

# Evolving Visual Sonar: Depth From Monocular Images

Martin C. Martin<sup>★</sup>

*Mobile Robotics Laboratory, Carnegie Mellon University, Pittsburgh, PA*

---

## Abstract

To recover depth from images, the human visual system uses many monocular depth cues, which vision research has only begun to explore. Because a given image can have many possible interpretations, constraints are needed to eliminate ambiguity, and the most powerful constraints are domain specific. As an experiment in the automatic discovery and exploitation of constraints, Genetic Programming was used to find algorithms for obstacle detection. The algorithms are designed to be a replacement for sonar, returning the location of the nearest obstacle in a given direction.

The evolved algorithms worked surprisingly well. Errors were largely transient. The algorithms generalized to both novel views of the office environment and to unseen obstacles. They were combined with a simple reactive wandering program originally written for sonar. The result exhibited good performance in an office environment, colliding only with obstacles outside the robot's field of view.

Time to collision results and failure modes are presented. Code is available for download.

*Key words:* Genetic Programming, Robotics, Visual Navigation, Monocular Vision

---

## 1 Introduction

Video cameras are attractive sensors that provide a large amount of data at high speed and low cost. These properties make them particularly attractive in robotics. However, current robotics navigation algorithms, from potential fields (Arkin, 1998) to simultaneous localization and mapping (SLAM) (Thrun, 2002), need the distance to the first obstacle in various directions.

To extract distances, stereo vision is typically used. Stereo vision returns a dense depth map, but requires precise synchronization and calibration to find the corresponding line to within a few pixels. It also requires significant CPU power. Yet a full depth map is not needed, as sonar sensors allow reliable navigation. Can we produce a kind of “visual sonar,” a simple algorithm that produces a few depth estimates per image?

Monocular images clearly contain sufficient information, as people can understand them and psychologists have elucidated many monocular depth cues. Yet it is far from obvious how to compute depth from such an image. Monocular images were central to early work in computer vision, where the problem was commonly recognized as under constrained, in that many different physical scenes can result in the same image. For example, two scenes that differ only by having one object closer to the camera and smaller can produce exactly the same image.

The solution was to incorporate *constraints*, such as that chimneys sit on top of houses, rather than hover in the air in front of them. Marr considered

---

\* Now at Icosystem Corporation, 10 Fawcett St., Cambridge, MA 02138

*Email address:* [martin@metahuman.org](mailto:martin@metahuman.org) (Martin C. Martin).

*URL:* [www.metahuman.org/martin](http://www.metahuman.org/martin) (Martin C. Martin).




Fig. 1. Evolved programs return the pixel coordinates of the lowest obstacle in a given column of the image. In the above image, the white bar covers ground pixels, while the black bar covers obstacles. The cross between these two shows the exact transition point. The five other crosses show the ideal response for five other columns.

constraints “the critical act in formulating computational theories” of vision (Marr, 1978). And the second half of the influential book Duda et al. (1973) focused on finding the constraints necessary for scene understanding. Yet today computer vision researchers look for domain independent techniques since these have the most general applicability.

Previous work in monocular vision (summarized below) relies on the researcher to discover domain constraints and embody them in algorithms. To create algorithms for new environments requires someone with a background in Robotics. This severely restricts the practical application of the algorithms. This paper provides a way out of the tension between domain dependence and independence, by proposing a general, domain independent tool which searches a set of algorithms to discover which constraints to exploit.

The main contribution of this work is the *automatic* discovery of domain constraints. This paper uses genetic programming to create obstacle detection

algorithms. In particular, we created a function that takes in a  $320 \times 240$  grayscale image, as well as the horizontal location of a column in the image, and returns the lowest non-ground pixel in the image. Thus, we define an obstacle as anything other than the ground, be it a chair, a wall or a person. See Figure 1. We used the simplest, most common form of genetic programming, described in Koza (1992). The only significant addition was a node for iterating a window over an image. Moving a window over an image is a common component of many vision algorithms. This work extends earlier efforts (Martin, 2002) by extending the previous analysis to explore issues of interest in pattern recognition.

The remainder of this section reviews previous work in both monocular vision and the evolution of visual routines. Section 2 describes the details of the evolutionary computation setup. Section 3 describes the performance of the system on the training data, as well as the results of using the evolved vision algorithm to control a robot. Finally, section 4 analyzes the results and discusses their implications.

### *Previous Work*

A few research groups have tackled monocular vision for robotics. Polly the robot (Horswill, 1993) looked for the lowest obstacle in a given column of an image. This exploits a documented monocular cue: If most obstacles touch the floor and the floor is roughly flat, then obstacles higher in the image are further away (Ooi et al., 2001).

To find the lowest obstacle, Polly’s creator assumed that the carpet was al-

ways textureless, and thus declared any area with visual texture to be an obstacle. Although algorithmically quite simple, it allowed Polly to give over 100 tours of the 7th floor the MIT Artificial Intelligence Lab, logging hundreds of hours without incident (Horswill, 1994). Lorigo (1996) assumed the bottom of the image was always ground, and looked for areas that differed significantly, declaring them obstacles. Ulrich et al. (2000) took ground to be the portion of a previous image that the robot has already traversed.

The work in Xie (1996) used a stereo pair of cameras, but without matching windows between the two images. Instead, it projected features from one image into the other, at the location expected if the feature were on the ground plane. When features fail to match up, this indicates they are not on the ground plane, suggesting that they are an obstacle. While this reduces the computational burden, it still requires careful calibration of the cameras and their positions relative to the ground plane. If either the cameras or the objects are moving, it also requires temporal synchronization of the cameras.

The beebot (Camus et al., 1999) used the divergence of the optical flow field to compute time to impact, steering where this was lowest. It has wandered the lab at 20 cm/s for as long as 26 minutes without collision. Sobey (1994) used a monocular camera combined with a zigzag motion to estimate the range to objects, and estimate the safe distance of travel in the present direction.

Others have evolved visual routines before. Johnson et al. (1994) evolved programs to find hands in binary images of people. The images were generated by having people stand in front of a blue screen, and marking every non-blue pixel as part of the person. Nodes included the centroid and bounding box of the silhouette, horizontal edge detectors and the ability to find the leftmost,

rightmost or average of a set of detected edges. Their results were promising, correctly classifying up to 93% of images. That performance was better than the best algorithms they were able to write by hand.

Cliff et al. (1993) evolved programs that allowed the Sussex gantry robot to move toward a white square while avoiding a white triangle in an otherwise black world. Images from the CCD camera were first reduced to three numbers, the average brightness over three circles. Smith (1998) used a 16 pixel one-dimensional camera on a robot soccer field. Several people have evolved convolution masks, such as interest operators and edge detectors, that return a result for every pixel (Bala et al., 1996, 1997; Ebner, 1998; Harris et al., 1996; Harris, 1997; Benson, 2000; Guarda et al., 1998). Teller et al. (1997) evolved programs to classify images which contained functions to return the min, max, mean and variance of image pixels over a rectangle. They also made use of an array of memory, as a way of maintaining state.

## 2 Methods

A supervised learning framework was employed. Images of three different environments were obtained, and the location of the lowest non-ground pixel determined manually. Evolutionary Computation was then given these images, and the output compared to these manually determined responses. The best evolved programs were then used to control a robot, whose performance was measured.

Images were recorded with a black and white analog video camera. Processing was performed by an off board dual 700 MHz Pentium III computer. Evolved

algorithms took in an image, that is, a  $320 \times 240$  array of grayscale values, and the horizontal location within an image of a column. They returned an estimate of the lowest non-ground pixel in that column of the image, see Figure 1.

The camera was mounted on the robot and pitched 51 degrees down from horizontal. The training set was gathered while the robot navigated using sonar, not vision. The view from the camera, a series of  $320 \times 240$  grayscale images, was passively recorded. A detailed description of the navigation algorithm can be found in Martin (2001).

Three datasets were used. The first two consisted of images from two hallways in different buildings, Newell Simon Hall (NSH) and the Field & Mobile Robotics Building (FMRB). The third data set consisted of every other image from the first two. Both environments sported glossy, textureless gray walls, which are notoriously difficult for local depth estimation techniques such as stereo and optical flow. The robot made two or three turns in each environment, and its shadow is clearly visible in most frames. In NSH there were several people visible at different times. In that same environment, the carpet had a large red stripe at one end of the hallway. In FMRB, the overhead lights were burned out in one part of the hallway, so that the intensity of the carpet varied widely.

To establish ground truth, the author determined and recorded the ideal response of the vision subsystem, i.e., the location of the lowest non-ground pixel, in each of six columns in each image. The number six was chosen somewhat arbitrarily to match the typical number of sonar sensors in the cameras roughly 90 degree field of view. The three training sets had approximately 70

Objective	Given an image and a horizontal position within it, return the first non-ground pixel in that column.
Architecture of individuals	Two trees, each with a single iterate node at the root.
Fitness cases	Six columns in each of 65-71 images. 390-426 total.
Raw fitness	The percentage of fitness cases correct. A fitness case was correct if the pixel location estimated by the evolved program was within 10 pixels of the authors determination of the ideal answer.
Standardized fitness	100% minus raw fitness.
Miscellaneous Parameters	51 generations, population size of 4,000, tournament selection with tournament size of 7, ramped-half-and-half with min size 6 and max size 9.

Table 1

Summary of the evolutionary run

images each. The total number of fitness cases was therefore approximately  $70 \times 6 = 420$  per data set.

Each evolved program was run six times on each image, once for each column where training data was assigned. Fitness was the number of “correct” responses. A column was judged correct if the value returned was within 10 pixels of the manually determined ideal response.

The vision algorithm was evolved using Genetic Programming (Koza, 1992). Genetic programming is a genetic algorithm in which individuals are programs, represented as parse trees. The set of node types is chosen by the programmer

and is dependent on the problem domain. Random programs are created by choosing a random node type for the root, then choosing a random node type for each of its arguments, continuing recursively. A fixed number (say 4000) are created this way and each one is assigned a scalar fitness. Individuals are then randomly chosen, with the fitter individuals more likely to be chosen, and are either copied verbatim into the next population (replication) or a random subtree is chosen in two different individuals and swapped (crossover). Once 4000 individuals are created this way, the original individuals are discarded, the new individuals are assigned fitness and the process repeats.

A new type of node, dubbed the *iterate* node, was created, which moved a rectangle vertically over the image. The node's three children were the size of the rectangle, the horizontal location (in pixels) of the column in which to iterate, and a subtree to execute at each location. This last argument could incorporate the results of various image operators applied to the window.

The full list of node types (i.e. function and terminal sets) is summarized in Table 2. In addition to the *iterate* node and rectangle sizes, the list includes standard arithmetic operators, various parameters of the image, flow control, nodes to set and read five floating point registers, and both the average, and average of squared value, of common image operators. The registers were inspired by Teller et al. (1997).

Each individual was composed of two trees. Each tree had a single *iterate* node, always at the root, and this rule was enforced during creation and crossover of individuals. During evaluation, the first tree was run, then the second tree, which could use the results of the first. Using two trees allows two pass algorithms, for example estimating properties of areas likely to be floor (such as

root	<code>iterate-up, iterate-down</code>
rectangle sizes	<code>2×2, 2×3, 3×2, 3×3, 2×4, 4×2, 4×4, 5×5, 2×6, 6×2, 3×6, 6×3, 6×6, 7×7, 2×8, 8×2, 3×8, 8×3, 8×8</code>
arithmetic	<code>+, -, ×, % (protected division), square, sin, cos, tan, atan, atan2, min, max</code>
parameters	<code>x-obstacle</code> , the horizontal pixel location in which to find the obstacle; <code>window-area</code> , the area of the window in pixels; <code>image-max-x</code> (always 319); <code>image-max-y</code> (always 239); <code>first-window</code> , +1 if this is the first location of the window for this iteration, -1 otherwise; <code>x</code> and <code>y</code> , the center of the rectangle in pixels; random constants
flow control	<code>prog2; prog3; break</code> , which halts the execution of the tree, returning immediately without any more iterations; <code>if-le</code> , the if less than or equal to operator
memory	<code>set-a ... set-e, read-a ... read-e</code>
results of first tree	<code>iter2-a ... iter2-e</code> (only available in the second tree)
image operators	average and average-of-squared over the window: raw, truncated median, median corner, Sobel magnitude, and four directional Moravec interest operators

Table 2

The set of node types that genetic programming combined into programs.

the bottom middle of the image), then comparing these to the desired column. Iterate nodes do not return a value in the traditional sense. Instead, the final values of the five registers are used as five return values from the tree. Thus,

the second tree can use the values of any of the five registers of the first tree. The result of the entire individual was simply the final value of the  $a$  register of the second tree.

The other parameters used during evolution are summarized in Table 2. The miscellaneous parameters are those used by (Koza, 1994), with the exception that we used a larger population. Early, informal experiments with a population size of 2,000 showed poor performance, while the larger population size worked quite well. On a dual 700 MHz Pentium III the times to assign fitness to a single individual varied widely but averaged 725 msec. The time for an entire run of 51 generations also varied widely, averaging approximately 20 hours.

### 3 Results

All three experiments achieved a fitness of greater than 85%. They did this despite burned out lights and other effects that caused the carpet’s average intensity to vary from zero to 145 out of 255; despite large gradients caused by imaging artifacts; despite moir patterns of image noise; and despite the shadow of the robot.

On the FMRB data set, in all but two of the runs, the best individual in the initial population achieved a fitness of 21.6%. The other two runs achieved 29.3% and 36.4% in the initial population. The runs on the other two datasets were similar.

The best individual from all FMRB runs achieved 91.8%. Some examples from the training set are shown in Figure 2. While hard to see, the white squares




Fig. 2. Results of the best evolved individual on test (i.e. not training) data. The evolved program achieved the correct answer over 90% of the time.

are the size of the window chosen by evolution. The other two conditions did almost as well, achieving 90.5% on the NSH data set, and 84.8% on the combined data set. In all three experiments the errors were transient, that is, an error in one image would disappear in the next image, even when the same objects were visible. As is discussed below, this means the errors are from noise in the imaging process and are easily filtered out. The exception of the stripe of red carpet in Newell Simon Hall, which was uniformly considered an obstacle when near the camera.

### *Validation On The Robot*

All three individuals generalized surprisingly well when run on new video of the same hallway. They were relatively insensitive to the pitch and the horizontal location of the column in the image. They worked on a range of iris settings and camera tilts and didn't overfit to column location. They detected obstacles that weren't present during training, such as chairs or people, with only a little less fidelity than they detected walls. They were also fast, running at about 10 Hz on a 700 MHz Pentium III.

If most obstacles touch the ground, then given rough estimates of the camera's height and tilt angle, these image locations can be translated into distances. A hand written, reactive algorithm used these distances to avoid obstacles in real time. The algorithm was originally developed using sonar, and ran with the distances from evolved programs with very little tweaking. A detailed description of the algorithm can be found in Martin (2001). The three different evolved programs, trained on the different data sets, all used the same avoidance algorithm. The vision subsystems were run on twelve columns in the image, twice as many as used in training, as six columns allowed rather large objects to pass between the columns.

The robot was then run in the same environment(s) it was trained in. Videos of this online validation can be found at [www.metahuman.org/martin](http://www.metahuman.org/martin). These routes included views of parts of the environment never seen during training. Obstacles were present that were not present during training, such as chairs, trash cans and people. It avoided these obstacles well. Most errors were also conservative, i.e., declaring that a pixel was an obstacle when it was, in fact,

ground.

Most errors were also transient, lasting only a frame or two even when the camera was stationary. As mentioned previously, the exception was the red carpet in the NSH data set, which was handled poorly even on the training data. It was classified as an obstacle when nearby, causing the navigation to consistently treat it as a wall. In 5 validation runs, the mean distance and time to collision for the FMRB individual was 133 meters, 20:01 minutes, and the longest was 260 meters, 39:32 minutes. These are comparable to the published performance of the hand-written systems.

Not surprisingly, the vision subsystem from the combined data set performed a little worse in each environment than the vision system developed from data only in that space. It did not have a new class of error, but instead was more likely to make one of the errors made by the other vision subsystems. In the FMRB hallway, its average distance and time to collision were 72 meters and 10 minutes respectively.

### *Structure of Best Evolved Individuals*

The best individual from the FMRB condition was simplified. The resulting program was quite readable, see Figure 3. The two trees used two very different techniques. The second tree, whose  $a$  register holds the obstacle location, was essentially a decision tree, although one of the decisions has a non-linear boundary in the space spanned by two image operators. The other tree, which detected obstacles at the bottom of the image, was a recurrent mathematical function involving both a single image operator and the location of the column.

**First tree:**

$b$  (initial value) :=  $3.40282 \times 10^{38}$

Iterate-Up,  $3 \times 3$  window, in column *desired-x*:

$$b := \frac{b(1+h) - (1+h+h^2+h^4)\text{desired-}x}{h^5};$$

Where  $h$  is *moravec-horizontal*.

**Second tree:**

Iterate-Down,  $3 \times 3$  window, centered on the desired column:

**if**  $y \leq 9$  **then**

$a := y$ ;

**if**  $\text{iter2-}b > 9$  **then**

$a := y$ ;

**if**  $\text{median} > 35.4444$  **then**

{

**if**  $\text{gradient} > 413.96$  **then**

$a := y$ ;

**if**  $\text{gradient} > 239$  **and**  $(\text{gradient}/239)^4 > \text{moravec-slash}$  **then**

$a := y$ ;

}

Fig. 3. The best individual from all runs on the FMRB data set. This individual has been simplified without changing its behavior. The meanings of the various identifiers are listed in Table 2. The best individuals evolved on the other data sets had more complex structure, making greater use of memory.

The second tree uses the non-directional gradient magnitude as a primary indication of an object, although at intermediate values (between 239 and 414) it uses a non-linear comparison with a directional gradient. However, gradients are ignored when the median image brightness is less than 35. On closer inspection, this resolved a problem with the image digitizer which would round any pixel of brightness less than 32 to 16. This produced areas of brightness 16 next to pixels of brightness 33, creating artificial gradients. Ignoring areas darker than 35 eliminates these spurious gradients.

The other best-of-experiment individuals were similarly simplified. Although

all individuals were composed of two trees, most individuals ignored the results of the first tree. Thus, one tree would most likely have been sufficient. Evolution had exploited a number of techniques. The best NSH individual was a sequence of if-then conditions similar to a decision tree but involving non-linear combinations of up to five different image operators. In all cases, the bottom row of the image was handled using different code than the rest of the image. The mechanisms for this varied; on the FMRB data set it used two different trees for the two conditions, whereas in the NSH and combined experiments a multitude of `if-else` statements were used to run different code at different locations. Interestingly, the raw image statistic was never used, although the median filtered statistic was. All directional gradients of the image intensity were used except the vertical.

## 4 Discussion

This work explored a simple methodology for automatically discovering algorithms that exploit domain constraints. Genetic programming was used to search through a space of monocular vision algorithms. In all three environments, evolution produced algorithms which achieved over 85% fitness. Even on a new stream of images, all errors lasted only a single frame except for the red carpet in the NSH environment. These algorithms were then combined with a reactive obstacle avoidance algorithm originally developed for sonar. The combined system worked surprisingly well, generalizing to areas of the environment not seen during training. Performance was comparable to that of existing hand coded systems from the literature. The code used for this work may be downloaded from [www.metahuman.org/martin](http://www.metahuman.org/martin).

Evolution generated a number of different algorithm styles, from recurrent mathematical expressions to a decision tree with non-linear decision surfaces. In all cases, the bottom row of the image was handled using different code than the rest of the image. This division into two cases—boundaries between obstacles and floor, vs. obstacles that go to the bottom of the image—was discovered automatically. Nothing in the representation suggested the two different cases. Instead, this reflects a natural dichotomy in the images: at the very bottom of the image a program must detect the presence or absence of an obstacle, but in the middle of an image it must detect the *transition* between ground and non-ground, which is likely to have a significant gradient. This division was not noticed by any of the three groups mentioned in the introduction who created Polly style algorithms by hand, yet in retrospect it is obvious and natural.

This work is also the first example of the evolution of a complete vision system for a pre-existing environment. Earlier evolutionary work, such as that by Cliff et al. (1993); Smith (1998) used simple environments constructed especially for the robot, and as a result, could process at most 16 pixels.

This work can also be seen as automating the construction of Polly style algorithms. The design of Polly’s vision system required someone with knowledge of robotics and computer vision, whereas the current work only needs someone to mark obstacle location in an image with a simple GUI.

Genetic programming uses what amounts to a scripting language for a representation. This leads to very different results than general function approximation such as neural networks. For example, if a given input is largely unrelated to the output, genetic programming will often ignore the input altogether,

whereas a neural network will simply fit to the noise. In searching the space of expressions, genetic programming is using a representation that mathematicians and computer scientists feel is natural.

The vision sensor is cheap, a simple monocular,  $320 \times 240$  grayscale camera such as a Webcam. This contrasts sharply with stereo vision platforms such as the Triclops which cost thousands of dollars due to need for precision synchronization and calibration. In contrast, our calibration could be performed by any teenager with a protractor and yardstick.

This simplicity comes from exploiting domain knowledge. Computer vision is currently focused on domain independent methods, yet these have limits. This work provides a way out of the tension between domain dependence and independence, by proposing a general, domain independent tool which analyzes a domain to discover which constraints to exploit.

## References

- R.C. Arkin, 1999. Behavior-Based Robotics, second ed. The MIT Press, Cambridge, MA.
- J.K. Bala, K. DeJong, J. Huang, H. Vafaie, and H. Wechsler, 1996. Visual routine for eye detection using hybrid genetic architectures. In Intl. Conf. on Pattern Recognition, Vienna, Austria.
- J.K. Bala, K. DeJong, J. Huang, H. Vafaie, and H. Wechsler, 1997. Using learning to facilitate the evolution of features for recognizing visual concepts. *Evolutionary Computation*, 4(3):297-312.
- K. Benson, 2000. Evolving automatic target detection algorithms that logically

- combine decision spaces. In Proc. of the 11th British Mach. Vision Conf., pp. 685-694, Bristol, UK, Sept. 2000.
- T. Camus, D. Coombs, M. Herman and H.T. Hong, 1999. Real-time Single-workstation Obstacle Avoidance Using Only Wide-field Flow Divergence. *J. Computer Vision Research*. 1(3).
- D. Cliff, P. Husbands and I. Harvey. *Evolving Visually Guided Robots*. In Proc. of SAB92, the Second Intl. Conf. on the Simulation of Adaptive Behaviour. MIT Press, 1993.
- R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons.
- M. Ebner, 1998. On the Evolution of Interest Operators using Genetic Programming. EuroGP98: the First European Workshop on Genetic Programming, Paris, France, R. Poli et al. (ed.), 14-15 April 1998, pp. 6-10.
- A. Guarda, C. Le Gal, and A. Lux. Evolving visual features and detectors. In Intl. Sympo. on Comp. Graphics, Image Proc. and Vision, Rio de Janeiro.
- C. Harris and B. Buxton, 1996. Evolving edge detectors. Technical Report Research Note RN/96/3, University College London, London, UK, Jan. 1996.
- C. Harris, 1997. An Investigation into the Application of Genetic Programming Techniques to Signal Analysis and Feature Detection. Ph.D. Dissertation, University College London, London, UK.
- I. Horswill, 1993. Specialization of Perceptual Processes. Ph.D. Dissertation, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, MA May 1993.
- I. Horswill, 1994. Visual collision avoidance by segmentation. In IROS '94: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 902-909.
- M.P. Johnson, P. Maes and T. Darrell, 1994. Evolving Visual Routines, in

- Artificial Life IV, Proc. of the Fourth Intl. Workshop on the Synth. and Simul. of Living Systems., R. Brooks and P. Maes (eds.), Bradford Books.
- J. R. Koza, 1992. Genetic Programming. MIT Press.
- J.R. Koza, 1994. Genetic Programming II. MIT Press.
- D. Marr and K. Nishihara, 1978. Visual Information Processing: Artificial Intelligence and the Sensorium of Sight. *Technology Review*, 81(1), pp. 2–23.
- M. C. Martin, 2001. The Simulated Evolution of Robot Perception, Carnegie Mellon University, Robotics Institute, Ph.D. Dissertation.
- M.C. Martin, 2002. Genetic Programming for Robot Vision. 7th Int'l. Conf. on the Simulation of Adaptive Behaviour, Edinburgh.
- L. Lorigo, 1996. Visually-Guided Obstacle Avoidance in Unstructured Environments. MIT Master's thesis, Artificial Intelligence Laboratory.
- T.L. Ooi, B. Wu and Z. J. He, 2001. Distance determined by the angular declination below the horizon. *Nature*, vol. 414, 8 Nov. 2001, pp. 197-200
- P.J. Sobey, 1994. Active Navigation with a Monocular Robot. *Biological Cybernetics*, 71(5), pp. 433–440. Springer-Verlag.
- T. M. C. Smith, 1998. Blurred vision: Simlation-reality transfer of a visually guided robot. In *EvoRobot98: Proc. of the 1st European Workshop on Evolutionary Robotics*, Paris, France, April 1998.
- S. Thrun, 2002. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel eds., *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann.
- A. Teller and M. Veloso, 1997. PADO: A new learning architecture for object recognition. In *Symbolic Visual Learning*, Oxford Press, 1997, pp. 77-112.
- I. Ulrich and Illah Nourbakhsh, 2000. Appearance-Based Obstacle Detection with Monocular Color Vision. In *Proc. AAAI National Conference on Arti-*

ficial Intelligence, Austin, TX, July/August 2000.

M. Xie, 1996. Matching free stereo vision for detecting obstacles on a ground plane. *Machine Vision and Applications*, 9(1), pp. 9–13. Springer-Verlag.