Norsk Regnesentral
NORWEGIAN COMPUTING CENTER

Norsk Regnesentral
Norwegian Computing Center
(+47) 22 85 25 00
nr@nr.no · www.nr.no

Note

SatTrafikk

Project report 2011

Note no SAMBA/05/12

Authors Siri Øyen Larsen, Arnt-Børre Salberg and Line Eikvil

Date February 2012

Contract JOP.11.11.2 (Norwegian Space Centre)

Norsk Regnesentral

Norsk Regnesentral (Norwegian Computing Center, NR) is a private, independent, non-profit foundation established in 1952. NR carries out contract research and development projects in the areas of information and communication technology and applied statistical modelling. The clients are a broad range of industrial, commercial and public service organizations in the national as well as the international market. Our scientific and technical capabilities are further developed in co-operation with The Research Council of Norway and key customers. The results of our projects may take the form of reports, software, prototypes, and short courses. A proof of the confidence and appreciation our clients have for us is given by the fact that most of our new contracts are signed with previous customers.

Norsk Romsenter

Norsk Romsenter (Norwegian Space Centre, NSC) is a government agency under the Ministry of Trade and Industry. NSC promotes the development, co-ordination and evaluation of national space activities as well as supports Norwegian interests in the European Space Agency (ESA). Earth observation involves all activities related to collection of information on the Earth's surface or atmosphere from instruments on board satellites. The Norwegian Space Centre's application programme supports users, research communities and businesses in testing the potential of Earth observation from satellites. Priority is given to the development of applications having public benefit.

Statens Vegvesen Vegdirektoratet

Statens Vegvesen Vegdirektoratet (The Norwegian Public Roads Administration (NPRA)) is responsible for the planning, construction and operation of the national and county road networks, vehicle inspection and requirements, driver training and licensing. It is also authorized to grant subsidies for ferry operations. The objective of the NPRA is to develop and maintain a safe, ecofriendly and efficient transport system. This is being done on a sound, professional basis by interacting with politicians, users and other interested parties.

Norsk Regnesentral
Norwegian Computing Center
Postboks 114, Blindern
N-0314 Oslo, Norway

Besøksadresse
office address
Gaustadalléen 23A, B
N-0371 Oslo, Norway

Telefon · telephone
(+47) 22 85 25 00
Telefaks · telefax
(+47) 22 69 76 60

Bankkonto · bank account
8200.01.48888
Foretaksnr. · enterprise no.
NO 952125001 VAT

Internett · internet
www.nr.no
E-post · e-mail
nr@nr.no

Title	**SatTrafikk project report 2011**
Authors	**Siri Øyen Larsen, Arnt-Børre Salberg and Line Eikvil**
Date	February 2012
Publication number	SAMBA/05/12

Abstract

The SatTrafikk project is a research project conducted by the Norwegian Computing Center and funded by the Norwegian Space Centre and the Norwegian Public Roads Administration, where the main objective is to develop a future system for automated road traffic counts from satellite imagery. The project was started in 2007 as a follow-on to the European Space Agency project "Road Traffic Snapshot" (2006-2007), which demonstrated the feasibiliy of this kind of vehicle detection.

In order to apply the SatTrafikk vehicle detection software in operational scenarios, several key components must work together. Up until 2010, modules for the detection of roads, vehicles, and clouds and cloud shadows, were separately developed. In 2011 we have integrated these modules into a processing chain with the aim of operationalizing the derivation of traffic statistics from satellite imagery. In this report we describe the newest methodological improvements and expansions of the automatic detection algorithms. The entire processing chain has been validated on five QuickBird and WorldView-2 scenes from Mid- and Northern Norway, with satisfactory results.

Keywords	Remote sensing, image analysis, pattern recognition, vehicle detection, road traffic statistics, very high resolution satellite images
Target group	Road traffic authorities
Availability	Open
Project number	220 497
Research field	Earth observation
Number of pages	26

Contents

1 Introduction

Road networks are resources of major importance for the society. Operation and development of road networks is a central activity for several public institutions, such as the Norwegian Public Roads Administration (Statens Vegvesen Vegdirektoratet, SVV). Traffic statistics is a key parameter for this activity. The primary source of traffic statistics today is ground based counts generated using various types of equipment mounted in or close to the road. Such equipment, e.g., radar, induction loops or pressure sensors, counts the number of cars passing a given location on the road during a period of time. Important statistics describing traffic can be derived from these counts, most importantly the so-called AADT, i.e., the Annual Average Daily Traffic, which is the average number of vehicles passing a given location during one day, taken as an average over a year. In Norway AADT is estimated using ground based vehicle counts in combination with statistical tools developed at the Norwegian Computing Center (Norsk Regnesentral, NR).

For fairly large parts of the Norwegian road network AADT is still unknown. The reason is that installation and operation of measurement equipment for ground based counts are both difficult and expensive, hence there are relatively few counting locations as seen in a geographical scope. Especially, AADT is missing for most roads with low traffic density on national basis.

Over the last few years, very high resolution satellite sensors have opened up for alternative means of traffic monitoring. Vehicle counts based on automated satellite image analysis can provide useful additional information to traditional traffic surveillance. A significant advantage of satellite based technology is that it does not require installation or operation of equipment in the road, thus maintenance demands and operating costs are no longer an issue. Moreover, a satellite image can cover large geographical areas and in principle this allows for AADT estimation of all the roads in the region, as opposed to only a few roads, as one is restricted to using fixed ground based measurements.

In 2006-2007 NR conducted the European Space Agency project "Road Traffic Snapshot" (see http://dup.esrin.esa.it/projects/summaryp92.asp) in cooperation with SVV and Institute of Transport Economics (Transportøkonomisk Institutt). Fundamental algorithms for vehicle detection using satellite images were demonstrated in the project. In the following project "SatTrafikk", starting in the summer of 2007 and funded by SVV and the Norwegian Space Centre (Norsk Romsenter, NRS), the detection algorithm was further developed and optimized. In 2008 the methods were validated on a large data set containing a variety of road conditions from different parts of the country. While the original methods were developed using satellite images covering the urban area of Oslo, this larger data set contained images from Sennalandet, Bodø, Kristiansund, Østerdalen, Eiker and Sollihøgda. These images revealed different types of challenges when it comes to automatic vehicle detection in different parts of the country, and especially there are differences between urban and rural areas.In 2008 we also conducted research on the expected quality of AADT estimates (derived from the statistical basis curve method), given the availability of one or a few satellite images per year. For roads with relatively large AADT as seen in a national context (i.e., AADT >20,000 vehicles) the results were promising (absolute error less than 20% given two satellite images a year), with the precondition that the vehicle detection algorithm is fairly accurate. For roads with smaller AADT (<20,000 vehicles) the corresponding average error was around 25%. With AADT less

than 1,000 vehicles, a larger error is expected, although, as traffic statistics hardly exist for such roads, there was no data evidence to verify this.

On the background of the above mentioned experiences with vehicle detection in different environments, combined with SVV's need for improved geographical coverage of traffic statistics, the focus of the SatTrafikk project in 2009 was directed towards relatively low traffic density roads, mainly located in rural areas. Large parts of the previously suggested methods were replaced or modified to meet the specific challenges related to vehicle detection under these conditions, and to optimize the detection strategy in general. After 2009, the vehicle detection algorithm may be considered as robust in cloud free images and manually based road detection. In 2010, algorithms for automatic road detection and automatic detection of clouds (with corresponding cloud shadows) were developed. These modules are essential for operational use of SatTrafikk since clouds will often be present in optical images, and satellite images are not delivered with sufficient precise geo-referencing.

In order to apply the SatTrafikk vehicle detection software in operational scenarios, several key components must work together. Up until 2010, modules for the detection of roads, vehicles, and clouds and cloud shadows, were separately developed. In 2011 we have put these modules into a processing chain together with necessary help modules for image and vector import etc. This involved various technical and organizational challenges, as well as new methodological problems, especially regarding the fact that the preconditions for vehicle detection changed when replacing the manually delineated road mask with the automatically produced one.

The operationalization process and technical documentation is described in a separate document [4], while the current report describes the methodological aspects that we have worked with during 2011.

2 Automatic road detection

In 2010 the basis methodology for the automatic road detection was developed [3]. The algorithm is based on using GIS data of the roads as prior information of the position of the road in the satellite image. Even though the 2010 road segmentation algorithm worked very well for many road segments, the precision needed to be improved before it could be used as part of a vehicle detection system.

2.1 Improvement of the road segmentation

The road segmentation is based on a mathematical model (called snake) of the road, initialized by the road GIS vector data. The vector data is needed to find the correct road, and to find the location of the road in the image, but cannot be used directly, as the vector and image data are not accurately co-registered. To simplify the snake based segmentation algorithm, the image is re-sampled along lines perpendicular to the road vectors. This transforms the road into a nearly straight line. The position of the midline of the road is found in multi-spectral resolution (using the blue image band), and then the position is adjusted using region-growing and various refinement techniques of the segmentation result in the panchromatic image. We will here describe the new improvements of the algorithm.

2.1.1 Geo-referencing the satellite image

To improve the precision of the road detection, we wanted to include elevation information into the process of geo-referencing the satellite image, and thereby improve the match of the road GIS data to the road shown in the satellite image. Our initial attempt was based on geo-referencing the satellite image using a digital elevation model (DEM) and the rational polynomial coefficients (RPC) delivered along with the satellite image. The RPC constitute a replacement model of the camera in that they are a generic set of equations that map object coordinates into image coordinates. However, since the quality of the 25m DEM covering Norway was not good enough, the geo-referenced satellite image was severely distorted (Figure 2-1), and an alternative approach for co-registering the vectors to the image was needed.

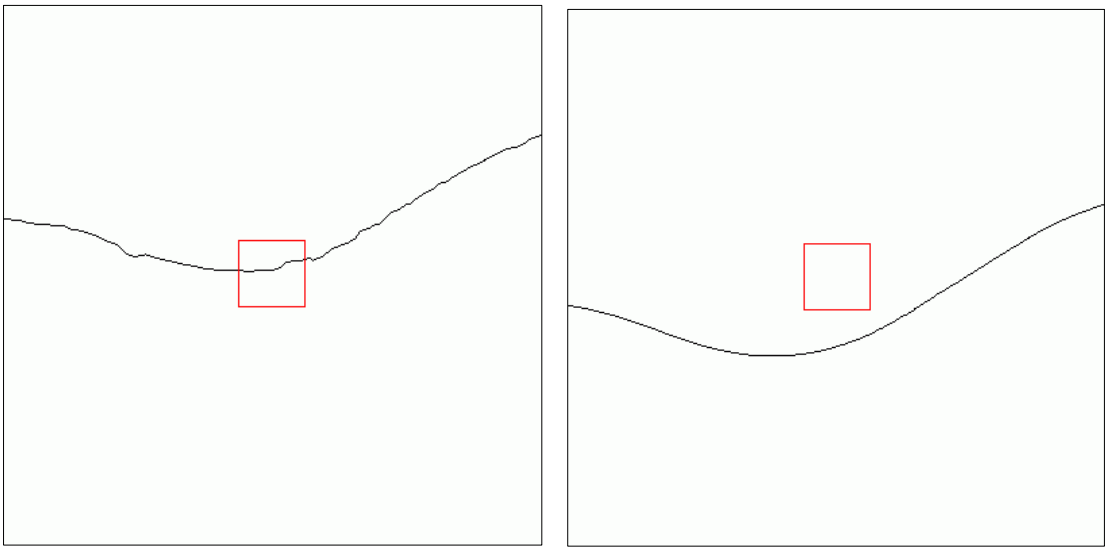

Figure 2-1: Example of road detection in DEM-based geo-referenced satellite image. Left: Road segment after geo-referencing using DEM. Right: Road segment without geo-referencing. The distortion introduced by the DEM-based geo-referencing is clearly shown left image.

Since the road GIS-data are delivered with height information in addition to geographical position, we reverse the geo-referencing approach and *use the RPC model to map the GIS-data into the satellite image*. In this case the satellite image is not resampled (still in ortho-ready projection). Moreover, the approach is not based on the national 25m DEM, and no DEM-related distortion is introduced in the mapped GIS-data. However, the height information in the GIS-data is based on GPS, and this measure may be inaccurate in some cases.

2.1.2 Controlling the snake model

The snake model basically consists of two terms that describes the shape of the snake. One term that is related to the prior information about the model (called internal force), which in this case is the GIS-data, and one term that is related to the image data (called external force). The behavior of the snake is controlled by two parameters that weight the internal and external forces.

The segmented road had a tendency to be displaced from the correct position at bus stops (bright area next to the road), in cloudy (bright) areas, etc. Therefore, we modified the parameter α, which controls the external forces of the snake, in order to make the movements of the snake less (image) data dependent, i.e., to trust the vector data more (see Figure 2-2). This could be done as the vector positions after geo-corrections were more accurate. Furthermore, while the cloud mask was previously not used to assist road detection, we now incorporate this information by setting α=0 in locations that (according to the cloud mask) are covered by clouds, i.e., where there is no information in the image that can guide the snake.

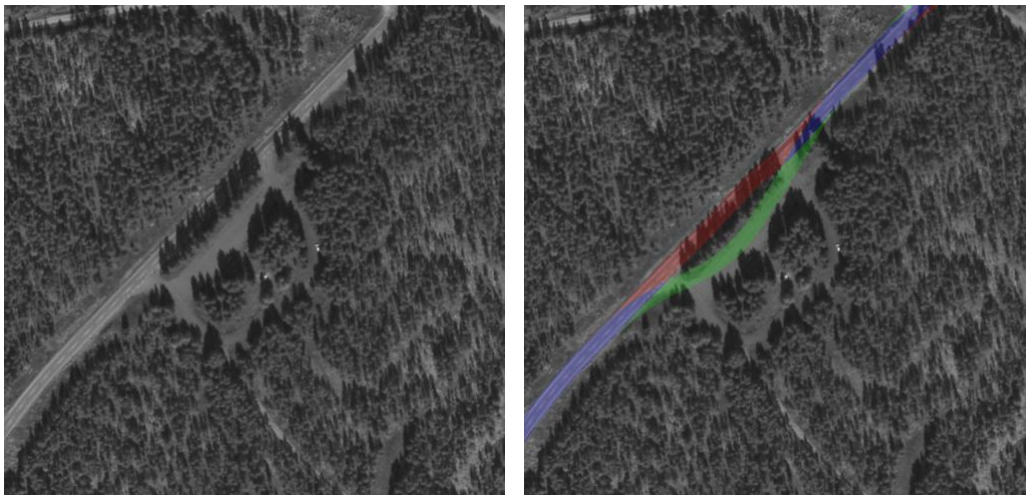

Figure 2-2 In this image we see that the view of the road is almost totally blocked by dark trees, while there is a bright passage to the right of the road. If the α-parameter of the snake is too high, the algorithm will rely more on the image data than the vector data, and the road mask will be drawn towards the bright area on the side in stead of the real road. In the right image the new (red) and old (green) road masks (blue indicate overlap between the two) are overlaid on the image.

2.1.3 Panchromatic adjustment

The panchromatic adjustment step has also been improved, mainly because the region growing step, which was previously performed using IDL's built-in function REGION_GROW, have been replaced by our own implementation. The initialization of the region growing is unchanged: the initial region is computed by resizing the one pixel wide center line image found in multi spectral resolution to panchromatic resolution, hence to a ~4 pixel wide center line, and then performing a morphological closing operation. The main concept of region growing also remains unchanged, i.e., neighboring pixels are included if the intensity value of the new pixel differs from the mean region intensity by less than a preset constant times the

standard deviation of the region pixels. However, while IDL's REGION_GROW function uses the global statistics of the region, we have implemented a version where only local region pixel values (in a 49×49 pixels neighborhood) are considered. The main problem with IDL's REGION_GROW for the purpose of road segmentation is that the standard deviation of pixels on a global level is very high, since it captures both very bright areas (e.g. clouds) and very dark areas (e.g. shadows), thus the new region is grown far too large. Furthermore, the panchromatic adjustment step also incorporate distance thresholding, to ensure that pixels close to the midline is included in the road mask, and to limit the width by discarding pixels very far from the midline. We discovered that the segmented road in some cases was too narrow, especially for the WorldView-2 images, and therefore we translated the threshold distance measures to metres and changed the distance thresholding implementation to become sensor independent.

2.1.4 Road GIS- data errors

It should also be mentioned that we have spent a lot of work detecting and fixing bugs in the vector data that, when untreated, leads to strange (bad) results, such as:

1. Missing data

2. Wrong direction vectors

3 Improvements of the automatic vehicle detection

Initial reports documented that the performance of the vehicle detection was very good [1]. However, by including automatic road detection, allowing the use of road segments in hazy and cloud shadowed areas, and allowing clouds to cover parts of the road, further adjustments and improvements of the vehicle detection algorithm was necessary.

The proposed segmentation strategy is based on a so-called Laplacian of Gaussian filter which is used to search through the image for elliptically shaped "blobs", i.e., regions of relatively constant intensity that is brighter or darker than the local background. The size of the elliptical main lobe of the filter corresponds to the size of typical vehicles, and the orientation corresponds to the orientation of the road, hence the vehicles, in the image. Interesting locations in the image can be extracted from the image response to the filter. Details describing the approach can be found in 2009 project report [1]. We will here explain the main adjustments and improvements that have been made in the 2011 project.

3.1 Contrast threshold

Even if the filtering technique is robust towards local contrast changes, in that it only considers pixels in a neighborhood whose size is equivalent to the size of the filter, we experienced that we failed to capture some of the vehicles with low contrast. Such vehicles are typically found in areas with a lot of haze and cloud shadows, or they are dark vehicles appearing on a very dark background, which is the case when the asphalt is new. The reason why the method failed to capture these vehicles was that the estimated blob contrast threshold - which refers to the absolute contrast - was too strict. The contrast relates to the intensity values, which were normalized over the road pixels of the global panchromatic image. However, it is better to normalize the image locally, since this makes the span of intensity values smaller, thus the contrast enhancement becomes better (see Figure 3-1). This is especially important when the image contains a lot of clouds, with very high intensities, and at the same time, very dark, shadow areas, because then the difference between a local and a global normalization is larger.

We also performed a study to optimize the contrast threshold, and found that 0.6 and 1.0 were suitable thresholds for the absolute value of the estimated contrast of dark and bright blobs, respectively.

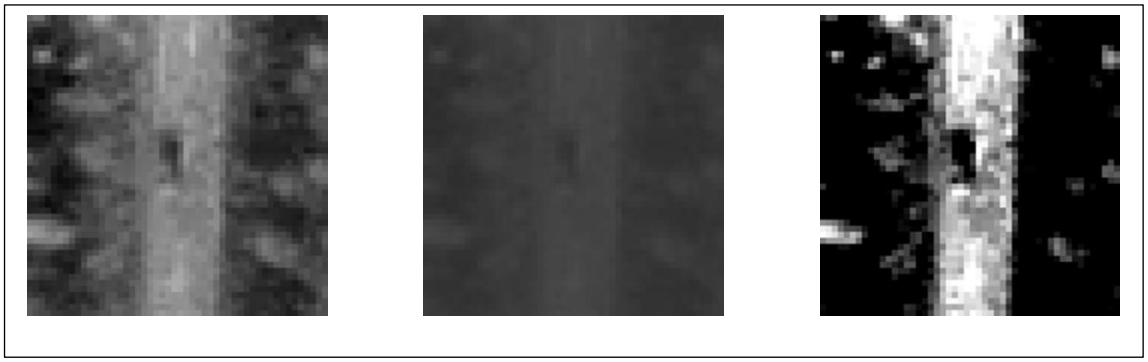

Figure 3-1. Subsection of image, showing a car in a shadowed area. Left; displayed using linear histogram stretch over the full range of values. Middle; globally normalized image (before subsectioning), displayed using linear histogram stretch on the range [-2.0, 1.0] Right; locally normalized image, displayed using linear histogram stretch on the range [-2.0,1.0].. The range of histogram values of the middle and right images are [-1.49,-0.93] and [-6.06,5.76], respectively.

3.2 Generalization

3.2.1 Adaptation to other satellite sensors

The vehicle detection method was developed using QuickBird images only. In 2010, we bought two WorldView-2 scenes, which have a resolution of 0.5m in the panchromatic band (as opposed to 0.6m for QuickBird). Since the goal is to achieve a fully automatic solution for vehicle detection, the algorithm must be able to read sensor/resolution information from the input data and adjust parameters for blob detection (scale space search) etc. accordingly, so that the vehicle detection becomes independent of image resolution. It is also more intuitive to measure object features in the metric system, rather than relative to pixel size.

3.2.2 Adaptation to non-connected road segments

The stand-alone implementation of the vehicle detection module was developed and validated with manually delineated road masks, which are far more accurate than the automatically generated ones. The introduction of automatic road detection has produced new challenges for vehicle detection as some vehicles may be cut in half or completely left out, and non-road areas included in the automatic mask yield more clutter and more non-vehicle objects. Furthermore, the manual road detection was performed in (more or less) cloud-, haze-, and shadow-free areas only, and avoiding road segments with more than two lanes. On the contrary, in the latest version we have decided to include the entire road, except areas covered by clouds, as defined by the cloud mask (which, again, is automatically constructed, hence, inevitably, it is not completely perfect). However, shadow and/or haze are present in most of the images we have analyzed this time, and even dominating the area in some of them. We have also adapted the method to allow analysis of an arbitrary number of non-connected road segments in the same image, as opposed to the previous version, where only one connected piece of road could be handled.

The only manual interaction taken at this point is the construction of masks covering roads with more than two lanes, in order to exclude analysis of these areas, since the system is not prepared for this, and the road vector data available as of today does not contain information about the number of lanes.

3.3 Integrating the cloud and cloud shadow mask

In 2010, we developed a method for detection of clouds and cloud shadows, which produces a cloud mask that classifies the image by labeling the pixels that belong to cloud or cloud shadow. As already mentioned in Section 2.1.2, the cloud mask yields useful information to help guide the snake for the road detection. In the vehicle detection module, the cloud mask is exploited to avoid areas completely covered by clouds (where vehicle detection does not make sense, as the ground can not be observed). A naive approach to avoid clouded areas would be to simply multiply the road mask with the opposite cloud mask. However, as illustrated in Figure 3-2, the cloud may enter the road from the side, and cover only a part of the road in the width direction. For the AADT estimate to make sense, we must report the length of the observed road in the image, thus, we should exclude all parts where the road can not be observed in its entire width. The solution we have developed is as follows: We represent the road by points along the midline. Then, for each cloud pixel inside the road mask, we calculate the midline point nearest in distance to the cloud pixel, and exclude this point. Finally, we analyze each connected line of midline points not excluded by the cloud mask. If the length of the line segment is below a preset threshold (currently set to 100 m), the corresponding road segment should not be considered for vehicle detection.

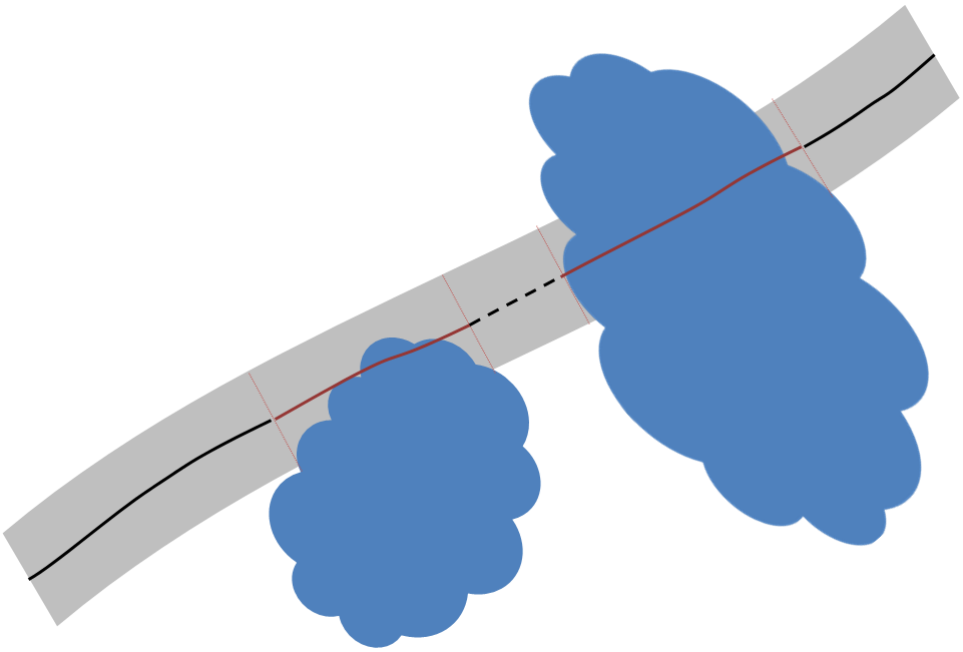

Figure 3-2 Sketch showing how we automatically estimated the observed road. Solid black lines indicate the midline of the road, solid red lines indicate the midline of road that should not be analyzed due to cloud cover, while the dotted black line is the midline of observed road that should not be analyzed because the road segment is too short.

We have also added a feature which indicates the cloud mask value for each object (this should be a number between 0(=cloud free, no shadow), and 1(=shadow)). This feature is currently not used for object classification, as its discriminative power in a kNN-classifier is poor. However, it may be useful for some sort of pre- or post-classification, as the object intensity contrast is expected to be lower in shadow and/or haze (the method can easily be modified so that it also produces a haze mask). This has not been studied so far.

3.4 Classification

As before, we use object feature extraction and a kNN-classifier to categorize the road objects after segmentation. However, since the road mask is produced automatically, and therefore less precise, in addition to the other changes in the system, the training data base had to be updated. We also found it useful to introduce some more classes (in the previous version, only two classes - car and no-car - were used), especially because we now also need to separate vehicles of different size. Furthermore, we have developed a new method for linking the objects after classification, in order to avoid double counts of the same vehicle, since vehicles are often composed of two-three objects (Figure 3-3). It should also be noted that the tree shadow pre-classification method (described in the project report from 2009 [1]), is no longer used, as we discovered that it lacked robustness and the generality needed to hold for a larger variety of image conditions.

Since the automatically produced road mask is less precise than the manually delineated mask, there are more non-vehicle objects to be classified. There are especially many objects with more or less rectangular or oval shape at the side of the road, in between vegetation or in the ditch alongside the road.

The class division scheme is shown with example objects in the following illustration:

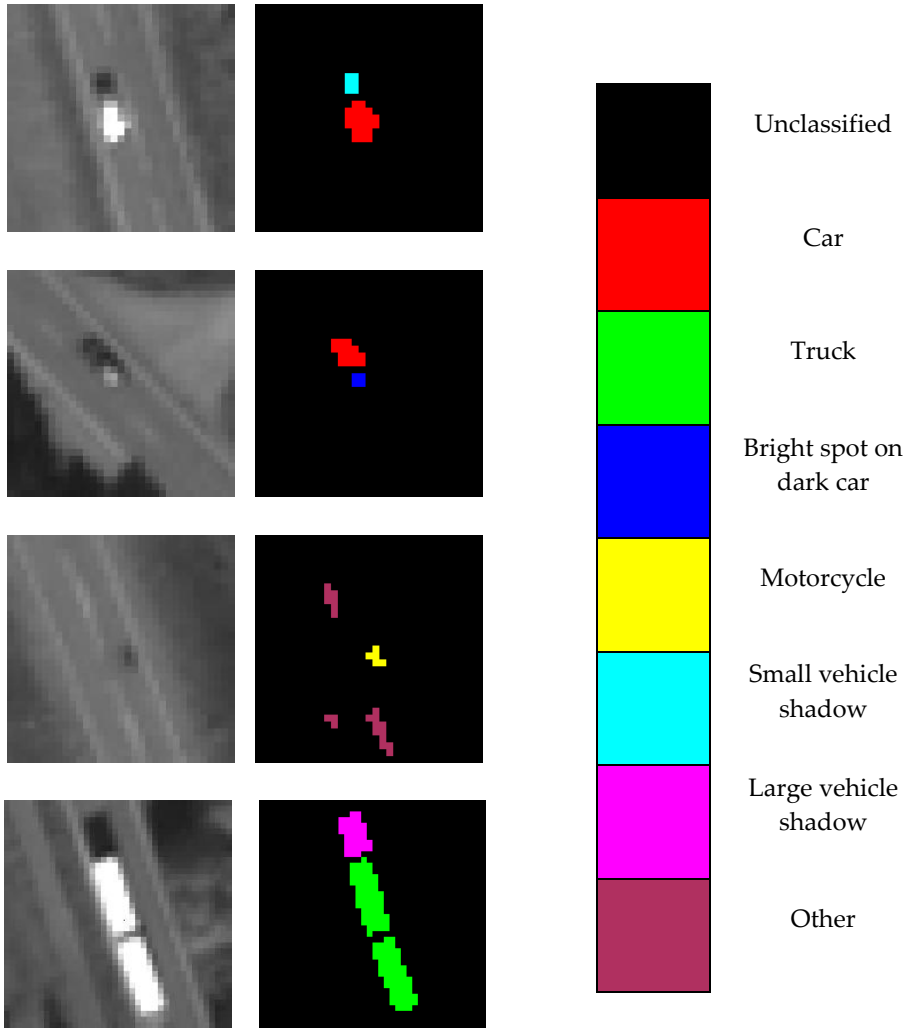

Figure 3-3 Class division scheme for classification of road objects.

As before, we scaled the features (to have zero mean and unit variance) prior to classification, and we treat classification of bright and dark objects separately. We have also used feature selection methods to seek the optimal set of features for both cases, and we also found it advantageous to use $K=5$ (instead of $K=3$, as in the previous version).

The following features were chosen for bright objects:

1. LoG amplitude

   The amplitude of the image response to the $\nabla^2 G$ filter at the scale of the detected blob (for detailed explanation, see [1]).

2. longitudinal (short side) contrast

   The mean contrast in the longitudinal directions, defined as the intensity difference (in the rotated panchromatic image) between the blob center and points lying 1.5 times the

length of the major principal half axis of the main lobe of the filter, to the left/right (i.e., contrast between the object center and points lying in front of/behind the object, related to the direction of the road).

3. mean pan intensity

   The mean intensity (in the panchromatic image) of the pixels of the object region.

4. pan standard deviation

   The standard deviation of the intensity values of the pixels of the object region.

5. local pan mean

   The mean intensity of the pixels of the object region after local intensity normalization. The local intensity normalization is performed by subtracting the mean and dividing by the standard deviation of the intensities within the road mask.

6. sobel gradient

   The mean absolute value of the Sobel gradient of the pixels of the object region.

7. area

   Object area measured in square metres.

8. spread

   A measure of the spread of the object region, defined using the region's second order central moments: $(\mu_{20} + \mu_{02})/(\text{area}^2)$.

For dark objects, the features are:

1. LoG amplitude

   As above.

2. longitudinal (short side) contrast

   As above.

3. mean pan intensity

   As above.

4. sobel gradient

   As above.

5. perimeter

   The object boundary is found using dilation of the object region with a 3x3 flat structuring element, and the perimeter is calculated by multiplying the number of boundary pixels with the pixel resolution (in metres).

6. distance from midline

   The distance (in metres) between the centroid of the object region and the closest point on the midline of the road.

7. blob contrast

   The estimated contrast of the elliptical blob found using the Laplacian filters (for detailed explanation, see [1]).

8. object width

   Width of the bounding box of the object region - after rotation so that the orientation of the region is aligned with the coordinate axes.

3.4.1 Training the classifier

We have developed a module for semi-automatic marking of objects (training), where only those objects close to a vehicle have to be labeled, while the remaining objects are assumed to belong to the class "Other". This means that we first have to define a vehicle buffer mask, and we do this in ENVI by inspecting the image and marking each vehicle location with a point ROI, and then we grow a buffer mask around the vehicle locations using the built-in ENVI functionality "Create buffer zone from ROIs" (which can be found in the ROI Tool menu's "Options" tab). We then construct the vehicle buffer mask, using an integer distance kernel, and maximum distance 30 pixels. The next step is to loop through all the objects overlapping the vehicle buffer zones and register the class. We have developed a simple ENVI/IDL script which makes this easy (see Figure 3-4). Each object is presented to the operator in the panchromatic image, where it is contoured, and the operator may also display the segmentation image, and/or a non-contoured panchromatic image on the side and link the images, in cases where it is difficult to make a choice of class in the first case. It should be noted that the marking is nevertheless a manual process, and since the resolution is relatively poor (compared to the size of the objects we are going to detect), one should not expect complete agreement between two different human operators.

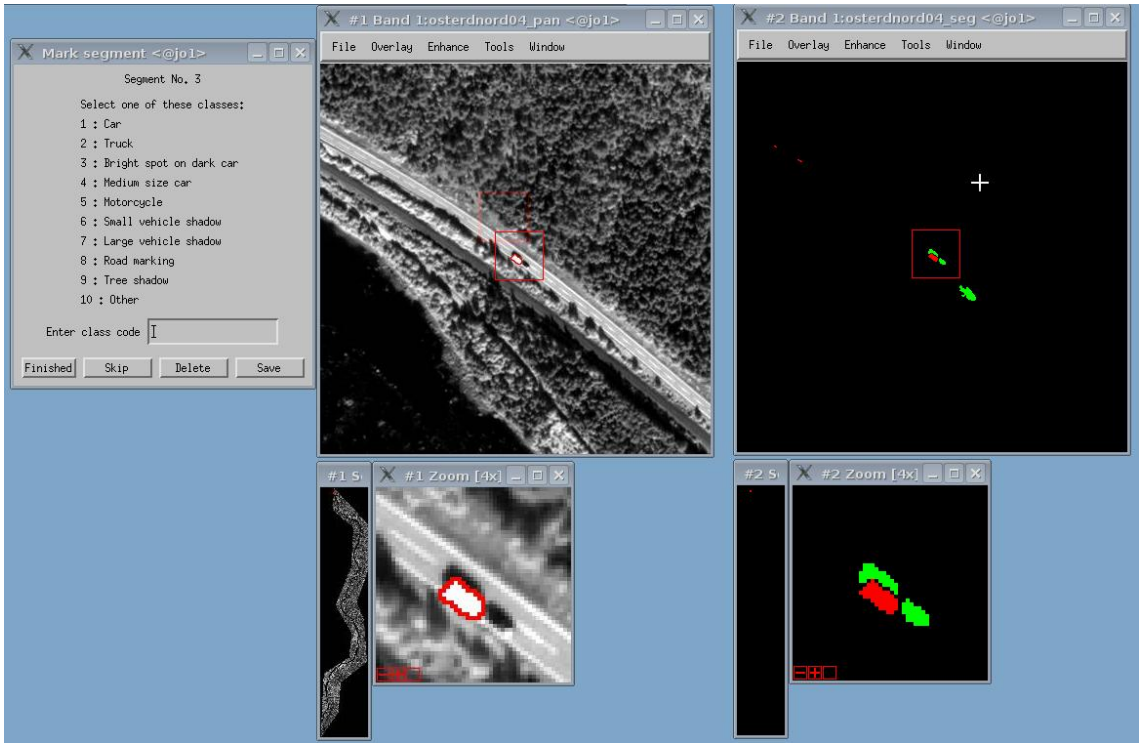

Figure 3-4. Graphical user interfacing for marking objects (training the classifier).

3.5 Object linking and vehicle type registration

After classification, each vehicle may be represented by one or more road objects - typically, the vehicle and its shadow, but we also have several examples where either the vehicle or the shadow object is lost, either in the segmentation or the classification. Therefore the objects should be linked, in order to ensure that each vehicle is counted once, and only once.

In previous work, using two classes only (vehicle and no-vehicle), the final vehicle image was constructed by dilating the bright objects classified as vehicles in the direction of the expected shadow, and then add the dark vehicle objects, in order to ensure overlap of the segments. This technique was no longer suitable when using more classes (see Figure 3-3). Moreover, predicting the size and shape of vehicle shadows is very challenging, regarding the fact that the resolution is poor related to the size of the vehicles, the height of the vehicle is not known exactly (it varies especially between cars and trucks), and, most importantly, the segmentation of the vehicle may fail to capture the exact shape and size of the vehicle (which do not influence the vehicle count, but have a lot to say for shadow shape prediction). We have therefore developed a new strategy for object linking, where the only assumption is that vehicle objects close to each other in the length direction of the road belong to the same vehicle. It should be noted that if the traffic is standing still (e.g. in rush hour queues), this would not be a suitable technique. However, the system we have developed is not meant for such cases, and it is fair to assume that there is at least one vehicle length between two subsequent vehicles.

For each object classified as vehicle or vehicle shadow, we construct a rectangular box with the same center, width and length as the object, but oriented in the same direction as the road at this location. We then extend the box so that it becomes twice as long, and 1.2 times as wide (Figure 3-5). Next, overlapping boxes are treated as one vehicle, and the vehicle type is registered. If a vehicle consists of objects with non-consistent vehicle types (e.g., car+large vehicle shadow, or car+truck), we choose the largest type.

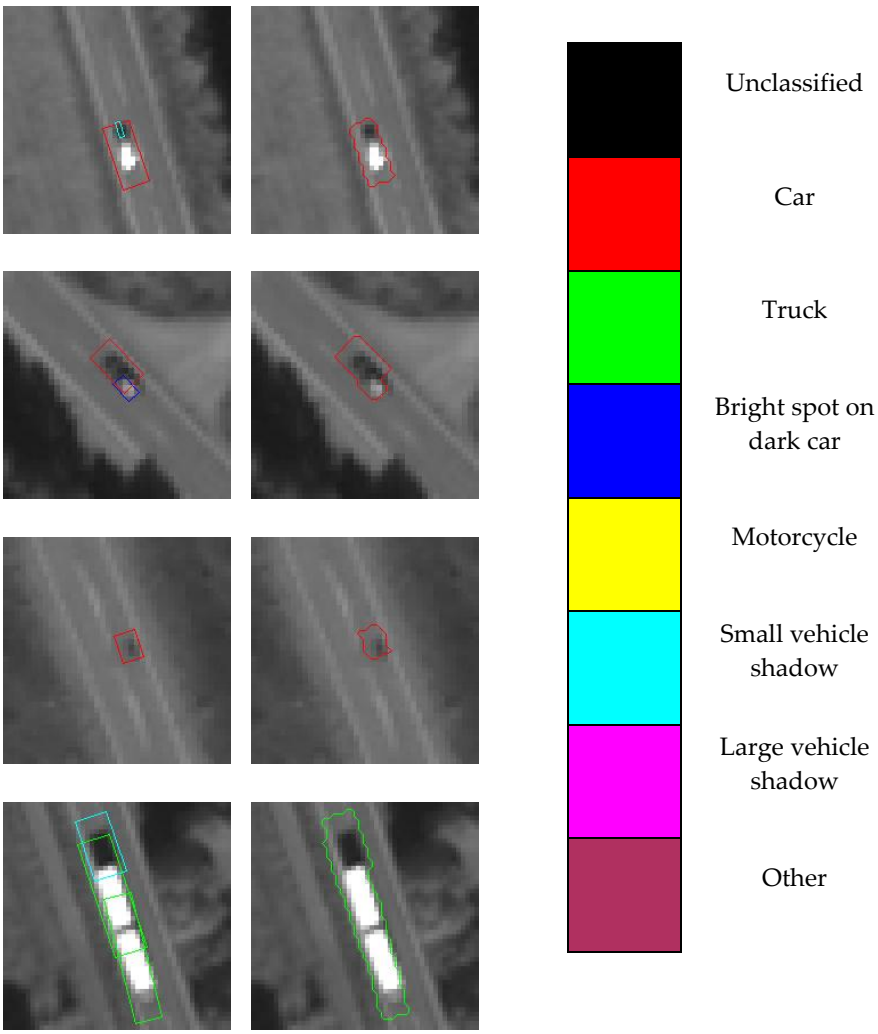

Figure 3-5. Object linking. Left column of images show rectangular boxes capturing the classification result. Right column illustrates the final vehicle result.Note two classification errors: the motor cycle is classified as car, and the vehicle shadow of the truck is classified as small vehicle shadow. In the latter case, the final vehicle is still classified as truck, since the other objects in the vehicle group are classified as truck.

3.6 NVRH estimation

The number of vehicle passing a point in the road per hour (NVRH) is given by:

$$\frac{number\ of\ observed\ vehicles\ in\ the\ image}{length\ of\ observed\ road\ in\ the\ image\ (km)} \cdot vehicle\ speed\ \ km\ \ h\ .$$

The number of observed vehicles is counted as described in the previous section, while the length of the observed road in the image is found automatically from the extracted midline of the road (Section 3.3). Since the midline is a digitized version of the reality, consisting of (densely spaced) linear line segments, we can not simply use the total length of the line segments, because this would yield a higher number than the true road length. Instead, we "smooth" the midline, using every fourth point only, and calculate the length. The vehicle speed can be approximated using the road speed limit.

4 The SatTrafikk-system

In 2008-2010, the SatTrafikk-project has developed methodology for:

- Automatic construction of road mask from satellite image and GIS-data (road detection)
- Automatic detection of clouds and cloud shadows
- Automatic detection of vehicles
- Estimation of the number of vehicles on a given road segment

In addition, two new components were developed in 2011:

- Vehicle type classification
- NVRH estimation

In 2011, the individual components were connected into a processing chain from satellite data to NVRH. The processing chain is referred to as the *SatTrafikk-system* (Figure 4-1).

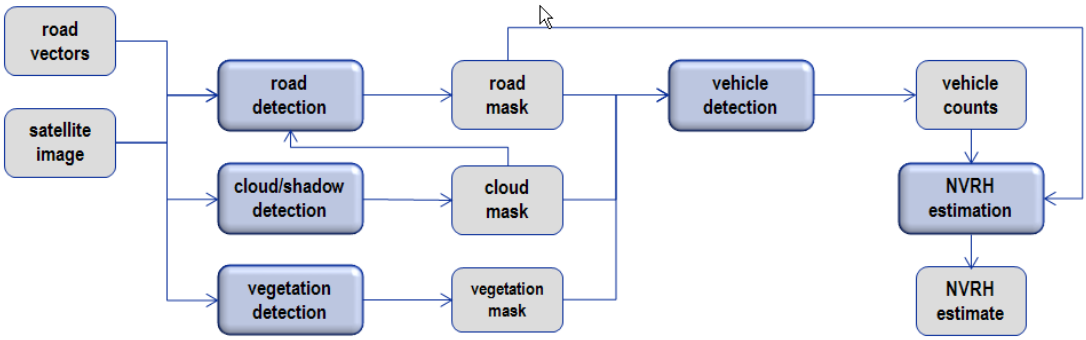

Figure 4-1 SatTrafikk processing chain.

5 Validation and discussion

The methods were validated on a total of five QuickBird and WorldView-2 scenes, details of which can be found in Table 5-1. In each case, the entire road in the image was analyzed, except those parts where the road has more than one lane in each direction (according to manually constructed "ignore area" mask), and areas covered by clouds (as defined by the automatically constructed cloud mask).

Table 5-1 Validation data.

Location	*Road*	*Date*	*Time (UTM)*	*Image area (km^2)*	*Observed road length (km)**
Østerdalen north	RV3	Aug 10, 2004	10:39	59	31.6
Østerdalen south	RV3	Aug 10, 2004	10:39	94	43.1
Østerdalen north	RV3	Sep 6, 2009	10:29	59	28.8
Mosjøen	EV6	Jun 10, 2010	11:30	95	31.0
Nordkjosbotn	EV6	Jun 8, 2010	11:02	55	12.7

*Estimated from the automatically generated road mask.

The length of the observed road has not been validated, except for the Østerdalen north 2004 image, where the true number should be 30.9km (this is found manually by comparing the image with a map, and finding the corresponding start and end road coordinates (given in metres along each road segment with a given ID number) in the image and on the map. As we see from Table 5-1, the number we have calculated is too high, indicating that the mid line should be smoothed further (perhaps using a smoothing spline) before estimating the length (cf. Section 3.6).

The validation was performed manually as follows: for each scene we made a "ground truth database" of vehicles that should be found. Then, the algorithm was applied, and for each vehicle in the ground truth database, the operator registered 1) whether or not it was detected, 2) whether or not the vehicle type registration was correct, and 3) whether or not it was counted twice. The remaining detections were registered as false. We also tried to interpret the surrounding image area of the false detections, and noted some possible explanations why the given object was falsely registered as a vehicle by the automatic method.

Results for each of the five images are presented in Table 5-2 through Table 5-6. The overall detection rate is 85.4%, and false detection rate is 8.0% (not counting numbers in parentheses - see tables for explanation). Example false detections are showed in Figure 5-1.

Counting vehicle type, we see that 11.2% of the detected vehicles have been classified as the wrong type. It is interesting to note that vehicle type mis-classifications occur much more often for trucks than for cars (2% of the detected cars and 35% of the detected trucks are classified as a different type), while all the detected motorcycles (100%) are mis-classified (as cars).

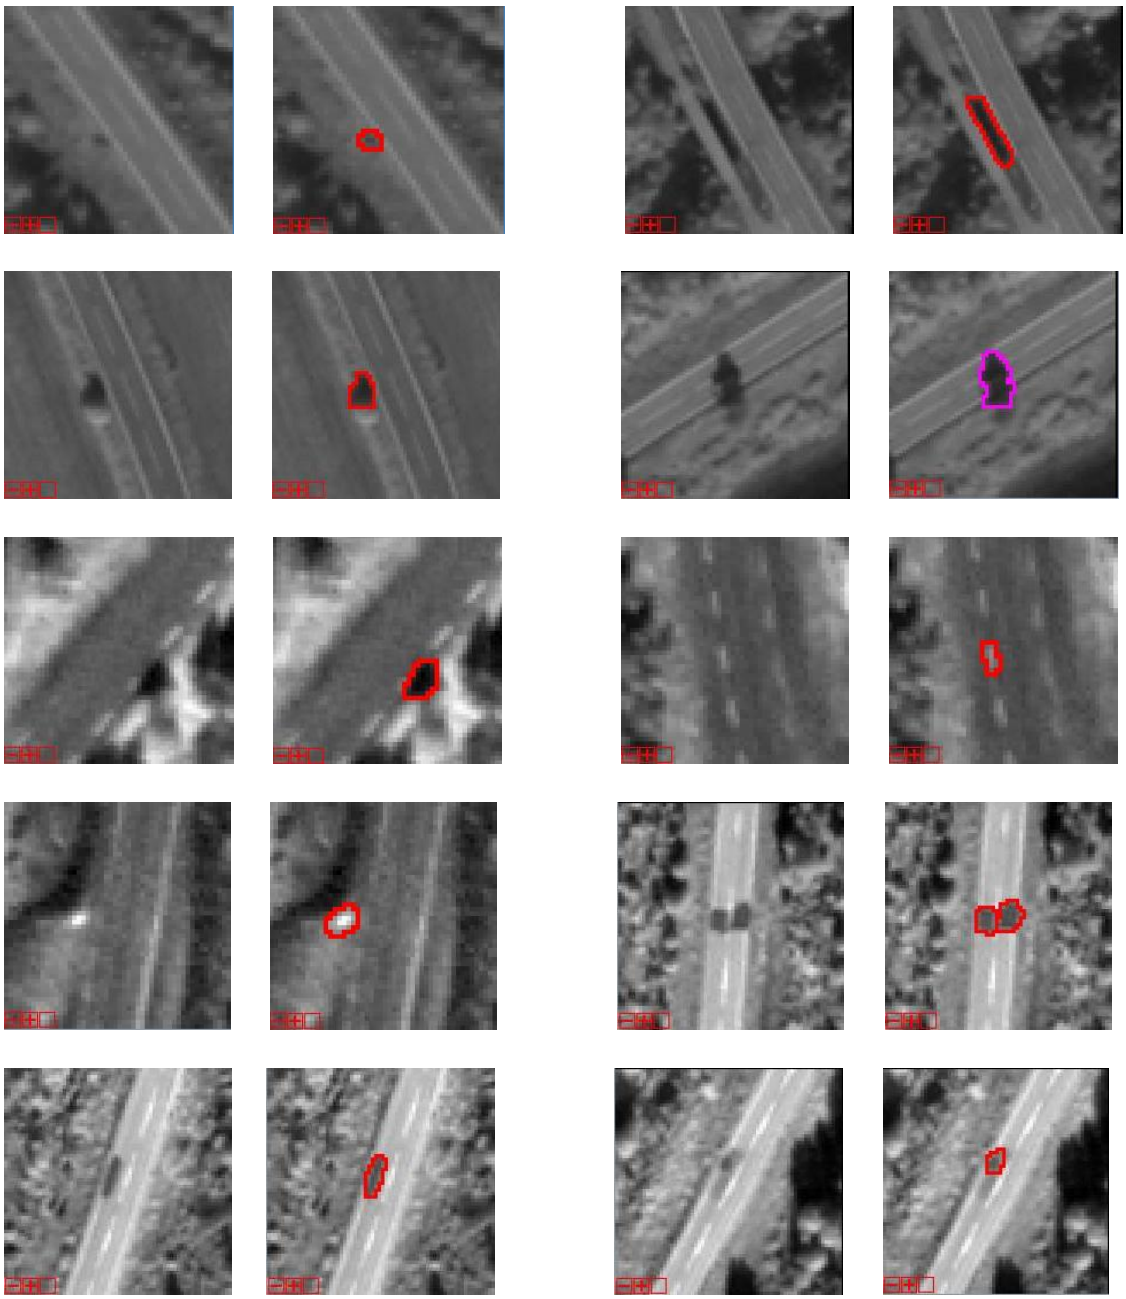

Figure 5-1 Example false detections.

Table 5-2. Validation results for the Østerdalen north 2004 image.

		vehicle type			
		car	truck	mc	**total**
Østerdalen north 2004	true vehicle	30	11	2	**43**
	found vehicle	30	11	1	**42**
	vehicle type error	0	2	1	**3**
	double count	0	1	0	**1**
	false detection	5	0	0	**5**
	segmentation ok	30	11	2	**43**

Table 5-3 Validation results for the Østerdalen south 2004 image. Numbers in parentheses refer to objects that are false detections wrt. the groundtruth, but at the same time true vehicle detections, although the vehicles in question are parked near the road, but not on it, hence, they should not be counted (Figure 5-2).

		vehicle type			
		car	truck	mc	**total**
Østerdalen south 2004	true vehicle	100	24	5	**129**
	found vehicle	91	22	5	**118**
	vehicle type error	2	7	5	**14**
	double count	2	0	0	**2**
	false detection	9 (+1)	1 (+1)	0	**10 (+2)**
	segmentation ok	94	23	5	**122**

Table 5-4 Validation results for the Østerdalen north 2009 image.

		vehicle type			
		car	truck	mc	**total**
Østerdalen north 2009	true vehicle	33	6	0	**39**
	found vehicle	28	5	0	**33**
	vehicle type error	1	2	0	**3**
	double count	0	0	0	**0**
	false detection	2	1	0	**3**
	segmentation ok	32	6	0	**38**

Table 5-5 Validation results for the Mosjøen image. Numbers in parentheses do not include vehicles thathave no overlap with the road mask (Figure 5-3).

		vehicle type			
		car	truck	mc	**total**
Mosjøen	true vehicle	57 (56)	14 (12)	0	**71 (68)**
	found vehicle	40	11	0	**51**
	vehicle type error	1	6	0	**7**
	double count	0	1	0	**1**
	false detection	7	0	0	**7**
	segmentation ok	54	11	0	**65**

Table 5-6 Validation results for the Nordkjosbotn image.

		vehicle type			
		car	truck	mc	**total**
Nordkjos-botn	true vehicle	19	13	0	**32**
	found vehicle	16	8	0	**24**
	vehicle type error	0	3	0	**3**
	double count	0	0	0	**0**
	false detection	0	0	0	**0**
	segmentation ok	18	12	0	**30**

As we have seen before, the detection rate varies considerably between the images, and the performance is closely related to the image conditions. Not surprisingly, the best performance is seen in the Østerdalen north 2004 image (detection rate 97.7%), which is a totally cloud-, haze- and shadow-free image, with clear and orderly conditions. In the opposite end of the scale we have the Mosjøen image (detection rate 71.8%), which contains a lot of clouds, haze and fog, making the image grainy, much of it also combined with low contrast shadow areas.

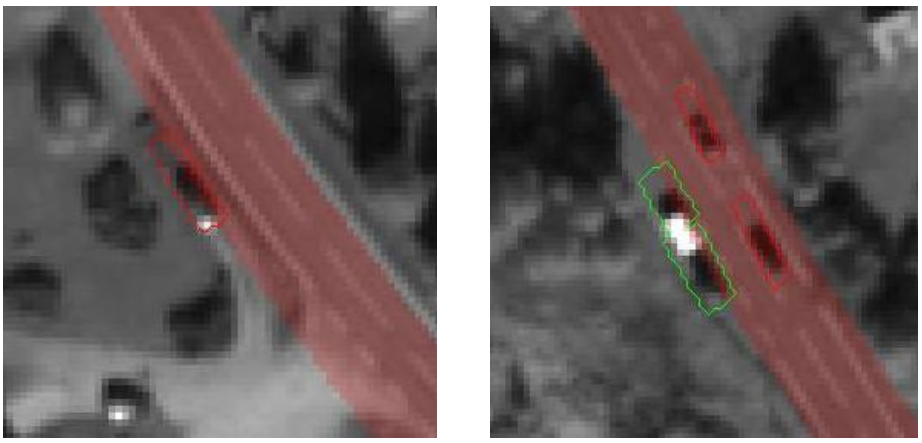

Figure 5-2 "False" detections. The car in the left image is parked next to the road, and therefore not a part of the groundtruth, yet it has been detected. A similar situation is presented in the right image, with the car and car trailer on the rest stop.

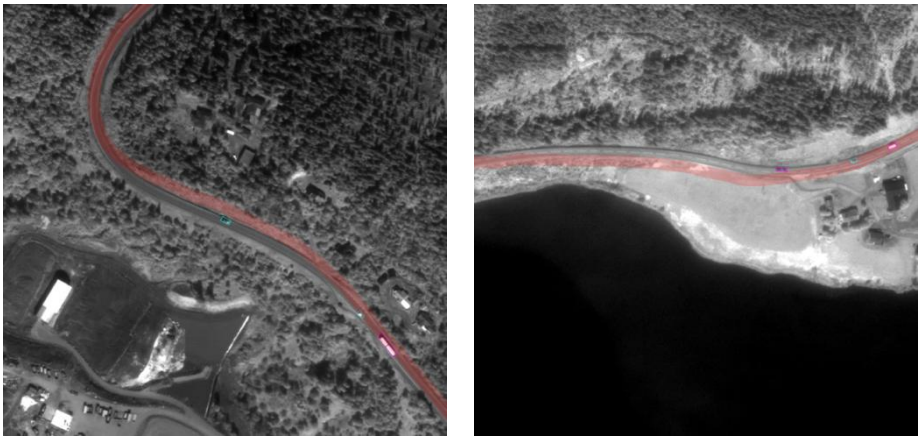

Figure 5-3 Examples of missed vehicles in situations where the road mask is completely off the road (no overlap between vehicle and road mask). The road mask is overlaid in red on the panchromatic image. Groundtruth vehicles are outlined in cyan (cars) and magenta (trucks).

6 Summary and conclusions

In 2011, we have integrated the separate detection modules developed in earlier years of the SatTrafikk project, in a common system (processing chain). In addition to the implementation of scripts that bind the modules together, several methodological and technical adjustments of the individual modules were necessary in order to make them work together.

The most important methodological improvements include:

- the construction of a geo-referencing module to utilize elevation information in order to improve the process of geo-referencing,
- modifications of the snake routine (in the road detection module) to make it less image data dependent, and as a special case of this, in the presence of clouds, the road detection should trust the vector data only,
- a new implementation of the region growing routine in the road detection module, which is based on local image values instead of global, and
- adjustments of the contrast calculations in the vehicle detection module,
- the development of a new class scheme and object linking strategy, and
- the development of a module for NVRH estimation.

Previous years, sub images (containing road segments of roughly 1-3 km) were extracted from the total scene, in order to avoid cloud and shadow areas, multiple lane roads, as well as to avoid heavy processing. In 2011, on the other hand, we wanted to focus on operationalization, with as little human interaction as possible, thus validation was performed on the original satellite data (.TIL file covering entire scene, automatically imported into ENVI). The only manual interaction performed was the construction of masks covering multiple lane roads, in order to avoid such. Moreover, the method has been organized to handle this automatically once the road vector data can be delivered with information about the number of lanes.

The precision of the road mask is a crucial quantity for high vehicle detection performance. For reliable operation it is necessary that the accuracy of the road mask is about 0.5m. This is an extreme situation compared to other remote sensing applications. It is expected that when information about bus stops, multiple traffic lanes, etc., are delivered, the performance of the road segmentation will be sufficiently high for most situations, and the vehicle detection performance will improve.

References

[1] Larsen, S. Ø., and A. B. Salberg (2009). SatTrafikk project report 2009, NR Note no SAMBA/55/09, Norwegian Computing Center.

[2] Salberg, A. B., Eikvil, L., Johansen, K., and Fredriksen, S. (2010). SatTrafikk 2010 - Operasjonelt scenario. NR Note no SAMBA/58/10.

[3] Salberg, A. B. and Eikvil, L. (2010). SatTrafikk project report 2010, Automatic road detection, Automatic cloud and cloud shadow detection. NR Note no SAMBA/57/10.

[4] Salberg, A. B., Larsen, S. Ø., Myrhaug, K., Solberg, R., Eikvil, L., and Johansen, K. (2011). Plan for operasjonalisering av SatTrafikk-systemet i 2012-2013. NR Note no SAMBA/55/11.