

Active Exploration and Keypoint Clustering for Object Recognition

Gert Kootstra, Jelmer Ypma and Bart de Boer

Abstract— Object recognition is a challenging problem for artificial systems. This is especially true for objects that are placed in cluttered and uncontrolled environments. To challenge this problem, we discuss an active approach to object recognition. Instead of passively observing objects, we use a robot to actively explore the objects. This enables the system to learn objects from different viewpoints and to actively select viewpoints for optimal recognition. Active vision furthermore simplifies the segmentation of the object from its background. As the basis for object recognition we use the Scale Invariant Feature Transform (SIFT). SIFT has been a successful method for image representation. However, a known drawback of SIFT is that the computational complexity of the algorithm increases with the number of keypoints. We discuss a growing-when-required (GWR) network for efficient clustering of the keypoints. The results show successful learning of 3D objects in real-world environments. The active approach is successful in separating the object from its cluttered background, and the active selection of viewpoint further increases the performance. Moreover, the GWR-network strongly reduces the number of keypoints.

I. INTRODUCTION

THE real world poses many challenging problems for artificial systems. Consider for instance the problem of recognizing objects in the real world. Many object recognition systems that are successful in controlled laboratory environments have problems with the uncontrolled and unpredictable properties of the real world. Whereas, for instance, illumination and background can be controlled in an artificial setting, this is not true for real-world environments. Natural systems deal with these problems by using active perception [8]. Instead of passively observing an object, many animals, including humans, explore the object to control the visual input (see fig. 1). The use of active perception is also very important for artificial systems [1, 21]. In this paper, we discuss an active approach to 3D object recognition in the real world by an autonomous robot. By actively changing viewpoint, the robot observes an object from different angles, making it possible to learn to recognize the object from any given viewpoint. Moreover, the system selects the viewpoint that is expected to be most informative for recognition. Fur-



Fig. 1. A 18 months old child exploring an unknown object. This enables the child to observe the object from different viewpoints. In the meanwhile, it makes it possible to discriminate the object from the background.

thermore, exploration of the object makes it possible to separate the object from its background, something that is non-trivial when passively observing an object on a highly cluttered background [16].

Like many current approaches to object recognition, our model describes the objects by a set of local interest points [9, 12, 24]. Description in terms of local interest points has the advantage that the representation is more robust to occlusions, clutter and noise. It is also less sensitive to changes in viewpoint. In our method we use the Scale Invariant Feature Transform (SIFT) for the detection and description of interest points [14]. Our approach is, however, not restricted to SIFT, but can also be used with other local image detectors and descriptors. In the rest of the paper, we use the term *keypoints* to refer to points of interest detected by SIFT.

Interest points have been successfully used for three-dimensional (3D) object recognition [4, 13, 18, 22]. These studies have demonstrated the ability to learn to recognize objects from multiple viewpoints and subsequently recognize these objects in cluttered scenes. However, learning in these studies takes place in well-controlled environments: the object is usually put on a turntable which carefully rotates the object, while taking pictures of the object with fixed lighting conditions – with the exception of [18] – and against a uniform background. This setup reduces the amount of noise and uncertainty and makes it trivial to separate the foreground from the background. It is therefore not representative for real-world environments. A real-world environment is usually highly uncertain and cluttered with many distracting features. In this paper, we present a method to learn objects in uncontrolled real-world environments, using active vision. We use a mobile robot to actively explore the objects and their environment.

The use of active vision to simplify perceptual tasks has been advocated by Ballard, who referred to it with the term *animate vision* [1]. In our approach, we make use of active

Manuscript received September 14, 2007.

Gert Kootstra is with the Artificial Intelligence Institute of the University of Groningen, Grote Kruisstraat 2/1, 9712 TS, The Netherlands (corresponding author, phone: +31 (0)50 363 6502, fax: +31 (0)50 363 6687, email: G.Kootstra@ai.rug.nl)

Jelmer Ypma was with the Artificial Intelligence Institute of the University of Groningen, The Netherlands (email: jelmer@ai.rug.nl).

Bart de Boer was with the Artificial Intelligence Institute of the University of Groningen and is currently with the Institute of Phonetic Sciences, University of Amsterdam, The Netherlands (email: B.G.deBoer@uva.nl).

vision in multiple ways. Firstly, we use it to separate the object from its background, similar to [5, 16]. We use a method that can be described as *what-moves-together-belongs-together*. The robot observes the object while circling around it, a behavior that is comparable to rotating an object in your hand (see fig. 1). While doing this, interest points belonging to the object will show little displacement on the camera image, since the object is near the center of rotation. Interest points in the background, on the contrary, show relatively large displacements, with the possible exception of points on the floor close to the object. The amount of displacement of an interest point is used to classify whether it belongs to the object or to the background.

Secondly, we use active vision to find stable interest points. By changing viewpoint, we can actively test whether an interest point is recognizable from nearby viewpoints. If so, we can classify the interest point as stable. This process will filter out points that are sensitive to rotation, translation and other affine transformations. This does away with the necessity to use affine invariant interest point detectors (e.g., [17]), which are not only computationally expensive, but have also been shown to perform worse on recognizing non-planar 3D objects than SIFT [18]. A similar approach to ours was taken in [11], where a behavior, inspired by insects, was adopted to find reliable visual landmarks.

Finally, we use active vision to gather more evidence for recognition. This is especially important in ambiguous situations. If from one viewpoint it is not possible to recognize an object, a more promising viewpoint can be selected. Although, as a human observer, we might have the impression that ambiguous situations are quite rare, we must remember that ambiguity strongly depends on the quality of the sensory system, as can be seen in [19]. In 3D object recognition, gathering more evidence improves recognition [23]. Some viewpoints will be more informative than others. We therefore propose a probabilistic method to select the viewpoint that is expected to be most informative as the next viewpoint. Viewpoint selection is also used by [3, 20], with the difference that we use a one-shot learning method and perform recognition in a real-world environment.

In addition to the use of active vision, we propose a method to reduce the number of keypoints in the keypoint database. One of the reasons that SIFT is so successful in object recognition is that it uses a large number of keypoints to represent one object [14]. This makes the system very tolerant to noise, and solves the problem of occlusions. There is, however, an important drawback, namely that a significant amount of computation in the recognition process is devoted to matching the observed keypoints with the keypoint database. Nearest-neighbor search methods like *kd*-tree search [6] that are efficient in low-dimensional spaces, do not do better than exhaustive search in the high-dimensional space of the SIFT features. An improvement in computation time can be achieved by an approximate best-bin-first method [2]. But even then, the computation time increases

with the number of stored keypoints, while the success in finding the nearest neighbor decreases. It is therefore very useful for 3D object recognition to reduce the number of stored keypoints in an efficient way.

In this paper, we use a growing-when-required (GWR) network [15] for efficient clustering of keypoints. When performing 3D object recognition, many of the acquired keypoints look very similar. There are several reasons for this. First of all, these are keypoints belonging to the same point on the object seen from different angles. Secondly, there are similarly looking points on repeating structures on the same object, and finally, we see ambiguous keypoints on different objects. Our GWR-network clusters these similar keypoints to attain efficient database use.

II. METHODS

We use the SIFT detection and description, as described in [14], as the basis for the 3D object recognition. We use a prior smoothing of each octave of $\sigma=1.2$ instead of $\sigma=1.6$ as proposed in [14], since this yields better performance in our experiments. Our method for matching the observed keypoints with the database is somewhat different. First, we focus solely on the individual matching of keypoints, and therefore do not use the geometric matching of sets of keypoints as used in [14]. And second, we use a threshold on the distance to the nearest neighbor, instead of a best to second-best ratio to determine a match, since this yielded better performance in our experiments. Details on our matching and recognition process are described further on.

In this section we will first discuss our active approach to object recognition, then our method to select the next viewpoint, and we will end with the method to cluster the SIFT keypoints.

A. Active Vision

One of the contributions of our study to improve 3D object recognition in real-world environments is the use of active vision. We make use of a mobile robot that explores objects by circling around them, observing them every 10 degrees. By actively changing viewpoint, the robot gathers new information that we use in two different ways: to detect stable keypoints and classify them as object or background, and to explore the object in order to gather more evidence to resolve ambiguous situations. Both methods are described in the following paragraphs.

A keypoint is considered stable if it is originally observed at an angle of θ degrees, and subsequently matched in the previous or next image, at $\theta \pm 10$ degrees. A keypoint \mathbf{k}_i is matched to its nearest neighbor in the previous image, \mathbf{k}_n , if the Euclidean distance between both is less than 0.6, where \mathbf{k} is the 128 dimensional feature vector of the keypoint. This filters out all keypoints that are only recognizable from one specific angle.



Fig. 2 A number of images from our database.

In the next step, we segment the stable keypoints belonging to the background from those belonging to the object. Each keypoint \mathbf{k}_i has a position (x_i, y_i) at which it is observed in the image. Since the object is in the center of rotation, object keypoints will move little when the robot is exploring the object, whereas the displacement will be relatively large for keypoints in the background. Furthermore, since the robot moves on a flat surface, keypoints will only move in the horizontal direction. Allowing some fluctuations, we classify a stable keypoint as an object point when

$$(|x_i - x_n| < x_T) \wedge (|y_i - y_n| < y_T) \quad (1)$$

where we use $x_T = 12$ and $y_T = 4$. Otherwise, the stable keypoint is classified as background. The successful use of this simple classification model nicely illustrates the power of active vision to simplify perceptual tasks¹. The robot explores the objects from 36 different angles and stores the stable object keypoints along with the object ID and pose. Doing so, the appearances of objects in a cluttered environment are learned.

Once the object database is in place, objects can be recognized. Based on the set of observed keypoints, \mathcal{O} , and the keypoint database, \mathcal{D} , we determine the activation of every model, $m_{ID,\theta}$, for object ID , and pose θ . The activation of a model is based on the set of observations, $\mathcal{S}_{ID,\theta} \subseteq \mathcal{O}$, that support the model.

$$\mathcal{S}_{ID,\theta} = \bigcup (\mathbf{p}_i \in \mathcal{O} \mid o_n = ID \wedge \alpha_n = \theta)$$

where o_n and α_n are respectively the object ID and pose of the nearest neighbor, \mathbf{k}_n , of \mathbf{p}_i in the keypoint database. Every supporting observation p_i in \mathcal{M} gives an activation a_i of

$$a_i = \exp(-|\mathbf{p}_i - \mathbf{k}_n|) \quad (2)$$

The total activation of model $m_{ID,\theta}$ given the observed keypoints, \mathcal{O} , and the keypoint database, \mathcal{D} , is given by

$$A_{ID,\theta}^\delta = \frac{\sum_{i \in \mathcal{S}_{ID,\theta}} a_i}{\sqrt{|\mathcal{D}_{ID,\theta}|}} \quad (3)$$

where $|\mathcal{D}_{ID,\theta}|$ is the number of keypoints in the keypoint database that are associated with object ID in pose θ and δ is the current viewpoint. Equation (3) gives the activation of a specific pose of an object. This activation increases with the number of supporting observations relative to the number of keypoints in the database associated with that object/pose. This makes that fewer matched observations are needed for objects that have relatively few interest points. However, the square root in the denominator causes the probability to increase with the number of object keypoints given the same ratio of matched observations to database keypoints. This reflects the idea that there is more confidence when there are more object keypoints.

Finally, the robot can actively gather more evidence for recognition. By rotating around the object, an action that is equivalent to the rotation of an object in your hand, the robot gathers more information about the object under consideration by viewing it from different angles. When driving around the object, we accumulate the evidence by

$$A_{ID,\theta}(t) = \sum_{\delta \in E} A_{ID,\theta}^\delta \quad (4)$$

where $A_{ID,\theta}(t)$ is the accumulated activation for object ID and pose θ at time t and $E = \{\varphi_0, \dots, \varphi_i\}$ is the set of viewpoints from where the observations are made. The change of viewpoint helps to disambiguate object and is therefore expected to result in more robust recognition of 3D objects. In the next section, we will discuss how the next viewpoint is selected by our model.

Finally, the activations of all poses for a given object together give the activation of an object model

$$A_{ID}(t) = \sum_{\theta \in \Theta} A_{ID,\theta}(t) \quad (5)$$

B. Next Viewpoint Selection

We use the active capabilities of the robot to explore the objects and gather more information from different viewpoints. In order to select the next viewpoint, we use a probabilistic approach. In this approach, the next viewpoint φ_{t+1} is the angle from where we expect the maximum activation of an objects-pose model. I.e.,

$$\varphi_{t+1} = \arg \max_{\gamma \in \Theta - \Phi} E(A_{ID,\theta}(t+1)) \quad (6)$$

where Θ is the set of all possible viewpoints, and $\Phi = \{\varphi_0, \dots, \varphi_{t-1}\}$ is the set of all previous viewpoints. The expected activation of the object-pose model when viewed from viewpoint γ at time $t+1$ is given by

$$E(A_{ID,\theta}(t+1)) = A_{ID,\theta}(t) + E(A_{ID,\theta}^\gamma \mid O_{ID,\theta}) P(O_{ID,\theta}) \quad (7)$$

$$E(A_{ID,\theta}^\gamma \mid O_{ID,\theta}) = \sqrt{|\mathcal{D}_{ID,\theta+\gamma}|} \quad (8)$$

¹ Slightly better results could be obtained if one takes advantage of the fact that keypoints in the background move in the same direction as the robot, whereas visible object keypoints are in front of the center of rotation, and therefore move in the opposite direction.

$$P(O_{ID,\theta}) = \frac{A_{ID,\theta}(t)}{\sum_{i=0}^N A_{o_i,\alpha_i}(t)} \quad (9)$$

In words, the expected new activation of the model is based on the old activation. This value is increased with the expected extra activation gained from the new viewpoint given that we are looking at object ID and pose θ , $E(A_{ID,\theta}^y | O_{ID,\theta})$, multiplied by the probability that we are actually looking at object ID and pose θ , $P(O_{ID,\theta})$. Equation (8) can be inferred from eq. (3), when we assume that we observe *all* keypoints belonging to object ID at pose $\theta + \gamma$. And the probability $P(O_{ID,\theta})$ is the activation of the model divided by the total activation of all object-pose models. By selecting the viewpoint that optimizes the expected activation of one of the object-pose models, we select the most informative viewpoint as the next.

C. Keypoint Clustering

As explained in the introduction, 3D object recognition with SIFT has the main disadvantage that the computational time needed increases with the number of keypoints stored in the database. We therefore use a growing-when-required (GWR) network [15] to efficiently cluster keypoints that are highly similar. A GWR-network is a clustering method, very similar to a growing-neural-gas (GNG) network [7]. Both networks are based on Kohonen's self-organizing maps (SOM) [10]. A SOM is an efficient method to cluster high-dimensional data. The disadvantage, however, is that the number of clusters needs to be set in advance. This makes a SOM highly inappropriate for object recognition with SIFT, since the number of clusters depends on the number of unique keypoints. A GNG-network is an adaptation of a SOM which can dynamically change the number of nodes, i.e., clusters, in the network. However, the drawback of a GNG-network is that new nodes are only added after a number of inputs. This is not desirable for object recognition, since we would like to add a node in the network when we observe a completely new keypoint. A GWR-network does just that, it adds nodes if this is required.

The GWR-network as described in [15] uses edges between nodes. This is based upon the SOM, and results in a topological preservation of the network in the sense that connected nodes in the network correspond to neighboring points in the input space. Usually, the edges in a GWR network are used in learning to move the neighboring nodes of the winning node closer to the presented input. This is undesirable for object learning, since the presentation of an input not only changes the representation of the corresponding keypoint, but also of neighboring keypoints. In the end, this will result in changing keypoints so much that they are not recognizable anymore. We therefore omitted the edges from the GWR-network.

For the description of our implementation of the GWR-

network, we follow the notation and description in [15].

Let K be the set of observed keypoints when learning the objects, A be the set of nodes in the network, \mathbf{w}_n be the weight vector of node n (of the same dimensionality as the SIFT keypoints), and t_n be the activation counter. Furthermore, each node holds a record, R_n , of all associated objects and poses. We initialize the network with $A = \{n_1\}$, where the weight vector of n_1 is initialized a randomly picked keypoint from K , and $t_1 = 0$. Then, for each keypoint \mathbf{k} from K we do:

1. \mathbf{k} and the object ID and pose θ to which the keypoint belongs are input to the network.
2. Select the best matching node $s \in A$, such that

$$s = \operatorname{argmax}_{n \in A} |\mathbf{k} - \mathbf{w}_n|$$
3. Calculate the activity of the winning node

$$a_s = \exp(-|\mathbf{k} - \mathbf{w}_s|)$$
4. Calculate the firing counter h_s

$$h_s = 1 - (1 - e^{-\alpha_b t_s / \tau}) / \alpha_n$$
 where $\alpha_b = 1.05$, $\alpha_n = 1.05$ and $\tau = 3.33$
5. if $(a_s < a_T) \wedge (h_s < h_T)$, add a new node r
 - $A = A \cup \{r\}$
 - $\mathbf{w}_r = \mathbf{k}$
 - $R_r = \{\langle ID, \theta \rangle\}$
 where $a_T = 0.8$ and $h_T = 0.4$.
6. Else, adapt the weights of the winning node
 - $\mathbf{w}_s = \mathbf{w}_s + \eta \cdot h_s \cdot (\mathbf{k} - \mathbf{w}_s)$
 - $R_s = R_s \cup \{\langle ID, \theta \rangle\}$
 where $\eta = 0.05$.
7. $t_s = t_s + 1$

When the presented keypoint is sufficiently similar to the winning node, it is clustered with that node, and the description of the node is altered to better represent all associated keypoints. If, on the other hand, the presented keypoint differs from the existing nodes, and the firing counter of the nearest node is below the threshold h_T , the presented keypoint is added as a new node. In this way, the GWR-network clusters similar keypoint, thus creating a smaller database for recognition.

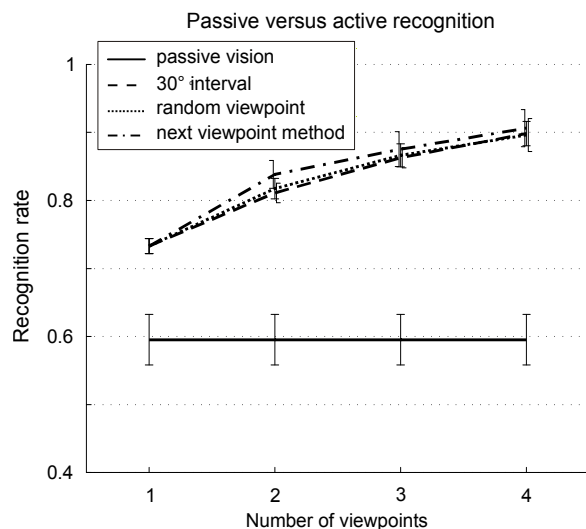


Fig. 3. The recognition rates for the passive approach and three active methods: a fixed 30° interval, random viewpoint selection and our next viewpoint method. The active methods are plotted as a function of the number of viewpoints used to explore the object. The plot of the passive method is flat, since it does not use exploration. The error bars give the 95% confidence intervals.

A record is kept of all objects and poses that correspond to the nodes in the network. This allows for supporting all objects containing similarly looking keypoints when such a keypoint is observed. This is in contrast with [14], where important evidence is discarded by choosing only keypoints that match uniquely with one object.

III. RESULTS

We used seven objects placed in a cluttered environment for our image database (see fig. 2). A mobile robot equipped with a CCD camera was used to take images from 36 different viewpoints around the objects. The image database consists of four different sets, where the orientation of the objects is respectively 0°, 90°, 180° and 270° with respect to the environment, resulting in a different background for the objects. In the experiments, training was done on one single set, while the other three sets were used to test the performance. This resulted in 12 different cross-validation tests.

In our first experiment, we tested the performance of our active approach to 3D object recognition and compared it with a passive approach that does not use robust keypoints filtering and multiple viewpoints. We furthermore compared the performance of our next viewpoint selection method to an approach where the next viewpoint is a simple 30° interval, as well as to random viewpoint selection. Fig. 3 shows the mean recognition rates over the 12 tests. For the active approaches, the performance is plotted as a function of the amount of evidence. Since the passive approach does not accumulate evidence, it is drawn as a horizontal line. The error bars show the 95% confidence intervals.

The active approaches clearly outperform the passive approach. Already with one viewpoint, the use of active vision to select stable object keypoints gives significantly better



Fig. 4 The mean recognition rates for standard SIFT, the GWR-network and standard SIFT using 36% of the keypoint database. Standard SIFT uses on average 6429 keypoints. The GWR-network has on average 2396 keypoint clusters (37% of standard SIFT). For the GWR-network and SIFT using 36% keypoints, the data is acquired from 10 experimental runs. The error bars give the 95% confidence intervals.

performance than passively considering all visible keypoints with respectively 73% and 60% success (t-test: p value $< 10^{-4}$). With increasing accumulation of evidence, the recognition rate rises from 73% to about 90% (t-test: p value $\ll 10^{-4}$). Our next viewpoint selection method shows a systematically better performance than the fixed interval and random viewpoint methods. Although the performance is not significantly better for the third and fourth viewpoint, it is significantly better for the second viewpoint with p values of 0.03 and 0.02 for the comparison with respectively the interval method and the random selection method. A significant increase in performance for the second viewpoint is important since this enables the system to recognize objects with fewer observations.

In the second experiment we compared the performance of the GWR-network with the standard SIFT method, both using active vision with a fixed interval of 30°. The learned keypoints are presented to the GWR-network in random order. We therefore performed ten different experimental runs to test the performance of the GWR-network. The standard SIFT method has on average 6429 keypoints in the database. The GWR-network has 2396 keypoint clusters (37% of standard SIFT), making it 2.7 times faster than standard SIFT. We also compared the GWR-network with standard SIFT that used only 36% of the keypoints, approximately the same amount of keypoints as the GWR-network used. These keypoints are selected randomly from the keypoint database. Again we performed 10 experimental runs. Fig. 4 shows the result of this experiment. We see that the standard SIFT method performs best. The GWR-network performs slightly (but significantly) worse. On the other hand, the GWR-network performs significantly better than standard SIFT using approximately the same number of keypoints. This

shows that the GWR-network effectively clusters the keypoints, using only 37% of keypoints with a minimum loss of performance.

IV. DISCUSSION

Our experiments show the successful use of object exploration for 3D object recognition. Exploration is used in different ways, (1) to detect stable keypoints, (2) to segment the object from the background, and (3) to select informative viewpoints. The active vision approach performs significantly better than when passively observing. The problem of the passive method is that it not only learns to associate the keypoints of the object itself with the object, but also the background keypoints. By selecting the next viewpoint based on the optimization of the expected activation of object models, the system highly increases its recognition performance with subsequent viewpoints.

Our second experiment tested the use of a GWR-network for clustering keypoints. Although the performance of the GWR-network in terms of recognition rates is slightly less than standard SIFT, it is computationally much more efficient. The GWR-network uses 2.7 times fewer keypoints, while performing significantly better than SIFT using the same number of keypoints. Reducing the amount of keypoints is important for object recognition using SIFT, especially with a growing number of objects. Our results using the GWR-network are promising, although there is room for improvements.

One of the problems with the GWR-network is that keypoints that are presented early to the network are badly represented in the final network, in contrast to keypoints that are presented later. This might account for the fact that the GWR-network performs worse than standard SIFT using the complete keypoint database. The keypoints presented first are simply lost. However, the fact that the network performs significantly better than SIFT using the same number of keypoints, shows that it is capable to effectively cluster the keypoints. Further research needs to be done to deal with the mentioned problem.

In this study, we did not use additional methods to improve the recognition rate. A good way to boost recognition is to use a geometric fit between sets of keypoints, for instance the geometric verification method described in [14]. This method can be used both with our active vision method and with the GWR-network. We expect a similar increase in performance for our methods.

Summarizing, we showed the successful use of active vision to simplify complex recognition tasks. We furthermore, demonstrated the possibility to reduce the number of keypoints for SIFT by using a GWR-network. Both methods make implementing object recognition in the real world more feasible.

REFERENCES

- [1] D. H. Ballard, Animate Vision, *Artificial Intelligence*, vol. 48, pp. 57-86, 1991.
- [2] J. Beis and D. G. Lowe, Shape indexing using approximate nearest-neighbour search in high-dimensional spaces, presented at Conference on Computer Vision and Pattern Recognition, Puerto Rico, 1997.
- [3] H. Borotschnig, L. Paletta, M. Prantl and A. Pinz, Appearance-based active object recognition, *Image and Vision Computing*, vol. 18, pp. 715-727, 2000.
- [4] V. Ferrari, T. Tuytelaars and L. v. Gool, Simultaneous Object Recognition and Segmentation from Single or Multiple Model Views, *International Journal of Computer Vision*, vol. 67, pp. 159-188, 2006.
- [5] P. Fitzpatrick, First Contact: an active vision approach to segmentation, presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, Nevada, 2003.
- [6] J. H. Friedman, J. L. Bentley and R. A. Finkel, An algorithm for finding best matches in logarithmic expected time, *ACM Transactions on Mathematical Software*, vol. 3, pp. 209-226, 1977.
- [7] B. Fritzsche, A Growing Neural Gas Network Learns Topologies, presented at Advances in Neural Information Processing Systems (NIPS'94), Denver, 1995.
- [8] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin, 1979.
- [9] C. Harris and M. Stephens, A Combined Corner and Edge Detector, presented at The Fourth Alvey Vision Conference, Manchester, UK, 1988.
- [10] T. Kohonen, The self-organizing map, *Proceedings of the IEEE*, vol. 78, pp. 1464-1480, 1990.
- [11] M. Lehrer and G. Bianco, The turn-back-and-look behaviour: bee versus robot, *Biological Cybernetics*, vol. 83, pp. 211-229, 2000.
- [12] D. G. Lowe, Object recognition from local scale-invariant features, presented at International Conference on Computer Vision, Corfu, Greece, 1999.
- [13] D. G. Lowe, Local Feature View Clustering for 3D Object Recognition, presented at IEEE Conference on Computer Vision and Pattern Recognition, Kauai, Hawaii, 2001.
- [14] D. G. Lowe, Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision*, vol. 60, pp. 91-110, 2004.
- [15] S. Marsland, J. Shapiro and U. Nehmzow, A self-organising network that grows when required, *Neural Networks*, vol. 15, pp. 1041-1058, 2002.
- [16] G. Metta and P. Fitzpatrick, Early integration of vision and manipulation, *Adaptive Behavior*, vol. 11, pp. 109-128, 2003.
- [17] K. Mikolajczyk and C. Schmid, An affine invariant interest point detector, presented at the 7th European Conference on Computer Vision, Copenhagen, Denmark, 2002.
- [18] P. Moreels and P. Perona, Evaluation of Features Detectors and Descriptors based on 3D Objects, *International Journal of Computer Vision*, vol. 73, pp. 263-284, 2007.
- [19] S. Nolvi, Adaptation as a more powerful tool than decomposition and integration, presented at the Workshop on Evolutionary Computing and Machine Learning, 13th International Conference on Machine Learning, Bari, Italy, 1996.
- [20] L. Paletta and A. Pinz, Active object recognition by view integration and reinforcement learning, *Robotics and Autonomous Systems*, vol. 31, pp. 71-86, 2000.
- [21] R. Pfeifer and C. Scheier, *Understanding Intelligence*: MIT Press, 1999.
- [22] F. Rothganger, S. Lazebnik, C. Schmid and J. Ponce, 3D Object Modeling and Recognition Using Local Affine-Invariant Image Descriptors and Multi-View Spatial Constraints, *International Journal of Computer Vision*, vol. 66, pp. 231-259, 2006.
- [23] S. D. Roy, S. Chaudhury and S. Banerjee, Active recognition through next view planning: a survey, *Pattern Recognition*, vol. 37, pp. 429-446, 2004.
- [24] C. Schmid and R. Mohr, Local greyvalue invariants for image retrieval, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, pp. 530-535, 1997.