Image analysis and processing for an automatic vehicle identification system

L. Coetzee*
Computing, *i*comtek, CSIR, P.O. Box 395, Pretoria, 0001
and E.C. Botha†
Department of Electrical, Electronic and Computer Engineering
University of Pretoria, Pretoria 0002, South Africa.

Abstract

This paper describes the image analysis and processing of an automatic vehicle identification system. The envisioned system will identify vehicles based on video or still images of unique identifying marks, namely metallurgical fingerprints, as proposed by De Kock in his South African patent [1]. Two candidate image analysis algorithms for automatic vehicle identification are selected. The selection criteria and the process for identification of these algorithms are described. The algorithms were tested on a dataset of real-world vehicle images. The experimental approach and the results obtained from applying these algorithms to the automatic vehicle identification problem form the main theme of our paper. Based on the results, the recommended system uses a nearest-neighbour classifier with features as extracted with a wedge-ring detector from the frequency domain of the image on a histogram preprocessed dataset.

1 Introduction

This paper describes the image analysis, processing as well as results as obtained from the NOVIS (**N**ational **O**n-line **V**ehicle **I**dentification **S**ystem) project[1]. It is envisioned that the final NOVIS system will be able to automatically verify vehicles based on unique features extracted from various locations in and on the vehicle e.g. engine block, welding signatures on doors and other areas of significance (i.e. the so called metallurgical fingerprints).

In Section 2 we provide sample images to better illustrate the concept of metallurgical fingerprints. In the following section (Section 3) we describe the selection criteria and the process used for the identification of the candidate algorithms. We follow this with a section containing our experimental approach, the results as obtained from the candidate algorithms as well as a discussion of the results (Section 4). Section 5 contains conclusions and a recommendation.

2 Data set description

Two different datasets were used during the different phases. An initial 100-vehicle dataset (900 images) was used in a human perception experiment to determine whether vehicles can be discriminated based on the nine locations initially identified [2, 3]. Five of these locations seemed to be useful for discrimination and were subsequently used in the experimental process (as described in Section 4) in conjunction with a dataset consisting of five samples of each location on each of the eight vehicles [4].

Figures 1 to 5 depict typical images of these five locations that seemed useful for discrimination.

3 Candidate algorithm selection process

In order to select appropriate image processing and analysis algorithms for the matching or verification problem, a set of selection criteria has to be adhered to based on inspection of the data.

*E-mail:louis.coetzee@csir.co.za.
†Corresponding author. E-mail: liesbeth.botha@eng.up.ac.za, Fax: 021 888 2694, Tel: 021 888 2547.
[1]The support of this work by DACST Innovation Project #23216 is gratefully acknowledged.

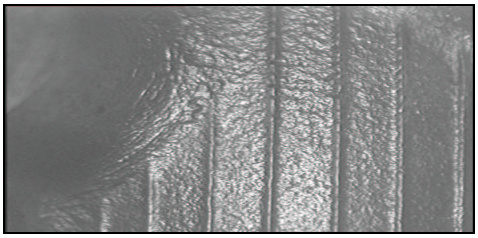

Figure 1: Location 001 – Left hand side on head

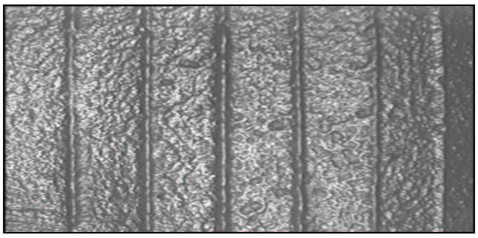

Figure 2: Location 002 – Right hand side on head

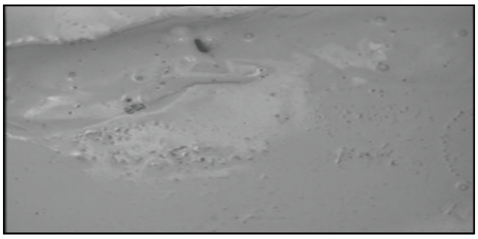

Figure 3: Location 003 – Weld on chassis

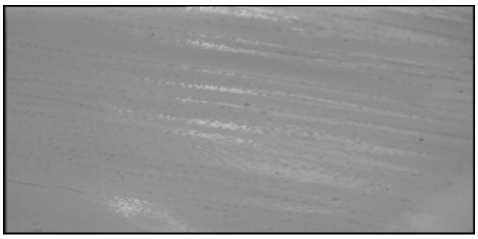

Figure 4: Location 004 – Glue on chassis

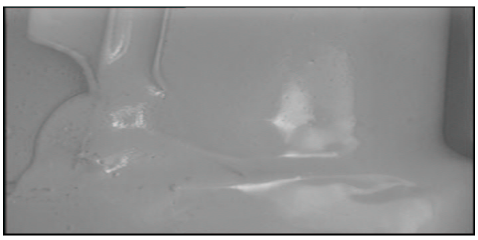

Figure 5: Location 005 – Left hand side front door

The following characteristics of the data (the images) have to be taken into account:

1. The images of the same location, but from different vehicles have widely varying degrees of similarity, depending on the location. This means that the intra-subclass differences can be small or large, and thus they have a large dynamic range.

2. The images at different locations on the same vehicle differ substantially, e.g. the images of the "bonnet" and of the "glue on the chassis" are totally different images. This means that we are dealing with large inter-subclass differences.

3. The images of the same car and same location, taken at different times and by a different photographer are very similar in a qualitative sense, but they differ somewhat in scale, in-plane and out-of-plane rotation, translation and lighting. These differences are small, due to the efforts to keep the conditions as similar as possible. However, these small qualitative differences give big quantitative differences. For instance, a single pixel change in translation will result in a huge inter-image distance if a simple inner product match is to be calculated.

4. Image data files are by nature large, and many calculations are required for matching.

Keeping the above considerations in mind, the following criteria are arrived at for selection of the algorithms:

- Characteristics 1 and 2 above imply that the features used for matching must not be image-specific, i.e. they must work for all kinds of different images. They must, however be sufficiently separable so that images with small differences can be distinguished.

- Characteristic 3 implies that the features used for matching must be invariant to the variations in terms of scale, translation and rotation encountered between images of the same sub-class.

- Characteristic 4 implies that the algorithm must not in itself be computationally expensive, as the images are data-intensive to start with.

As noted above, the images to be matched are fairly similar, but that several ways of transformation of the images can occur (translation, rotation, lighting, etc.). Due to the expected transformations of the images, feature-based classification was considered to be the most appropriate. The recommendation is to conform to the classic pattern recognition paradigm of (i) pre-processing, (ii) feature extraction and (iii) classification. Based on image inspection, no noise reduction seemed possible as there is no apparent and structured noise visible. A variety of feature extraction mechanisms exist. *Local* features were deemed inappropriate as the images differ too much between sub-classes and no properly defined small set of features exist (or can easily be described) for images from each of the locations. *Global* feature extraction, where features are calculated from the whole image, using standard algorithms and procedures is more appropriate. Two candidate algorithms were identified namely:

1. Fourier coefficients, and specifically a particular decimation of the Fourier space, namely a *wedge-ring structure* [5] and
2. A principal component analysis, which will be referred to as *eigenimages* [6, 7].

These approaches conform to the criteria set out above. Both are invariant to small transformations in the intra-subclass images and both should be able to distinguish close matches from different sub-classes. The performance is empirically investigated on the experimental study on the captured images as presented in Section 4.

4 Experimental approach and results

The final envisioned vehicle classification system will perform *verification*, but we conduct *recognition* experiments in order to determine the best classifier to use in the verification paradigm. Two datasets were generated from the original data:

1. A histogram-equalized dataset, where histogram-equalization was applied to each individual image (referred to as the **histogram dataset**) [8].
2. An average-adjusted dataset, where the average of each image was adjusted to 128 with pixels with grayscale values below zero clipped to zero and those above 255, clipped to 255 (referred to as the **average dataset**).

Figure 6 illustrates the different classification schemes applied to each of the datasets (as described above). Experiments were performed on the spatial domain of the images on each dataset using a **correlation classifier** [9]. In the feature-based approaches, features were extracted from the frequency domain using a wedge-ring detector [5, 10]. Classification results on the wedge-ring features were obtained using **nearest-neighbour** [9], **linear** [9] and **neural-net** [11] classifiers. Principal Component Analysis (PCA)was used as an alternative to the frequency domain analysis to extract features [6, 7]. Using the PCA-based features, experimental results were obtained for **nearest-neighbour** and **neural-net** classifiers.

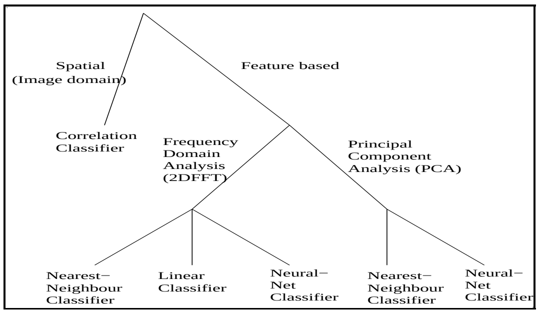

Figure 6: Experimental approaches

As described in Section 2, five samples of each of the five locations for each of the eight vehicles were available. Separate experiments were conducted for each of the locations in order to obtain the best classifier and parameters for each. For each location the experimental process was repeated five times: in each repetition a different sample was used as the test sample (and the remaining four samples as training samples). Furthermore, different feature vector dimensions were experimented with during the feature-based analysis combined with variations in the number of hidden neurons and various starting positions in weight space for the neural-net classifier.

A summary of the results are presented below. The presented results for a location is an averaged value obtained from the five trials at that location (variation in picking the test samples). The value used in the averaging process reflects the maximum classification rate obtained for a test sample at a location, regardless of the weight initialization, number of hidden neurons or feature vector dimension.

4.1 Histogram dataset

4.1.1 Spatial-based classifier

Table 1 presents the results using the correlation classifier on the histogram dataset. Location 003 had the poorest recognition rate, while location 005 achieved 100%.

Location:	001	002	003	004	005
Average%:	97.5	97.5	87.5	97.5	100

Table 1: Average results for *correlation classifier* using the histogram dataset

4.1.2 Feature-based classifiers

Tables 2 to 4 present the results for the nearest-neighbour, linear and neural-net classifiers using the histogram dataset and features as extracted from the frequency domain with the wedge-ring detector.

Location:	001	002	003	004	005
Average%:	100	97.5	100	100	100

Table 2: Average results for *nearest-neighbour classifier* using the histogram dataset with *wedge-ring features*

Location:	001	002	003	004	005
Average%:	100	97.5	100	97.5	100

Table 3: Average results for *linear classifier* using the histogram dataset with *wedge-ring features*

Location:	001	002	003	004	005
Average%:	100	97.5	100	100	100

Table 4: Average results for *neural-net classifier* using the histogram dataset with *wedge-ring features*

Recognition results using the histogram dataset with wedge-ring features as extracted from the frequency domain were excellent for all three classifiers. Location 002 was somewhat unstable, compared against the results for the other locations.

Tables 5 and 6 present the results for the nearest-neighbour and neural-net classifier using the histogram dataset and features extracted using principal component analysis.

Location:	001	002	003	004	005
Average%:	85	65	92.5	97.5	92.5

Table 5: Average results for *nearest-neighbour classifier* using the histogram dataset with *PCA features*

Location:	001	002	003	004	005
Average%:	90	95	100	100	100

Table 6: Average results for *neural-net classifier* using the histogram dataset with *PCA features*

Results obtained from the PCA features using the nearest-neighbour classifier were worse than those obtained from the neural-net classifier. Overall, the results using PCA features were poorer than those obtained using the wedge-ring features.

4.2 Average dataset

4.2.1 Spatial-based classifier

Table 7 presents the results using the correlation classifier on the average dataset. Locations 001 and 002 achieved perfect recognition rates, but the other locations' recognition rates are inadequate.

4.2.2 Feature-based classifiers

Tables 8 to 10 present the results for the nearest-neighbour, linear and neural-net classifiers using the average dataset and features as extracted from the frequency domain with the wedge-ring detector.

Very good recognition rates for all classifiers and all locations (except location 002) were obtained.

Location:	001	002	003	004	005
Average%:	100	100	35	55	65

Table 7: Average results for *correlation classifier* using the average dataset

Location:	001	002	003	004	005
Average%:	100	97.5	100	97.5	100

Table 8: Average results for *nearest-neighbour classifier* using the average dataset with *wedge-ring features*

Location:	001	002	003	004	005
Average%:	100	97.5	100	100	100

Table 9: Average results for *linear classifier* using the average dataset with *wedge-ring features*

Location:	001	002	003	004	005
Average%:	100	97.5	100	100	100

Table 10: Average results for *neural-net classifier* using the average dataset with *wedge-ring features*

Tables 11 and 12 present the results for the nearest-neighbour and neural-net classifier using the average dataset and features extracted using principal component analysis.

Recognition rates dropped slightly compared to those obtained using the other method of preprocessing.

4.3 Discussion

A number of conclusions can be drawn from the presented results. Locations 001 and 002 appear to be somewhat problematic for the classifiers. A 100% recognition rate was obtained only occasionally for these locations. In contrast, the experiments using locations 003, 004 and 005 obtained excellent results, very often 100%. The difficulty associated with locations 001 and 002 can most likely be contributed to the high information content associated with the texture (high frequency components) of these locations. Locations 003, 004 and 005 were less **"busy"** or textured, which allowed for the extraction of more meaningful features.

Regarding preprocessing techniques, the histogram dataset appears to be more reliable than the average dataset. Information was lost using the average dataset because of the clipping imposed on the pixel values (below zero and above 255 grayscale levels).

Viewing the performance of the classifiers using the PCA-based features, little difference in classification results are noticed between the nearest-neighbour and neural-net classifiers. In general, the results were slightly poorer compared to those using features extracted from the frequency domain. The correlation classifier obtained good results, but the computational complexity renders it unsuitable. Regarding the results obtained from the frequency domain, very little difference is observed between the various classifiers. Keeping the above in mind, a nearest-neighbour classifier is the classifier of choice, thus avoiding the complexity of retraining the linear and neural-net classifiers when new vehicles are added to the database.

Analyzing the various results as presented earlier, as well as the arguments as presented above, it is recommended to use a nearest-neighbour classifier with features extracted from the frequency domain using a wedge-ring feature extractor on histogram equalized images.

5 Conclusion

This paper presented selection criteria for identifying candidate algorithms for solving the vehicle verification problem. Using these algorithms in conjunction with two different preprocessing techniques a vast number of experiments were conducted to identify the most suitable approach for the NOVIS project. Analyzing the results, it was recommended to use a nearest-neighbour classifier with features extracted from the frequency domain using a wedge-ring feature extractor on histogram equalized images.

References

[1] S. De Kock. Method of identifying a vehicle, August 2000. South African Patent 2000/3961.

[2] S. de Kock, L. Langenhoven, and M.B. Cortie. Development of a prototype national

Location:	001	002	003	004	005
Average%:	77.5	87.5	97.5	100	100

Table 11: Average results for *nearest-neighbour classifier* using the average dataset with *PCA features*

Location:	001	002	003	004	005
Average%:	92.5	90	97.5	100	100

Table 12: Average results for *neural-net classifier* using the average dataset with *PCA features*

on-line vehicle identification system (NOVIS) based on metallurgical finger printing and digital imaging. First Progress Report. Technical Report C2955M, MINTEK, Randburg, South Africa, April 2000.

[3] E.C. Botha. Development of a prototype national on-line vehicle identification system (NOVIS) based on metallurgical finger printing and digital imaging. NOVIS Second Progress Report. Image Processing Algorithms. Technical Report LGI 2000/006, LGI, Pretoria, South Africa, Sep. 2000.

[4] L. Coetzee and E.C. Botha. Development of a prototype national on-line vehicle identification system (NOVIS) based on metallurgical finger printing and digital imaging. NOVIS Fourth Progress Report. Experimental Analysis of Image Processing Algorithms. Technical Report 2001/001, RE@UP, Pretoria, South Africa, Sep. 2000.

[5] L. Coetzee and E. C. Botha. Fingerprint recognition in low quality images. *Pattern Recognition*, 26(10):1441–1460, 1993.

[6] D.W. Purnell. Automatic face recognition in a heterogeneous population. Master's thesis, University of Pretoria, Pretoria, South Africa, Apr. 1998.

[7] D.W. Purnell, C. Nieuwoudt, and E.C. Botha. Automatic face recognition in a heterogeneous population. *Pattern Recognition Letters*, 19:1067–1075, 1998.

[8] D. H. Ballard and C. M. Brown. *Computer vision*. Prentice Hall, Englewood Cliffs, New Jersey, 1982.

[9] R. O. Duda and P. E. Hart. *Pattern classification and scene analysis*. Wiley, New York, 1973.

[10] L. Coetzee. Fingerprint recognition. Master's thesis, Department of Electrical and Electronic Engineering, University of Pretoria, Pretoria, South-Africa, 0002, 1992.

[11] E. Barnard and R. A. Cole. A neural-net training program based on conjugate-gradient optimization. *Oregon Graduate Center Technical Report No. CSE 89-014*, July 1989.