Submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence

An earlier version of this paper appeared in the *Workshop on Perception for Mobile Agents* (in conjunction with IEEE CVPR'99), pp. 74–81, June 1999

Real-time Recognition of Self-Similar Landmarks*

Daniel Scharstein, *Member, IEEE,* and Amy J. Briggs, *Member, IEEE*

Department of Mathematics and Computer Science
Middlebury College
Middlebury, VT 05753, USA
[schar,briggs]@middlebury.edu

Abstract

Mobile robot localization and navigation require robust and efficient methods for utilizing perceptual information. In this paper we propose a design for simple visual landmarks that can be unobtrusively added to an environment for navigation without a global geometric map of the workspace. We describe self-similar intensity patterns together with an efficient and reliable algorithm for their detection in real-time that can handle a wide range of affine transformations. An implementation of the method is available at www.middlebury.edu/~schar/landmark.

Index terms: Vision-based mobile robotics, pattern recognition, artificial landmarks, landmark design, real-time detection, sensor-based navigation

*Support for this work was provided in part by the National Science Foundation under POWRE grant EIA-9806108, VT-EPSCoR grant OSR-9350540, and by a grant to Middlebury College from the Howard Hughes Medical Institute.

1 Introduction

Navigation for mobile robots requires efficient methods for utilizing perceptual information. Simple landmarks that are strategically and unobtrusively added to the environment can provide reliable visual cues for efficiently obtaining localization information. In this paper we propose self-similar landmark designs and an algorithm for their detection that works under a wide range of viewing conditions in real-time.

1.1 Navigation using artificial landmarks

Many techniques have been employed for sensor-based navigation and localization. Industrial mobile robots have traditionally navigated by following painted lines on the floor [1] or tracking buried wires or infrared beacons. The disadvantage of these approaches is that they require substantial engineering of the environment. Recently many researchers have employed landmarks – either artificial [2, 3, 4, 5, 6, 7] or extracted from the environment [8, 9, 10, 11, 12] – to guide the motion of a mobile robot in indoor environments.

The most commonly used approach with artificial landmarks is heuristic: landmarks are designed and placed so that landmark detection under normal circumstances is straightforward. For example, some systems use simple patterns containing unique identifiers mounted on the ceiling [2]. The problem with the heuristic approach is that it only works under certain restricted conditions that are enforced for the sake of speed: the patterns must be viewed from a narrow range of distances and angles and will not be recognized if partially occluded.

Another proposed approach employs projective invariants [13]. The idea is to use smooth planar curves for recognition imaged under perspective viewing conditions that are invariant under changes in viewpoint. While interesting in theory, the method is not very robust and lacks practical applicability.

In this paper, we develop new landmark designs together with an efficient algorithm for their detection that can handle a wide range of affine transformations. In contrast to existing approaches that require analysis of a $2D$ image, our method finds matches along individual scanlines, without any preprocessing, making it suitable for real-time applications.

1.2 Related work

The approach taken in this paper and by a number of other research groups is to use artificial landmarks that can be easily and unobtrusively added to the environment. Becker *et al.* [2] use simple landmarks attached to the ceiling of the environment. The landmark design consists of a square grid of black and white tiles in a pattern that encodes the landmark orientation and unique identification. Because the landmarks are mounted on the ceiling at a fixed distance from the camera and will always be unobstructed, the problem of detecting landmarks in an image becomes straightforward. A shortcoming of this type of landmark design is that detection is sensitive to variations in lighting conditions. Taylor and Kriegman [7] use barcode landmarks consisting of five vertical stripes colored black,

gray, or white. Their recognition algorithm is to extract vertical edges and locate groups of six consecutive edges with the appropriate cross-ratios. The method requires the patterns to be unoccluded and placed at a specific orientation, and it relies on specialized hardware for real-time image processing. Lin and Tummala [4] propose more complicated three-dimensional landmarks consisting of a black circular disk on a white background with a smaller white disk mounted 12cm in front of the black disk. Projections of a circular disk in the image plane can be recognized using an elliptical Hough transformation technique. The method assumes that the center of the disks is set to the same height as the camera and that the pan-axis of the camera is parallel to the pattern.

Traditionally, vision-based robot navigation has proceeded from three-dimensional maps of the environment, constructed, for example, using stereo vision techniques [14]. More recently, landmarks have been used to navigate without a full environment model. Techniques for mobile robot navigation based on landmarks include those that are primarily reactive [9], those planned within a geometric environment map enhanced with perceptual landmarks [15, 10], and those based on a topological description of landmark locations without a global map [16, 11, 12].

Our landmarks are designed for a navigation system in which a topological map of landmark locations is first constructed during an exploratory phase and then used for navigation without relying on a global geometric map. Sensors are used for landmark detection and local obstacle avoidance. In this paper we focus on the important problem of designing landmarks that can be quickly and reliably detected under a wide range of viewing conditions with few assumptions about camera or landmark orientation, lighting conditions, or distances to the landmarks.

1.3 Outline of the paper

Section 2 introduces self-similar functions for the design of an optimally recognizable intensity pattern. We define a match measure that can be maximized over all translations to detect certain self-similar patterns in an image. An algorithm for finding such patterns in an image under general affine transformations is presented in Section 3. In Section 4 we analyze the expected shape of the local maxima of the match function. We then present our experimental results with real images in Section 5. We discuss ways of achieving real-time performance in Section 6, and conclude in Section 7.

2 Landmark Design

The key idea proposed here is to use self-similar intensity patterns as visual landmarks.

2.1 Self-similar functions

We define a function $f : \mathbb{R}^+ \rightarrow \mathbb{R}$ to be *p-similar* for a scale factor p, $0 < p < 1$, if $f(x) = f(px)\ \forall x > 0$. The graph of a p-similar function is self-similar; that is, it is identical to itself scaled by p in the x-direction (see Figure 1). Note that a p-similar function is also

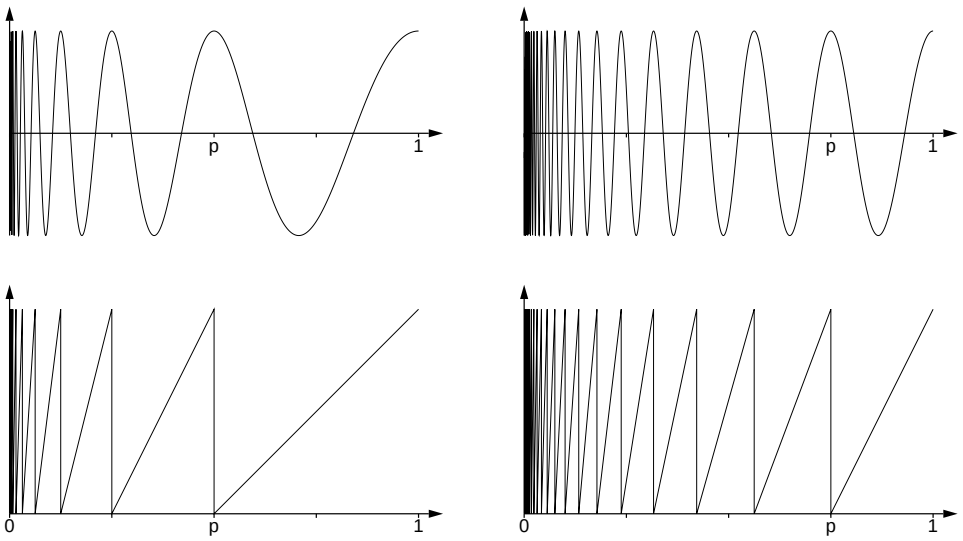

Figure 1: Examples of p-similar functions for $p = \frac{1}{2}$ (left) and $p = \frac{3}{4}$ (right).

p^k-similar, for $k = 2, 3, \ldots$ The property of p-similarity is invariant to scale:

$$f(\alpha x) = f(p\,\alpha x) = f(\alpha\, px). \tag{1}$$

This makes self-similar intensity patterns attractive for recognition, as the distance to the camera does not matter.

A p-similar pattern can be detected by comparing the observed intensity pattern to a version of itself scaled by p. We can accommodate patterns of limited spatial extent by restricting the comparison to a window of width w. Let $d_{p,w}(f)$ be the average absolute difference between the original and scaled functions over w:

$$d_{p,w}(f) = \frac{1}{w}\int_0^w |f(x) - f(px)|\, dx. \tag{2}$$

Then f is *locally* p-similar over w if and only if $d_{p,w}(f) = 0$. A simple method for detecting a locally p-similar pattern in a one-dimensional intensity function $I(x)$ would then be to minimize the above measure over translations $I_t(x) = I(x + t)$:

$$t_{\text{match}} = \arg\min_t d_{p,w}(I_t).$$

The minimal value of 0 is achieved only if I is p-similar over w at translation t. Unfortunately, all constant functions are p-similar for any p. Thus $d_{p,w}(I_t)$ would also be minimal in regions of constant intensity. To exclude locally constant functions, we must detect patterns that are self-similar *only* for scale p (and $p^2, p^3, \ldots$), and not for other scales. This can be achieved by *maximizing* the *mismatch* at scale $p^{1/2}$ (and $p^{3/2}, p^{5/2}, \ldots$). Let us assume without loss of generality that the range of observable intensities is $[0, 1]$. A maximal mismatch for scale $\sqrt{p}$ is then given if $|f(x) - f(\sqrt{p}\,x)| = 1$, or locally, if

$$d_{\sqrt{p},w}(f) = \frac{1}{w}\int_0^w |f(x) - f(\sqrt{p}\,x)|\, dx = 1. \tag{3}$$

In this case we say that f is (locally) $\sqrt{p}$*-antisimilar.*

We can combine equations (2) and (3), and revise our method for detecting self-similar patterns that are only p-similar for a given scale p. We will maximize the match function

$$m(t) = d_{\sqrt{p},w}(I_t) - d_{p,w}(I_t) \tag{4}$$

over all translations t. Note that the range of $m(t)$ is $[-1, 1]$, since the range of both terms on the right-hand side of (4) is $[0, 1]$. A locally constant function will yield $m(t) = 0$. Similarly, a random intensity pattern that is neither p-similar nor $\sqrt{p}$-similar will yield a response close to zero. A significant positive response is only expected for intensity patterns that are p-similar but not $\sqrt{p}$-similar.

2.2 An optimally recognizable pattern

Given the match measure m defined in (4), we now need to find an intensity function that yields the optimal response $m = 1$. That is, we seek a p-similar, $\sqrt{p}$-antisimilar function $s_p(x)$. To derive this function, let us switch to the domain of 1-periodic functions, i.e., functions F such that $F(z) = F(z + 1)\ \forall z$. A change of domain that maps 1-periodic functions F to p-similar functions f is given by

$$f(x) = F(\log_p x) = F(\ln x / \ln p), \tag{5}$$

since

$$\begin{aligned} f(px) &= F\left(\frac{\ln px}{\ln p}\right) = F\left(\frac{\ln p + \ln x}{\ln p}\right) \\ &= F\left(1 + \frac{\ln x}{\ln p}\right) = F\left(\frac{\ln x}{\ln p}\right) = f(x). \end{aligned}$$

Note that in general a translation of F by t corresponds to a scaling of f by p^t. Thus, we are looking for a function $S : \mathbb{R} \to [0, 1]$ such that

$$S(x) = S(x + 1)\ \forall x, \quad \text{and} \tag{6}$$

$$|S(x) - S(x + \frac{1}{2})| = 1\ \forall x. \tag{7}$$

A simple function satisfying both conditions is

$$S(x) = \begin{cases} 0, & x - \lfloor x \rfloor < \frac{1}{2} \\ 1, & x - \lfloor x \rfloor \geq \frac{1}{2} \end{cases} \quad = \lfloor 2(x - \lfloor x \rfloor) \rfloor, \tag{8}$$

i.e., a "square-wave" intensity function (see Figure 2). Note that there are infinitely many functions satisfying conditions (6) and (7). Equation (6) leaves an interval of width 1 unconstrained, say, $[0, 1)$. Equation (7) then leaves half of this interval unconstrained, say, $[0, \frac{1}{2})$. Equation (7) additionally restricts the function's values to be either 0 or 1, since the range of the function is limited to $[0, 1]$. Thus, any function $g : [0, \frac{1}{2}) \to \{0, 1\}$ can be used to generate a function S satisfying equations (6) and (7). Of these, we have chosen $g(x) = 0$.

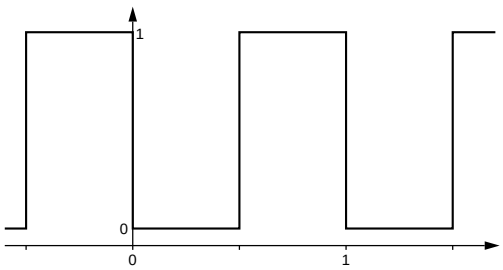

Figure 2: The square wave, a function $S : \mathbb{R} \to [0,1]$ such that $S(x) = S(x+1)$ and $S(x+\frac{1}{2}) = 1 - S(x)\ \forall x$.

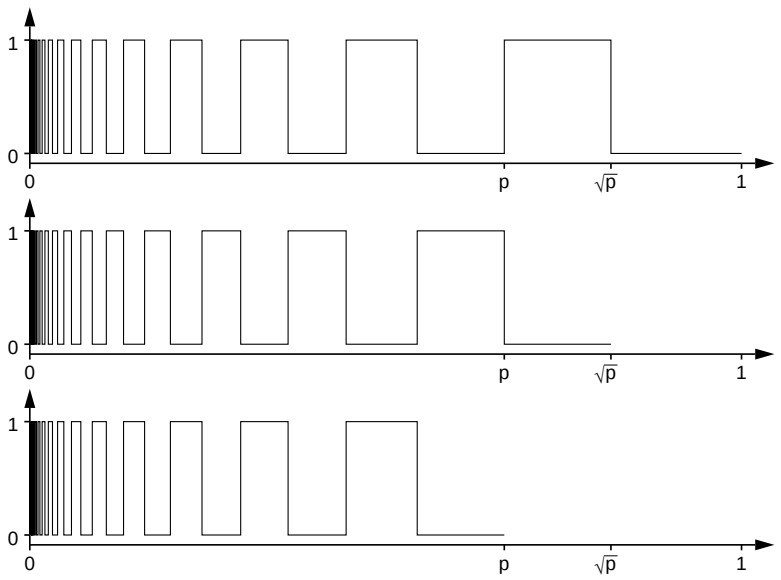

Figure 3: Self-similar square wave $s_p(x)$ for $p = \frac{2}{3}$ (top); maximal mismatch at scale $p^{1/2}$ (middle); and match at scale p (bottom).

Given our square wave function S defined in (8), we convert back to the domain of p-similar functions by setting

$$s_p(x) = S(\log_p x) = \lfloor 2(\log_p x - \lfloor \log_p x \rfloor) \rfloor \tag{9}$$

with the desired properties of $d_{p,w}(s_p) = 0$ and $d_{\sqrt{p},w}(s_p) = 1$ for any w (see Figure 3).

2.3 p-Similar Landmarks

We now have all the components for landmark design and recognition. The key step for moving to two dimensions is to use a pattern that is p-similar in one direction (say, horizontally), and constant in the other direction. If such a pattern is then sampled along any non-vertical line, the resulting intensity function is still p-similar because of the scale in-

variance of self-similarity (equation (1)). This allows us to detect two-dimensional p-similar patterns that have undergone an affine transformation by examining isolated scanlines.

Formally, let

$$L(x, y) = s_p(x) \tag{10}$$

be our two-dimensional landmark pattern, where $s_p(x)$ is the self-similar square wave function (equation (9)) for a fixed p. An affine transformation yields

$$A(L(x, y)) = L(ax + by + c, dx + ey + f) = s_p(ax + by + c).$$

Sampled at $y = y_0$, we get $s_p(ax + (by_0 + c)) = s_p(ax + t)$. Thus, the problem of finding an affine transformation of the two-dimensional pattern L has been reduced to finding a translation t of the one-dimensional pattern s_p.

3 Algorithm

The idea underlying the recognition algorithm is to find locations (x_m, y_m) in the image at which a scanline is locally p-similar and $\sqrt{p}$-antisimilar. To do this, we adopt the matching function m from equation (4) for scanlines in an image $I(x, y)$:

$$\begin{aligned} m_y(x) \;=\; & \frac{1}{w}\int_0^w |I(x+\xi, y) - I(x+\sqrt{p}\,\xi, y)|\, d\xi \\ & -\frac{1}{w}\int_0^w |I(x+\xi, y) - I(x+p\xi, y)|\, d\xi. \end{aligned} \tag{11}$$

The value of $m_y(x)$ depends on the intensities along the line (x, y) to $(x + w, y)$, and is constrained to the interval $[-1, 1]$. If the pattern s_p is present at (x, y), then $m_y(x) = 1$. It is easy to see that the pattern cs_p with reduced contrast $c < 1$ (i.e., difference between maximal and minimal intensities) will only yield a response $m_y(x) = c$. Other (non p-similar) intensity patterns will yield responses close to or below zero.

An algorithm for finding affine transformations of a landmark with intensities $L(x, y) = s_p(x)$ (for a known p) in an image can be formulated as follows:

```
for every k-th scanline y
    for all x
        compute m_y(x)
    mark all strong local maxima of m_y as matches.
```

This simple algorithm requires only $O(nw/k)$ operations for an n-pixel image. The computation of $m_y(x)$ can be adapted to discrete images by replacing the integrals in equation (11) by summations, and determining inter-pixel intensities using linear interpolation. The two parameters, k, the spacing of scanlines to search, and w, the window size, only depend on the smallest expected size of the landmark pattern and can be fixed. Typical values for an image of size 640×480 are $k = 6$ and $w = 45$.

The interesting question is: what constitutes a "strong" local maximum? A simple answer would be to require that a local maximum be greater than a fixed threshold $c_{\min}$ (corresponding to a minimum contrast). More information, however, can be gained from observing the shape of $m_y(x)$ in the vicinity of a local maximum.

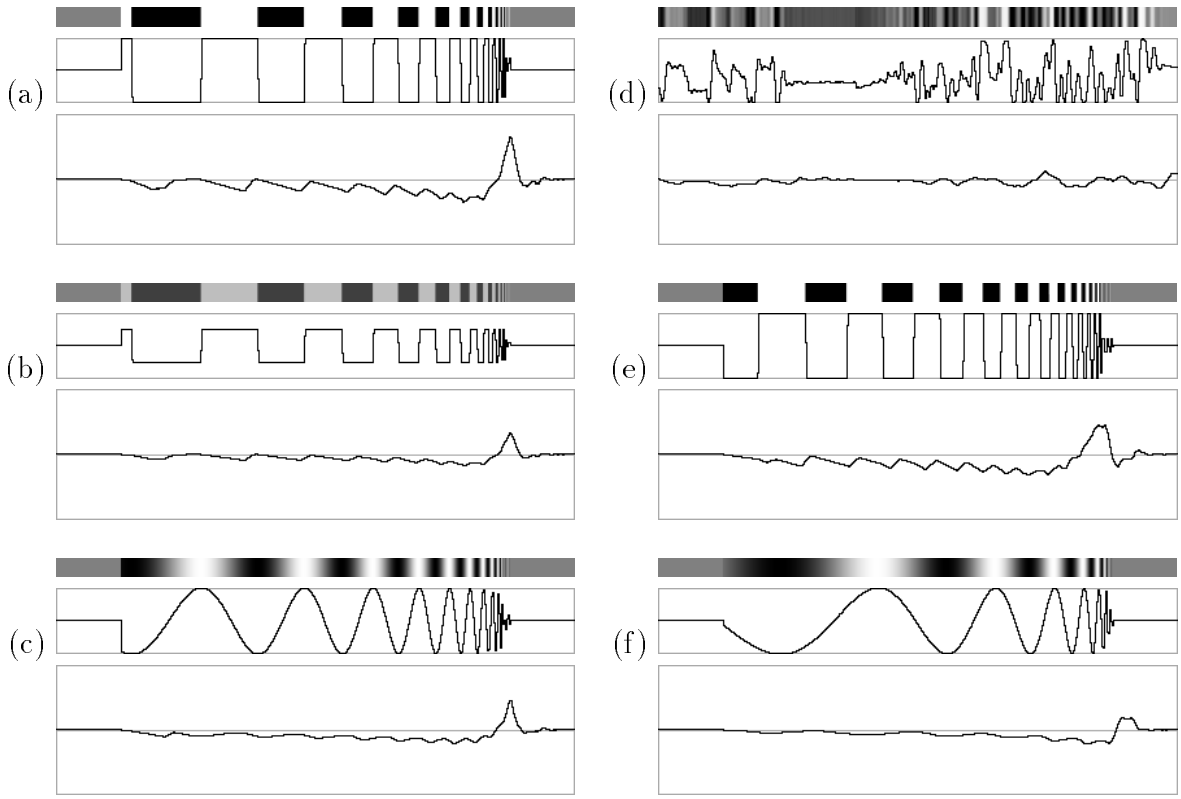

Figure 4: Various intensity patterns $I(x)$ and the response of the match function $m(x)$ for $p = \frac{2}{3}$ and $w = 50$. (See Section 4 for details.)

4 Finding matches reliably

In analyzing the shape of m[1] we have to take into consideration the fact that real images are discrete and have finite resolution. Sampling effects will significantly alter the results from the continuous case. We will therefore use synthetic images for our analysis, and then test our findings on real images in Section 5.

Figure 4 shows several intensity profiles of scanlines that were synthesized from continuous functions using 10-fold oversampling. Below each intensity plot is a plot of $m(x)$ for $p = \frac{2}{3}$ and $w = 50$. The length of each scanline is 400. The patterns on the left are locally $\frac{2}{3}$-similar at $x_m = 350$; (a) and (b) are square waves s_p with full and half contrast respectively, and (c) is a p-similar sine wave $\sin_p(x) = \sin^2(\pi \log_p x)$. All three patterns result in clear peaks at the correct location x_m; however, the value of the maximum $m(x_m)$ is less than expected. The observed values are (a) 0.66, (b) 0.33, and (c) 0.46, while the expected values are 1, $\frac{1}{2}$, and $\frac{2}{\pi} \approx 0.64$. The differences are due to sampling and interpolation, in particular at locations of discontinuities or strong change. The patterns on the right, on the other hand, are *not* $\frac{2}{3}$-similar: (d) is a "random" intensity pattern taken from a real image, (e) is a $\frac{3}{4}$-similar square wave, and (f) is a $\frac{1}{2}$-similar sine wave. The match function

[1] We drop the subscript y in this section as we are only considering single scanlines.

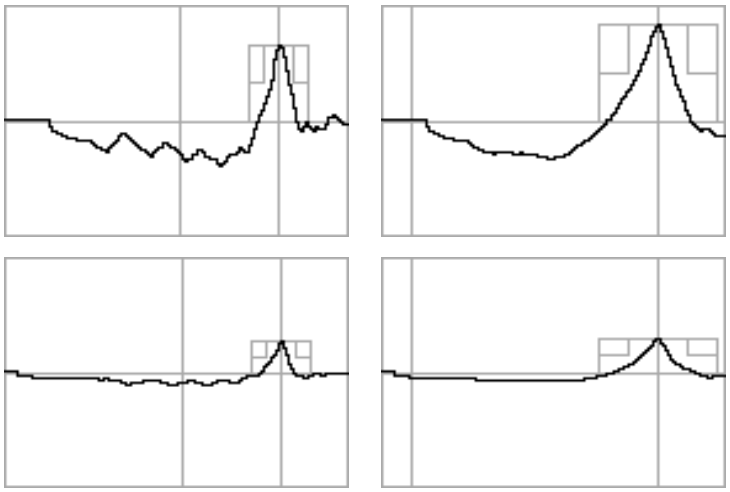

Figure 5: A shape constraint for detecting sharp local maxima. The two vertical lines indicate the window size: $w = 40$ (left), and $w = 100$ (right).

for the random pattern (d) has no strong peaks, but there are distinctive local maxima in the match functions for (e) and (f). They are, however, much more rounded and not as "sharp" as the peaks on the left.

We have performed a large number of experiments to determine the influence of the various parameters (window size, contrast and shape of the intensity function, and self-similarity factor p) on the shape of the local maxima of m. We found that the shape remains fairly constant over a wide range of parameter values, but scales in width and height. In particular, the height of the peak is proportional to the intensity contrast of the pattern, and the width of the peak at the level of the y-axis is about $1/5$ of the window size. The sides of the peak above the y-axis are typically slightly concave.

These observations can be used to derive a simple test for "sharpness": given a local maximum $v = m(x_m)$ greater than a threshold $c_{\min}$, test whether $m(x_m \pm \delta_x) < v/2$, where δ_x is approximately $w/10$. This test is illustrated in Figure 5. The top row shows the match function m for the square wave pattern s_p with full contrast as in Figure 4 (a); the bottom row shows m for a sine wave pattern $\sin_p$ similar to Figure 4 (c), except with half contrast. Each function is shown for $w = 40$ (left) and $w = 100$ (right). The two vertical lines show the window size w; the small boxes have width δ_x. Note that the shape of the peak is qualitatively the same in all four images.

5 Experiments

Our first experiment illustrates the robustness of the proposed method. Figure 6 shows two images of cluttered scenes containing many self-similar patterns at different distances and orientations. The algorithm correctly identifies all p-similar patterns for the specified similarity factor $p = \frac{2}{3}$, while ignoring p-similar patterns for different p. The parameters used were $w = 50$, $k = 6$, $c_{\min} = 0.1$, $\delta_x = 5$.

The second experiment demonstrates the use of the method for robot navigation. To

Figure 6: Two images containing many landmark patterns. All $\frac{2}{3}$-similar patterns have been detected. Clutter as well as p-similar patterns for $p \neq \frac{2}{3}$ do not result in matches.

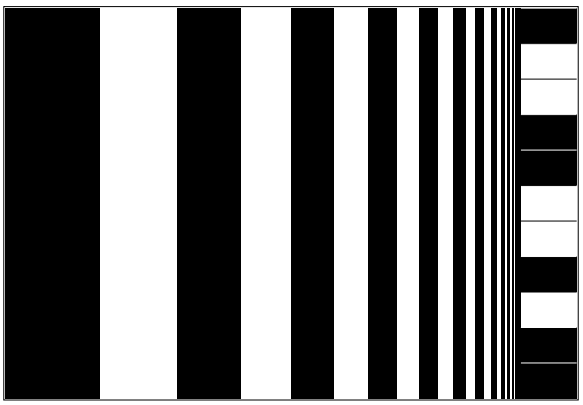

Figure 7: Our self-similar landmark pattern with barcode.

distinguish different landmarks, we add a simple binary barcode to the right side of the self-similar pattern (see Figure 7). The recognition algorithm is then extended to group all matches corresponding to a single landmark pattern as follows: First, the position and orientation of each landmark pattern in the image is determined by fitting a straight line to each set of three or more individual matches detected on consecutive scanlines. The exact vertical extent of each pattern is then estimated from the intensity distribution to the left of this line. Once the locations of the patterns are known, their barcodes can be decoded easily. Figure 8 shows a Pioneer 2 mobile robot equipped with a camera, and two sample pictures taken by the robot. All landmarks are detected correctly, and all landmark numbers are decoded except for one that was too far away. The parameters used for this experiment are $w = 45$ and $k = 6$.

The final experiment demonstrates the robustness of the method to different orientations. The algorithm, as described so far, only detects horizontally self-similar patterns in one direction. In practice, this restricts the range of allowable rotations to about ± 45 degrees, since several scanlines have to intersect a significant portion of the landmark pattern. While this range is perfectly sufficient for many applications (including robot navigation), it can easily be extended to 360 degrees — at the expense of a 4-fold increase in running time — by searching four rotated copies of the image. Figure 9 shows 12 images from a sequence of 60 images of a landmark pattern undergoing a rotation of 360 degrees. The pattern is detected correctly in 59 of the 60 images with the default setting of parameters.

6 Achieving real-time performance

The recognition algorithm as described above is fairly fast: typical running times for a straightforward implementation on a 450 MHz Pentium II machine are 0.33 seconds for 640×480 images and 0.08 seconds for 320×240 images (using parameters $w = 45$ and $k = 6$). If a rotation range of 360 degrees is allowed, the running times increase slightly less than fourfold to 1.1 and 0.27 seconds, respectively.

A dramatic speed-up can be achieved by further restricting the set of pixels for which $m_y(x)$ is computed. Recall that we are already considering only every k-th scanline. We

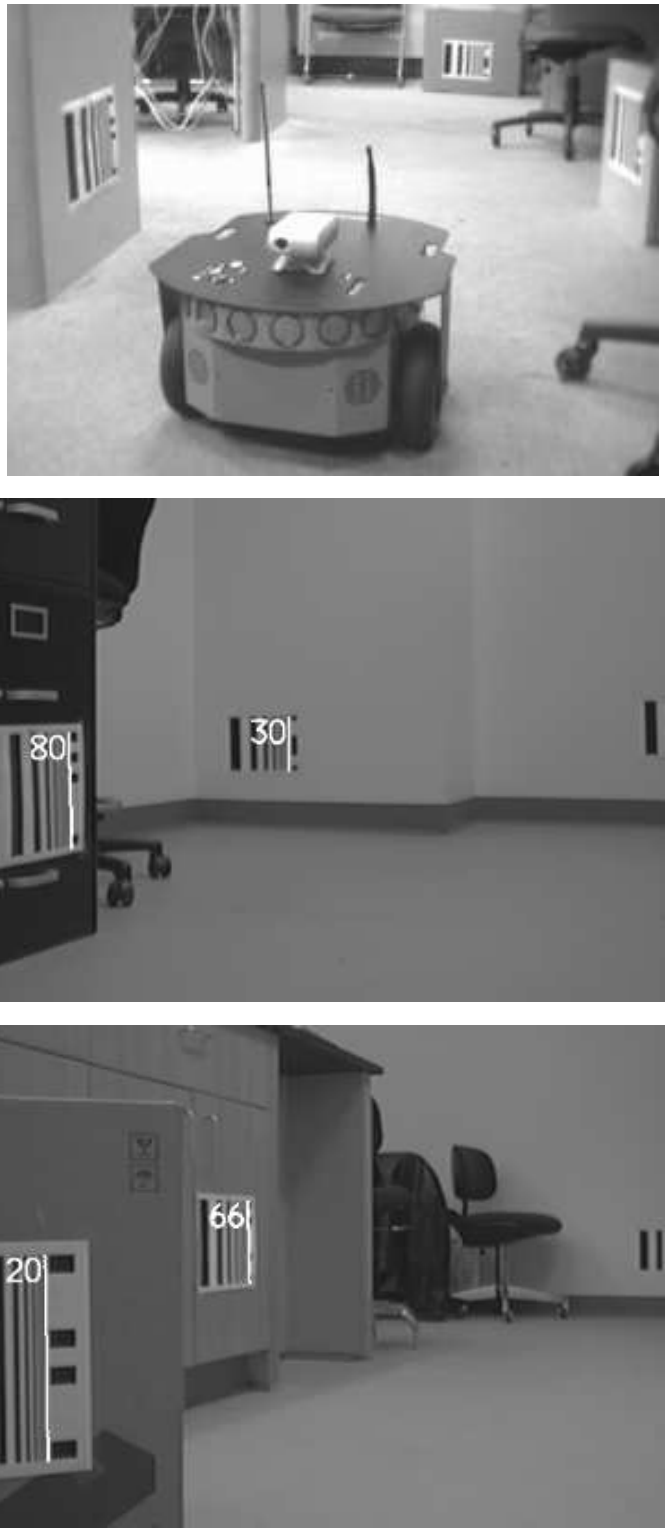

Figure 8: The mobile robot, and two images taken with its camera. All landmarks have been found, and all but one have been identified by their barcodes.

Figure 9: 12 images from a 60-image sequence demonstrating the robustness of the method for rotation.

can start by looking at only every $2k$-th scanline, and, only if a match is found, look at its neighboring scanlines as well. Since isolated matches are discarded anyway in the grouping process, no landmark will be missed. This technique can yield a speed-up of up to 2 if few or no landmark patterns are present in the image. Another opportunity for restricting the search is to make use of the fact that peaks corresponding to matches have a certain minimum width at a given height h (e.g., $h = c_{\min}/4$). Given a lower bound l for this width, it is possible to scan for peaks by looking only at every l-th pixel. A conservative

	Image size	
Allowable rotation	320×240	640×480
± 45 degrees	0.0123 s (81 fps)	0.0272 s (36 fps)
360 degrees	0.026 s (38 fps)	0.088 s (11 fps)

Table 1: Running times and frame rates of the optimized implementation for different image sizes and different allowable rotations.

bound for l is typically given by δ_x, i.e., half the allowable peak width used in the test for sharpness discussed in Section 4. Typical values for this number are 4 or 5 pixels. We have found, however, that even with much higher values for l (e.g., 10), very few peaks are overlooked (typically only those resulting from patterns with very low contrast), so further speedup can be achieved with minimal impact on robustness.

The new running times resulting from the above modifications are summarized in Table 1. Note that real-time performance (more than 30 frames per second) is achieved for full-size images with a rotation range of ± 45 degrees, as well as on half-size images with a rotation range of 360 degrees. Even in the slowest scenario, a full rotation range on 640×480 images, a rate of more than 10 frames per second can be achieved.

Given that landmarks can be recognized at frame rate, the method is ideally suited for real-time applications such as robot navigation. In our current system, the robot navigates an environment augmented with landmarks by traversing the edges of a landmark visibility graph. This graph is constructed during an initial exploration phase, and dynamically updated as the robot performs subsequent navigation tasks. The high speed of the recognition algorithm is crucial both for landmark detection and for reliable navigation: First, to record which landmarks are visible from a newly reached position, the robot performs a single 360-degree scan, during which images are continuously acquired and searched for landmarks. Second, to travel from one landmark to another the robot uses a visual servoing routine that adjusts the robot's motion to keep the landmark centered in its field of view. Because of the short latency between image acquisition and landmark detection, this servoing routine can also be employed for more challenging tasks, such as following a moving target. Many other applications are possible as well.

7 Conclusion

Artificial landmarks that can be unobtrusively added to an indoor environment can serve as practical and inexpensive aids for mobile robot navigation and localization. When detected quickly and reliably they can make task execution more robust by reducing uncertainty due to control and sensing errors. In this paper we have proposed self-similar intensity patterns as landmarks and have described an efficient algorithm for their detection. Our method robustly detects landmarks under a wide range of affine transformations, which are

a sufficient approximation of projective transformations in practice, given that the size of a landmark is typically small compared to its distance to the camera. The method operates on single scanlines without any preprocessing and runs in real-time without specialized hardware. The patterns and the algorithm have real-world applications not only in mobile robot navigation, but also in other situations where fast, reliable tracking of a pattern is required. One important emerging such application is *video insertion*, i.e., the real-time substitution of a pattern in a video stream [17] (for example, the insertion of regional advertisements on billboards in the national broadcast of a sports event [18]).

In summary, we have proposed the first practical method employing full affine invariants for landmark detection in real-time applications. An implementation of the algorithm is publicly available at `http://www.middlebury.edu/~schar/landmark/`.

Acknowledgements

We are grateful to Darius Braziunas, Huan Ding, Cristian Dima, Cuong Nguyen, and Sorin Talamba for their insights and their assistance in implementing the algorithm and in conducting the experiments reported here.

References

[1] E. McVey, K. Drake, and R. Iñigo, "Range measurements by a mobile robot using a navigation line", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 1, pp. 105–109, Jan. 1986.

[2] C. Becker, J. Salas, K. Tokusei, and J.-C. Latombe, "Reliable navigation using landmarks", in *Proceedings of IEEE International Conference on Robotics and Automation*, June 1995, pp. 401–406.

[3] M. Kabuka and A. Arenas, "Position verification of a mobile robot using standard pattern", *IEEE Journal of Robotics and Automation*, vol. RA-3, no. 6, pp. 505–516, Dec. 1987.

[4] C. Lin and R. Tummala, "Mobile robot navigation using artificial landmarks", *Journal of Robotic Systems*, vol. 14, no. 2, pp. 93–106, 1997.

[5] J. Salas, J. L. Gordillo, and C. Tomasi, "Visual routines for mobile robots: Experimental results", *Expert Systems with Applications*, vol. 14, pp. 187–197, 1998.

[6] K. Tashiro, J. Ota, Y. C. Lin, and T. Arai, "Design of the optimal arrangement of artificial landmarks", in *Proceedings of IEEE International Conference on Robotics and Automation*, June 1995, pp. 407–413.

[7] C. J. Taylor and D. J. Kriegman, "Vision-based motion planning and exploration algorithms for mobile robots", in *1994 Workshop on the Algorithmic Foundations of Robotics, A. K. Peters*, Goldberg, Halperin, Latombe, and Wilson, Eds., 1995, pp. 69–83.

[8] C. Fennema, A. Hanson, E. Riseman, J. R. Beveride, and R. Kumar, "Model-directed mobile robot navigation", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 20, no. 6, pp. 1352–1369, 1990.

[9] D. P. Huttenlocher, M. E. Leventon, and W. J. Rucklidge, "Vision-guided navigation by comparing edge images", in *1994 Workshop on the Algorithmic Foundations of Robotics, A. K. Peters*, Goldberg, Halperin, Latombe, and Wilson, Eds., 1995, pp. 85–96.

[10] B. Nickerson, P. Jasiobedzki, D. Wilkes, M. Jenkin, E. Milios, J. Tsotsos, A. Jepson, and O. N. Bains, "The ARK project: Autonomous mobile robots for known industrial environments", *Robotics and Autonomous Systems*, vol. 25, pp. 83–104, 1998.

[11] C. Owen and U. Nehmzow, "Landmark-based navigation for a mobile robot", in *Proceedings of Simulation of Adaptive Behaviour*. 1998, MIT Press.

[12] R. Sim and G. Dudek, "Mobile robot localization from learned landmarks", in *Proceedings of IEEE/RSJ Conference on Intelligent Robots and Systems (IROS), Victoria, BC*, Oct. 1998.

[13] S. Carlsson, R. Mohr, T. Moons, L. Morin, C. Rothwell, M. Van Diest, L. Van Gool, F. Veillon, and A. Zisserman, "Semi-local projective invariants for the recognition of smooth plane curves", *International Journal of Computer Vision*, vol. 19, no. 3, pp. 211–236, Aug. 1996.

[14] O. Faugeras, *Three-Dimensional Computer Vision*, MIT Press, Cambridge, MA, 1993.

[15] A. Lazanas and J.-C. Latombe, "Landmark-based robot navigation", *Algorithmica*, vol. 13, no. 5, pp. 472–501, May 1995.

[16] A. Kosaka and J. Pan, "Purdue experiments in model-based vision for hallway navigation", in *Proceedings of the Workshop on Vision for Robots in IROS'95, Pittsburgh, PA*, 1995, pp. 87–96.

[17] K. Hanna and P. Burt, "Video merging employing pattern key insertion", US Patent 5,566,251, Oct. 1996.

[18] G. Medioni, G. Guy, H. Rom, and A. François, "Real-time billboard substitution in a video stream", in *Proceedings of the 10th Tyrrhenian International Workshop on Digital Communications, Ischia, Italy*, 1998.