Eurographics Workshop on Visual Computing for Biology and Medicine (2017)
I. Hotz, D. Merhof, and C. Rieder (Editors)

Automatic thrombus detection in non-enhanced computed tomography images in patients with acute ischemic stroke

P. Löber[1], B. Stimpel[1], C. Syben[1], A. Maier[1], H. Ditt[2], P. Schramm[3], B. Raczkowski[3], and A. Kemmling[3]

[1]Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
[2]Siemens Healthcare GmbH, Forchheim , Germany
[3]Institute of Neuroradiology, University Medical Center Schleswig Holstein, Lübeck, Germany

Abstract

In case of an ischemic stroke, identifying and removing blood clots is crucial for a successful recovery. We present a novel method to automatically detect vascular occlusion in non-enhanced computed tomography (NECT) images. Possible hyperdense thrombus candidates are extracted by thresholding and connected component clustering. A set of different features is computed to describe the objects, and a Random Forest classifier is applied to predict them. Thrombus classification yields 98.7% sensitivity with 6.7 false positives per volume, and 91.1% sensitivity with 2.7 false positives per volume. The classifier assigns a clot probability $\geq$ 90% for every thrombus with a volume larger than 100 mm^3 or with a length above 23 mm, and can be used as a reliable method to detect blood clots.

CCS Concepts

•***Computing methodologies*** → *Classification and regression trees;* •***Applied computing*** → *Health care information systems;*

1. Introduction

Stroke is a serious neurological disease, and constitutes a major cause of death and disability worldwide, with 15 million people suffering a stroke per year [RDW16]. The main reason for a stroke is thrombosis, i.e., the formation of a blood clot inside a blood vessel supplying the brain. Identifying and removing the clot is crucial for a successful stroke recovery.

However, clot search is only conducted if a stroke is assumed, but in some cases a patient comes into hospital for various symptoms and is not treated as a stroke patient in the first place. In such a case, a routine computed tomography (CT) scan of the brain may be acquired, but no active search for a thrombus is performed. An automated detection system in non-enhanced computed tomography (NECT) images would decrease the probability to miss an obstruction, save time, and improve the clinical outcome.

To the best of our knowledge, no efforts have been made so far to automatically detect vascular occlusion in NECT images. Instead, several methods have been proposed to detect occlusion in CT angiography (CTA) data [MMD02, MAG14]. Nevertheless, studies have shown that a thrombus can be discovered in NECT images through its abnormal high density structure [LPG*92, GFBV83]. Semi-automatic methods, which involve region growing methods with manually defined seed points, have proven to allow for accurate thrombus assessment [RJR*10, KYC*08].

In this paper, we extend the semi-automatic methods and present a new approach that detects blood clots in NECT images completely automatically.

2. Material and Methods

This section describes the approach that is used for automatic occlusion detection in NECT images. First, preprocessing is applied to narrow the search space and extract possible candidates. Next, several features are computed, which involve a probabilistic vessel map that has to be registered to the patient. Finally, sampling methods are applied and classification with a Random Forest (RF) [Bre01] is performed.

2.1. Candidate Extraction

Candidate extraction consists of several steps. First, the skull is removed by seeded region growing that includes only brain tissue voxels $\leq$ 200 Hounsfield units (HU). Initial seed points are found in the center of the middle axial slice.

Next, brain center and the *dorsum sellae* — an anatomical landmark belonging to the skull base — are determined and used later as reference positions for location based features. Therefore, the brain volume is computed in each single slice, beginning at the top axial slice. The search stops when the volume starts to decrease and continues to decrease for three consecutive slices. The center of gravity (CoG) of this slice is then defined as brain center. Subsequently, for every slice below the brain center, the algorithm calculates the CoG

DOI: 10.2312/vcbm.20171245

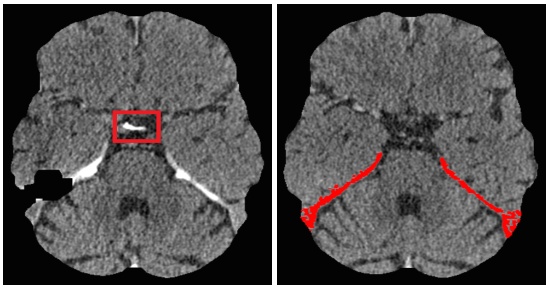

(a) (b)

Figure 1: *Example of skull base (a) and connective tissue segmentation (b).*

and searches in a small square with 50 × 50 voxels around it for voxels ≥ 200 HU. The topmost slice with a voxel fulfilling this criteria is then defined as skull base voxel.

Additionally, the skull base serves as seed point for a region growing method to segment connective tissue in the brain, which is removed from the region of interest (ROI). The region growing is performed with a 3D 6-connectivity in inferior vertical direction for voxels ≥ 60 HU. Skull base and connective tissue segmentation can be seen in Fig. 1.

Finally, thrombus candidates are extracted by thresholding, in such a way that only voxels ranging from 55 - 110 HU remain. The threshold is chosen similar to the ones used in [RJR*10] and [KYC*08]. To enable analysis of objects, connected voxels are clustered using 3D connected component labeling [RK82] with 6-connectivity. To remove noise and irrelevant objects, clusters with a volume being smaller than 1.5 mm^3 are discarded. In addition, clusters larger than 500 mm^3 can be discarded since a thrombus doesn't grow to this size [KYC*08].

2.2. Feature Generation

35 features are calculated to describe a thrombus. The full list can be found in Table 3. The features can be separated into intensity-based features, geometrical features, and location-based features.

Intensity-based Features

A thrombus can be distinguished from surrounding tissue by its high density. Therefore, intensity characteristics of the candidate and its surrounding tissue are computed. A cluster C is represented by maximum(C), mean(C), and variance(C) of its voxel intensities. The hull H around each cluster includes all direct neighbor voxels of a candidate. Based on the intensities of the hull voxels, the following features are computed: minimum(H), maximum(H), mean(H), variance(H), and maximum(H) − minimum(H). Additionally, mean(H) − mean(C) is computed, which compares hull intensities with cluster intensities.

The last group of gray value features delivers morphological information of the lesion and its surroundings. Therefore, multi-scale images are computed by convolving the image with Gaussian kernels. Three different standard deviations 0.5 mm, 1.0 mm, and 1.5 mm are applied. The resulting intensities at the voxel position of the candidate's maximum intensity prior to filtering are then used as features.

Geometrical Features

To describe the geometrical properties of a thrombus, the cluster volume in mm^3 is computed as well as the extents in horizontal (x), vertical (y), and slice (z) direction, and the diagonal length of the cluster's bounding box. Furthermore, the two-dimensional shape is analyzed in the slice where the candidate reaches the maximum slice volume. In this slice, the 2D volume of the cluster and its extents in x- and y- direction are computed, as well as the diagonal length of the two-dimensional bounding box.

Since a thrombus usually has a thin elongated form, the rectangularity and the aspect ratios are calculated to capture these shape characteristics [Ros00]. The rectangularity is defined by the ratio V/V_m of the object volume V and minimal bounding box volume V_m; and the aspect ratios $a_1 = \frac{extent(x)}{extent(y)}$, $a_2 = \frac{extent(x)}{extent(z)}$, and $a_3 = \frac{extent(y)}{extent(z)}$ put the different extents into relation.

Location-based Features

A thrombus can only appear in regions containing blood vessels. To describe the location of a cluster, the horizontal, vertical, and slice offset of the candidate's CoG to the brain center, as well as to the skull base are computed. Moreover, the Euclidean distances of the CoG to these locations are computed. These information are helpful since vascular occlusion are more likely to appear at the beginning of an artery, i.e., close to the skull base.

In addition, two features are calculated to emphasize regions where a thrombus can manifest. Therefore, a probabilistic vessel map and a reference brain image are utilized. The reference brain (atlas) is an image in a reference coordinate system, generated by the Montreal Neurological Institute as an average over multiple MRI scans of the brain [FEB*11]. The probabilistic vessel map is obtained by segmenting and averaging over numerous data sets by Kemmling et al. [KWB*12]. Alignment of the atlas and the vessel map, as well as of the atlas and the current data set, is ensured by diffeomorphic non-rigid registration [CHF02]. The values of the registered vessel probability map are then summed up over all cluster voxel positions, and the mean value is used as a feature. Additionally, the mean Euclidean distance of cluster voxels to the vessel map's centerlines is calculated.

2.3. Sampling and Classification

Since the candidate extraction technique covers a large part of the brain and is only based on intensities, many clusters remain, but in case of a stroke usually only one thrombus appears. Other possible negative clusters are, e.g., calcifications, artifacts, or bone that has not been removed completely. As a consequence, the amount of negative samples (majority class) exceed the number of positive samples (minority class) by a factor of ∼ 250. This highly imbalanced data set leads to a poor detection rate if no sampling method is used. To compensate this, a sequence of random under- and over-sampling is applied to the training data [HG09]. First, randomly selected negative samples are discarded such that ten times more majority samples than minority samples remain. Next, randomly

selected positive samples are duplicated until a balanced training set is achieved. The sampling ratios are determined empirically.

As last step, training and classification is performed with a RF classifier. Therefore, the RF implementation of the OpenCV C++ library [Bra00] is used. The model parameters that have to be determined are: number of trees, tree depth, minimal number of samples required at a tree node for it to be split, and number of randomly selected features at each node.

3. Experimental Setup

Thin-sliced reconstructions of NECT scans of the brain are available for 79 patients. Average slice width is 0.8 mm. Ground Truth (GT) segmentation of thrombi to train the RF are obtained semi-automatically with a region growing method, for which a seed point and upper and lower threshold are determined manually in every image. In each image one thrombus is present. The average thrombus volume is 73.8 mm^3, ranging from 17.6 to 188.8 mm^3; average thrombus length is 22.1 mm, ranging from 9.0 to 51.9 mm; and mean thrombus cluster intensity is 63 HU, ranging from 52 to 75 HU.

Classifier performance is evaluated and optimal RF model parameters are determined using nested cross validation (CV), which ensures a clear separation between training and validation data for an unbiased evaluation [VS06]. We split the 79 volumes into 8 disjunct folds and perform 8-fold nested CV. In nested CV, an inner loop is used to perform tuning of the parameters while an outer loop is used to evaluate the performance. In both loops, one fold is used for testing, and the remaining folds for training. Note that sampling methods are only applied to the training data.

4. Results

4.1. Classification Performance

Classification of the thrombus candidates with a RF generates a probabilistic value $p \in [0,1]$. Thus, a cutoff value has to be determined to make a final class assignment. The results of the nested CV for different cutoff values can be seen in the confusion matrix in Table 1. With $p \geq 0.25$, only one thrombus is missed, and 527 false positives (FPs) are generated. With $p \geq 0.50$, 7 thrombi are missed, and 217 FPs remain.

The performance measures for different cutoffs are listed in Table 2. The highest sensitivity with lowest FPs is achieved with a cutoff $p \geq 0.25$, yielding sensitivity of 0.987 and FP/volume of 6.67. If the median cutoff $p \geq 0.5$ is chosen, a sensitivity of 0.911 and an average number of FP per volume of 2.747 is achieved. With $p \geq 0.75$, sensitivity is 0.759 while only ~ 1 FP per volume is produced. With $p \geq 0.90$, more than 50% are detected, and only 0.392 FP/volume. Specificity and accuracy are close to 1.0 in each case.

Furthermore, the relationships between the probability a thrombus candidate is assigned, and the thrombus volume and the thrombus length, are illustrated in Fig. 2. All thrombi with a volume ≥ 100 mm^3 or a length ≥ 23 mm are assigned a probability higher than 90%. The smaller the thrombus, the more likely a lower probability is assigned. However, also smaller candidates can achieve a high probability. The only samples that achieve a probability ≤ 0.5 are small lesions with a volume ≤ 25 mm^3 or a length ≤ 11 mm.

Table 1: *Confusion matrix obtained after nested CV with different cutoff values p.*

	Pos.	Neg.
GT Pos.	78	1
GT Neg.	527	19130

(a) $p \geq 0.25$

	Pos.	Neg.
GT Pos.	72	7
GT Neg.	217	19440

(b) $p \geq 0.50$

Table 2: *Performance measures for different cutoffs.*

	Cutoff probability			
	0.25	0.50	0.75	0.90
sensitivity	0.987	0.911	0.759	0.519
FP/volume	6.671	2.747	1.025	0.392
specificity	0.973	0.989	0.996	0.998
accuracy	0.973	0.989	0.995	0.997

4.2. Feature Importance

In RF, the feature importance (FI) indicates how much the feature decreases the prediction uncertainty, and can be calculated for every feature. The result is displayed in Table 3. The most important features are extent in x-direction, diagonal length, and volume of the candidate, followed by the two features that are derived from the vessel probability map. 10 Features fall below a FI value of 1.0 and can be omitted. Table 4 shows the performance of the classifier with different feature selection. Classifying without the features derived from the vessel probability map slightly reduces the sensitivity and increases the number of FP per volume. Omitting the 10 features with lowest importance yields the same sensibility, and the number of FP per volume is only slightly increased.

5. Discussion

The goal of this work is to detect vascular occlusion. Hence, the most important measure is the sensitivity, i.e., the probability of detection. Our implemented system is able to detect all but one thrombus, proving that the goal can be fulfilled with the proposed method. The high values of specificity and accuracy show that most of the negative samples are correctly classified as such, but these measures don't possess high importance, as the negative samples are produced by the same system in the previous steps.

Although a sensitivity close to 1.0 can be achieved, a higher detection rate has the downside of producing more FPs. Therefore, a trade-off has to be made. In case of a clinical application, detecting positive samples is more important than producing a false alarm, because a FP has no consequence for the patient, whereas a missed thrombus can have severe impact on the clinical outcome. However, a false alarm takes up time of a radiologist and has a negative impact on the clinic's cost efficiency, and, therefore, the rate should not be too high. Thus, the cutoff $p \geq 0.50$ might be a reasonable trade-off, producing a detection rate of 91.1% and 2.7 FP per volume.

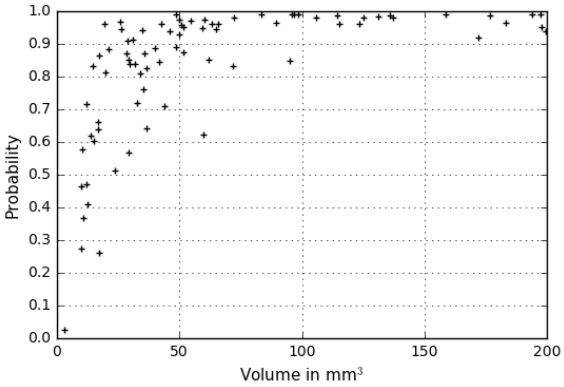

(a) *Relation between output probability and 3D volume.*

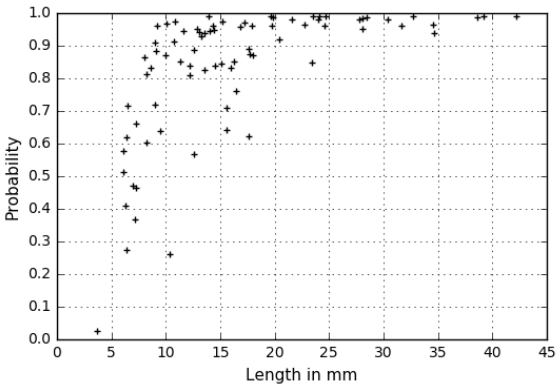

(b) *Relation between output probability and diagonal length.*

Figure 2: *Relation between output probability and thrombus volume (a), and between output probability and thrombus length (b).*

The high FP rate for high sensitivities is due to the fact that the extraction step covers a large part of the brain, while only bone and connective tissue are excluded. This leaves many negative clusters in the image, and makes sampling methods necessary to compensate for the imbalanced data set. Typical FP clusters are structures inside the middle cerebral artery (MCA) (Fig. 3a), or the basilar artery (Fig. 3b), where the vessel map exhibits high probability values. These FPs are hard to avoid, since a thrombus can likely emerge at these locations. Other typical FP samples are large hyperdense clusters referring to noise (Fig. 3c). All FPs can be discarded quickly after visual assessment through a radiologist.

Negative clusters and, consequently, the number of FPs, could be reduced with an increased value for the lower threshold limit during candidate extraction, but proper threshold selection is a difficult task. A higher threshold requires a higher density for the blood clot in order to extract it properly. Even with the chosen lower bound of 55 HU, some lesions are not fully segmented after candidate extraction. The one thrombus that is missed is such an under-segmented case. On the other hand, lower thresholds have shown to produce too many leakage effects into surrounding tissue, and the cluster is consequently not recognized as thrombus. The chosen lower bound extracts all samples without leakage effects.

Furthermore, FPs could be reduced with a higher cutoff value, but in this case, more of the small clots would be missed. However, all large thrombi are detected with a high probability. This is important since large occlusion cause the most severe damage. Examples of such certain classifications are shown in Fig. 3d and 3e.

The high probability for large clots is verified by the feature importance calculation. Next to volume and length, the extent in x-direction has proven to be the most important feature, because most thrombi possess an elongated structure along the MCA. Another important factor is the the vessel map, which provides two valuable

Table 3: *List of feature importances.*

Feature	Feature Importance
extent(x)	11.55
Diagonal length	9.69
3D volume	8.06
Mean vessel probability map value	5.14
Mean distance to vessel map	4.90
2D volume	4.62
Max(H) − Min(H)	4.42
Gaussian at $\sigma = 0.5$	4.42
extent(z)	4.19
extent(x) in maximum slice	4.08
offset(y) to brain center	3.57
Min(H)	3.19
Rectangularity	3.13
Mean(C)	3.00
Max(C)	2.83
Mean(H) − Mean(C)	2.72
extent(y)	2.20
Mean(H)	1.87
Variance(C)	1.79
Aspect ratio a_1	1.71
2D diagonal length in maximum slice	1.54
3D distance to skull base	1.49
3D distance to brain center	1.47
offset(z) to skull base	1.36
extent(y) in maximum slice	1.11
Gaussian at $\sigma = 1.0$	0.88
Max(H)	0.84
offset(x) to brain center	0.74
Aspect ratio a_3	0.66
offset(x) to skull base	0.59
Aspect ratio a_2	0.50
offset(y) to skull base	0.48
Variance(H)	0.45
offset(z) to brain center	0.41
Gaussian at $\sigma = 1.5$	0.22

Table 4: *Performance comparison with different feature selection ($p \geq 0.5$): I) All features are used; II) All features without vessel probability map; III) Only the 25 features with FI above 1.0.*

	I	II	III
sensitivity	0.911	0.861	0.911
FP/volume	2.747	4.013	2.975

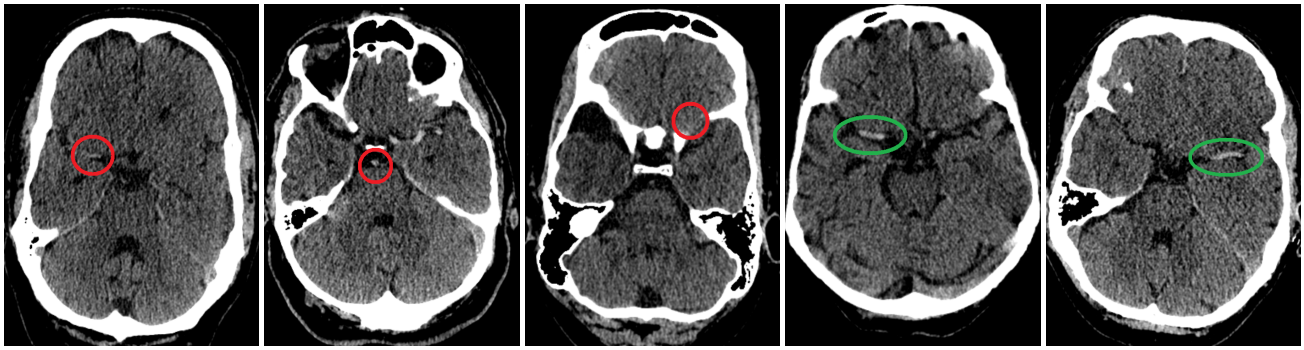

(a) *False Positive inside the MCA.* **(b)** *False Positive inside the basilar artery.* **(c)** *False Positive referring to noise.* **(d)** *Right MCA occlusion.* **(e)** *Left MCA occlusion.*

Figure 3: *Examples of false positive and true positive findings.*

features. These outweigh the other location-based features, which might be due to the fact that they describe the location more precise than the reference positions in the center of the brain and the skull base.

Although the results look promising, our method has limitations. First, the used data is already in the standardized atlas coordinate system. Some processing steps rely on the slice orientation and therewith require thin slices. With commonly used thick slice protocols resampling could become a problem if the orientation differs. In addition, our classification is performed on a candidate level within CT, but not on a patient level. If binary classification for presence of thrombus should be evaluated on a patient level, a negative control cohort of patients without occlusion is required.

6. Conclusion

We presented a fully automatic detection and classification method for vascular occlusion in NECT images. The proposed approach achieves a high detection rate, including all large MCA occlusion, while the number of FPs is in a reasonable range. A trade-off between sensitivity and number of FPs per volume can be made to adapt the method to the needs of the application. The method could prove valuable in clinical environment in order to detect blood clots and, respectively, to reduce the probability to miss a stroke. More sophisticated candidate extraction techniques such as an adaptive thresholding method might improve the performance in future works. With more training data, the method can be extended to deep learning approaches, which have gained huge successes in classification tasks in recent years [KSH12].

References

[Bra00] BRADSKI G.:. *Dr. Dobb's Journal of Software Tools* (2000). 3

[Bre01] BREIMAN L.: Random forests. *Machine Learning 45*, 1 (2001), 5–32. 1

[CHF02] CHEFD'HOTEL C., HERMOSILLO G., FAUGERAS O.: Flows of diffeomorphisms for multimodal image registration. *Proceedings of 2002 IEEE International Symposium on Biomedical Imaging* (2002), 753–756. 2

[FEB*11] FONOV V., EVANS A., BOTTERON K., ALMLI C., ET AL.: Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage 54*, 1 (2011), 313–327. 2

[GFBV83] GACS G., FOX A., BARNETT H., VINUELA F.: CT visualization of intracranial arterial thromboembolism. *Stroke 14*, 5 (1983), 756–762. 1

[HG09] HE H., GARCIA E.: Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering 21*, 9 (2009), 1263–1284. 2

[KSH12] KRIZHEVSKY A., SUTSKEVER I., HINTON G.: Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* (2012), 1097–1105. 5

[KWB*12] KEMMLING A., WERSCHING H., BERGER K., KNECHT S., ET AL.: Decomposing the Hounsfield unit: Probabilistic segmentation of brain tissue in computed tomography. *Clinical Neuroradiology 22*, 1 (2012), 79–91. 2

[KYC*08] KIM E., YOO E., CHOI H., LEE J., ET AL.: Thrombus volume comparison between patients with and without hyperattenuated artery sign on CT. *American Journal of Neuroradiology 29*, 2 (2008), 359–362. 1, 2

[LPG*92] LEYS D., PRUVO J., GODEFROY O., RONDEPIERRE P., LECLERC X.: Prevalence and significance of hyperdense middle cerebral artery in acute stroke. *Stroke 23*, 3 (1992), 317–323. 1

[MAG14] MAIORA J., AYERDI B., GRAÑA M.: Random forest active learning for AAA thrombus segmentation in computed tomography angiography images. *Neurocomputing 126* (2014), 71–77. 1

[MMD02] MASUTANI Y., MACMAHON H., DOI K.: Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis. *IEEE Transactions on Medical Imaging 21*, 12 (2002), 1517–1523. 1

[RDW16] RAFFELD M., DEBETTE S., WOO D.: International stroke genetics consortium update. *Stroke 47*, 4 (2016), 1144–1145. 1

[RJR*10] RIEDEL C., JENSEN U., ROHR A., TIETKE M., ET AL.: Assessment of thrombus in acute middle cerebral artery occlusion using thin-slice nonenhanced computed tomography reconstructions. *Stroke 41*, 8 (2010), 1659–1664. 1, 2

[RK82] ROSENFELD A., KAK A.: *Digital Picture Processing*, 2 ed. Academic Press, San Diego, CA, 1982. 2

[Ros00] ROSIN P.: Measuring shape: ellipticity, rectangularity, and triangularity. *Proceedings of 15th International Conference on Pattern Recognition 1* (2000), 952–955. 2

[VS06] VARMA S., SIMON R.: Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics 7*, 1 (2006), 91. 3