Real Time Automatic License Plate Recognition in Video Streams

Fredrik Trobro

September 23, 2007

Abstract

In recent years there has been an increased commercial interest in systems for automatic license plate recognition. Some of the existing systems process single images only, some even requiring vehicles to stop in front of a gate so that a still image of good quality can be taken. This thesis presents an approach that makes it possible to process 25 images per second on a standard PC from 2004, allowing identification decisions to be based on information from several images instead of just a single one. Thus, the impact of a single bad image is reduced, and vehicles can be allowed to drive past the camera without hindrance.

In order to reach the necessary processing speed, a simplified Stauffer-Grimson background estimation algorithm has been used, enabling the system to only search for license plates on objects that are in motion. The main method for finding license plates has been a computational-wise economical connected component labeling algorithm. A basic pixel-by-pixel comparison OCR algorithm has also been implemented. A real life test running for twelve days has shown the complete system to have a rate of successful identification at 80 %.

Supervisors: Håkan Ardö and Kalle Åström, Lund University

Contents

Figure 1: The diesel tank that was the inspiration for this master's thesis.

1 Introduction

As of the last decade, theft by opportunity-seeking criminals travelling by car has become an increasing problem in the southern Swedish countryside. Expensive tools are a prime target for these touring criminals, but also unguarded fuel tanks. The father of the author of this master's thesis is a farmer and has more than once had the unfortunate experince of diesel theft from an unguarded tank (Figure 1) on the farm. Since the offenders are unlikely to be travelling by foot, the idea came to mind that a system for automatic license plate recognition could be used for surveillance of the area around the diesel tank. The system could have a list of authorized vehicles, and send a picture message to the owner's mobile phone in case an unauthorized vehicle approaches the diesel tank.

A monitoring system depending on real time automatic licence plate recognition would of course be easy to circumvent by just removing or covering the licence plates of a vehicle, but it was deemed unlikely that offenders would expect such a system and take precautions. Against well-informed and prepared offenders the best security system in the world would probably not suffice.

The fact that the diesel tank in Figure 1 lies detached on open ground makes automatic license plate recognition extra difficult, since vehicles can approach from any direction. A multitude of cameras would be needed around the diesel tank. Instead, it was decided to implement the system at another farm, guarding that farm's sole driveway (Figure 2). At the time of this writing, the system has been up and running at that location for almost three years and has made possible the identification of several unauthorized vehicles. The camera might

Figure 2: The driveway as seen by the implemented automatic license plate recognition system.

also have had a deterring effect, since no offence has been committed at the location during these three years.

1.1 Background

Due to the many possible commercial applications of automatic license plate recognition, much work has been done in the area. Already in 1976, the Police Scientific Development Branch in the UK started developing a system that was up and running in 1979. In 1981 the first arrest due to a stolen car being detected by this system was made [Wik07].

However, since most previous work has been done by private corporations, much of the underlying theory is kept secret. But from the publically available articles it can be deduced that the different solutions for automatic lincense plate recognition generally consist of two parts:

1. Finding license plates in images
2. Reading text from license plates

The second part of the problem, reading text, is really just a subset of the vast field of optical character recognition (OCR) and will be dealt with briefly in this thesis. The focus of this thesis will instead lie on the first part of the problem, finding license plates in images. One notable solution for that part of the problem is the use of a Hough transform to find the edges of license plates [KG95]. The Hough transform has been shown to give good results, but it reacts on any sharp straight lines in an image, e.g. those found in windows on houses. Another approach is to search for areas with high contrast gradients, such as those caused by black text on a white background [Dra97].

1.2 Aim of the Thesis

The aim of this thesis is to create a system for real time automatic license plate recognition using standard off-the-shelf hardware. The system must be usable in real-life conditions, it must have a high rate of successful identifications and run in real-time. The system must not be depending on external triggering signals to tell when a vehicle is in front of the camera, and must not require vehicles to stand still in front of the camera. An attempt will be made to utilize information from all frames containing a certain license plate, instead of just choosing a single image to detect the license plate in. Therefore it will be desirable to process 25 video frames (images) per second.

The system will be implemented in Sweden, and therefore optimized for Swedish license plates. The system will also be optimzed for license plates from countries in the vicinity of Sweden, as long as the success rate for Swedish license plates is not reduced.

One conceivable approach for license plate recognition would be to simply do OCR on entire images, and then determine if any recognized text represents a license plate. It would however be very hard to implement an OCR system that on the chosen hardware would be able to process 25 frames per second while successfully recognizing text in any size, angle or location. Also, personalized license plates (license plates containing arbitrary text) would be hard to distinguish from arbitrary text on the sides of vehicles etc. In consequence of this, this thesis will focus on finding license plates in images by other means. Once the license plate edges have been located, the OCR part of the problem will be much easier since size, angle and location of the text will be known.

1.3 Related Work

In order to reduce the area in which the search for license plates will be performed, the Stauffer-Grimson algorithm [Sta99] for background estimation was used. This method keeps track of several different background color values for each pixel, thus it can handle leaves swaying in the wind etc by classifying both the leaves and the sky behind them as background. The variance of color values in a pixel determines how far a pixel can deviate from a background color while still being classified as background. This allows the method [Sta99] to have acceptable performance in noisy video, and in dark areas of an otherwise bright scene. The background model is updated after each processed frame, so that changes in the background will be reflected in the background model after a period of time.

Previous work in the area of license plate recognition has utilized Hough transforms [KG95] or searched for high contrast gradient concentrations [Dra97]. Connected component labeling has been chosen for the system described in this thesis because of its relatively low computational cost, and because it enables the system to find license plates of varying sizes and orientations in an image.

The Hough transform [GW02] has a very high computational cost, with the hardware chosen in this thesis it would not be possible to process anyway near 25 frames per seconds. Processing 25 frames per second has been one of the main goals in this thesis, since it allows identification to be based on many images of the same license plate instead of just one, thereby increasing chances for success.

The approach described in [Dra97] (searching for high contrast gradient concentrations) depends on triggering signals from an external sensor to determine when a vehicle is visible in the scene. Also, the license plates in the images are required to match a predetermined size. These requirements makes an automatic license recognition system more complex to set up than if the more flexible approach of connected component labeling is used.

2 Hardware

When discussing real-time performance and rate of successful identifications for the automatic license plate recognition system, the hardware configuration must be taken into consideration. In this section follows a specification of the hardware that has been used.

2.1 Computer

CPU: AMD Athlon XP 2800+

CPU-frequency: 2.07 GHz

RAM: 512 MB

Graphics: VIA/S3G KM400/KN400 (embedded on motherboard)

OS: Microsoft Windows XP Professional, Service Pack 2

2.2 Camera

Manufacturer: Monacor

System: 625 lines, PAL/CCIR

Sensor: CCD, 1/3", color

Resolution: 420 lines

Sensitivity: 1 lux

Shutter: Electronic, auto

SNR: >50 dB

Lens: 12.0 mm/F2.0

Viewing angle: 26°

Features: Back Ligth Compensation, Automatic White Balance, Automatic Gain Control

Figure 3: An example of the interlacing effect. Odd and even rows are not captured at the same time, there is a 20 millisecond offset between them causing obects in fast motion to appear garbled.

3 Methods

The image resolution was set to 352x288 (a quarter of the maximum possible resolution) in order to avoid interlacing (Figure 3) and increase speed. Higher resolution would make it easier for the system to find license plates and read the text, but the increased computational cost would make it impossible to process 25 frames per second on the chosen hardware. Lower resolution would make license plates indecipherable. To compensate for the low resolution, a zoom-lens was used (26° viewing angle).

3.1 Background Estimation

As an initial way of eliminating areas that do not contain a moving vehicle, the Stauffer-Grimson algorithm [Sta99] was used. But the chosen CPU was not powerful enough to do this in real-time on a video stream with the resolution 352x288 and also search after license plates, therefore the method was simplified. The first alteration was to just use every fourth pixel in both the horizontal and the vertical dimension. Thus, a single pixel decides for an area of sixteen pixels whether it should be classified as background or foreground, reducing the computational cost by a factor of sixteen.

The second simplification of [Sta99] was the use of a constant, instead of the variance, when deciding the size of the interval of color values that were classified as background. That way, there was no longer any need for calcu-

Figure 4: The marked area is classified as foreground by the simplified Stauffer-Grimson algorithm. Noise has been removed by morphological opening.

lating a variance. Only two equations where then needed when updating the background model after each processed frame:

Equation 1 describes the updating of weights (ω) for all K distributions in a single pixel. The weights are then normalized so that their sum is 1. The learning rate (α) is a constant value. $M_{k,t}$ is defined as 1 for the model that matched X_t in this pixel and 0 for all other models. Equation 2 describes the updating of the mean value (μ) for the matching model. Here ρ in [Sta99] has been replaced by the constant value α in order to speed up calculations.

$$\omega_{k,t} = (1-\alpha)\omega_{k,t-1} + \alpha(M_{k,t}) \tag{1}$$

$$\mu_t = (1-\alpha)\mu_{t-1} + \alpha X_t \tag{2}$$

The simplifications of [Sta99] reduces its accuracy, but that is of little concern as long as the license plates are still classified as foreground. Much noise was introduced in the background model, but it could be removed by morphological opening [GW02] (Figure 4).

The sole purpose of using background estimation has been to eliminate areas in which moving vehicles cannot be found. By only searching for license plates in areas classified as foreground, the computational cost for each frame is greatly reduced. Buffering of frames allows the system to spend longer time on frames filled with movement, and then catch up with the real-time video stream when the scene is static again.

3.2 Morphology

In order to remove foreground areas so small that they are unlikely to contain a vehicle, and also to remove noise and small outgrowths on larger areas, the binary foreground estimation is first eroded three times, then dilated four times (morphological opening, [GW02]). The reason for doing one more dilation than erosion was that this increases the chance for the entire license plate to be classified as foreground. Implementation complexity was reduced by morphing several times with a single pixel element instead of morphing once with a larger structure element.

3.3 Connected Component Labeling

In this thesis, the main algorithm for finding license plates has been a simple connected component labeling algorithm that examines four-way pixel connectivity [Coh93]. The goal of this algorithm is to find segments (areas) of similarly colored pixels, surrounded by pixels of contrasting color. For instance, a white license plate surrounded by a black border.

For every pixel in the image, the level of red, green and blue color is compared to the corresponding color levels of the four pixels above, below, to the left and to the right of the examined pixel (four-way connectivity). If two adjacent pixels have level differences in each color smaller than a predetermined threshold, they are considered to be connected and thus are given the same segment label.

A dictionary of segment labels is created, so that when two segments are found to be connected a translation of the second segment label to the first segment label can be added to the dictionary. That way each pixel will be given a label only once, but the label can afterwards be translated if the segment turned out to be part of a larger already labeled segment.

3.4 Heuristics

More than once in this thesis, heuristics (empirically evolved rules not theoretically proven) have been used to determine whether an object could be a license plate or not. A license plate is assumed to ...

- ... never be bigger than one third of the image width or one fifth of the image height. Zooming in the camera more than that would not be reasonable, since the location of license plates in the scene cannot be predicted with a high enough certainty.
- ... never be smaller than 4 % of the image width or 2 % of the image height. Even with an image resolution of 720x576, it would not be possible to interpret the characters on a license plate smaller than that.
- ... have the shape of a parallelogram.
- ... have four distinct corners.
- ... have an angle between its top and bottom border of less than 4°. When a license plate is viewed head on, it has the shape of a rectangle. But when viewed from slightly to the side and above, it has more the shape of a parallelogram. But not truly a parallelogram since the farther short

side will be slightly shorter than the nearest, and thus the two long sides will not be truly parallel. Because of that, and because of possible image noise, the license plate long sides are allowed to be at an angle of up to 4° relative each other.

3.5 Thresholding

Thresholding is one of the most simple methods in image analysis, and yet it is also one of the most powerful. The trick is to find a good value to use as threshold. In this thesis, thresholding is used when a rough location of a supposed license plate has already been acquired. The mean luminance value for an area just containing the supposed license plate is then calculated, since the luminance values for the license plate background and the luminance values for the license plate borders will almost always be on opposite sides of that mean value, even if the license plate has a gradient shadow. After creating a binary image with the mean luminance value as threshold, the license plate borders are easy to find.

3.6 Cut-out

When a license plate has been located, it is cut out and transformed into the shape of a horizontal rectangle of predetermined size. Pixel values are resampled by means of bilinear interpolation [GW02]. Any skew is thus corrected and the characters are lined up horizontally. That, and the fact that all transformed cut-outs have the same width and height in pixels, makes OCR and comparison of license plates in different video frames much easier.

3.7 Plate Matching

If information from several video frames is to be used for the identification of a license plate, license plates in consecutive frames must be matched since there might be more than one of them visible in a frame. For the primary matching method, a simple coordinate comparison was chosen. Coordinates for license plates in the current video frame are compared to coordinates for license plates found in the previous frame.

If a license plate in the current frame is not close enough to the location of a license plate in the previous frame, the contrast is enhanced for the cut-out of the new license plate and the sum of absolute pixel differences between the new cut-out and the mean images for license plates found in recent frames is calculated. If that sum is not low enough for any of the old license plates, it is concluded that a new license plate has been found. But just in case, the OCR result will later be compared to that of the old license plates in order to find matching plates.

3.8 Mean Image

Instead of performing OCR directly on the license plate in a video frame, a grayscale mean image is formed from cut-out license plates in consecutive frames. That way, small skewing errors and video noise will be leveled out. A single failed cut-out will also have less impact.

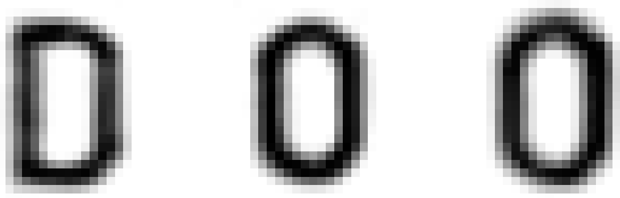

Figure 5: The letters D, O and the numeral zero, from photos of older Swedish license plates.

The mean image is updated on-line (Equation 2) so that there is no need for keeping all the license plates from previous frames in memory.

3.9 Contrast Enhancement

The grayscale mean images are normalized by means of contrast enhancement, so that a pixel-by-pixel comparison OCR method can be used. After empirical studies, it was decided to carry out the contrast enhancement according to the following formula:

- The darkest 5 % of the pixels are made completely black.
- The pixels that fall into the 5 % to 15 % range in the luminance histogram are given new luminance values in the 0 % to 25 % range of the maximum possible luminance.
- The pixels that fall into the 15 % to 35 % range in the luminance histogram are given new luminance values in the 25 % to 75 % range of the maximum possible luminance.
- The pixels that fall into the 35 % to 90 % range in the luminance histogram are given new luminance values in the 75 % to 100 % range of the maximum possible luminance.
- The brightest 10 % of the pixels are made completely white.

After the contrast enhancement, characters in the license plate are black and stand out clearly from the background, which is bright white even if it was e.g. green or yellow in the video frame. The characters are also thin enough to be clearly recognizable.

In some places in the world, e.g. Malaysia, license plates consist of white text on a black background [YTKA]. For such license plates, the system would have to detect that grayscale images should be inverted before adding them to a mean image. But since no such license plates were ever sighted during development or testing of this system, that feature was not implemented.

3.10 Optical Character Recognition

Since the focus of this thesis was on finding license plates, not on the OCR part of the problem, a simple pixel-by-pixel comparison method was selected for the OCR. Reference images were created for all letters and numerals that can appear on standard Swedish license plates, which in fact does not include all letters in the Swedish alphabet. The letters I, Q, V, Å, Ä and Ö never appear on standard Swedish license plates, since from a distance they might easily be confused with other characters. They might however appear on personalized license plates, but no such license plates were sighted during the development of this system and therefore those characters were not included in the system.

On older Swedish license plates it is diffucult to tell apart the letters D and O. Even worse, the letter O and the numeral zero look identical (Figure 5). Some work was therefore put into adjusting the reference character images in order to increase the chances of distinguishing between D and O. More than one reference image of each character was included in the system. Also included were reference images depicting the tax sign found in the center of the rear license plate on Swedish vehicles, and the small T-sign on taxi car license plates. Those non-characters are not included in the result of the license plate recognition, their reference images rather help the system to avoid confusing non-characters with real characters.

There is however no way to distinguish between the letter O and the numeral zero in older Swedish license plates (Figure 5). Except for their positions. Standard Swedish license plates always consist of three letters to the left, a gap, and then three numerals to the right. Thus, whenever a license plate containing six characters is found, it is assumed to be a standard Swedish license plate. The first three characters are interpreted as letters, and the last three characters are interpreted as numerals. There is then a risk of misinterpreting foreign or personalized license plates, but tests showed that the system's rate of successful identifications increased due to the assumption about standard Swedish license plates since in Sweden the standard Swedish license plates far far outnumber foreign or personalized license plates. For license plates containing more than six characters, the system allows any combination of letters and numerals.

One benefit of the otherwise deficient OCR method of pixel-by-pixel comparison is that the quality of the character matchings can easily be evaluated. It's simply a matter of adding together all pixel differences between the unknown character and the chosen reference image. The best match will be the reference image resulting in the lowest sum of pixel differences. That sum can then be compared to the difference sum of the best match made in previous frames. When the final result of the license plate recognition is to be presented, the best character matches in each position from all frames in which the license plate appeared will be shown. Thus, the system will have increased the chances of successful license plate recognition by using information from many different video frames, which was one of the aims of this thesis.

3.11 Corner Tracking

As a last resort for tracking license plates between consecutive frames, corner tracking (with a bit of optical flow calculations added) is used. In other words, corner tracking is normally not used. It is only used as a way of finding the new location for a license plate that was found in the preceding frame but could not be found in the current frame by any other means than corner tracking. The reason for this conservative use of corner tracking is that during testing it proved to give less exact coordinates for license plates than the other methods.

For each corner of a license plate, a small window of pixels around it in the previous frame is fit into different positions around the expected location of the corner in the current frame. The position in which the sum of pixel differences is lowest is chosen as the most likely position of the corner in the current frame.

A method for corner detection is then used to get precise coordinates for the best corner in the vicinity of that position. The method depends on the license plate background to be brighter than its surroundings, so the method is not always successful, but adequate as a last resort. A small and square binary mask is created for each of the four license plate corners. The quadrant that will be inside the license plate, if a mask is placed with its center on top of its corresponding license plate corner, is painted white. The other three quadrants are painted black. That is, the mask for the top left license plate corner will have its bottom right quadrant painted white, etc.

The masks are then applied on the video frame by moving them in one-pixel-steps around the suspected location of their corresponding license plate corners. For each mask in each position, the luminance for all video frame pixels under the white quadrant are added together. The luminance values for all video frame pixels under the three black quadrants are then subtracted from that sum. The position in which the sum reaches its maximum is chosen as the most likely position for the license plate corner corresponding to that mask.

4 System Overview

A system for automatic license plate recognition was implemented on the previously listed hardware, using the previously described methods arranged according to the flowchart in Figure 6. For each step in the flowchart, a supposed license plate must pass all the tests in order to be allowed to continue to the next step. Objects failing a test are discarded and not investigated further. Here follows a description of the steps in the flowchart:

1. Background estimation. In order to reduce the computational cost, license plates are only searched for in areas classified as foreground.

2. Connected components. Areas enclosed by sharp boundaries (e.g. license plates) are located by using the connected component labeling algorithm.

3. Heuristics. The shape and size of areas found in step 2 is examined, to determine which areas that merit further investigation.

4. Thresholding. When image resolution is low, the gap between letters and license plate top and bottom borders is sometimes too small to be seen in the image. The license plate will thereby be split into several different areas by the connected components algorithm. As a countermeasure, the area found in step 2 is expanded, so that it hopefully encompasses any other areas belonging to the same license plate. To promote further analysis, a binary image is created from the expanded area by using the mean luminance value of that area as a threshold. Doing so increases chances that the licence plate as a whole can be separated from its background, even if partial areas of the license plate have gradual transitions to the background instead of sharp edges.

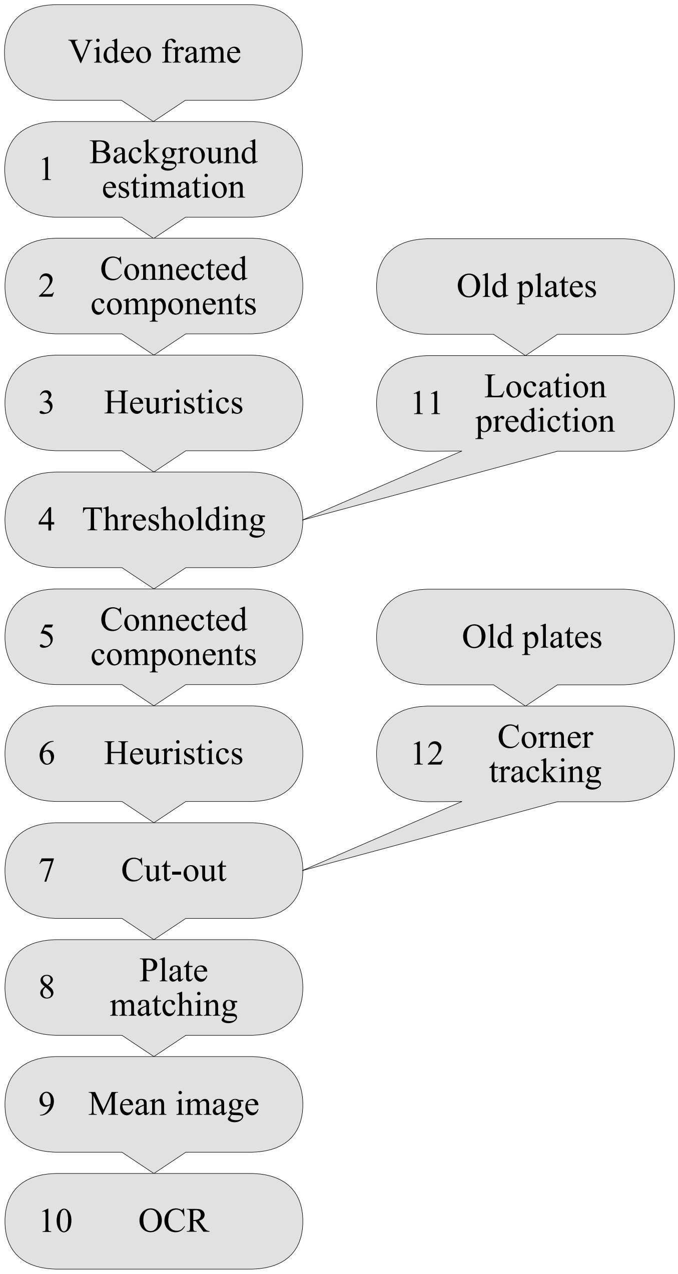

Figure 6: System overview.

5. Connected components. This time the connected components algorithm is not utilized to directly find license plates. Instead, it is used on the binary image from step 4 to find areas that are not part of the license plate. Any area connected to the borders of the binary image (the expanded area from step 4) is considered to not be part of the license plate. The remaining areas will then constitute the whole license plate, even if some of the areas are not connected to each other.

6. Heuristics. Another set of heuristics is examined, this time involving distance between the four corners of the supposed license plate, and the parallelism of its top and bottom borders.

7. Cut-out. Having passed all tests, the license plate is now ready to be cut out. And, by means of bilinear interpolation [GW02], it is made completely horizontal and given a predefined width and height.

8. Plate matching. All license plates found in a video frame are compared against those found in previous frames. For all matches, the information from the current video frame will be added to the information collected in previous frames.

9. Mean image. Those license plates that were found to be matching license plates found in previous frames are incorporated into the corresponding mean images by using Equation 2. License plates not seen in previous frames are used as foundations for new mean images. All mean images are contrast enhanced so that the characters stand out as much as possible.

10. OCR. Mean images from step 9 are interpreted by means of OCR. For each license plate, the degree of certainty in individual character recognitions are compared to the highest degrees of certainty for characters in corresponding positions in previous frames. The best matches in each position are presented on screen as the system's interpretation of the license plate. However, license plates seen for the first time in the current frame must have at least four characters with a high degree of certainty in order to be accepted as license plates.

11. Location prediction. For license plates that were found in any of the last five video frames but not in the current frame, an attempt is made to predict their location by assuming that they move at constant speed and retain their size and color. Their last known movement vector is added to their known or predicted location in the previous frame. The new location is then inserted into step 4 of the flowchart, where the old license plate's last known size and luminance threshold are used.

12. Corner tracking. For old license plates that were put through the location prediction in step 11 but still not found in the current frame, corner tracking is used in order to find the four license plate corners in the vicinity of the location predicted in step 11. Those four coordinates are then used as corners for the cut-out in step 7.

Figure 7: Example image from the reference movie.

5 Results

During development, a 16 minutes long film from a busy street in Lund (Figure 7) was used for testing the automatic license plate recognition. The system was optimized until a 98 % rate of successful identification (license spelled out correctly) was reached for that film.

5.1 Problems

The 98 % success rate could however not be maintained at the site of the implemented system (Figure 2). That is to be expected, the system will have become optimized for the film that was used as reference during development. In addition to that, the following causes for the reduced success rate were identified:

1. Image resolution. The reference film (Figure 7) had a resolution of 720x576 pixels, but at the site of the implemented system (Figure 2) a resolution of 352x288 pixels was used. The resolution was reduced in order to avoid interlacing (Figure 3) and to increase processing speed. In the reference film vehicles were apporaching the camera almost completely head-on, therefore interlacing was less of a problem and could be handled by a de-interlacing filter.

2. Headlights blinding the camera. The camera was equipped with a system for automatic gain control, thus darkness causes the camera to greatly increase its sensitivity to light. When a vehicle then approaches the camera, only the bright headlights against a dark background are visible in the video feed from the camera, making it impossible to discern the license plate even for a human eye.

3. Shadow edges across license plates (Figure 8). During a day, the sun might come into a position that causes many license plates to be partly covered

Figure 8: A license plate partly in shadow, making it difficult to detect.

Figure 9: A Serbian license plate, lacking clearly defined edges.

in shadow. The connected components algorithm will then treat the dark and bright halves as two different areas. The thresholding or the contrast enhancement might turn the dark half completely black.

4. License plates lacking clearly defined edges (Figure 9) will not be detected at all by the connected components algorithm.

5. Dust or snow whirled up by vehicles (Figure 10). Whenever a cloud of dust or snow is trailing a vehicle, the vehicle's rear license plate will be more difficult to detect.

Problem number one was alleviated by zooming in the camera on the point where most license plates appear, at the cost of not detecting license plates mounted high above the ground. Problems number two, three and four might be solved by using infrared lighting and cameras, provided that license plates reflect more infrared light than their surroundings and that vehicle headlights do not produce enough infrared light to blind an infrared camera. Due to economic constraints the infrared approach was not tested in practise.

Figure 10: Dust and snow whirled up by a vehicle, obscuring its rear license plate.

Problem number two is less severe in Sweden during summer, since daylight then prevails from very early in the morning to very late in the evening. Reversely, during winter daylight is scarce in Sweden and problem number two becomes severe.

Problem number five is alleviated in rainy weather, since then there is no snow and dust is bound to the ground. Problem number three is also alleviated in rainy/cloudy weather, since then there are no sharp shadows.

5.2 Test Period

In order to test what the implemented system is capable of under good conditions, results where monitored during a period of twelve mostly rainy summer days in 2007 (Table 1).

Type	Not detected	Plate detected, failed OCR	Plate detected, successful OCR	Total
Car going in	3	2	29	34
Car going out	3	1	30	34
Truck going in	0	2	4	6
Truck going out	3	2	1	6
Total	9	7	64	80

Table 1: Results from test period (June 25, 2007 to July 6, 2007).

The rate of successful identification was 80 % for this period. There were no false detections, all images saved by the system did contain a license plate.

There is a notable difference in the success rate for cars (87 %) versus that for trucks (42 %). One reason might be that one of the trucks in the test period had Polish license plates (Figure 11). Since the amount of trucks during the test period was less than 20 % of the amount of cars, the Polish truck had a large impact on the results and the degree of certainty is much lower for trucks

Figure 11: Polish license plate mounted in a difficult position on a truck.

than for cars. Nevertheless, license plates on trucks are often mounted in places that make them hard to detect (Figure 11).

6 Conclusions

This master's thesis has shown that there are clear benefits to be had from running an automatic license plate recognition system on every frame in a video stream instead of just on isolated still images. Combined with the rather simple methods of connected component labeling and least-pixel-difference OCR, this system was nevertheless able to match a more advanced still image system utilizing gradient concentrations and neural network OCR [Dra97]. Both systems reached a rate of succesful identification at 80 %.

As a comparison, according to [Dra97] the best system for automatic license plate recognition existing in 1995 was the one by Parsytec and Saab, having a success rate of 86 %. However, in 2006 the European Parking Association reported that "Today's technology can reach over 99.5 % success rate and read plates at speeds over 250 km/h." [Ass06].

7 Future Work

Searching for areas with high contrast gradients [Dra97] might be a wiser approach than connected component labeling [Coh93] or the Hough transform [KG95] since not all license plates have clearly defined edges (Figure 9). Also, sharp shadow edges across license plates (Figure 8) reduce accuacy for the connected component labeling approach. On the other hand, infrared lighting and cameras might make connected component labeling the perfect choice, provided that license plates reflect more infrared light than their surroundings. Infrared might also solve the serious problem of cameras being blinded by car headlights during nighttime, or strong sunlight during daytime.

A better method for OCR than pixel-by-pixel comparison should be used, e.g. some kind of neural network approach.

Acknowledgements

The author would like to thank his parents, Inger and Göran Trobro, for help with installing and maintaining the system at the test location. The author would also like to thank his supervisors, Håkan Ardö and Kalle Åström at Lund University, for help during development.

References

[Ass06] European Parking Association. Case study - license plate recognition. `http://www.parking-net.com/ShowCaseItem-5750.aspx`, April 2006. [Online; accessed 28-July-2007].

[Coh93] Harvey A. Cohen. One-pass gray-scale image segmentation. In *Proc DICTA-93*, pages 572–578, Sydney, dec 1993.

[Dra97] Sorin Draghici. A neural network based artificial vision system for licence plate recognition, 1997.

[GW02] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Prentice Hall, Upper Saddle River, N.J., 2nd edition, 2002.

[KG95] V. Kamat and S. Ganesan. An efficient implementation of the hough transform for detecting vehicle license plates using dsp's. *rtas*, 00:58, 1995.

[Sta99] Chris Stauffer. Adaptive background mixture models for real-time tracking. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 246–252, 1999.

[sus05] susning.nu. Registreringsskylt. `http://susning.nu/Registreringsskylt`, May 2005. [Online; accessed 3-August-2007].

[Wik07] Wikipedia. Automatic number plate recognition — wikipedia, the free encyclopedia. `http://en.wikipedia.org/w/index.php?title=Automatic_number_plate_recognition&oldid=139307319`, 2007. [Online; accessed 28-June-2007].

[YTKA] Keem Siah Yap, Yong Haur Tay, Marzuki Khalid, and Tahir Ahmad. Vehicle licence plate recognition by fuzzy artmap neural network.

Appendix: Recognition Examples

The following images have been automatically saved by the implemented license plate recognition system. The images are shown in pairs, to the left an image of the vehicle and to the right an image with the result of the license plate recognition. The result is shown in the bottom left corner of the image to the right, it is the final mean image of the vehicle license plate followed to the right by a collection of characters that the system believes represent the contents of the license plate.

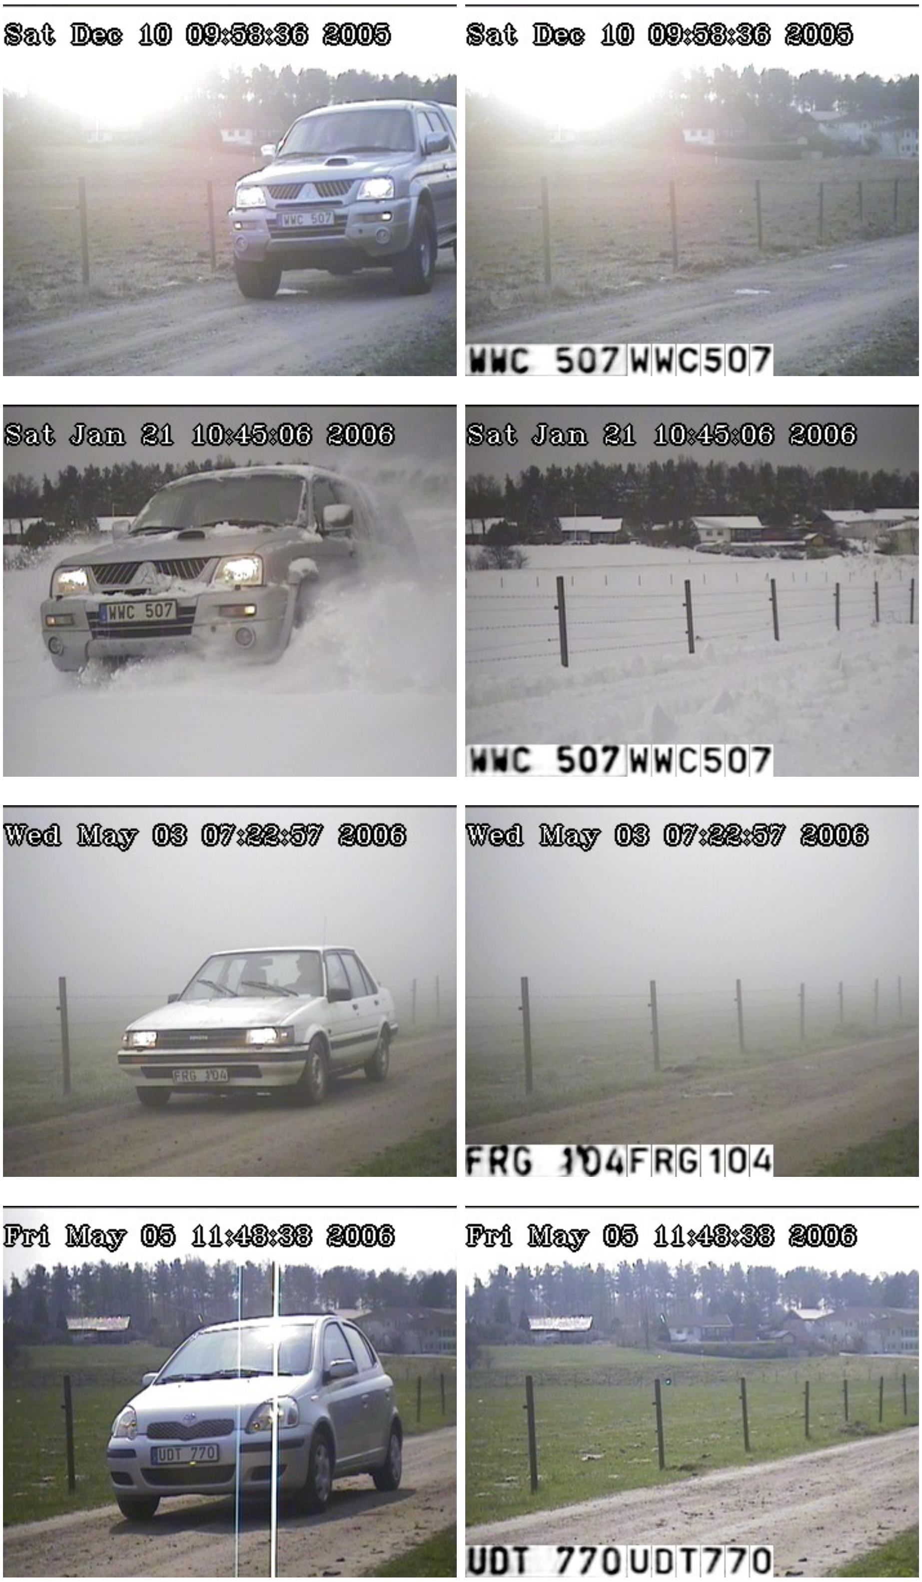
Sat Dec 10 09:58:36 2005
Sat Dec 10 09:58:36 2005
WWC 507
WWC 507 WWC507
Sat Jan 21 10:45:06 2006
Sat Jan 21 10:45:06 2006
WWC 507
WWC 507 WWC507
Wed May 03 07:22:57 2006
Wed May 03 07:22:57 2006
FRG 104
FRG FRG104
Fri May 05 11:48:38 2006
Fri May 05 11:48:38 2006
UDT 770
UDT 770 UDT770

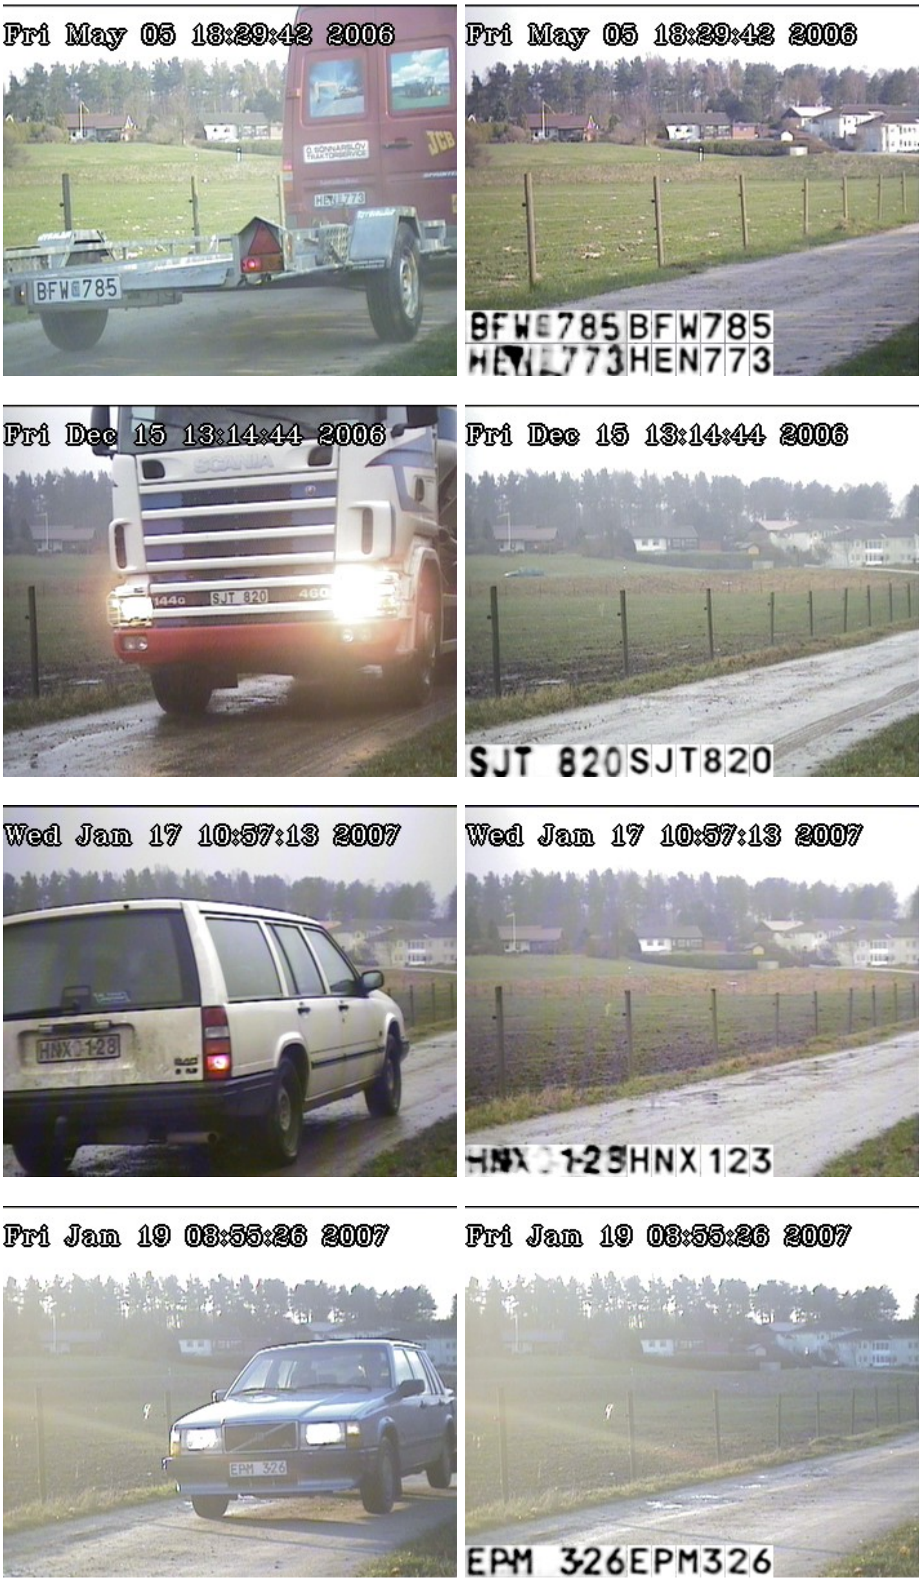

Fri May 05 18:29:42 2006
BFW 785
Fri May 05 18:29:42 2006
BFW785
HEN773
Fri Dec 15 13:14:44 2006
SJT 820
Fri Dec 15 13:14:44 2006
SJT 820 SJT820
Wed Jan 17 10:57:13 2007
Wed Jan 17 10:57:13 2007
HNX123
Fri Jan 19 08:55:26 2007
EPM 326
Fri Jan 19 08:55:26 2007
EPM 326 EPM326


Sat Jan 20 15:49:03 2007
Sat Jan 20 15:49:03 2007
UHO 461 UHO461
Tue May 01 11:26:35 2007
MITSUBISHI
Tue May 01 11:26:35 2007
WWC 507 WWC507
Tue Jun 12 17:54:12 2007
Tue Jun 12 17:54:12 2007
EPM 326 EPM326
Tue Jun 19 12:00:35 2007
SCANIA
SCANIA
Tue Jun 19 12:00:35 2007
CPP 736 CPP736


Tue Jun 19 13:20:35 2007
CPP 736
Tue Jun 19 13:20:35 2007
CPP 736CPP736
Tue Jun 26 20:23:20 2007
PL
LRY E950
Tue Jun 26 20:23:20 2007
LRY E950LRYE950
Wed Jun 27 12:14:33 2007
MERCEDES-BENZ
Wed Jun 27 12:14:33 2007
FXU868