

An Automated Tachograph Chart Analysis System

A. Antonacopoulos and D.P. Kennedy

PRImA Group, Department of Computer Science, University of Liverpool
Peach Street, Liverpool, L69 7ZF, United Kingdom
<http://www.csc.liv.ac.uk/~prima>

Abstract. This paper describes a new system that analyses tachograph charts. These circular charts are legal records of information on the different types of driver activity (driving, other duty, standby and rest) and vehicle data (speed and distance travelled). As each driver of each passenger and goods vehicle over a certain capacity must use a chart for every 24-hour period, there is a significant need for automated analysis (currently, tachograph charts are analysed manually).

The system starts by determining the shape parameters of the chart (location of the centre and radius). The position of the start of the 24-hour period (radius from centre to 24 hour tick mark) is then estimated. Finally the driver activity trace (recorded in a circular manner) is extracted, converted into a linear representation and recognised. Results from the evaluation of the system against professionally prepared ground truth indicate at least 94% accuracy in reading the driving time even on difficult (with scratches and marks) charts.

1 Introduction

Tachographs are devices that are fitted by law to all goods and passenger vehicles above a certain capacity (3.5t in the EU). These devices record onto circular paper-based charts the speed of the vehicle, different types of driver activity and the distance travelled (and sometimes additional information, such as engine speed and fuel consumption) over a 24-hour period [1].

Tachograph charts are legal documents that can, at any time, be inspected by the police for breach of driving regulations (relating to the supervision of the social welfare conditions for road traffic). Tachograph charts can also be used as evidence in court. Therefore, the accuracy by which the information is extracted from them is crucial.

There is a wealth of information contained on a tachograph chart. Apart from the information that the driver has to enter (start and end location and mileage reading), the other parameters that are recorded automatically can be cross-referenced to reveal inconsistencies. For instance, if the tachograph was tampered with to hide some part of the driving trace, there would be an inconsistency with the distance and speed traces.

It is common practice for transportation companies to gather the charts from their drivers at regular intervals and have them analysed by service bureaux. Given

the multitude of vehicles around the world (most countries have the same tachograph standards) there is a significant market for the analysis of these charts.

Currently, the most widespread tachograph chart analysis method involves the mounting of the chart under a magnifying lens and the visual examination of the information. A semi-automated system using a turntable under a fixed linear CCD has been developed to read relatively coarse information. This approach has the significant disadvantage of being very limited in the information it can extract (does not allow cross-referencing with other information on the chart) and it requires the manual placing of each chart on the turntable.

Ideally, human intervention should be kept to a minimum. The system proposed in this paper uses a commercial scanner to obtain images of the charts. A stack of charts can be placed in a sheet feeder, thus eliminating the need for accurate placement by hand. The scanned image of each chart is then analysed to extract the required information.

The characteristics of the chart and of the scanned image are described in sections 2 and 3, respectively. The system and its processes are detailed in Section 4. The paper concludes with the presentation of results and a discussion in Section 5.




Figure 1. A tachograph chart.

2 Chart Nature and Characteristics

A tachograph chart is made out of a circular thin card coated with wax. Sandwiched between the card and the wax is a layer of special ink. When the chart is inserted into the tachograph, three styli record different types of information by scratching the wax off the card. Where the ink is exposed to air it darkens to become black. The chart completes a full rotation inside the tachograph in 24 hours. Each driver has to use a new chart for each 24-hour period.

An example of a tachograph chart can be seen in Figure 1. Starting from the centre, four main areas with recorded information are present. At the innermost part of the chart, the handwritten information is completed by the driver. It records their name, starting point, destination, dates of departure and arrival, kilometre recordings before and after the trip, and total distance covered (in kilometres).

The next band of information records the distance travelled in up and down strokes of 5Km each (see Figure 2). This trace is useful in calculating the total distance travelled and in detecting when the vehicle stopped (the stroke becomes horizontal).

The third band of information is the one most often checked first for breach of driving regulations. It records the type of activity (mode) performed by the driver at any given point in time during the 24 hours. The different activities are represented by varying the thickness of the circular trace (see Figure 3). The two thickest trace bands denote driving time and other working time (other tasks at the workplace), respectively. The next trace band (thinner than the two above) denotes stand-by time—co-driver times, sleeping in cabin during the trip, and waiting times. Finally, the thinnest band denotes resting time. It should be noted that this very thin rest band does not always appear at the bottom of the mode trace. Instead it could be placed anywhere within the limits of the width of the duty band (e.g., it could appear starting from the middle of the width of the driving mode trace). The most important band here is the driving time, which is the first to be checked for calculating the total hours spent driving during 24 hours. The other bands are also useful for determining the number and duration of breaks.




Figure 2. Example of distance trace.

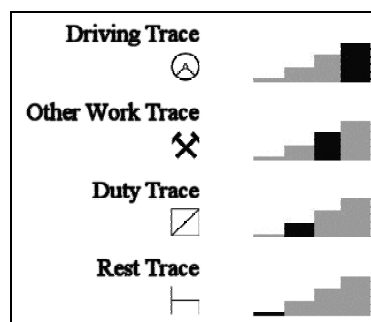


Figure 3. Explanation of mode trace.

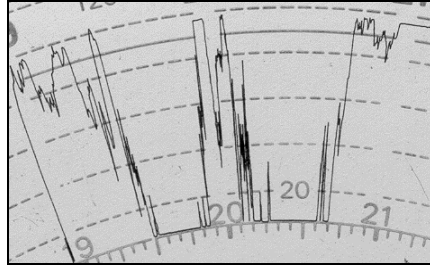


Figure 4. Example of speed trace.

The area of the chart furthest from the centre (see Figure 4) records the speed at which the vehicle travelled, displaying details of acceleration and deceleration.

Apart from the recorded information, the following characteristics can be observed. First, there is a significant amount of pre-printed information comprising digits and tick marks corresponding to the hours in the day, and various symbols aimed at explaining to the human operator the relative thickness of each of the activities of the mode trace.

Second, the hole in the centre of the chart has a shape that indicates a particular direction. In fact, the longest axis of the hole coincides with the notional axis passing through the tick marks of the 12 hour point on one side and of the 24 hour point on the other. Figure 5 illustrates the landmarks of the tachograph chart. Although different manufacturers will vary the pre-printed information on the chart, the main layout is basically the same (the layout illustrated in the figure is the one used in the vast majority of charts). It can also be seen that there is a semicircle printed running between the 12 and 24 hour points. This line could be useful in establishing the direction of the 12–24 axis.




Figure 5. Chart layout model showing landmarks.




Figure 6. Example of scratches on a chart.




Figure 7. Example of frayed edge.

The final physical characteristic to be noted is the colour of the tachograph chart. While the background of the chart (wax) is a very light grey, the pre-printed information can appear in different colours. In the majority of cases (over 90%) the pre-printed information appears in medium green, which has distinct grey level (see below). Other colours (red or blue) have similar lightness.

The nature of the chart and that of the tachograph apparatus in general, give rise to a number of problematic situations that impact on the performance of the system. First, the fragile wax coating of the chart is easily damaged by rough handling (the nature of the work does not permit any assumptions regarding careful handling of the charts). By the same principle that information is recorded on the chart (ink darkens when exposed to air), unwanted scratches and marks can appear. These scratches are the same colour as the useful information and frequently intersect with it (see Figure 6). Bad handling can also cause damage to the inner (hole) and the outer edge of the chart. Figure 7 illustrates a frayed edge. An indistinct edge can cause difficulties in establishing the correct location of the chart (see below).




Figure 8. Examples of mode trace overlapping with pre-printed information.




Figure 9. A (smoothed) histogram of a typical chart.



Figure 10. Shadow introduced by scanning

Second, a potential problem can be caused by variations in temperature in the vehicle. In cold weather, the wax is harder and the stylus does not remove it completely off the trace. This makes the corresponding trace appear lighter and non-uniform, as the ink has not darkened to the same degree across the trace. Thin gaps in the trace may be confused with standby mode, especially when the vehicle is frequently stopping (e.g., a delivery vehicle in an urban environment).

Finally, either through normal wear or through manipulation, a stylus may become slightly bent. The corresponding trace will not, therefore, be recorded at the intended location on the chart. Instead it may overlap with some of the pre-printed information. This can be a problem when parts of the mode trace overlap with the explanatory symbols for the same trace (see Figure 8).

3 Image Characteristics

From the above description of the nature of the tachograph chart it can be concluded that, in order to preserve potentially differentiating information, the scanned image should be greyscale. As can be seen in Figure 1, the different types of information (and noise) are present in the form of different greylevels.

It is interesting to note that the structure of the greylevel histogram of a chart is very stable. This structure can be seen in Figure 9 (the reader can refer to Figure 1). Starting from the lightest region on the right, the rightmost peak corresponds to the background generated by the scanner cover. The highest peak to its left corresponds to the background of the chart itself (wax), while the low hump to the left of that corresponds to the pre-printed information. The recorded (etched) information appears as the very low peak to the far left (black) of the histogram.

A number of characteristic artefacts are noted in the image. First, it can be observed that although the useful information appears to be well separated in greylevel from the rest of the data, scratches and marks appear in the same greylevels too. Furthermore, due to the cold temperature effects mentioned in the previous section, some trace regions may not appear as uniformly black as expected in general.

The scanning process itself introduces some of the artefacts in the image. Due to the chart being made out of card (thicker than normal paper), a shadow appears at the outer edge (perimeter) and inside the central hole (see Figure 10). This shadow complicates the estimation of the exact location of the chart and its centre (see below).

The scanning resolution required is determined to be at least 400 dpi. At that resolution the resulting image (roughly cropped to a rectangle around the chart) is about 2000 by 2000 pixels. At 400 dpi the speed trace and the rest mode trace are represented as very thin lines (frequently 1-pixel wide). As the primary goal of the system is to extract the driving mode trace, losing some information from the rest trace is not very serious. Moreover, as the speed trace is also not exploited in the current implementation, a subsampled image can be used during some operations for faster processing (see below).

4 The System

The input to the system is a greyscale (8-bit at 400dpi) image, containing a tachograph chart. This image can be the result of batch-scanning a stack of charts using a high-speed scanner with an appropriate sheet-feeder. The chart could land anywhere in the scanner, although scanner software enables the *approximate* cropping of the image to a square containing the chart. In the image, no assumption is made about the orientation of the 12-24 axis of the chart.

In order to extract the driving mode trace (the goal of the current system), the circular chart has to be accurately located in the image and its centre must be established. Following that, the orientation of the 12-24 axis is estimated, in order to determine the time of the day at which each of the driving intervals occurred. The circular band describing the mode trace is read after ‘unwinding’ the chart, obtaining a linear

representation. Each of these stages is explained in more detail in each of the following subsections.

4.1 Chart localisation

Chart localisation involves three processes: histogram analysis, tracing of the outside edge of the chart and estimation of the parametric form of the chart (circle radius and location of the centre).

From the (smoothed) image histogram, the greylevel range corresponding to the image background (scanner cover) can be determined, as explained in Section 3 above. Having that information, an attempt is made to trace the outside edge of the chart. This is not a straightforward way to recognise the circular shape, as explained above, as the circular shape of the chart is not accurately represented in the image due to edge imperfections and shadow effects.

As the tracing is an image-based operation and the shape of the chart is a basic circle, the image is safely subsampled to a third of its original size. The tracing follows the boundary between what is identified as the image background and what is not [2]. At the end of this process, to adjust for the shadow, an attempt is made to 'push in' each edge point toward the centre of the disk. Each edge point can move inward so long as the next pixel toward the centre is at least as dark as the original edge pixel. As an attempt to maintain the shape of the original circle and avoid spurious 'inlets', each original pixel is not allowed to go in more than a fixed length (usually 1 to 2 pixels) further than its neighbours.

The identified boundary pixels can be used as the basis for estimating the parameters of the circular shape of the chart. The Hough transform is used for this estimation [3]. To improve efficiency, only one every 50 boundary points (experimentally determined) are used as input. In addition, the Hough transform is performed in two stages. In the first stage, a coarse accumulator is used. A smaller fine-resolution accumulator is then used in the second application of the transform, centred around the parameters identified at the end of the first stage.

At the end of the localisation process, the centre and radius of the circular chart are available for the identification of the 12-24 axis and the subsequent reading of the mode trace.

4.2 Location of the 12–24 axis

Although not essential for calculating the total driving time during a 24-hour period, the location of the 12-24 axis enables the registration of each time interval on the chart with the actual time of day it occurred.

A first approximate location of the 12-24 axis is based on the assumption that it coincides with the orientation of the central hole (see Section 2 and Figure 5). To the best of the authors' knowledge, this assumption always holds (it is not stated as a fact since there is no information on the rules governing the actual manufacturing process of the charts). To achieve this first estimate, the inside edge of the central hole is traced and the edge point that lies further from the centre is noted. The line through

that point and the centre should coincide with the required axis. In fact, the error of this process is within the $\pm 7.5^\circ$ range, which equates to ± 30 minutes with the current image parameters.

A more accurate estimate is obtained by starting with the above information and complementing in the following way. A strip is read in a circular manner along the perimeter of the chart (full-size image). This strip is wide enough (10% of the radius) to contain the numerals denoting the hours. The resulting strip is a long narrow rectangle, where the pixel intensities along the columns are summed to produce a 1-dimensional signal. This signal is then averaged and a standard 1-D Sobel filter $[-1, 0, 1]$ is passed twice over it. The resulting signal contains both positive and negative peaks. Peaks whose magnitude is less than a third (experimentally determined) of the maximum magnitude are removed and the remaining positive and negative peaks are concatenated.

A characteristic pattern (a group of successive negative-positive-negative peaks) corresponds to single digit numerals and a different pattern (a group of successive negative-positive-negative-positive-negative peaks) to double digit ones. Other combinations of peaks are considered to be irrelevant. It should be noted that the peak in the centre of each of the two patterns above denotes the presence (location) of the hour tick mark.

The next task is to identify the 24 hour tick mark. For this purpose, any tick mark that gives a well-defined signature in the 1-D signal is identified and the corresponding angular difference between that and the approximate estimate (from the hole alignment is used). As the hole estimate is accurate to within 30 minutes, the true hour of the day can be established for the chosen tick mark. From that established tick mark, the accurate position of the 24 hour tick mark can then be calculated. The resulting axis is the one passing through the resulting tick mark and the identified centre of the chart.

4.3 Reading the mode trace

Having established the centre of the chart and its start and end (radius to 24 hour tick mark), the chart is 'unwound' to form a long rectangle. This 'unwinding' is performed in the full-size image by reading pixels along straight lines (Bresenham's algorithm [4] is used) as consecutive radii from the centre of the chart toward the perimeter.

As discussed in Section 3, recorded information appears as the darkest regions the image. In particular, the mode trace lies approximately (it is not certain, as explained in Section 2) in the middle of the chart, i.e. the half radius area. Suitable regions are identified as bounding boxes (parts of the mode trace can be isolated)

Within each of the bounding rectangles identified above, the greylevel values of the pixels are summed along each column. Columns corresponding to driving have a very low count, while columns corresponding to rest (or no activity) have the highest count. Each column is then grouped with its neighbours (if appropriate) to form bands representing the same activity: driving (lowest count), duty (medium count), rest (high count) and nothing (highest count).

As each column corresponds to a 14-second interval, the start time and duration of the different mode bands are then established.

5 Results and Discussion

The performance of the system was evaluated by analysing a representative (leaning toward difficult) data set of 108 charts varying from average difficulty (55%) to difficult (34%) to very difficult (11%). An example of a chart of average difficulty can be seen in Figure 11, whereas an example of a very difficult (mainly due to scratches) chart can be seen in Figure 12. For each of the charts there was corresponding ground truth prepared by bureau professionals. The accuracy of the system was determined by examining the discrepancies between the total duration of an activity during 24 hours as identified by the system and as read by the professionals.

The overall accuracy of the system for reading the *driving* mode is 94%. More specifically, charts of average difficulty were read with 97.9% accuracy, while difficult and very difficult charts were read with 94.4% and 90% accuracy respectively.

With regard to the *duty* mode, the overall accuracy was 97.2%. This can be broken down to 99.3% accuracy for charts of average difficulty, and 95.2% and 97.1% accuracy for difficult and very difficult charts, respectively.

All the above results are regarded as very good by the professionals, as 5 to minute discrepancies are not crucial in everyday circumstances. Some facts that shed light to the reading of charts by a human operator will help put the performance of the system into context. The device used by human operators has a resolution of 1 minute, whereas each mode trace pixel in the automated system (with the current settings) corresponds to 14 seconds. In practice, human operators 'average out' regions of sudden stops and starts, rendering precise estimates of accuracy rather difficult.

The main reason for error can be isolated as the inaccurate location of the centre of the chart. The precise location of the centre is crucial to all the processes that perform circular reading. In some cases, there is clearly room for further improvement, especially in making the initial tracing and centre estimation operations more robust.

Another issue to be raised is the appearance of the rest node trace. As mentioned in Section 3, this appears as 1-2 pixels wide in the 400dpi image. Moreover, as it is so thin, it tends to be represented by lighter greylevels. If the rest trace was deemed to be necessary to read accurately, the image would have to be scanned at a higher resolution (at the cost of memory and efficiency).

While the system is complete, it is still an experimental prototype. Further work will concentrate on two fronts. First, other options for improving accuracy will be investigated and integrated into the existing setup. Second, engineering aspects such as optimising the system for speed and, interfacing with and controlling the high-speed scanner will be pursued.

In conclusion, this paper has presented a system that automates the analysis of tachograph charts. The charts are scanned using a commercially available scanner, as opposed to the proprietary setup of a previous semi-automated method. The information of most interest is that recorded in the work mode trace, and the system concentrates on this. However, there is a wealth of information present on the chart

which can be read from the same image, for its own worth or for cross referencing and validating the mode trace.

6 Acknowledgement

The authors are grateful to Tachograph Analysis Consultants (TAC) Ltd of Liverpool UK, for providing the tachograph charts and ground truth.

References

- [1] VDO Kienzle, *Manual for the Evaluation and Use of the Original Kienzle Tachograph Chart*.
- [2] T. Pavlidis, *Algorithms for Graphics and Image Processing*, Springer, Berlin, 1982.
- [3] M. Sonka, V. Hlavac and R. Boyle, *Image Processing, Analysis and Machine Vision*, 2nd ed, PWS – International Thomson Publishing, 1999.
- [4] D.F. Rogers, *Procedural Elements for Computer Graphics*, McGraw-Hill, 1985.




Figure 11. A chart of average difficulty.




Figure 12. A 'very difficult' chart to analyse.