Vanishing Point Detection without Any A Priori Information

Andrés Almansa, Agnès Desolneux, Sébastien Vamech

► To cite this version:

Andrés Almansa, Agnès Desolneux, Sébastien Vamech. Vanishing Point Detection without Any A Priori Information. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2003, 25 (4), pp.502-507. 10.1109/TPAMI.2003.1190575 . hal-00170785

HAL Id: hal-00170785

<https://hal.science/hal-00170785>

Submitted on 1 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Vanishing point detection without any a priori information

Andrés Almansa, Agnès Desolneux, Sébastien Vamech

CMLA - ENS Cachan, 61 avenue du président Wilson

94235 Cachan cedex, France

{almansa,desolneux}@cmla.ens-cachan.fr, vamech@dptmaths.ens-cachan.fr

This work was partially funded by the French Space Agency (CNES) and Cognitech Inc. We would like to specially thank Jean-Michel Morel for his encouragement, support, ideas, and for introducing us to Gestalt-based image analysis. A. Almansa would also like to thank Lenny Rudin and Pascal Monasse for useful discussions, references, and motivation to address this problem, as well as for adapting and optimizing this method for Cognitech's video processing software. Finally we gratefully acknowledge the very pertinent observations of the reviewers which greatly contributed to improve the quality of this paper.

Abstract

Even though vanishing points in digital images result from parallel lines in the 3D scene, most of the proposed detection algorithms are forced to rely heavily either on additional properties (like orthogonality or coplanarity and equal distance) of the underlying 3D lines, or on knowledge of the camera calibration parameters, in order to avoid spurious responses. In this work we develop a new detection algorithm that relies on the Helmholtz principle recently proposed for computer vision by Desolneux, Moisan and Morel [6], [7], both at the line detection and line grouping stages. This leads to a vanishing point detector with a low false alarms rate and a high precision level, which doesn't rely on any a priori information on the image or calibration parameters, and doesn't require any parameter tuning.

Index terms—*Vanishing point, perceptual grouping, Gestalt theory, Helmholtz principle*

I. INTRODUCTION

Sets of parallel lines in 3D space are projected into a 2D image obtained with a pin-hole camera to a set of concurrent lines. The meeting point of these lines in the image plane, is called a *vanishing point*, and may eventually belong to the line at infinity of the image plane in the case of 3D lines parallel to the image plane. Even though concurrence in the image plane doesn't necessarily imply parallelism in 3D (it only implies that all 3D lines intersect the line defined by the focal point and the vanishing point), the counterexamples for this implication are extremely rare in real images, and the problem of finding parallel lines in 3D is reduced to finding vanishing points in the image plane.

The usefulness of precise measurements of vanishing points, among other geometric primitives, was demonstrated for instance in [4], in the framework of forensic applications of single view metrology. Since the seminal work of Barnard [13], however, automated computational methods for vanishing points detection in digital images have been based on some variation of the Hough transform in a conveniently quantized Gaussian sphere. Several refinements of these techniques followed, but most recent works suggest that this simple technique often leads to spurious vanish-

ing points[12]. In order to eliminate these false alarms most authors considered some kind of joint Gestalt, which adds some other property to 3D parallelism like coplanarity and equal distance between lines [11] or orthogonality between the three main 3D directions [8], [12], [9]. In addition, knowledge of the intrinsic camera calibration parameters is commonly assumed [8], [3] by these methods, or they are designed mostly for omnidirectional images [3]. To the best of our knowledge, the question of reliably determining whether an image *actually contains* some vanishing points and its *number* has not yet been addressed systematically.

In this work we show that 3D parallelism alone is a significant Gestalt in many images of man-made environments and that it can be reliably detected with a low number of false alarms and a high precision level, without using any secondary property, or any a priori information on the image or calibration parameters, and without any parameter-tuning. We do not claim that secondary properties (like equal distance, or orthogonality) should not be used in any circumstance; this can be useful for some applications, and our technique may be extended to these situations. But in many applications a pure vanishing point detector is more useful, since it can be used to determine some calibration parameters of the camera (which are needed in other approaches relying on orthogonality for instance). The key improvements with respect to previous vanishing point detectors are the following: *(i)* The primitives that are accumulated in (an equivalent of) the Gaussian sphere are line segments, which are themselves detected with an almost-zero false alarms rate, by a refinement of the method presented in [5]. *(ii)* Our criterion to determine a meaningful vanishing point from a large vote in the Gaussian sphere is deduced from the Helmholtz principle [7], thus producing a low number of false alarms, without need for threshold-tuning. *(iii)* Finally a Minimum Description Length (MDL) criterion is used to further restrict the number of spurious vanishing points and to deal with the masking phenomenon.

This article should be considered as a more extended application of the general technique presented in [7] (in this same special issue), where the Helmholtz principle is explained in more detail. The refinements that are needed for the alignment detection algorithm are described in [2] and in [1, chapter 4]. These changes aim mainly at eliminating multiple responses for a single line segment, and at eliminating the precision parameter to obtain a fully parameterless method.

II. DETECTION OF VANISHING POINTS

As in the case of alignments we shall define a meaningful vanishing point in terms of the Helmholtz principle. Our objects in this case will be all the meaningful segments obtained by the method we described in the previous section. The common feature we shall seek for among these segments is a common point v_∞ met by all their supporting lines. Due to measurement errors, these lines will rather meet within a more or less small subset V of the image plane, which we shall call *vanishing region*. To consider all possibilities we need to consider a *finite* family of such regions $\{V_j\}_{j=1}^M$, such that it covers the whole (*infinite*) image plane, *i.e.* $\bigcup_{j=1}^M V_j = \mathbb{P}_2$. In [14], an intelligent such partition is proposed whereas most works use a partition of the image plane such that the projection of each vanishing region on the Gaussian sphere has a quasi-constant area [13], [8], [12], [9]. This partition has the advantage that it assigns the same precision to all 3D orientations, but it requires knowledge of the internal camera calibration parameters. However a practical application of the Helmholtz principle (Sections II-A and II-B) leads to a different partition of the image plane into vanishing regions (to be introduced in Section II-C), which shares some qualitative properties with the common Gaussian sphere partition.

A. Meaningful vanishing regions

Assume that a total of N segments were detected, with supporting lines $l_1, l_2, \dots, l_N$. We consider the event that at least k among these N lines meets a given vanishing region V_j . Under the assumption that all lines are independent with the same distribution, the probability of such an event is $B(p_j, N, k)$, where p_j is the probability that a line meets the vanishing region V_j . Moreover, since the M regions V_j are chosen to sample all possible vanishing regions, we make $N_T = M$ such tests. Thus the number of false alarms for a vanishing region V_j can be defined as:

$$\text{NFA}(V_j) := MB(p_j, N, k), \quad (1)$$

and as usual, the vanishing region is ε -meaningful if k is sufficiently large to have $\text{NFA}(V_j) \leq \varepsilon$.

In order to actually find the value of NFA and the minimal value $k(j, \varepsilon)$ of k such that V_j becomes meaningful, we need to know the probabilities p_j . This is the subject of the next section.

B. Probability of a line meeting a vanishing region.

Gratefully this geometric probability problem has been very elegantly solved in [10]: Choose a suitable measure for a random line G on the plane (it is shown that there is only one way to do this, up to a multiplicative constant, in a translation- and rotationally invariant way), and consider two convex sets K_1 and K_2 of the plane. Then the measure of all lines meeting both sets is:

$$\mu[G \cap K_1 \neq \emptyset \text{ and } G \cap K_2 \neq \emptyset] = \begin{cases} L_2 = \text{Per}(K_2) & \text{if } K_1 \subseteq K_2 \\ L_i - L_e & \text{if } K_1 \cap K_2 = \emptyset \end{cases} \quad (2)$$

where the *external perimeter* L_e is the perimeter of the convex hull of K_1 and K_2 and the *internal perimeter* L_i is the length of the “internal envelope” of both sets, which is composed of the internal bitangents to K_1 and K_2 and parts of their perimeters.¹

¹The proof of this result can be found in [10] and a sketch of the proof in [2] or [1]

This result can be directly applied to our problem of determining p_j in the case where the vanishing region $V_j \subseteq \Omega$ is contained in the (convex) image domain Ω . Since we can only observe line segments that intersect the image domain, the probability we are interested in is actually:

$$p_j = P[G \cap V_j \neq \emptyset \mid G \cap \Omega \neq \emptyset] = \frac{\mu[G \cap V_j \neq \emptyset \text{ and } G \cap \Omega \neq \emptyset]}{\mu[G \cap \Omega \neq \emptyset]} = \frac{\text{Per}(V_j)}{\text{Per}(\Omega)}. \quad (3)$$

For vanishing regions $V_j \cap \Omega = \emptyset$ external to the image domain, we just apply the second case of equation (2) and the probability becomes

$$p_j = \frac{L_i - L_e}{\text{Per}(\Omega)}. \quad (4)$$

The intermediate case where there is an intersection but no inclusion is treated as this second case with $L_i = \text{Per}(K_1) + \text{Per}(K_2)$.

C. Partition of the image plane into vanishing regions

In this section we address the problem of choosing a convenient partition of the image plane into vanishing regions. For this purpose we use the following criteria:

Equal probability. We try to build a partition such that the probability $p_j = P[G \cap V_j \neq \emptyset]$ that a random line G of the image meets a vanishing region V_j is constant for all regions. Without this equiprobability condition, certain vanishing regions would require many more meeting lines to become meaningful than others, *i.e.* they would not be *equally detectable*, which is not desirable.

² We can easily deduce from the results of the previous section that this equiprobability condition implies that the size of V_j increases dramatically with its distance from the image, which agrees

²For instance the partition into regions whose projection into the Gaussian sphere has constant area, does not necessarily satisfy this equal probability condition. This was observed by [8] in the case of uniformly distributed 3D lines. In this case, lines almost parallel to the image plane become much less probable than lines which are almost orthogonal. Despite the correction proposed in [8], this still leads to problems in the detection of vanishing points when the perspective effect is very low (distant vanishing points, or lines almost parallel to the image plane), as observed by [12].

with the fact that the *localization error* of a vanishing point increases with its distance from the image. Thus, with the equiprobability condition, we obtain the localization error of the vanishing points as a consequence of their detectability.

Angular precision. The size and shape of the vanishing regions should be in accordance with the angular precision of the detected line segments. In [12] the author addresses this problem by considering a localization error of 1 pixel at the ends of the segment, so the precision of the segment's orientation is $d\theta = \arctan \frac{1}{l}$, where l is the length of the segment. The supporting line of the segment should be rather considered as a “cone” with angle $d\theta$. When such a cone intersects a vanishing region the corresponding accumulator is updated by a value proportional to the angular fraction of the cone covered by the vanishing region. This fraction becomes 1 if the vanishing region is larger than the width of the cone. Here we can use Shufelt's concept, or a similar one whose only difference is that we threshold by not considering fractional intersections. If V_j covers the whole width of the uncertainty cone of the segment, then we consider that there is an intersection. Otherwise, the intersection is uncertain, and we say that the segment does not meet the vanishing region. This leads us to construct the vanishing regions in such a way that their size is comparable to the width of the corresponding vanishing cones.

Now we shall construct a partition of the plane into vanishing regions that closely satisfies both criteria above. The partition is composed of two families of vanishing regions. The first (“interior”) one consists of regions entirely contained in the image domain Ω , and the second (“exterior”) one consists of regions outside the image domain.

For simplicity, we shall approximate the image domain by its circumscribed circle, and consider the image domain Ω as a circle of radius $R = N/\sqrt{2}$. In order to meet the angular precision requirement, all exterior regions V will be portions of sectors of angle $d\theta$ lying between distances

d and d' from the image center O . Then, the probability $p_e(d, d') = \frac{L_i - L_e}{\text{Per}(\Omega)}$ that a random line meeting the image domain Ω does also meet V can be expressed after a simple trigonometric calculation as a function of d , d' and $d\theta$.

Concerning the interior regions, we chose a simple tiling of the circle Ω with square tiles. The side of each square is chosen to be equal to the side of the exterior tiles closest to the image domain, *i.e.* $2R \sin(d\theta)$. The perimeter of the interior regions is therefore equal to $8R \sin(d\theta)$, and the probability that a line meets an interior vanishing region is:

$$p_i = \frac{\text{Per}(V)}{\text{Per}(\Omega)} = \frac{4 \sin(d\theta)}{\pi} \quad (5)$$

This ensures that all interior regions have the same probability, and that their size is in accordance with the coarsest angular precision $d\theta$ of the line segments. Now we have to choose the values of d and d' to ensure that all exterior regions have the same probability $p_e = p_i$. To do so, we start with the first ring of exterior regions setting $d_1 = R$, and we choose d'_1 by solving the equation $p_e(d_1, d'_1) = p_i$ for d'_1 . Then we fill the second ring of exterior tiles by setting $d_2 = d'_1$ and solving the equation $p_e(d_2, d'_2) = p_i$ for d'_2 . We iterate this process until we get $d' > d_\infty$, where d_∞ is such that $\lim_{d' \rightarrow \infty} p_e(d_\infty, d') = p_i$. We can easily check that d_∞ is finite and satisfies:

$$4 \sin(d\theta) = 2d\theta + \frac{\pi}{2} - \beta - \frac{1}{\cos \beta} + \tan \beta \quad \text{where} \quad \beta = \arccos \left(\frac{R \cos(d\theta)}{d_\infty} \right) \quad (6)$$

Regions in the last ring will then be unbounded, with probability $\leq p_i$. They represent parallel lines in the image plane. Figures 1 and 2 show some examples of this partition of the image plane for different precision levels $d\theta$.

D. Final remarks

In this section we introduce some additional criteria to suppress spurious vanishing points and to eliminate the angular precision parameter $d\theta$.

Multi-precision analysis. The choice of a fixed value for the angular precision parameter $d\theta$ requires a compromise between detectability and localization error of vanishing points. We are interested in the highest possible precision level (smaller localization error in the vanishing point). On the other hand if the precision level is too fine with respect to the angular precision of the segments, the vanishing region will be hardly detected. The optimal level will approximately match the precision of the segments converging to this vanishing point, and our strategy will be to try to adjust the precision level automatically to this value. From simple calculations on the definition of the NFA, we observe that for a total $N = 1000$ lines, we need about 300 concurrent lines to be meaningful at precision $d\theta = \frac{\pi}{16}$, whereas only 15 concurrent lines are enough at precision $d\theta = \frac{\pi}{1024}$. But we would only need 7 concurrent lines, if the total number of lines was $N = 100$. This discussion motivates the procedure described below.

As in the case of alignments, instead of fixing a single angular precision level, we will consider multiple dyadic precision levels $d\theta = 2^{-s}\pi$ for n different values of s in a certain range $[s_1, s_n]$. In our experiments $s = 4, 5, \dots, 7$ showed to be the most useful range, but this can be adjusted to the range of precision levels of the extracted segments. According to the discussion above, at each precision level $d\theta$ we should only keep those segments with a precision level no coarser than $d\theta$. Coarser segments would significantly increase N_s (thus increasing the detection threshold k) without significantly increasing the number k of lines meeting the vanishing region. Now we can apply the previously described method for all precision levels. This procedure, however, may multiply the expected number of false alarms by a factor no larger than n . So, in order to keep the false alarms rate smaller than ε we modify equation (1) as follows:

$$\text{NFA}(V_{j,s}) = \frac{M_s}{n} B(p_s, N_s, k). \quad (7)$$

The vanishing region is considered ε -meaningful if k is large enough to obtain $\text{NFA}(V_{j,s}) \leq \varepsilon$.

With this definition the total expected number of false alarms from this multi-precision analysis can be easily shown to be no larger than ε . The problem is that a single vanishing point may be meaningful at several different precision levels, and we only want to keep the best explanation for it.

Local maximization of meaningfulness. When a huge number of segments meet a vanishing region $V_{j,s}$ they also meet some of the neighboring regions at the same precision level s , as well as all coarser regions $V_{j,s'} \supseteq V_{j,s}$ and some finer regions $V_{j,s''} \subseteq V_{j,s}$. Therefore, these neighboring regions too, are likely to become meaningful, but are not necessarily the best explanation. To choose among them the best explanation we introduce the following maximality concept: A vanishing region $V_{j,s}$ from a multi-precision family of partitions of the image plane is *maximal* if it is more meaningful than any other region intersecting it. More precisely, $V_{j,s}$ is maximal if:

$$\forall s' \in [s_1, s_n], \forall j' \in \{1, \dots, M_{s'}\}, \overline{V_{j',s'}} \cap \overline{V_{j,s}} \neq \emptyset \Rightarrow \text{NFA}(V_{j',s'}) \geq \text{NFA}(V_{j,s}) \quad (8)$$

where $\overline{A}$ denotes the closure of a set A . Note that the condition $\overline{V_{j',s'}} \cap \overline{V_{j,s}} \neq \emptyset$ includes both neighboring regions at the same level, as well as coarser regions containing $V_{j,s}$ and finer regions contained in it.³

Minimum Description Length. Figure 1 shows all the maximal 1-meaningful vanishing regions that are detected for the photograph of a building. Clearly the first three correspond to real orientation in the 3D scene, whereas the other three are an artificial mixture of different orientations. Observe that these mixtures are less meaningful than the original ones, because only a small portion of the segments in each direction can participate. Therefore, these artificial vanishing regions can be filtered out by an MDL criterion similar to the one we used for segments. Among all maxi-

³We used this condition instead of inclusion, because the equal probability constraint that we used to construct our partition means that regions at precision level $s + 1$ cannot always be completely included in a single region at the coarsest precision level s . In this situation this “non-empty intersection”-type condition is better suited.

mal meaningful vanishing regions we start a competition between them, based on the principle that each segment has to choose a single vanishing region which best explains its orientation. More precisely, a segment with supporting line l is assigned to the vanishing region $V_{j,s}$ such that $\text{NFA}(V_{j,s})$ is smallest among all regions $V_{j,s}$ met by l . Then we recompute $\text{NFA}(V_{j,s})$ for all meaningful segments using equation (7), with the only modification that instead of k we consider $k' \leq k$ which is the number of lines that do not only meet $V_{j,s}$, *but also have been assigned to the vanishing region* $V_{j,s}$. If the number of false alarms is still smaller than ε , then the vanishing region is a *maximal MDL meaningful*.

E. Algorithm

In order to avoid mutual exclusions, the MDL criterion is run iteratively. In the first iteration the V_{j_1} with lowest NFA is selected as MDL meaningful. Then k (and the corresponding NFA) is updated for the remaining meaningful V_j 's by discounting all segments meeting V_{j_1} . Thus the NFA can only increase and the number of meaningful regions decreases. In the i -th iteration the i -th meaningful region V_{j_i} with lowest NFA is selected as MDL meaningful, and the remaining meaningful regions are updated by discounting from k (and the corresponding NFA) all the segments meeting V_{j_i} . The iteration stops when $\text{NFA}(V) > \varepsilon$ for all remaining regions V .

Sometimes this procedure will still miss some weak vanishing points which are “masked” by stronger vanishing points composed of much more segments. These may not be perceived at first sight, but only if we manage to unmask it by getting rid of the “clutter” in one way or another. For instance we may focus our attention into the corresponding region, or we can hide the stronger features. This unmasking mechanism can be simulated by zooming into a certain region of interest as illustrated in Figure 2(d), or by continuing our MDL iteration as follows: When no more meaningful MDL vanishing regions exist, remove all line segments that meet the already detected

vanishing points $V_{j_1}, \dots, V_{j_i}$. Thus the total number N of segments will decrease and so will do $\text{NFA}(V_j) = MB(p, N, k)$ in equation (1). Thus some vanishing points may become meaningful again and we can restart the previous iteration. This iteration allows to distinguish a first group of features that are meaningful on its own right, from a second group which is only detected in the absence of the first group's masking effect.

The complexity of the algorithm we just described for N lines and M tested vanishing regions is $o(N\sqrt{M})$ for computing the line-cell intersections, plus the MDL which takes about $o(NM')$ per iteration where $M' \ll M$ is the number of meaningful regions that have been detected at that iteration. The number M is fixed and equal to 270360, 67828, 17195, when the finest angular precision level is respectively $\frac{\pi}{1024}$, $\frac{\pi}{512}$, $\frac{\pi}{256}$. An optimized version of this software was reported to run on a 1GHz Pentium processor in 0.16 seconds for $N = 64$, up to 1.88s for $N = 1024$. In an image of size 512×512 we usually detect a few hundred alignments, so the running time of the vanishing point detector part is negligible with respect to the alignment detection module which does the bulk of the work. The complexity of the latter module is $o(N^4)$ for an image of size N , and the running time for $N = 512$ is about 40 seconds, also on a 1GHz Pentium processor.

III. EXPERIMENTS AND DISCUSSION

Figures 1 and 2 show the results of applying our algorithm for vanishing point detection on several images.⁴ In most cases consisting of man-made environments the most relevant orientations are detected, without any false alarms. Figures 1(e) illustrates the need of the MDL criterion in order to filter out artificial vanishing points that may appear when the real vanishing points are

⁴In all our experiments we used $\varepsilon = 1$, for coherence with previous works. But we could also have used a much smaller value, since in all the examples presented here all real vanishing points have $\text{NFA} < 0.0001$. Furthermore $\varepsilon = 1$ means that we can expect in average one false vanishing point in a random image, which is quite high with respect to the reduced number of vanishing points we usually find in real images.

Fig. 1. (a-d) Detected line segments for a building image and the only three maximal MDL vanishing points that are detected. They correspond to the two horizontal orientations and to one vertical orientation. Note that no orthogonality hypothesis was used, thus it can be used a posteriori in order to calibrate some camera parameters. For each vanishing point we only display the segments that contributed to this point at the automatically chosen precision. (e) Before applying the MDL criterion some spurious vanishing regions remain. Note that they arise from mixtures of real vanishing regions, and that they are significantly less meaningful and less precise than the real vanishing regions. Therefore, during MDL, most segments vote for the real vanishing region instead of these “mixed” ones, so that after MDL their number of false alarms decreases and they are no longer meaningful.

Fig. 2. (a-c) Detected line segments for an image of a painting by Uccello and the only two maximal MDL vanishing points that are detected. Note that the vanishing points corresponding to the oblique wall and the staircase are missed. This is due to the fact that both the alignment detection and the vanishing point detection are global, and the less meaningful segments and vanishing points are masked by the more meaningful horizontal and vertical orientations. (d) *Illustration of the “masking” phenomenon.* When we select the wall subimage in the previous figure, more alignments are detected, and a new vanishing point that was masked in the global image becomes meaningful. This is due to two cooperating effects. First the masking phenomenon at the alignment detection level means that we detect in this subimage more meaningful segments than in the global image. Secondly, at the vanishing points detection level, the total number of segments is smaller, which means that the minimal number of concurrent lines for a vanishing region to become meaningful k is also small. A similar result can be obtained by restarting the MDL iteration a second time with the remaining segments after all MDL meaningful vanishing points have been detected and the contributing segments removed.

Fig. 3. Accidental vanishing points. When applied to images of man-made environments which actually contain vanishing points, the method very rarely detects accidental vanishing points. But this does happen in natural images in which we do not perceive such vanishing points. Here we show one of the worst such examples that we found in our experiments. In this case the detected vanishing points are probably not perceived because they are made up mostly of segments that are not perceived as straight lines in the first place. Many of these segments would be better explained as meaningful curved boundaries, and therefore will never give rise to vanishing points. Hence the false alarms in the vanishing points detection phase are here to some extent the result of some special kind of false alarms in the alignment detection phase. Further experiments on natural images showed this kind of false alarms of vanishing points (due to some false alarms in line segments which are actually curved boundaries) to be the most prominent one.

extremely meaningful. Note that after MDL (figures 1(b) to 1(d)) we only get the main three directions (two horizontal and one vertical).

Figure 2 illustrates the masking phenomenon. Here the less meaningful directions corresponding to the wall are “masked” by the many segments in the horizontal and vertical directions, but it can be “unmasked”. See the figure captions for a more detailed explanation.

Finally Figure 3 shows the limitations of the proposed method when applied to natural images not containing vanishing points (see caption for details). This and other similar experiments further enforce the conclusion in [7] on the importance of addressing the conflicts between Gestalts. Indeed if we were able to resolve the conflict between the alignment and the curved boundaries Gestalts we would eliminate many “false” line segments and thus further reduce the number of false alarms in the vanishing point detection phase. Our experiments suggest that this approach might be complementary to (and in certain cases better adapted than) other approaches to reducing spurious responses rather based on joint Gestalts at the vanishing point detection level.

It is quite difficult to build an experimental setup which allows to fairly compare our method with previously proposed ones. The reason is that our assumptions are quite different here, since we do not treat the same problem: whereas most previous works [8], [12], [9], [11] look for joint Gestalts that combine 3D parallelism with some other property, whereas here he try to push the pure partial Gestalt of 3D parallelism to its limits.

An exception is the recent work in [3], which only relies on 3D parallelism and has been shown to produce highly accurate vanishing points, but assumes knowledge of the camera calibration parameters and omnidirectional images, which is not exploited by our method. The importance of this knowledge is not thoroughly discussed in [3], but it was crucial in [8] in order to reduce spurious responses. The work in [3] relies on a Hough transform as in [8] in order to determine the number of vanishing points and is therefore prone to the same sensitivity to the internal calibration parameters. For this reason it can be considered as complementary to our method. In fact our method could be used either in the initialization its step to determine the number and approximate positions of vanishing points more reliably, or as a validation step to reduce the number of false alarms.

REFERENCES

- [1] Andrés Almansa. *Échantillonnage, interpolation et Détection. Applications en imagerie satellitaire*. PhD thesis, ENS Cachan, 61 av. du Président Wilson, 94235 Cachan, France, December 2002.
- [2] Andrés Almansa, Agnès Desolneux, and Sébastien Vamech. Vanishing points are meaningful gestalts. Technical Report CMLA-2001-24, CMLA, ENS Cachan, 2000.
- [3] Mathew E. Antone and Seth Teller. Automatic recovery of relative camera rotations for urban scenes. In *Proceedings of CVPR 2000*, 2000.
- [4] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *International Journal of Computer Vision*, 2001.
- [5] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Meaningful alignments. *International Journal of Computer Vision*, 40(1):7–23, 2000.
- [6] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Edge detection by helmholtz principle. *Journal of Mathematical Imaging and Vision*, 14(3):271–284, 2001.
- [7] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. A grouping principle and four applications. *This same special issue of PAMI*, 2001.
- [8] Evelyne Lutton, Henri Maître, and Jaime Lopez-Krahe. Contribution to the determination of vanishing points using hough transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4):430–438, 1994.
- [9] C. Rother. A new approach for vanishing point detection in architectural environments. In *British Machine Vision Conference*, pages xx–yy, 2000.
- [10] Luis A. Santaló. *Integral Geometry and Geometric Probability*, volume 1 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1976.
- [11] F. Schaffalitzky and A. Zisserman. Planar grouping for automatic detection of vanishing lines and points. *Image and Vision Computing*, 18(9):647–658, June 2000.
- [12] J.A. Shufelt. Performance evaluation and analysis of vanishing point detection techniques. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(3):282–288, March 1999.
- [13] S.T. Barnard. Interpreting perspective images. *Artificial Intell.*, 21:435–462, 1983.
- [14] T. Tuytelaars, M. Proesmans, and L. Van Gool. The cascaded hough transform. In *International Conf. on Image Processing, ICIP-97*, volume 2, pages 736–739, 1997.