Consistency Checking for Vision Tasks

M. Middendorf, H. H. Nagel, and A. Ottlik
Institut für Algorithmen und Kognitive Systeme,
Universität Karlsruhe (TH), 76128 Karlsruhe, Germany
e-mail: {markusm|nagel|ottlik}@iaks.uni-karlsruhe.de

Abstract

Video surveillance systems usually have to operate on thousands or ten-thousands of frames. Interactive frame-by-frame assessment of results, therefore, is time consuming and expensive. We present a simple and fast approach which allows automated cross-checking of image segmentations obtained by different algorithms or two versions of the same algorithm.

Image regions show up in a number of video surveillance tasks, such as face detection, vehicle tracking, or gesture recognition. We define a figure-of-merit which reflects the consistency of pixel-wise represented image regions. The practical use of this approach is demonstrated by computing the consistency of regions obtained by model-based vehicle tracking and motion-based image segmentation. As will be seen, algorithmic difficulties of either of these approaches are closely linked to significant changes of the defined figure-of-merit.

1. Introduction

Video surveillance tasks often lead to an implicit or explicit segmentation of images. *Recognition tasks* implicitly split the image into an area corresponding to the recognized object(s) and a background, *region tracking tasks* separate the image of the tracked object(s) from each other and from the background.

Whereas similar approaches within the same 'state space' can easily be checked against each other by defining a distance measure on the state vectors, this is not the case when different approaches are to be compared.

For this reason, a figure-of-merit has been developed in order to quantify the compatibility of image segments obtained by different approaches. Extensive experiments have been performed in order to establish the suitability of this figure-of-merit as a criterion for the decision whether both results are compatible with each other or not. Such algorithmic consistency tests become important in large scale experiments for parameter tuning or the evaluation of methodological modifications. Often, only small changes can be observed in such experiments. This necessitates the accumulation of large sets of observations to obtain a reliable basis for conclusions. Visual scrutiny by humans thus becomes tedious and prone to inspection errors – unless the majority of consistent results can be detected and skipped quickly in order to concentrate on the critical cases.

2. The Approaches to be Analysed for Consistency

The advantages of the approach are demonstrated both on the PETS 2000 image sequence (down-scaled 2:1) and on another, more complex traffic sequence recorded by a stationary camera at an innercity road intersection ('the intersection sequence'), see Figure 1. For this purpose, the results of two common computer vision tasks are analysed:

Model-based tracking: A model-based vehicle tracking approach is used, similar to the ones reported by, e. g., [2] or [4]. The contour of the vehicle model, projected into the image plane, encloses a coherent image segment. If more than one vehicle has to be tracked within one frame, one obtains multiple segments corresponding to the vehicle images, and one segment corresponding to the background. The center column of Figure 1 shows the projection of vehicle models overlaid onto the images.

Motion-based segmentation: Alternatively, we use a motion-based image segmentation as it has been proposed by [3]. Discontinuities in Optical Flow (OF) fields lead to a segmentation of the image. The right most column of Figure 1 illustrates the results of this approach using differently colored clusters.

No part of the approach will take advantage of the knowledge that these particular approaches are compared, i. e., any other approach leading to a pixel-wise segmentation of images might be used instead. We demonstrate the usability of the defined figure-of-merit by referring to these approaches since it enables us to closely relate significant changes of this criterion to observable conditions or failures of one of these approaches. Although both a tracking failure

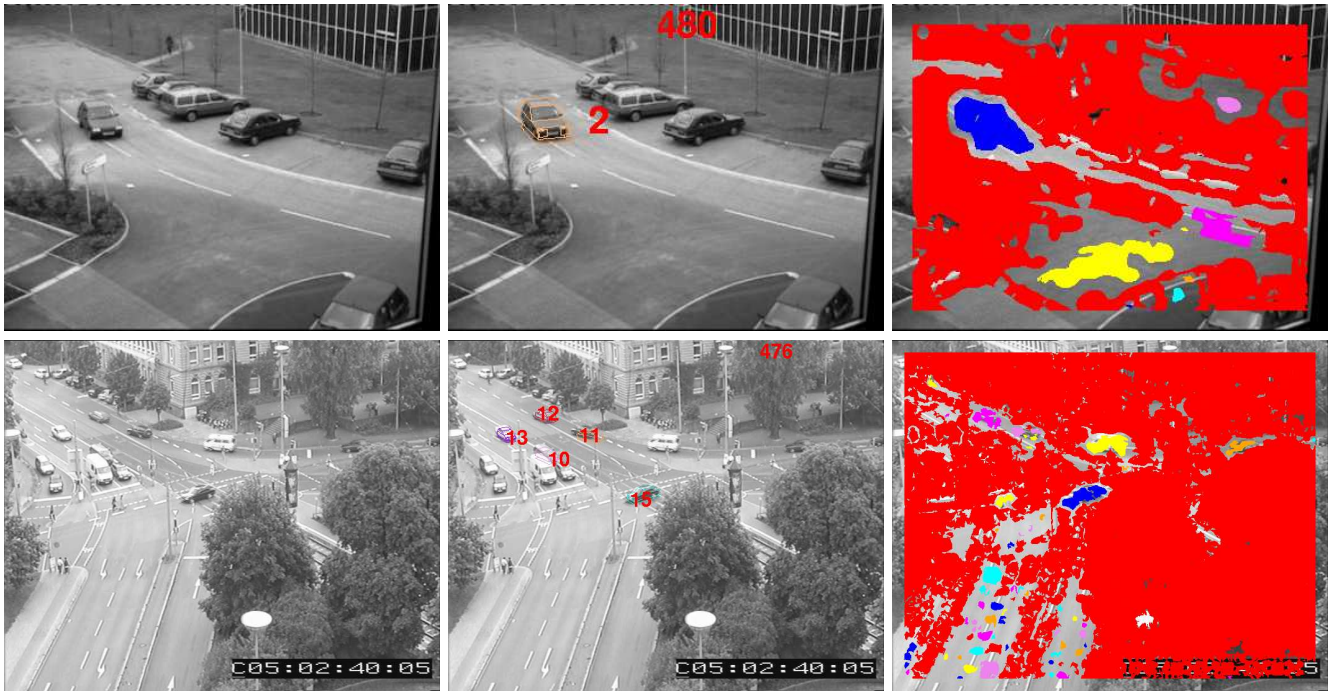

Figure 1: Representative frames from the down-scaled PETS 2000 sequence (frame #480, top) and from a more complex image sequence (smaller vehicle images, multiple objects) recorded at an innercity road intersection (frame #476, bottom): Original images (left), projection of 'fastback model' for tracked vehicles (center), and motion-based image segmentation (right). The border area of these image frames was not considered during motion-based segmentation, since gradient estimation, underlying the OF estimation, is not possible there.

and a segmentation error can be detected easily by human inspection, experience has shown that it is by no means trivial to detect them automatically. In the sequel, we thus study the question how well it may be possible to *detect algorithmically* whether or not either one of these approaches has failed.

3. A Consistency Test

Using the current state of a model-based vehicle tracking approach, we are in a position to compute the set of all pixels which are covered by the projection of a vehicle model into the image plane. In what follows, one has to distinguish between 'vehicles' and 'vehicle *candidates* (VC)'. The term 'vehicle candidate' refers to what the model-based tracking algorithm decides to be a vehicle, whereas 'vehicle' refers to a *real* vehicle. 'Vehicle *candidate* reference numbers (VCRN)' are automatically assigned to all VCs at the time when tracking is initialized.

Let '*Set of Vehicle Image Pixels* (SoVIP_i)' denote the set of all pixels which belong to the projection of the vehicle model assigned to VC #i. In addition, a '*Set of Optical Flow Segment Pixels* (SoOFSP_j)' will be defined for every image region j obtained by motion-based segmentation (OF-field segment). Let $|\cdot|$ denote the cardinality of an arbitrary set.

Using these definitions, the following figure-of-merit for consistency between an OF-field segment and the current estimate of a VC image obtained by model-based tracking can be defined:

$$\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j) = \frac{|\,\mathsf{SoVIP}_i \cap \mathsf{SoOFSP}_j\,|}{|\,\mathsf{SoVIP}_i \cup \mathsf{SoOFSP}_j\,|}\,. \quad (1)$$

The area corresponding to the *intersection* in the numerator of this ratio as well as the two underlying sets are illustrated in Figure 2. $\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j)$ will evaluate to zero when the intersection between SoVIP_i and SoOFSP_j vanishes. In the – unlikely – case of a perfect overlap between the VC image and an associated OF-field segment, the resulting figure-of-merit will be $\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j) = 1$.

In case an OF-field segment does not overlap the current VC image, we could assume that no relation exists between this OF-field segment and the projection of a vehicle model into the image plane. Experience has shown, however, that the projection of a vehicle model may gradually 'loose contact' with the corresponding vehicle image, for example due to intermediate (partial) occlusion of the vehicle by some opaque object in the scene like, e. g., a tree or a large board with road indicators. In such a case, $\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j)$

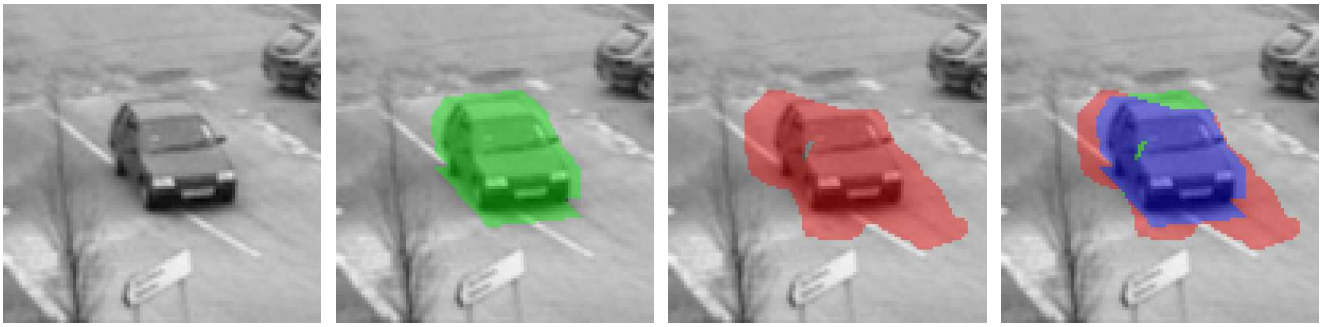

(a) Image of VC #2 in frame 480 (b) SoVIP$_2$ (c) SoOFSP$_3$ (d) SoVIP$_2$, SoOFSP$_3$ and their union

Figure 2: Panel (a) shows a window cropped frome frame 480 of the PETS 2000 image sequence illustrated in Figure 1 around the image of vehicle candidate (VC) #2. Panel (b) shows the same window, but overlayed by the projection (in *green*) of a fastback model positioned in the scene according to the current estimate of VC #2's position. Panel (c) illustrates the only OF-field segment (in *red*) that overlaps the projection of the vehicle model. Panel (d) is overlayed by the *union* of the projected fastback-model and the OF-field segment, with the *intersection* of these two overlays in *blue*. It can be seen that – in this (deliberately selected) example – both segments overlap surprisingly well.

should gradually decrease towards zero.

A value close to zero for $\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j)$ should be expected, too, when a vehicle slows down to a more or less complete stop. In such a case, the OF-vectors overlapping the vehicle image will exhibit very small values. As a consequence, they can no longer be separated reliably from the large OF-segment corresponding to the stationary parts of the depicted scene. In such a case, SoVIP will overlap – almost – completely the OF-segment corresponding to the stationary parts of the depicted scene. The cardinality of the intersection between SoOFSP and SoVIP thus will correspond roughly to the cardinality of SoVIP whereas the cardinality of the union of these two sets will be approximately equal to the much larger cardinality of the 'background segment'. Their ratio, therefore, will become very small.

4. Detailed Discussion of Examples

The considerations outlined in the preceding section have been used to generate test results based on the PETS 2000 image sequence.

4.1 PETS 2000 sequence, VC #2

Figure 3 illustrates the value of $\mathsf{consistency}_{\mathrm{OFS/VI}}(2,j)$ as a function of frame number, where j is the number of the 'best matching' motion-based image segment, selected automatically for every frame, i. e. $j = \arg\max(\mathsf{consistency}_{\mathrm{OFS/VI}}(2,j)$. For convenience, let

$$\mathsf{max\text{-}consistency}_{\mathrm{OFS/VI}}(i) = \max_j(\mathsf{consistency}_{\mathrm{OFS/VI}}(i,j)) \ . \tag{2}$$

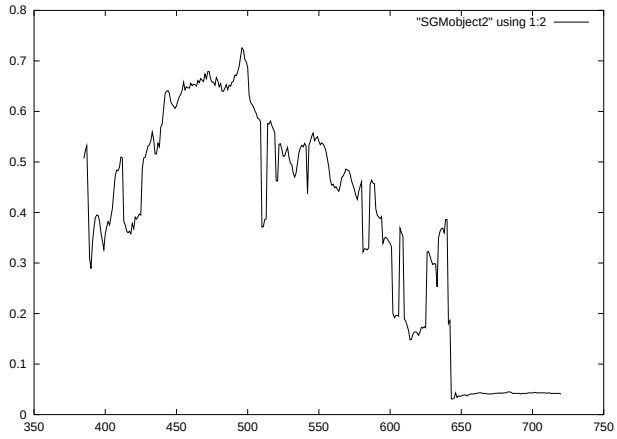

Figure 3: Plot of $\mathsf{max\text{-}consistency}_{\mathrm{OFS/VI}}(2)$ – see equ. (2) – as a function of the frame number for vehicle with VCRN #2 from the PETS 2000 sequence, i. e. the 'best matching' OF-field segment is selected automatically for each frame.

VCRN 2 is assigned to the vehicle that enters the field of view at about frame 370. $\mathsf{max\text{-}consistency}_{\mathrm{OFS/VI}}$ oscillates since time of the initialization of the model-based tracking process at frame 385 to about frame 430, and has quite low values. A close-up look at the image sequence (see Figure 4, left-most panel) reveals that during this time period, the vehicle image is partly occluded by slim branches, leading to an over-segmentation of the OF-field in *some* frames. This effect is further aggravated by the initial right-turn of the vehicle which leads to a changing appearence of the vehicle image. This violates the assumption underlying the OF-estimation, namely that grayvalues remain essentially

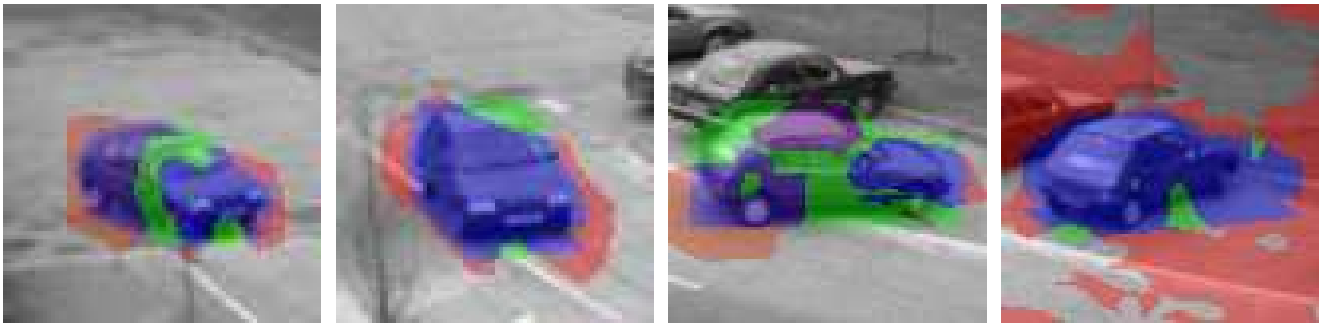

frame #400
2 segments

frame #470
1 segment

frame #602
4 segments

frame #660
1 segment

Figure 4: These clippings from frames of the PETS 2000 image sequence show *all* OF-field segments which overlap the projection of the vehicle model. Note that max-consistency$_{\mathrm{OFS/VI}}$ takes into account only the best matching segment.

constant.

From frame 430 to 500, segmentation results are satisfactory (Figure 4, 2nd panel), as is reflected by high values for max-consistency$_{\mathrm{OFS/VI}}$. After a short occlusion by another branch of the tree at about frame 510, the vehicle slows down and starts to turn left. This, again, causes an over-segmentation of the OF-field (3rd panel), leading to decreasing and oscillating values of max-consistency$_{\mathrm{OFS/VI}}$. When the vehicle comes to a stop, the motion-based segmentation naturally fails. This fact can be recognized by constantly low values of max-consistency$_{\mathrm{OFS/VI}}$ (an additional example is discussed in connection with Figure 7).

4.2 PETS 2000 sequence, VC #1 and #3

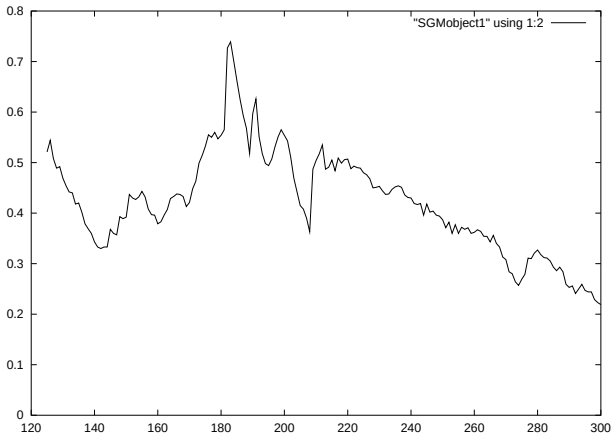

Figure 5: Plot of max-consistency$_{\mathrm{OFS/VI}}$ as a function of frame number for vehicle with VCRN #1 from the PETS 2000 sequence.

For VC #1 and #3, the values of max-consistency$_{\mathrm{OFS/VI}}$ (Figures 5 and 6) are much more stable than for VC #2. A

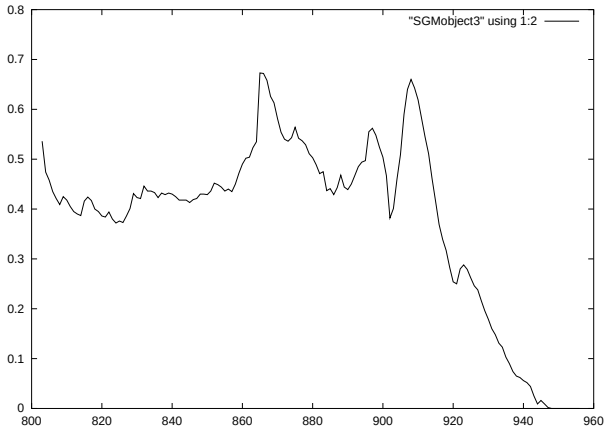

Figure 6: Plot of max-consistency$_{\mathrm{OFS/VI}}$ as a function of frame number for vehicle with VCRN #3 from the PETS 2000 sequence.

few observations, however, need to be discussed:

- The values of max-consistency$_{\mathrm{OFS/VI}}$ are only ≈ 0.5 on average: The motion-based segmentation algorithm 'separates' the image at motion discontinuities. Since the background, i. e. the image of the street, is more or less homogeneous, these discontinuities may be situated away from the 'real' vehicle image boundaries. The motion-based image segment which matches the projection of the vehicle model best, therefore, often includes many pixels from the background. Thus, even in the best case, SoVIP $\not\approx$ SoOFSP, but SoVIP $\subset$ SoOFSP and $|\mathrm{SoOFSP}| \gg |\mathrm{SoVIP}|$.

- The values of max-consistency$_{\mathrm{OFS/VI}}$ change significantly. This effect, too, is due to the structure of the background. The values of max-consistency$_{\mathrm{OFS/VI}}$ are higher if some grayvalue structure close to the vehi-

cle image limits the extent of the motion-based image segment, and they are lower if the background is homogeneous.

- For VC #1, the value of max-consistency$_{OFS/VI}$ slowly but steadily decreases after frame 220. This is due to the fact that the margin of pixels around the vehicle image which belongs to SoOFSP, but not to SoVIP, has about constant width, provided the background structure does not change significantly. When the size of the vehicle image – and thus |SoVIP| – decreases due to the perspective, the impact of these margin pixels increases, leading to decreasing values of max-consistency$_{OFS/VI}$.

4.3 Intersection sequence, VC #15

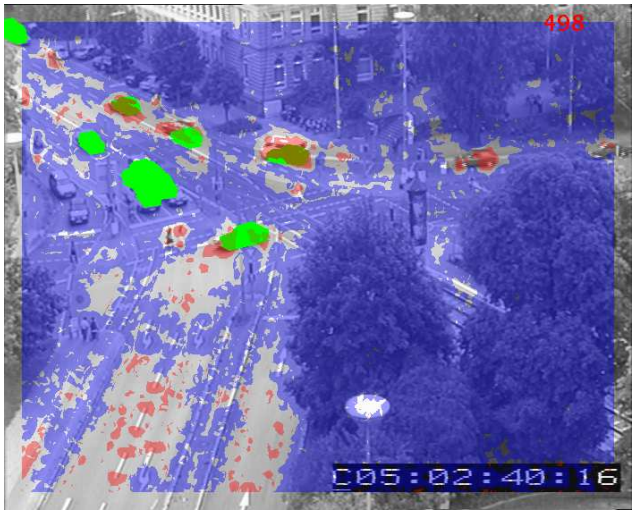

Figure 7: Frame 498 from the intersection image sequence: A fastback model (in *green*) is overlaid for each VC which has been initialised up to this frame. The *blue* overlay indicates the OF-segment corresponding to the background. The *red* overlays represent OF-field segments obtained by the motion-based segmentation.

In contradistinction to the PETS sequence – where only a single vehicle moves at any one time – the intersection sequence comprises several moving vehicles with much smaller images.

Frame 498 from this sequence has been reproduced in Figure 7, with all pixels overlayed in *green* which correspond to projections of vehicle candidates initialized up to this frame, i. e. all pixels in SoVIP for any initialised vehicle candidate. The overlay in *blue* here represents the largest coherent OF-segment which in this case corresponds to a considerable part of the the stationary background. Image areas without any overlay within the interior of Figure 7 are due either to discontinuities or to regions where the grayvalue variation is too small to estimate an OF-vector and, moreover, *this fact has been determined algorithmically*. The (smaller) overlays in *red* indicate additional OF-segments. These segments comprise segmented vehicle images and 'noise segments'. The latter are due to the fact that '(nearly) homogeneous' image areas may isolate smaller OF-field segments corresponding to the background from the large one.

In Figure 8, max-consistency$_{OFS/VI}$ has been plotted as a function of the frame number for the vehicle (more precisely, for VC #15) driving from right to left close to the center in Figure 7. When this VC had been ini-

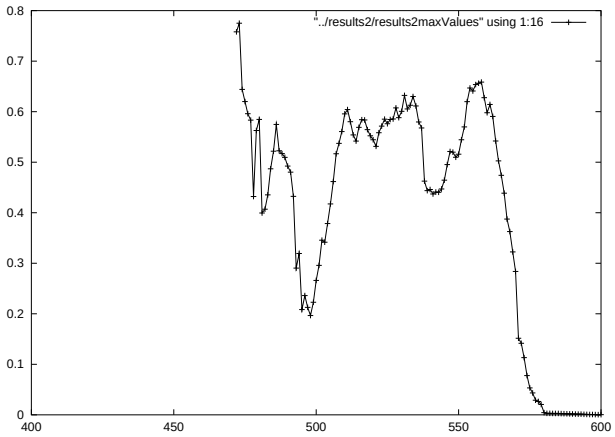

Figure 8: Plot of max-consistency$_{OFS/VI}$ as a function of the frame number for a vehicle with VCRN 15 from the intersection image sequence.

tialised for model-based tracking around frame 470, the OF-segment and the projected vehicle candidate image overlapped very well, resulting in a value around 0.75 or more for max-consistency$_{OFS/VI}$. Some phenomena like oversegmentation and the extension of OF-field segments well beyond the vehicle image proper have been discussed already in Subsections 4.1 and 4.2.

Starting around framenumber 490, a new phenomenon (see Figure 9) begins to dominate this ratio, namely the motion discontinuities related to the stationary mast. The trailing OF-segment decreases in size with increasing framenumber. At framenumber 498, the occluding mast has cut the hypothetical single OF-segment corresponding to this entire vehicle into two about equally small parts as illustrated by Figure 9.

When the vehicle continues to pass behind the mast, its visible leading part will become greater and the algorithm selects the OF-segment corresponding to this leading part for the determination of max-consistency$_{OFS/VI}$. This ratio will thus begin to rise again as can be seen in Figure 8 until

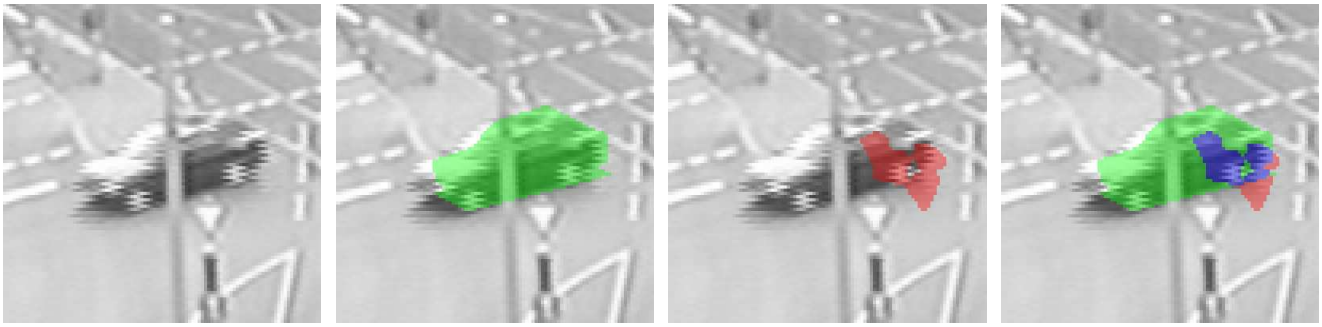

(a) Image of VC #15 in frame 498 (b) SoVIP (c) SoOFSP (d) SoVIP, SoOFSP and their union

Figure 9: Panel (a) shows a window cropped frome frame 498 of the intersection image sequence illustrated in Figure 1 around the vehicle corresponding to VC #15 when it is partially occluded by a mast. Panel (b) shows the same window, but overlayed by the projection of a fastback model positioned in the scene according to the current estimate of the VC position. Panel (c) illustrates the OF-field segment obtained by the motion-based segmentation approach. Panel (d) is overlayed by the *union* of the projected fastback-model and the OF-field segment, with the *intersection* of these two overlays colored *blue*.

it reaches a local maximum. From frame 557 on, the vehicle will be occluded increasingly by the foliage of trees. Eventually, the OF-segment corresponding to the vehicle will vanish, resulting in a zero value for $\text{max-consistency}_{\text{OFS/VI}}$.

5. Additional Experiments

Extensive experiments are summarized in Table 1. The model-based tracking approach used only a single vehicle type (a 'fastback'). The results refer to all vehicles of this type which appear within the intersection image sequence (and those that look like 'fastbacks' in at least a few frames). Column A ('Good results') assumes consistency between the VC's image and the largest overlapping OF-field segment if $\text{max-consistency}_{\text{OFS/VI}}$ exceeds 0.33 for a *number of consecutive* frames. The bottom line of the sixth column from Table 1 indicates that this case has been found only in about $12\ \%$ of all potential tracking steps.

Columns B through D show the results of an *interactive* assignment to the most common reasons for observed inconsistency values, whereas column E summarizes the remaining cases.

- In about $29\ \%$ of all potential tracking steps, the vehicle velocity estimated by model-based tracking is so small that one cannot expect to obtain large values for $\text{max-consistency}_{\text{OFS/VI}}$ (column B). These cases can be detected algorithmically by considering the estimated velocity.

- Another category comprises those cases where the vehicle is about to stop or start, i. e. its velocity is small and changes significantly, and – as a consequence – the OF-segmentation is unreliable (column C). This case, too, can be identified algorithmically.

- A fourth category is constituted by those cases where the vehicle velocity is estimated to be large enough in order to justify the expectation that $\text{max-consistency}_{\text{OFS/VI}}$ will assume a sizable value. OF-field segmentation problems, however, prevent the generation of a sufficiently large OF-segment (column D).

- The remaining cases include full or partial occlusions or the period where the vehicle leaves the field of view and is only partially visible.

6. Potential Improvements

It can be observed frequently that the OF-field segment which best[1] matches to the projected vehicle model switches from one relative position with respect to the vehicle image to another one. In many cases, this is caused by oversegmentation of the estimated OF-field which in turn is preferred in comparison with an undersegmentation of the OF-field. The value obtained for $\text{max-consistency}_{\text{OFS/VI}}$ in such a case naturally is much smaller and may give an inappropriate cue regarding the quality of 3D-model-based tracking.

One could consider, therefore, to take *not only the best matching* OF-field segment into account, but in addition the *second best matching, too*. In principle, this opens a Pandora's box of combinations. To prevent unwanted effects, it may be advisable to restrict these possibilities. The idea consists in supplementing the consistency test based on equ. (1) by a second, '*weaker*' one. It could look like the following:

[1] See equ. (2) in Subsection 4.1

VCRN	From Frame	To Frame	Terminated due to	# of frames	Categorisation of tracking result: Good results A	Vehicle stands B	Stopping/ Starting C	Segm. problems D	Other E
10	97	1300	Failure	1204	0	950	230	24	0
11	212	720	LFV (F)*	509	90	0	0	300	119
12	302	645	LFV*	344	160	0	0	184	0
13	311	2000	EIS	1690	98	945	180	467	0
15	472	580	LFV	109	90	0	0	0	19
17	525	2000	EIS	1476	25	850	150	451	0
18	590	870	LFV*	281	150	0	0	131	0
19	683	943	LFV*	261	166	0	0	95	0
21	876	1040	LFV (F)*	165	0	0	0	60	105
22	889	1220	Failure	332	70	0	0	70	192
23	977	1200	Failure	224	30	0	194	0	0
25	1195	1400	Failure	206	16	0	0	115	75
27	1277	1800	LFV*	523	58	0	0	455	10
29	1328	1430	LFV*	103	23	0	15	45	20
30	1362	1491	Failure	130	0	0	0	79	51
31	1384	1900	LFV	517	45	0	20	452	0
33	1440	1910	LFV*	471	40	0	0	431	0
35	1528	1580	LFV*	53	4	0	0	24	25
36	1546	1980	LFV*	435	10	0	0	400	25
43	1674	2000	EIS	327	30	0	0	277	20
46	1753	2000	EIS	248	20	0	0	208	20
				9608	1125	2745	789	4268	681
				100%	11,7%	28,6%	8,2%	44,4%	7,1%
without standing vehicles, 100%=6863 frames				–	16.4%	–	11.5%	62.2%	9.9%

Table 1: Explanation of abbreviations: VCRN = Vehicle *Candidate* Reference Number. LFV = Vehicle leaving field of view. EIS = End of image sequence. The entry 'LFV (F)' in column four indicates poor model-based tracking results that might also justify to consider the tracking to be a failure before the vehicle leaves the field of view. The fifth column '# of frames' shows the total number of frames for which the vehicle candidate identified by the VCRN in the first column has been tracked by the model-based approach under specified conditions *and* for which the corresponding vehicle is at least partially visible. It thus represents the maximum number of (potentially successful) tracking steps. The asterisk in column 4 indicates cases where the *vehicle* leaves the field of view, but where the *vehicle candidate* is still available for a number of frames. Columns A to F are explained in the text.

Let S_i, $i = 1, \ldots, n$ be the OF-field segments that overlap the projection of a vehicle candidate, where S_1 is the 'best matching' segment (according to $\mathsf{consistency}_{\mathrm{OFS/VI}}$) and S_n the 'worst matching' (but still overlapping) segment. Furthermore, let $M_k = \cup_{i=1\ldots k} S_i$, i. e. the union of the k best matching segments. The basic principle of the 'weak consistency test' is to compute $\mathsf{consistency}_{\mathrm{OFS/VI}}$ using $M_k, k \neq 1$ instead of M_1. This restricts the 2^n possible subsets of $S_1, \ldots, S_n$ to only n sets of segments that are taken into account.

It is expected that the detrimental effect of oscillating $\mathsf{consistency}_{\mathrm{OFS/VI}}$-values can be reduced without the introduction of too many new problems. It possibly even may allow to raise the threshold from the currently used value of 0.33 and thereby contribute to a higher discriminative power of the $\mathsf{consistency}_{\mathrm{OFS/VI}}$-test.

Figure 10 plots these new values for $\mathsf{consistency}_{\mathrm{OFS/VI}}$, taking into account the best matching OF field segment (solid), both the best and the second best matching segment (dashed), and a varying number of segments, whose union M_k gives the best match (dotted). From this plot, one can conclude:

- Usage of more than a single segment reduces the oscillations, but it may not prevent small values of $\mathsf{consistency}_{\mathrm{OFS/VI}}$ in all frames (note, e. g., frame 1270).
- The segments have to be chosen carefully, otherwise the results may become worse than before (frames 1300 ff.).
- In some frames, using more than two segments may be advisable.

The true problem of this approach is that one either has

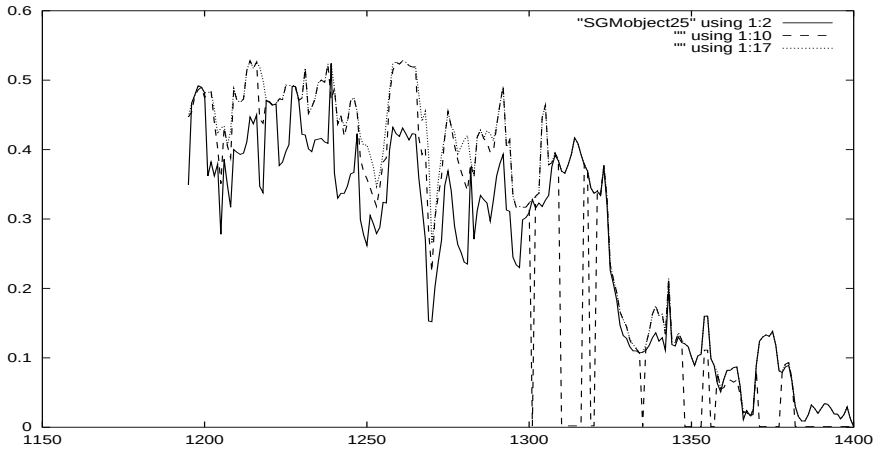

Figure 10: Plot of consistency$_{\mathrm{OFS/VI}}$ as a function of the frame number for a vehicle of type 'fastback' with vehicle *candidate* reference number **25** from the intersection image sequence, taking into account the best matching OF field segment (solid), both the best and the second best matching segment (dashed), and a varying number of segments, whose union M_k gives the best match (dotted).

to fix a number of segments in advance (and as one can see, already 'two' may be too much at some frame-times) or the algorithm uses an a-posteriori rule to decide which segments are taken into account. Furthermore, it does not improve results in case the segmentation is basically wrong. Finally, the considered unions of segments might not include the optimal solution, i. e. the union of, e. g., the second and third best matching segment (which is not considered in the algorithm sketched in this section) might give better results than any other set of segments.

7. Summary and Conclusions

The ratio of the cardinalities given by equ. 1 has been introduced as a criterion for quantification of compatibility between the estimated vehicle image and the estimated OF-field segment for the same vehicle. It turned out that the variation of this ratio as a function of the frame number could be associated with intuitively plausible observations for these two estimates. Numerous experiments – some of which have been discussed in the preceding sections – corroborated initial impressions.

It thus begins to appear justifiable to incorporate a plausibility test into a model-based tracking system in order to obtain an 'algorithmic handle' for the detection of tracking problems. OF-field segments usually result from a spatiotemporally *localized* computation. It thus is influenced by other aspects of the input image sequences than results of a 3D-model-based approach. The latter relies on a single initialisation and subsequently on a highly nonlinear forward-propagation of information. These differences with respect to sources and types of errors support the hypothesis that consistency between results obtained by such different approaches is a good cue that both processes are still 'on track'. Stimulated by the positive experience with initial experiments of the kind reported here, an analogous ratio has been investigated as a compatibility cue between two *tracking* results – one obtained by a model-based tracking approach similar to [2] and another one by a data-driven tracking approach based on [1]. For a sequence of limited duration which was available for tests, these compatibility cues agreed with intuition.

Acknowledgments

We thank R. Deriche and M. Rousson for stimulating discussions.

References

[1] T. Brox, M. Rousson, R. Deriche, and J. Weickert: *Unsupervised Segmentation Incorporating Color, Texture, and Motion.* In: Proc. 10th International Conference on Computer Analysis of Images and Patterns CAIP'03, 25-27 August 2003, Groningen, The Netherlands. LNCS, Springer-Verlag: Berlin·Heidelberg·New York/NY 2003.

[2] M. Haag and H.-H. Nagel: *Combination of Edge Element and Optical Flow Estimates for 3D-Model-Based Vehicle Tracking in Traffic Image Sequences*. International Journal of Computer Vision **35**:3 (1999) 295-319.

[3] M. Middendorf: *Untersuchungen zur Bestimmung des Optischen Flusses*. Doctoral Dissertation, 21 July 2003, Fakultät für Informatik der Universität Karlsruhe (TH), (in German).

[4] A. Pece and A. Worrall: *Tracking without Feature Detection.* In: Proceedings First IEEE Workshop on Performance Evaluation of Tracking and Surveillance (PETS 2000), March 31, 2000, Grenoble, France, pp. 29–37, 2000.