

Adaptive prototype-based fuzzy classification

Nicolas Cebron*, Michael R. Berthold

*Nycomed Chair for Bioinformatics and Information Mining, Department of Computer and Information Science, University of Konstanz,
78457 Konstanz, Germany*

Abstract

Classifying large datasets without any a priori information poses a problem especially in the field of bioinformatics. In this work, we explore the problem of classifying hundreds of thousands of cell assay images obtained by a high-throughput screening camera. The goal is to label a few selected examples by hand and to automatically label the rest of the images afterwards. Up to now, such images are classified by scripts and classification techniques that are designed to tackle a specific problem. We propose a new adaptive active clustering scheme, based on an initial fuzzy *c*-means clustering and learning vector quantization. This scheme can initially cluster large datasets unsupervised and then allows for adjustment of the classification by the user. Motivated by the concept of active learning, the learner tries to query the most “useful” examples in the learning process and therefore keeps the costs for supervision at a low level. A framework for the classification of cell assay images based on this technique is introduced. We compare our approach to other related techniques in this field based on several datasets.

Keywords: Fuzzy clustering; Classification; Active learning; Image mining; Cell assays; Noise handling

1. Introduction

The development of high-throughput imaging instruments, e.g. fluorescence microscope cameras, resulted in them becoming a promising tool to study the effect of agents on different cell types. These devices are able to produce more than 50,000 images per day; up to now, cell images are classified by a biological expert who writes a script to analyze a cell assay. As the appearance of the cells in different assays changes, the scripts must be adapted individually. Finding the relevant features to classify the cell types correctly can be difficult and time-consuming for the user.

The aim of our work is to design a classifier that is both able to learn the differences between cell types and is easy to interpret. As we are dealing with non-computer experts, we need models that can be grasped easily. We use the concept of clustering to reduce the complexity of our image dataset. Cluster analysis techniques have been widely used in the area of image database categorization.

Especially in our case, we have many single cell images with a similar appearance that may nevertheless be categorized in different classes. Another case might be that the decision boundary between “active” and “inactive” is not reflected in the numerical data that are extracted from the cell image. Furthermore, the distribution of the different cell types in the whole image dataset is very likely to be skewed. Therefore, the results of an automatic classification based on an unsupervised clustering may not be satisfactory, thus we need to adapt the clustering so that it reflects the desired classification of the user.

* Corresponding author.

E-mail addresses: nicolas.cebron@uni-konstanz.de (N. Cebron), michael.berthold@uni-konstanz.de (M.R. Berthold).

As we are dealing with a large amount of unlabeled data, the user should label only a small subset to train the classifier. Choosing randomly drawn examples from the dataset helps to improve the classification accuracy but needs a large number of iterations to converge. Instead of picking redundant examples, it would be better to pick those that can “help” to train the classifier.

This is why we try to apply the concept of active learning to this task, where our learning algorithm has control over which parts of the input domain it receives information about from the user. This concept is very similar to the human form of learning, whereby problem domains are examined in an active manner.

After introducing the Cell Assay Image Miner in Section 2, we give an overview of state of the art techniques in Section 3 that are related to our work. We shortly revise the fuzzy c -means (FCM) algorithm with noise detection in Section 4. A sampling scheme that makes use of the fuzzy memberships is proposed in Section 5. We show results in Section 6, before drawing conclusions in Section 7.

2. Cell assay image mining

In this section we introduce the Cell Assay Image Miner, a software to explore and categorize cell assay images. A typical cell assay image is shown in Fig. 1.

To identify interesting substructures in one image, the original image must be segmented in order to calculate the features for each cell individually. Unfortunately, the appearance of different cell types can vary dramatically. Therefore, different methods for segmentation have to be applied according to the different cell types. However, the individual cells in one image tend to look similar.

Currently, good results are obtained by an approach that detects a cell nucleus in an image based on a trained neural network. After this step, a region growing is performed in a similar manner to the approach described in [15]. The result of such a segmentation step is shown in Fig. 2.

After the image has been segmented, we can calculate the features on each small subimage of a cell individually. The feature extraction module calculates features of a cell image based on the histogram (first order statistics) or based on the texture (second order statistics). There are also modules for the calculation of Zernike moments [24] and a line feature module that samples points in an image along a vector. The histogram features comprise the mean, variance, skewness, kurtosis, and entropy of the histogram.

The 14 texture features from Haralick [12] represent statistics of the co-occurrence matrix of the gray level image. Four co-occurrence matrices from horizontal, vertical, diagonal, and antidiagonal directions are averaged to achieve rotation invariance. These features provide information about the smoothness, contrast, or randomness of the image—or more general statistics about the relative positions of the gray levels within the image.

Currently, the different feature modules are not integrated to form a combined feature vector. One possibility is to assign weights to each feature in order to control its influence on the classification. At present, we use the feature

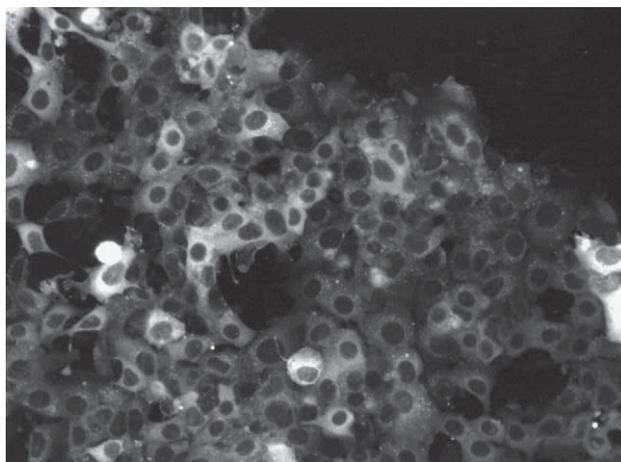


Fig. 1. Original cell image taken by a high-throughput screening microscope camera.

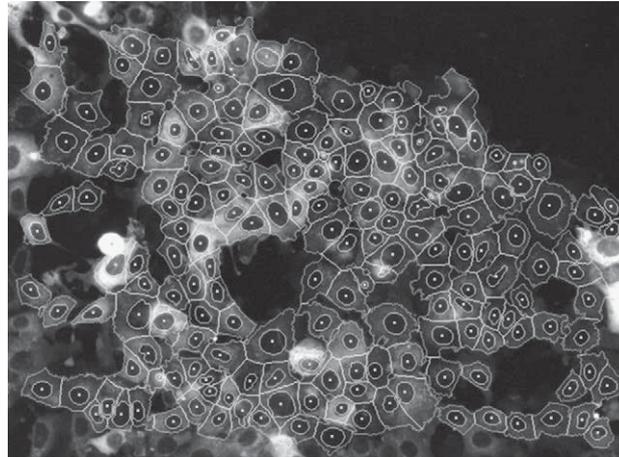


Fig. 2. Segmented cell image.

Key	Mask	000 0	000 1	D ASM	D Contrast	D Correla...	D Variance	D IDFM	D SumAv...	D SumEnt...	D Entropy	D DiffEntr...	D ICM1	D ICM2	D Mean
001 001_1				0.003	40.951	8,278,475.228	865.402	0.437	58.379	3,853,049,9...	7.219	3.357	-0.349	0.963	27.476
001 001_2				0.003	43.955	14,312,165....	1,442.186	0.412	74.935	3,850,449,5...	7.199	3.355	-0.395	0.976	35.807
001 001_3				0.006	24.457	3,298,290.137	567.331	0.542	46.609	3,858,187,6...	6.79	3.097	-0.358	0.955	21.262
001 001_4				0.001	95.142	67,850,007....	3,482.665	0.297	115.29	3,828,314,9...	7.798	3.702	-0.424	0.988	54.682
001 001_5				0.001	226.528	298,024,54....	11,173.029	0.197	203.959	3,791,608,8...	7.743	3.739	-0.502	0.996	98.043
001 001_6				0.001	269.863	104,461,02....	4,569.976	0.248	130.187	3,833,704,9...	7.377	3.529	-0.547	0.997	62.369
001 001_7				0.003	43.717	17,088,075....	1,576.672	0.394	78.399	3,854,109,0...	7.046	3.282	-0.439	0.984	36.641
001 001_8				0.003	56.669	20,702,452....	1,854.506	0.315	86.992	3,853,269,5...	6.729	3.192	-0.441	0.983	42.013
001 001_9				0.004	26.218	11,133,549....	1,424.536	0.449	76.741	3,855,797,7...	6.737	3.13	-0.391	0.969	36.068
001 001_10				0.002	34.029	27,050,971....	2,176.543	0.362	91.432	3,846,060,8...	7.428	3.486	-0.404	0.981	42.286
001 001_11				0.005	22.7	5,401,953.392	855.852	0.535	56.229	3,853,994,4...	7.074	3.239	-0.349	0.958	25.966
001 001_12				0.003	75.814	13,149,090....	1,301.413	0.448	64.114	3,842,784,0...	7.483	3.471	-0.403	0.98	30.975

Fig. 3. Table showing each cell with its corresponding mask and numerical features.

modules according to requirements of the cell assay images. In Fig. 3 we show a table with the single cell images and the Haralick features. The numerical features that we compute based on these images constitute our feature vectors. As we can see from these preprocessing steps, the number of datapoints may become very large; as we segment thousands of images into small subimages (approximately 200 small cell images per original image), we reach an order of millions of images. Our goal is to classify the original images by classifying each individual cell within.

At the beginning, we do not have any labeled instances, but we can make use of a biological expert who is able to provide a class label for each cell image that is shown to him. The problem is to classify the whole dataset with as few labeling steps as possible. We have a certain degree of freedom considering the misclassification as the whole image is classified by a majority decision over the small cell images. If a clear majority decision can be made, the image is not considered further. Borderline cases with equal distributions of classes are sorted into a special container to be assessed

manually by the biological expert. It becomes apparent that this approach allows for a rather high fault tolerance, as a human will have no objections to labeling a few images by hand rather than risk a misclassification.

In the next sections we propose a scheme that tackles this special setting by first clustering the whole unlabeled dataset unsupervised and then assigning class labels to the cluster prototypes. This classification can then be adjusted by the user; we propose a query function that tries to select the most useful examples by taking into account the fuzzy memberships.

3. State of the art

In many classification tasks it is common that a large pool of unlabeled examples U is available whereas the cost of getting a label for an example is high. The concept of active learning [6] tackles this problem by enabling a learner to pose specific queries, chosen from an unlabeled dataset. In this setting, we assume that we have access to a noiseless oracle that is able to predict the class label of a certain sample. Given an unlabeled dataset U , a labeled dataset L , and a set of possible labels C , we can describe an active learner as a tuple (f, q) . $f : L \mapsto C$ is the classifier, trained on the labeled (and sometimes also the unlabeled) data. The query function q makes a decision based on the currently labeled samples, which examples from U should be chosen for labeling. The active learner returns a new classifier f' after each pool query or a fixed number of pool queries.

For the sake of completeness, we mention also two other settings in active learning: in stream-based active learning [9] (an online version of pool-based active learning) a learner receives a stream of unlabeled examples and has to decide for each example whether to query its label or not. Especially the Query by Committee algorithm should be mentioned in this setting. It induces an even number of classifiers: whenever they disagree on an example, this example is selected for labeling.

The second setting is the selective sampling approach [1], where the learner is free to construct useful examples and then requests their label. Current research on theoretical foundations of active learning are rare, recently [7] gave lower and upper bounds for the number of labels needed with a greedy active learning strategy.

Many active learning strategies for different kinds of algorithms exist. In [6], a selective sampling is performed according to where the most general and the most specific hypotheses disagree. The hypotheses were implemented using feed-forward neural networks with backpropagation. Active learning with support vector machines (SVM) has also become very popular. The expensive learning process for the SVM can be reduced by querying examples with a certain strategy. In [20], the query function chooses the next unlabeled datapoint closest to the decision hyperplane in the kernel induced space. SVM with active learning have been widely used for image retrieval problems [18,21] or in the drug discovery process [22].

To model the underlying distribution of the given unlabeled data, we find it useful to use an approach that clusters the data. To date, research on approaches that combine clustering and active learning has been sparse.

In [19], clustering and active learning are combined in a possibilistic framework. The idea is to select the most representative samples to adjust the clustering in a coarse-to-fine strategy.

In [2], a clustering of the dataset is obtained by first exploring the dataset with a *farthest-first-traversal* and providing *must-link* and *cannot-link* constraints. In the second *consolidate*-phase, the initial neighborhoods are stabilized by picking new examples randomly from the dataset and again by providing constraints for a pair of datapoints.

In [11], an approach for active semi-supervised clustering for image database categorization is investigated. It includes a cost-factor for violating pairwise constraints in the objective function of the FCM algorithm. The active selection of constraints looks for samples at the border of the least well-defined cluster in the current iteration.

However, our approach differs from the others in the way that the data are preclustered before supervision enhances the classification accuracy. Thus, our scheme is able to explore and classify a large unlabeled dataset in a fast and accurate way.

4. FCM with noise detection

The FCM algorithm [3] is a well-known unsupervised learning technique that can be used to reveal the underlying structure of the data based on a similarity measure. Fuzzy clustering allows each datapoint to belong to several clusters, with a degree of membership for each one. We use the extended version from [8] for the added detection of noise.

Let $T = \vec{x}_i, i = 1, \dots, |T|$ be a set of feature vectors for the data items to be clustered, $W = \vec{w}_k, k = 1, \dots, c$ a set of c clusters. V is the matrix with coefficients where $v_{i,k}$ denotes the membership of \vec{x}_i to cluster k . Given a distance function d , the FCM algorithm with noise detection iteratively minimizes the following objective function with respect to v and w :

$$J_m = \sum_{i=1}^{|T|} \sum_{k=1}^c v_{i,k}^m d(\vec{w}_k, \vec{x}_i)^2 + \delta^2 \sum_{i=1}^{|T|} \left(1 - \sum_{k=1}^c v_{i,k}\right)^2 \quad (1)$$

$m \in (1, \infty)$ is the fuzzification parameter and indicates how much the clusters are allowed to overlap each other. The first term corresponds to the normal FCM objective function, whereas the second term arises from the noise cluster. δ is the distance from every datapoint to the noise cluster c . This distance can either be fixed or can be updated in each iteration according to the average interpoint distances. Objects that are not close to any of the cluster centers \vec{w}_k are therefore detected as having a high membership to the noise cluster. J_m is subject to minimization under the constraint

$$\forall i : 0 \leq \sum_{k=1}^{c-1} v_{i,k} \leq 1 \quad (2)$$

FCM is often used when there is no a priori information available and thus can serve as an overview technique.

5. From clustering to classification

Based on the prototypes obtained from the FCM algorithm, we can classify the dataset by first providing the class label for each cluster prototype and then by assigning the class label of the closest prototype to each datapoint.

Datapoints that are detected as noise are removed because they do not help to enhance the classification.¹ We will give reasons for doing so later.

In order to have enough information about the general class label of the cluster itself that represents our current hypothesis, we perform a technique known as *cluster mean selection* [10]. It helps us to determine the necessary number of cluster prototypes for the classification. Each cluster is split into subclusters; subsequently, the nearest neighbor of each cluster prototype is selected for the query procedure. If the class distribution within the current cluster is not homogeneous, we replace the prototype with the prototypes of the subclusters. We call this the *exploration phase*, as we are trying to get an overview of which kind of categories exist in the dataset.

A common problem is that the cluster structure does not necessarily correspond to the distribution of the classes in the dataset. The redefinition of cluster prototypes could increase the classification accuracy. We make use of the learning vector quantization (LVQ) algorithm for this task, which is described in the following section. Instead of randomly choosing prototypes for the LVQ, we use the prototypes obtained by the FCM algorithm.

5.1. Learning vector quantization

LVQ [17] is a so-called competitive learning method. The detailed steps are given in Algorithm 1. The algorithm works as follows: for each training pattern, the nearest prototype is identified and updated. The update depends on the class label of the prototype and the training pattern. If they possess the same class label, the prototype is moved closer to the pattern, otherwise it is moved away. The learning rate ε controls the movement of the prototypes. The learning rate is decreased during the learning phase, a technique known as *simulated annealing* [16]. The LVQ algorithm terminates if the prototypes stop to change significantly. One basic requirement in the LVQ algorithm is that we can provide a class label for each training point \vec{x}_i that is randomly sampled. We assume that the training set is unlabeled—however an expert can provide us with class labels for some selected examples. As we can only label a small set of examples, we

¹ For the Cellminer application one could show those examples as potentially interesting outliers to the user but for the construction of a global model they do not carry much information.

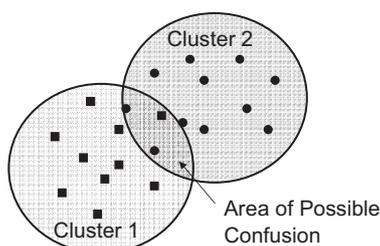


Fig. 4. Two clusters that overlap and the resulting area of possible confusion.

need to optimize the queries with a strategy to boost the classification accuracy while keeping the number of queries at a low level. In the next section, we propose a query function that attempts to solve this problem.

Algorithm 1. LVQ algorithm

- 1: Choose R initial prototypes for each class $m_1(k), m_2(k), \dots, m_R(k), k = 1, 2, \dots, K$, e.g. by sampling R training points at random from each class.
 - 2: Sample a training point \vec{x}_i randomly (with replacement) and let $m_j(k)$ denote the closest prototype to \vec{x}_i . Let g_i denote the class label of \vec{x}_i and g_j the class label of the prototype.
 - 3: **if** $g_i = g_j$ **then** {that is they belong to the same class}
 - 4: move the prototype toward the training point:
 $m_j(k) \leftarrow m_j(k) + \varepsilon(\vec{x}_i - m_j(k))$, where ε is the learning rate.
 - 5: **end if**
 - 6: **if** $g_i \neq g_j$ **then** {that is they belong to different classes}
 - 7: move the prototype away from the training point:
 $m_j(k) \leftarrow m_j(k) - \varepsilon(\vec{x}_i - m_j(k))$
 - 8: **end if**
 - 9: Repeat step 2, decreasing the learning rate ε to zero with each iteration.
-

5.2. Selection of examples based on fuzziness

The selection of new examples is of particular importance as it influences the performance of the classification. Assuming access to a noiseless oracle it is vital to gain as much information as possible from the smallest possible number of examples. The prior data distribution plays an important role, in [5] the authors propose to minimize the expected error of the learner:

$$\int_x E[(\hat{y}(x; D) - y(x))^2 | x] P(x) dx \quad (3)$$

where E denotes the expectation over $P(y|x)$ and $\hat{y}(x; D)$ the learner's output on input x given training set D . The idea is to weight the uncertainty of the classifier with the distribution of the data. If we act on the assumption that the underlying structure found by the FCM algorithm already inheres an approximate categorization, we can select further examples by querying datapoints at the partition boundaries.

We assume that the most informative datapoints lie between clusters that are not well separated from each other. We call these regions "areas of possible confusion". This coincides with the findings and results in [10,19]. Fig. 4 demonstrates this setting: There are two clusters; datapoints have been assigned the class label of their closest prototype. As we expect that the distance between similar images in the feature space is small, we can label datapoints close to the prototype with a high confidence, whereas the confidence is lower for points lying between different clusters.

To identify the datapoints that lie on the frontier between two clusters, we propose a new procedure that is easily applicable in the fuzzy setting. Rather than dynamically choosing one example for the labeling procedure (which would slow down the process), we focus on a selection technique that selects a small batch of N samples to be labeled. Note that a data item \vec{x}_i is considered as belonging to cluster k if $v_{i,k}$ is the highest among its membership values. If we consider the datapoints between two clusters, they must have an almost equal membership to both of them.

The selection is performed in two steps: Initially, all datapoints are ranked according to their memberships to cluster prototypes; subsequently, the most diverse examples are chosen from this pool of examples to avoid choosing points that are too close to each other. The ranking is based on the fuzzy memberships and can be expressed for each datapoint \vec{x}_i as follows:

$$\text{Rank}(\vec{x}_i) = 1 - (\min |v_{i,k} - v_{i,l}|) \quad \forall k, l = 1, \dots, c, \quad k \neq l \quad (4)$$

Note that we also take into account the class label of each cluster. Only if the clusters correspond to different classes is the rank computed.

After all datapoints are ranked, we can select a subset with high ranks to perform the next step: diversity selection. This prevents the active clustering scheme from choosing points that are too close to each other (and therefore are together not that interesting). We refer to the *farthest-first-traversal* [13] usually used in clustering. It selects the most diverse examples by choosing the first point at random and the next points as farthest away from the current set of selected instances. The distance d from a datapoint x to the set S is defined as $d(S, x) = \min_{y \in S} d(x, y)$, known as the *min-max distance*.

While taking into account samples at the decision boundaries between clusters, the current hypothesis should also be verified. A cluster mean selection step as mentioned in the exploration phase helps to consolidate the classification.

We summarize the procedure we have developed so far in the following section.

5.3. Adaptive active classification

Our adaptive active classification procedure is based on a combination of the techniques that have been mentioned above. All steps are listed in Algorithm 2.

The algorithm pursues two goals: 1. exploration of the dataset to get an initial classification and subsequently; 2. exploitation of the dataset to obtain a classification that corresponds more closely to the semantics of the expert. We start to cluster our dataset with the FCM algorithm, because we expect dense regions in the feature space that are likely to bear the same class label. Therefore, the FCM algorithm gives us a good initialization and prevents us from labeling unnecessary instances.

The noise detection in the clustering procedure serves the same purpose: Rare datapoints that represent borderline cases should not be selected, as these noise labels would influence the classification in a negative way. Furthermore, these samples would be useless for the classification. However, note that in this manner, we are able to present unusual and/or outlier cases to the user, that could be interesting to him.

After a batch of N examples has been selected from within each cluster and from the borders of the clusters, the user interaction takes place: the expert has to label each example. The newly labeled samples are then added to the current set of labeled samples L . After this step, the cluster prototypes can be moved based on the training set L .

Algorithm 2. Adaptive active clustering procedure

- 1: $L \leftarrow \emptyset$
 - 2: **while** Examples in Cluster have different class labels **do**
 - 3: Perform the FCM algorithm on current cluster
 with noise detection (unsupervised).
 - 4: Filter out datapoints belonging to noise cluster.
 - 5: Label cluster prototypes.
 - 6: Add the labeled prototypes to L .
 - 7: **end while**
 - 8: **while** Classification accuracy not satisfactory **do**
 - 9: $T \leftarrow$ Select m training examples at the borders.
 - 10: Select n examples from T with diversity selection.
 - 11: Ask the user for the labels of these samples, add them to L .
 - 12: Move the prototypes according to L .
 - 13: Decrease the learning rate ε .
 - 14: **end while**
-

The question is when to stop the movement of the prototypes. The simulated annealing in the LVQ algorithm will stop the movement after a certain number of iterations. However, an acceptable solution may be found earlier, which is why we propose further stopping criteria:

5.3.1. *Validity measures*

Can give us information of the quality of the clustering [23]. We employ the within cluster variation and the between cluster variation as an indicator. This descriptor can be useful for the initial selection of attributes. Naturally, the significance of this method decreases with the subsequent steps of labeling and adaptation of the cluster prototypes.

5.3.2. *Classification gradient*

We can make use of the already labeled examples to compare the previous to the newly obtained results. After the labels of the samples inside and between the clusters have been obtained, the cluster prototypes are moved. The new classification of the dataset is derived by assigning to each datapoint the class of its closest cluster prototype. By comparing the labels given by the user to the newly obtained labels from the classification, we can calculate the ratio of the number of correctly labeled samples to the number of falsely labeled examples.

5.3.3. *Tracking*

Another indicator for acceptable classification accuracy is to track the movement of the cluster prototypes. If they stop moving because new examples do not augment the current classification, we can stop the procedure.

5.3.4. *Visual inspection*

If the datapoints are linked to images (as in the setting we describe in Section 2), we can make use of them. Instead of presenting the numerical features, we select the corresponding image of the data tuple that is closest to the cluster prototype. We display the images with the highest membership to the actual cluster and the samples at the boundary between two clusters if they are in different classes.

6. Experimental results

In this section, we want to demonstrate the mode of action of our classification scheme on an artificial dataset. As the cell assay image data that we are working on are confidential, we have chosen a similar and comparable cell image dataset from the NISIS pap-smear competition. We also compare the active LVQ algorithm with active SVM [20] on the satimage dataset from the UCI repository [4].

6.1. *Artificial data*

Fig. 5 shows the two-dimensional test data in a scatterplot. The different gray tones correspond to the different classes in this dataset. This is a typical example for a dataset where the distribution of the classes is skewed. Fig. 6 clarifies the difference between random selection on the left side and examples chosen with ranking and diversity selection on the right side. The latter helps the LVQ algorithm to improve the classification accuracy more quickly as can be seen in Fig. 7, which shows the classification error in percent over the number of iterations of the LVQ algorithm.

Another issue that we want to take a look at is the benefit of batch sampling. One could argue that it is enough to determine the most interesting point at each iteration and then to move the prototypes. We perform a batch sampling that allows a diversity selection to be carried out, too. The benefit of batch sampling is demonstrated in Fig. 8, where we plot the error in percent for sampling just one datapoint at each iteration versus sampling multiple points in each iteration. In fact, the single sampling approach performs much worse than random selection in this case.

6.2. *Cell assay image data*

The task in the NISIS pap-smear competition is to classify pre-stages of cervical cancer in cells before they progress to invasive carcinoma. The data consist of 917 images of pap-smear cells, classified carefully by cyto-technicians and doctors. Each single cell image is described by 20 numerical features, and the cells fall into seven classes. A basic data analysis [14] includes linear classification results, in order to provide lower bounds on the acceptable performance

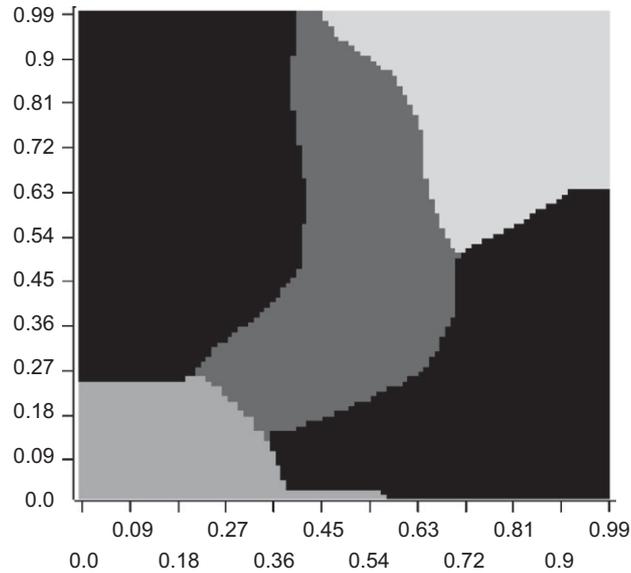


Fig. 5. Scatterplot with two-dimensional test data.

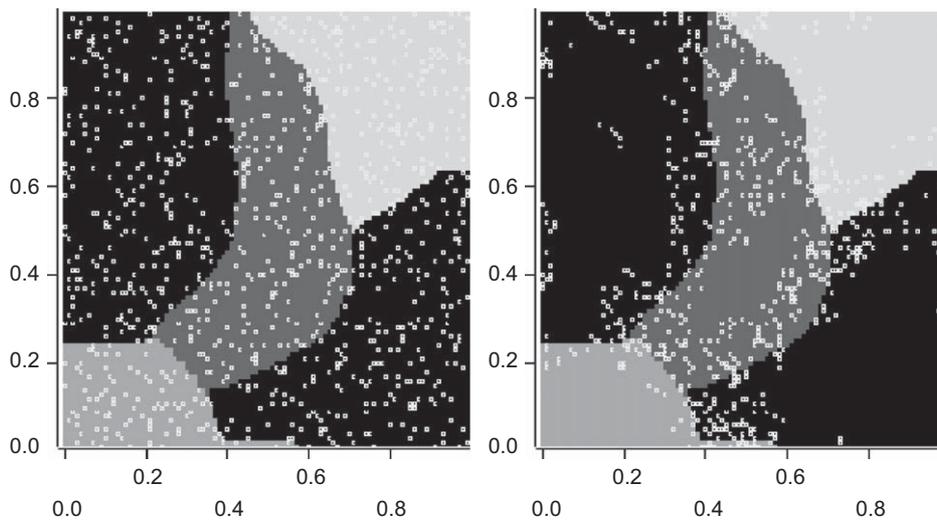


Fig. 6. Different selection techniques: random selection (left) and diversity selection (right).

of other classifiers. We compared our approach to an approach with an SVM with active learning [20], mentioned in Section 3. However, it must be noted that the active SVM is initialized differently by choosing random examples from each class. In our setting of cell assay image mining, where we have no labeled instances at the beginning, this step would not be possible, and a random initialization of the SVM would decrease the performance significantly. It must also be noted, that the performance of the active SVM depends heavily on the chosen kernel function. We used a polynomial kernel with which the active SVM performed best.

Fig. 9 shows the error rate of five test runs with our adaptive classification scheme. At the beginning, the classification error decreases significantly. After all classes have been found, it continues to decrease for further 10%.

As an exploration phase is missing in the work of [20], the variance of the classification error at the beginning is very high. With an increasing number of iterations, the classification error becomes more stable. After a number of approx. 200 training samples, the active SVM performs slightly ($\sim 5\%$) better (Fig. 10).

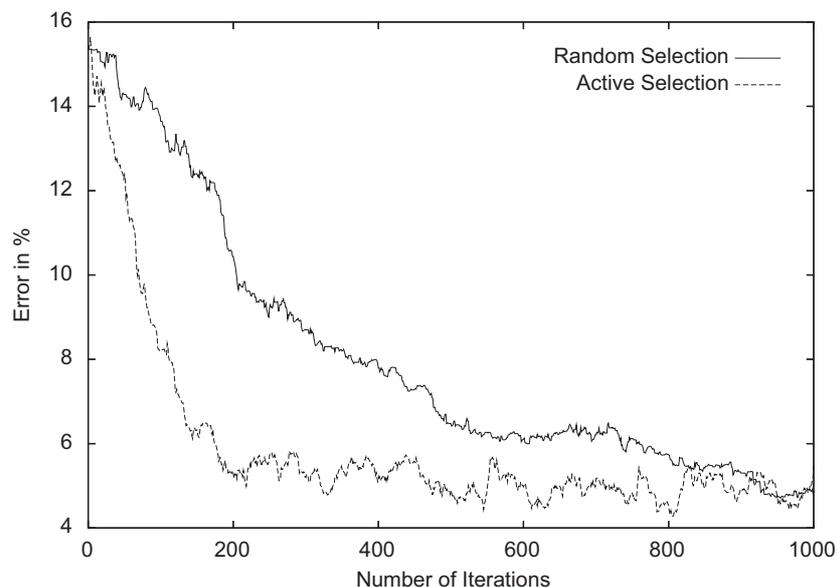


Fig. 7. Active vs. random selection.

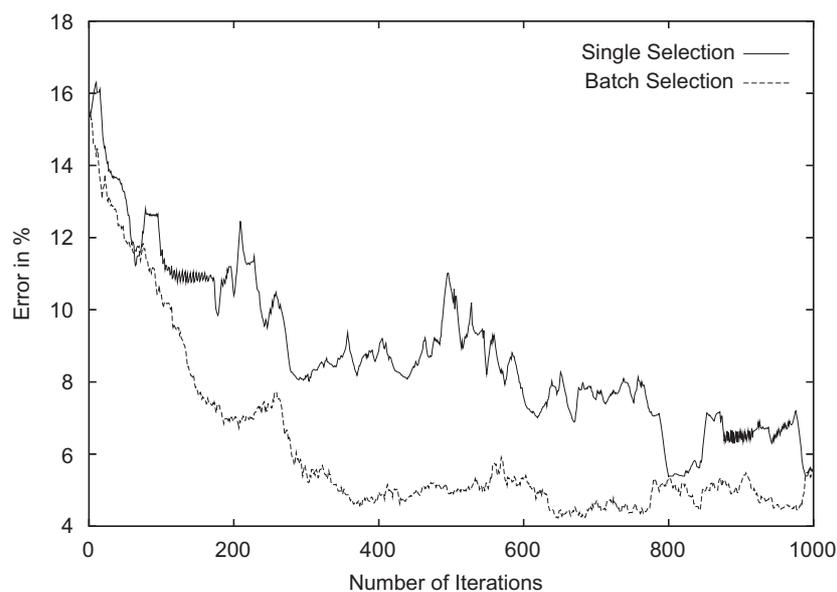


Fig. 8. Single Sampling vs. Batch Sampling.

We can observe from this result that our adaptive active classification scheme is able to tackle the problem of cell assay classification. Its performance is better than random selection and comparable with an SVM with active learning. The advantage of our active classification scheme is the better performance at the beginning, which is highly desired in our setting. The active SVM needs noticeably more time than our algorithm: the pure training time without user interaction takes 175 s for 200 samples, whereas our scheme needs 8.2 s.

6.3. Satimage data

We compared the classification error of our active LVQ algorithm with active SVM [20] on the satimage dataset [4] that contains 6435 cases split into six classes in a 36-dimensional feature space. Although this dataset does not inherit

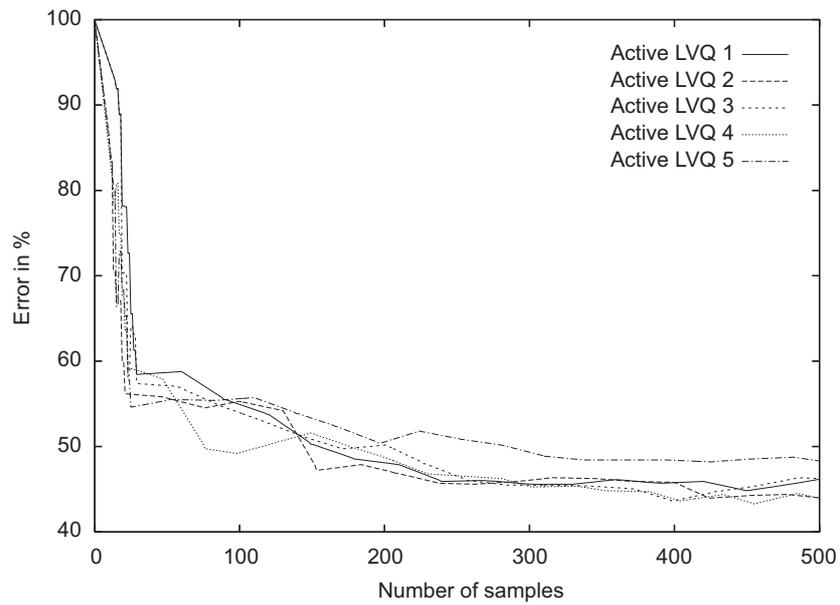


Fig. 9. Active LVQ on pap-smear dataset.

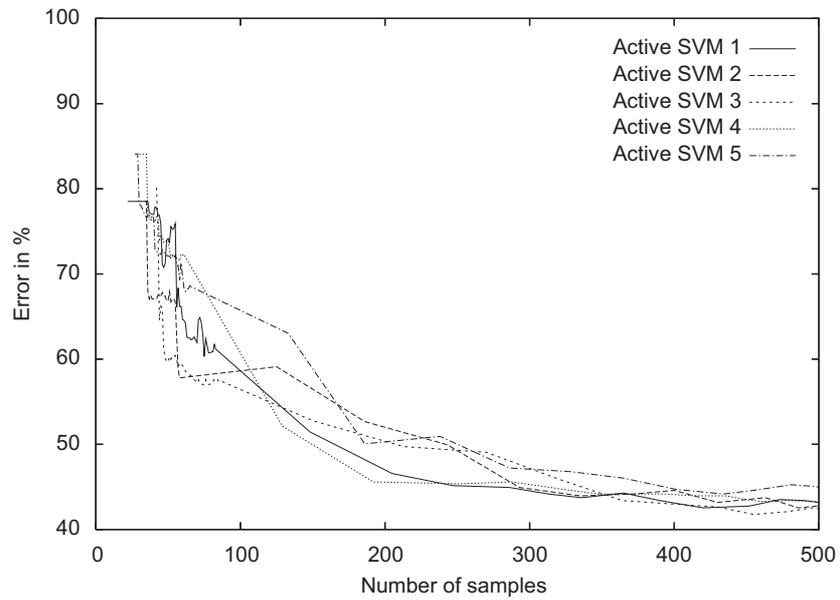


Fig. 10. Active SVM on pap-smear dataset.

the structure for which our scheme has been developed, our adaptive active classification scheme performs very stable in the first iterations, see Fig. 11.

As can be clearly seen, the active selection of datapoints in the learning process of the LVQ algorithm leads to a significantly faster convergence of the classification, especially at the first iterations. This corresponds totally to our objective of keeping user interaction at a low level.

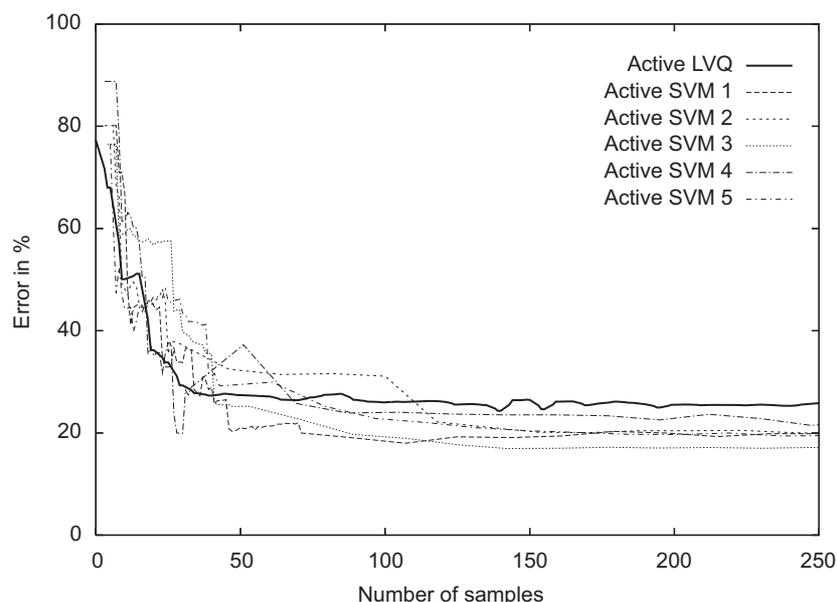


Fig. 11. Active LVQ vs. active SVM on the satimage dataset.

7. Conclusion

In this work, we have addressed the problem of classifying a large dataset when only a few labeled examples can be provided by the user. We have introduced a new adaptive active classification scheme that starts with the fuzzy c -means algorithm for an initial clustering. The classification of the dataset is obtained by labeling the cluster prototypes and assigning to all datapoints the label of the closest prototype. We have proposed to move the cluster prototypes, similar to the learning vector quantization (LVQ) method to obtain results closer to the expectation of the user. From the unlabeled pool of instances, new examples are chosen by a query function that makes use of the fuzzy memberships to the cluster prototypes combined with a diversity selection. Based on the labels of the selected examples at the borders between clusters and the labeled examples inside clusters, the prototypes are moved. We have shown that the misclassification rate can be improved more quickly. We have discussed an application in the mining of cell assay images, where the data often inherits the aforementioned properties.

Acknowledgments

This work was partially supported by DFG Research Training Group GK-1042 “Explorative Analysis and Visualization of Large Information Spaces”.

References

- [1] D. Angluin, Queries and concept learning, *Mach. Learn.* 2 (3) (1988) 319–342.
- [2] S. Basu, A. Banerjee, R.J. Mooney, Active semi-supervision for pairwise constrained clustering, in: M.W. Berry, U. Dayal, C. Kamath, D.B. Skillicorn (Eds.), *SDM*, SIAM, Philadelphia, PA, 2004.
- [3] J.C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms*, Plenum Press, New York, 1981.
- [4] C.L. Blake, D.J. Newman, S. Hettich, C.J. Merz, *UCI repository of machine learning databases*, 1998.
- [5] D. Cohn, Z. Ghahramani, M. Jordan, Active learning with statistical models, *Adv. in Neural Inform. Process. Syst.* 7 (1995) 705–712.
- [6] D.A. Cohn, L. Atlas, R.E. Ladner, Improving generalization with active learning, *Mach. Learn.* 15 (2) (1994) 201–221.
- [7] S. Dasgupta, Analysis of a greedy active learning strategy, In: *NIPS*, 2004.
- [8] R.N. Dave, Characterization and detection of noise in clustering, *Pattern Recognition Lett.* 12 (11) (1991) 657–664.
- [9] Y. Freund, H.S. Seung, E. Shamir, N. Tishby, Selective sampling using the query by committee algorithm, *Mach. Learn.* 28 (2–3) (1997) 133–168.

- [10] B. Gabrys, L. Petrakieva, Combining labelled and unlabelled data in the design of pattern classification systems, *Internat. J. Approx. Reason.* 35 (3) (2004) 251–273.
- [11] N. Grira, M. Crucianu, N. Boujema, Active semi-supervised fuzzy clustering for image database categorization, in: H. Zhang, J. Smith, Q. Tian (Eds.), *Multimedia Information Retrieval*, ACM, 2005, pp. 9–16.
- [12] R.M. Haralick, K. Shanmugam, I. Dinstein, Textural features for image classification, *SMC* 3 (6) (1973) 610–621.
- [13] D.S. Hochbaum, D.B. Shmoys, A best possible heuristic for the k -center problem, *Math. Oper. Res.* 10 (2) (1985) 180–184.
- [14] J. Jantzen, et al., Pap-smear benchmark data for pattern classification (<http://fuzzy.iau.dtu.dk/downloads/smeas2005/>), 2005.
- [15] T.R. Jones, A. Carpenter, P. Golland, Voronoi-based segmentation of cells on image manifolds, in: Y. Liu, T. Jiang, C. Zhang (Eds.), *CVBIA*, in: *Lecture Notes in Computer Science*, Vol. 3765, Springer, Berlin, 2005, pp. 535–543.
- [16] S. Kirkpatrick, C.D. Gelatt Jr., M.P. Vecchi, Optimization by simulated annealing, *Science* 220 (4598) (1983) 671–680.
- [17] T. Kohonen, The self-organizing map, *Neurocomputing* 21 (1–3) (1998) 1–6.
- [18] T. Luo, K. Kramer, D.B. Goldgof, L.O. Hall, S. Samson, A. Remsen, T. Hopkins, Active learning to recognize multiple types of plankton, *J. Mach. Learn. Res.* 6 (2005) 589–613.
- [19] H.T. Nguyen, A. Smeulders, Active learning using pre-clustering, in: C.E. Brodley (Ed.), *ICML*, ACM, 2004.
- [20] G. Schohn, D. Cohn, Less is more: active learning with support vector machines, in: P. Langley (Ed.), *ICML*, Morgan Kaufmann, Los Altos, CA, 2000, pp. 839–846.
- [21] L. Wang, K.L. Chan, Z.H. Zhang, Bootstrapping svm active learning by incorporating unlabelled images for image retrieval, in: *Proc. IEEE Comput. Soc. Conf. on Computer Vision and Pattern Recognition*, Vol. 1, 2003, pp. 629–634.
- [22] M.K. Warmuth, J. Liao, G. Rätsch, M. Mathieson, S. Putta, C. Lemmen, Active learning with support vector machines in the drug discovery process, *J. Chem. Inform. Comput. Sci.* 43 (2) (2003) 667–673.
- [23] M.P. Windham, Cluster validity for fuzzy clustering algorithms, *Fuzzy Sets and Systems* 5 (1981) 177–185.
- [24] F. Zernike, Diffraction theory of the cut procedure and its improved form, the phase contrast method, *Physica* 1 (1934) 689–704.