez C, Ezquerro N, Santana CA (2006) Constraining and summarizing association rules in medical data. *Knowl Inf Syst* 9(3):259–289
14. Perneckzy R, Drzezga A, Diehl-Schmid J, Schmid G, Wohlschlagler A, Kars S, Grimmer T, Wagenpfeil S, Monsch A, Kurz A (2006) Schooling mediates brain reserve in Alzheimer's disease: findings of FDG PET. *J Neurol Neurosurg Psychiatry* 77:1060–1063
15. Sese J, Kurokawa Y, Monden M, Kato K, Morishita S (2004) Constrained clusters of gene expression profiles with pathological features. *Bioinformatics* 20(17):3137–3145
16. Železný F (2003) RSD—a system for relational subgroup discovery through first-order feature construction—user's manual. 2003. v1.0
17. Železný F, Lavrač N (2006) Propositionalization-based relational subgroup discovery with RSD. *Mach Learn* 62(1-2):33–63
18. Walker PR, Smith B, Liu QY, Famili F, Valdes J, Liu Z, Lach B (2004) Data mining of gene expression changes in alzheimer brain. *Artif Intel Med* 31(2):137–154
19. World Health Organization (2005) ICD-10: International statistical classification of diseases and related health problems (Tenth Revision), 2nd edn. World Health Organization, Geneva, Switzerland
20. Wu X, Kumar V, Ross Quinlan J, Ghosh J, Yang Q, Motoda H, McLachlan GJ, Ng A, Liu B, Yu PS, Zhou Z-H, Steinbach M, Hand DJ, Steinberg D (2008) Top 10 algorithms in data mining. *Knowl Inf Syst* 14(1):1–37

Author Biographies



Jana Schmidt finished her diploma thesis in bioinformatics in 2007 and is a PhD student at the chair for Bioinformatics at the Technische Universität München since 2008. Her main research interests are machine learning and data mining techniques on medical datasets and process mining. She is also a paramedic at the Johanniter-Unfall-Hilfe.



Andreas Hapfelmeier received his diploma in bioinformatics in 2007 and is a PhD student at the chair for Bioinformatics at the Technische Universität München since 2008. His research interests focus on machine learning and data mining techniques on massive medical datasets.



Marianne Mueller received her diploma in Computer Science in 2004 from University Freiburg and is a PhD student at the chair for Bioinformatics at the Technische Universität München since 2005. Her current research interests focus on descriptive mining techniques on medical datasets.



Robert Pernecky received his MD degree in 2001 and his PhD degree in 2004 from Universität München, Munich, Germany. He was an assistant doctor at the Centre for Cognitive Disorders, Department of Psychiatry and Psychotherapy, Technische Universität München from 2002 to 2008. He is head of the Neurobiological Laboratory, Department of Psychiatry and Psychotherapy, Technische Universität München since the beginning of 2009. His research interests focus on the neurobiology of degenerative dementias, including clinical, neuroimaging, neurogenetic, and neurochemical studies. His scientific work has led to over 50 peer-reviewed journal articles and he was also awarded several research prizes for his studies.



Alexander Kurz is professor of psychiatry at Technische Universität München and head of the Centre for Cognitive Disorders at the Department of Psychiatry. He has a long record of dementia research and a broad spectrum of scientific interests including genetic risk factors of Alzheimer's disease, neuroimaging and biochemical markers of neurodegeneration, early-life determinants of cognitive impairment, psychological interventions for patients with early dementia, and dementia caregiver counseling.



Alexander Drzezga received a MD degree from Technische Universität München (TUM) and is a board certified nuclear medicine physician since 2003. Since 2003 he has been a senior physician (consultant) and since 2005 Assistant Professor for Nuclear Medicine of the Department of Nuclear Medicine at TUM. He is currently a Visiting Assistant Professor at Harvard Medical School, Massachusetts General Hospital/ Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA, USA. His research interests focus on clinical and preclinical imaging in neurodegenerative disorders, including functional and molecular imaging tools (PET, MRI), multimodal imaging approaches and small animal imaging technologies.



Stefan Kramer is a professor at the computer science department of Technische Universität München in Germany, where he leads the bioinformatics group. He has organized several conferences and workshops, co-edited special issues of journals, given invited talks and tutorials, and served on a number program committees in the areas of machine learning and bioinformatics. His research interests are in bioinformatics, cheminformatics, machine learning, and data mining.