

A HYBRID LEARNING SYSTEM FOR OBJECT RECOGNITION

Klaus Häming and Gabriele Peters

Lehrgebiet Mensch-Computer-Interaktion, FernUniversität in Hagen, Universitätsstr. 1, Hagen, Germany

Keywords: Machine learning, Reinforcement learning, Belief revision, Object recognition.

Abstract: We propose a hybrid learning system which combines two different theories of learning, namely implicit and explicit learning. They are realized by the machine learning methods of reinforcement learning and belief revision, respectively. The resulting system can be regarded as an autonomous agent which is able to learn from past experiences as well as to acquire new knowledge from its environment. We apply this agent in an object recognition task, where it learns how to recognize a 3D object despite the fact that a very similar, alternative object exists. The agent scans the viewing sphere of an object and learns how to access such a view that allows for the discrimination. We present first experiments which indicate the general applicability of the proposed hybrid learning scheme to this object recognition tasks.

1 INTRODUCTION

We already proposed a similar learning system for object recognition and object reconstruction, which was based on the reinforcement learning (Sutton and Barto, 1998) component only (Peters, 2006). Figure 1 shows a diagram representing the general idea of this earlier proposed system. In (Peters, 2006) we discussed an application of this architecture to the problem of creating a sparse, view-based object representation.

Our new hybrid learning system is based on the same system. But now we extend it by introducing a belief revision component and apply it to object recognition rather than to reconstruction.

A special property of the object recognition in question is, that it aims to work on arbitrary 3D-objects. It also aims to allow the discrimination of rather similar objects which differ in a subtle detail only. The latter property introduces the additional problem of identifying a distinguishing object part and also coping with situations in which these parts are not visible. A straight forward approach to solve this is to simply scan the whole object in a regular pattern until a suitable view is found. However, the goal is to keep the system away from collecting unnecessary data, after all.

An additional constraint we want to impose is to create a solution which is strictly view-based, without the need of additional (i.e., 3D) information. In this work, we therefore use a feature-based approach.

2 RELATED WORK

Before we introduce our experimental set-up, we want to point out differences of this approach to similar work on object recognition.

First, we do not train our system to detect a particular object or feature, as for example (Viola and Jones, 2001) do for face detection and (Gordon and Lowe, 2006) for the more general case of finding an arbitrary object in an image. We want our system to detect characteristic differences between objects as well as cases in which these differences are absent.

In (Schiele and Crowley, 1998), for every object the most distinctive view has been determined beforehand. The recognition routine used this knowledge to jump to these positions for the object detection. This is not a strictly view-based approach, since a global co-ordinate system is necessary to allow for the new positions to be assumed. We want to avoid a reliance on such supplementary information.

Apart from feature-based approaches, eigenspaces are common approaches, as described in (Murase and Nayar, 1995) and (Deinzer et al., 2006), where the latter focuses on integrating the cost of viewpoint selection into the learning framework, which we do not. Eigenspace approaches need whole images as input. We want to be able to develop the system towards robustness against changes in distance and partial occlusion. Hence, we prefer a feature-based approach, because it puts less constraints onto the image acquisition.

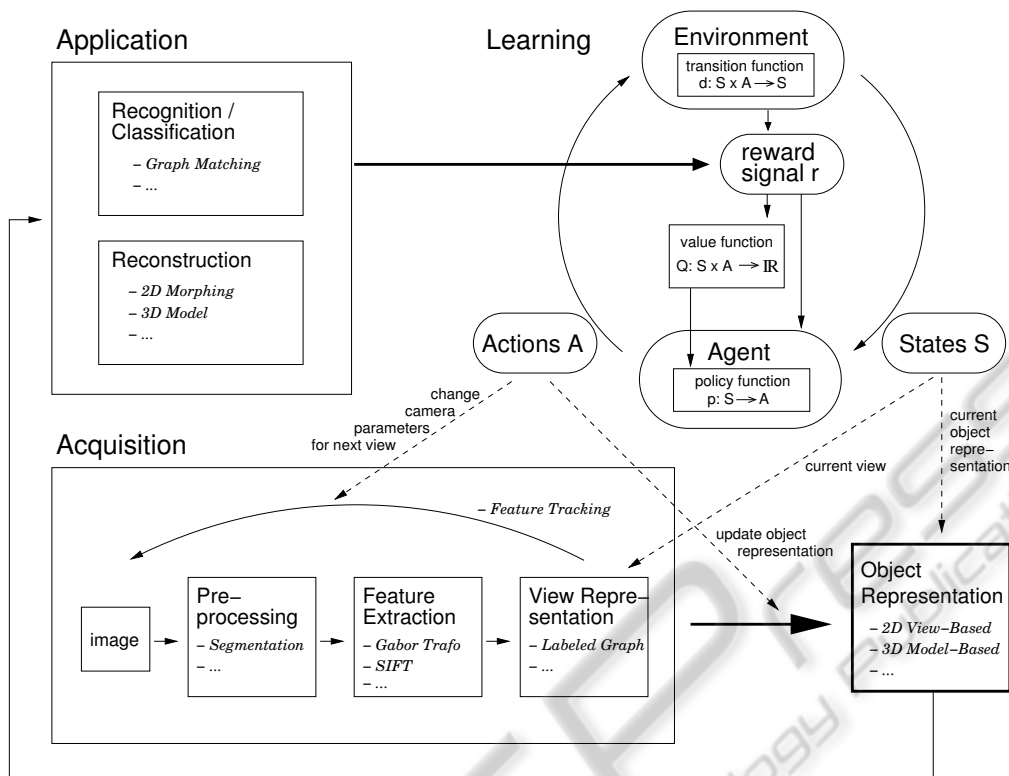


Figure 1: Earlier version of the learning system for computer vision. On the right upper part the reinforcement component is shown. On the left several application areas are listed, e.g., object recognition or object reconstruction. The system acquires images of objects (acquisition part) and incorporates information extracted from the images into the up-to-now learned object representation (shown in the lower right). With the current representation the application is performed. The result of the performance is given in form of a reward back to the learning component.

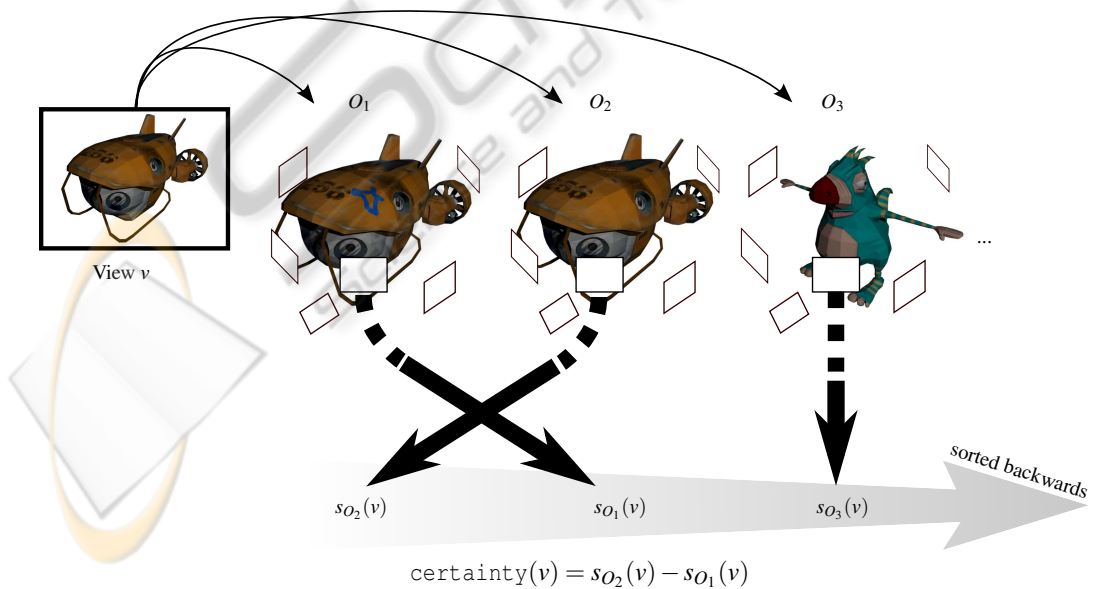


Figure 2: A certainty score derived from a comparison of a view against an object database. In this example, object O_2 is the most similar one to the given view v while O_1 is ranked second. The difference of their respective similarity score constitutes the certainty score. So, if the view is found to match O_1 and O_2 equally well, the certainty will be low and hence the current view rated as not being sufficiently discriminative between O_1 and O_2 . The two depicted submarines are examples of similar objects the agent learns to distinguish.




Figure 3: Extension of the earlier proposed reinforcement learning agent with the additional belief revision component. This induces two levels of learning. The lower level learning (or implicit learning) is represented by the reinforcement learning framework. The higher level learning (or explicit learning) is represented by the introduction of a ranking function (“OCF”) as a filter on the possible actions presented to the policy of the reinforcement learning framework. This way the agent is enabled to learn in rather large state spaces because the ranking function allows the recognition of states even if they are not exactly the same. The reward function r returns 100 whenever the agent reaches a goal state and 0 everywhere else. We use a ϵ -greedy policy with $\epsilon = 0.1$. The “decode”-module creates a symbolical state description from the state signal.

3 EXPERIMENTAL SET-UP

We simulate our object recognition experiments on images of 3D-models of objects. Some of these 3D-models are slightly modified copies of others. Figure 2 includes an example of such a 3D-model of an object and its modified version.

We provide the agent with a database of the objects against which it compares its visual input. Here, the visual input consists of features which are computed by the SURF feature detector (Bay et al., 2006). Based on a similarity score, this comparison yields a vote for one of the candidate objects together with a value representing the vote’s certainty. The basic principle behind the similarity score is to relate the number of *matched* features to the number of all detected features.

Whenever the agent perceives a view which does not allow the object’s recognition, because it can belong to a multitude of candidates, the certainty value will be low. On the other hand, if the current view allows the identification of the object, it will be high. Figure 2 details these ideas. The agent itself is a Q -learner (Sutton and Barto, 1998) that uses a two-level representation of its current belief as proposed in (Sun et al., 2001) and analogous to the one used in (Häming and Peters, 2010). The lower level of the learning system is implemented as a Q -table, while the higher level uses a ranking function (Spohn, 2009) to symbolically represent the agents perception. A schematic picture of this approach is shown in Figure 3.

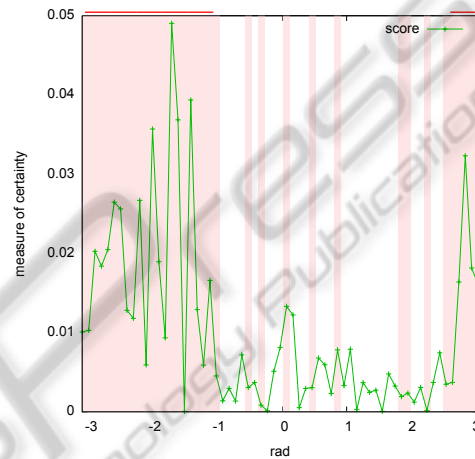
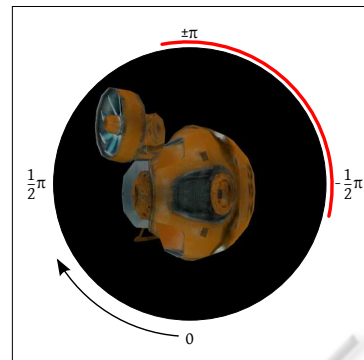


Figure 4: Example of how to extract a threshold to define a goal state. The diagram shows the course of the certainty score on a great circle around the object. The red markings show the positions from which the distinguishing modification on the submarine can be seen clearly enough to trigger a high certainty score. (The modification in this example is given by a blue, star-shaped mark on the surface of the submarine, which can be seen in Figure 2.) The background of the graph is colored light red where the agent was able to identify the correct model.

4 RESULTS

Assessing the certainty score while the agent surrounds an object on a great circle as depicted in the upper part of Figure 4, we can record the course of the certainty score along with the ability of the agent to identify the object as shown in the lower part of Figure 4. The identification of the object using a thus determined threshold defines the goal state. The agent is then given the task to learn how to find a position in relation to the object from which it can recognize it. The agent’s movements were restricted to a sphere around the object during these experiments. As it turns out, the agent is able to rapidly learn where to look at to identify an object, as Figure 5 reveals.

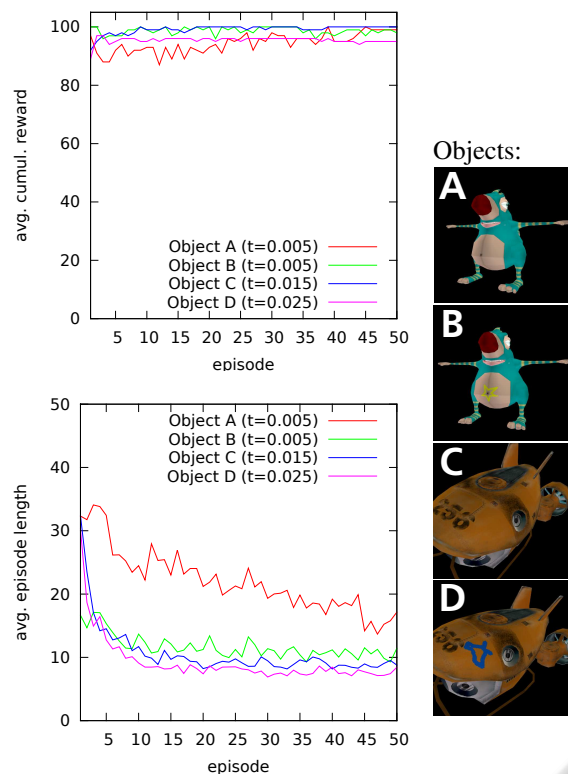


Figure 5: Example learning curves. The top graph shows the averaged cumulative rewards during the course of 50 episodes. The bottom graph shows the averaged number of steps it took the agent to reach the goal. The objects the agent had to recognize are shown on the right. Additional to the submarine models, two dragon models were used which differ in the presence or absence of a yellow star surrounding the belly button. The threshold values used for the certainty score are given in the diagram legends. The results were averaged over 100 runs.

5 CONCLUSIONS

We presented a hybrid learning system which consists of two different machine learning components, namely a reinforcement learning component and a belief revision component. This system was applied to an object recognition problem. We demonstrated in a first experiment, that the agent is able to learn how to access such views of an object that allow for a distinction of the object from a very similar but different object. As an indicator for this ability we regard the promising learning curves and decreasing episode lengths depicted in Figure 5. In the current state of development, our system, of course, still exhibits weaknesses. For example, the threshold-based goal state identification is not robust enough to be universally applicable. In particular, it turned out to depend on the distance of the camera to the object. Summariz-

ing, we have reason to assume the general applicability of our hybrid learning approach to object recognition tasks.

ACKNOWLEDGEMENTS

This research was funded by the German Research Association (DFG) under Grant PE 887/3-3.

REFERENCES

Bay, H., Tuytelaars, T., and Van Gool, L. (2006). Surf: Speeded up robust features. In *9th European Conference on Computer Vision*, Graz Austria.

Deinzer, F., Denzler, J., Derichs, C., and Niemann, H. (2006). Integrated viewpoint fusion and viewpoint selection for optimal object recognition. In Chanteler, M., Trucco, E., and Fisher, R., editors, *British Machine Vision Conference 2006*, pages 287–296, Malvern Worcs, UK. BMVA.

Gordon, I. and Lowe, D. G. (2006). What and where: 3d object recognition with accurate pose. In Ponce, J., Hebert, M., Schmid, C., and Zisserman, A., editors, *Toward Category-Level Object Recognition*. Springer-Verlag.

Hämig, K. and Peters, G. (2010). An alternative approach to the revision of ordinal conditional functions in the context of multi-valued logic. In *20th International Conference on Artificial Neural Networks*, Thessaloniki, Greece.

Murase, H. and Nayar, S. K. (1995). Visual learning and recognition of 3-d objects from appearance. *Int. J. Comput. Vision*, 14(1):5–24.

Peters, G. (2006). A Vision System for Interactive Object Learning. In *IEEE International Conference on Computer Vision Systems (ICVS 2006)*, New York, USA.

Schiele, B. and Crowley, J. L. (1998). Transinformation for active object recognition. In *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, page 249, Washington, DC, USA. IEEE Computer Society.

Spohn, W. (2009). A survey of ranking theory. In *Degrees of Belief*. Springer.

Sun, R., Merrill, E., and Peterson, T. (2001). From implicit skills to explicit knowledge: A bottom-up model of skill learning. In *Cognitive Science*, volume 25, pages 203–244.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press, Cambridge.

Viola, P. and Jones, M. (2001). Rapid object detection using a boosted cascade of simple features. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 1:511.