Learning Strategies to Select Point Cloud Descriptors for 3D Object Classification: A Proposal

J. Garstka and G. Peters

University of Hagen, Germany

Abstract

We propose a reinforcement learning approach for an adaptive selection and application of 3D point cloud feature descriptors for the purpose of 3D object classification. The result of the learning process is an autonomously learned strategy of selection of descriptors with the property that the successive application of these descriptors to a 3D point cloud yields high classification rates among a large number of object classes. The order in which the descriptors are applied to an unfamiliar point cloud depends on the features calculated in previous steps of the descriptor sequence, i.e., the sequence of descriptors depends on the object to be classified, thus it is highly adaptive. Our approach starts with a given number of descriptors and object classes, but it is able to adapt dynamically to changes in the environment. For example, further descriptors can be added during the learning process, and new object classes are created autonomously if necessary.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Object recognition I.2.6 [Artificial Intelligence]: Learning—Reinforcement Learning I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—3D/stereo scene analysis

1. Introduction

Due to a wide range of applications such as scene understanding, navigation or applications in robotics like grasping or scene manipulation, the classification and recognition of 3D point clouds has been a fundamental part of research in visual computing for the last few years. Additionally, the appearance of cheap 3D cameras like the Microsoft Kinect made these fields of application available to a broader public. Most of the current algorithms compare objects pairwise by matching the descriptions of whole objects or of local features. There are many state of the art 3D point cloud feature descriptors with different recognition rates and time complexities, e.g., [RBMB08,TSDS10,RC11,HPPLG11]. However, in general the computational costs of calculation and comparison for local feature descriptors are high. Furthermore, there are a few algorithms which describe point clouds globally, e.g., [Kei99, PR99, PRM*00, SKO00, VS02]. But global descriptors are too imprecise for an accurate assignment of a 3D point cloud to an object class in most cases. As with many other problems, there is not one best feature descriptor in the domain of 3D object recognition and classification [Ale12]. This raises the question of selecting and applying descriptors depending of the situation at hand.

This proposal provides a concept which may offer an answer to this question. We present a method, where we use reinforcement learning (RL) to learn the successive selection and application of several point cloud descriptors to obtain high classification rates.

2. Our approach

Our approach is based on a RL framework with a finite number of actions and states [SB98]. More precisely we use Q-learning. It utilizes an adaptive set of global and local feature vectors for some object classes. It is initialized with the feature vectors of a small set of already classified 3D point clouds.

2.1. Environment

A *state* consists of a set object classes and a set feature descriptor algorithms. Furthermore a state can comprise results from feature descriptors. It is not intended, that this option is used for all descriptors, but it can be used for basic shape attributes (e. g., flat, longish or uniformly), which can be obtained by PCA, to the state. The intention to embed the elementary results of different, probably predominant global

DOI: 10.2312/egp.20141065

descriptors into the state, is to enable the RL agent to use different actions due to different descriptor results.

An *action* consists of the application of a 3D point cloud descriptor. More precisely, it consists of three steps: 1) the selection of the descriptor, 2) the calculation of the feature vector(s) and 3) the elimination of candidates from the object classes by comparing the current feature vector(s) with all feature vectors learned so far in the previous steps. Each descriptor is applied only once.

2.2. Learning Process

Without any restrictions the reinforcement learner terminates naturally, if the number of potential candidates of suitable object classes is zero or all 3D point cloud descriptors have been used. But the natural termination is not desirable, since we propose a time limit how long a single object classification should take maximally. Without this limitation the RL framework would probably learn to use a descriptor with high accuracy like PFH [RBMB08] or SHOT [TSDS10], but it would take a very long computational time to compare the feature vectors for all objects in the object classes. Moreover, it makes no sense to wait until the set of potential object classes is empty. Thus, the learning process terminates at the latest when only one object class remains. At this point we have to clarify the conditions under which the RL agent gets an immediate reward. The goal of our approach is to determine the quality of the states during the learning and classification process. Thus, there will be no rewards for the achievement of any intermediate state. The only states at which an evaluation of the result is possible, are the terminal states. When a terminal state is reached where only one object class remains and the object class does match, the reward is linearly interpolated in $r \in [0,1]$, with $r = 1$ at $t = 0$ and $r = 0$ at $t = timelimit$. Otherwise, if there are more object classes left or the object class does not match, the reward is $r = -1$.

Initially, the RL starts with a random policy for maximum exploration. In this phase, we start with the initially known and classified objects, since the decision whether the final object class fits or not is straightforward. If the policy gets more stable, the exploration is reduced in favor of exploitation. This method is called ε-greedy, meaning that most of the time those actions are selected, which maximize the expected reward, but with probability ε a random action is selected. In this way it is possible that the system adapts to changes of the environment over time. This balance of exploration and exploitation allows the system to grow online and allows us to add new descriptors to the system. The ε-greedy strategy automatically leads to the occasional use of those new feature descriptors and an adaption of the policy π.

A big advantage of our approach compared to classical object classification approaches (e. g., neural networks based ones) consists in the fact that the number of object classes can increase dynamically during the learning and classification process. To compensate for possible faulty decisions of the RL agent, the classification of an unclassified object is repeated automatically multiple times with an increasing ε-value to allow more random selections of feature descriptors. In case, the object cannot be assigned to one of the object classes at hand, a new class is created, which means that the agent has learned a new object class autonomously.

3. Conclusion

This proposal suggests a system which learns a strategy to select and apply 3D point cloud descriptors with the goal to classify a point cloud with high accuracy, namely among a large number of object classes and within a preset time limit. The proposed approach is based on reinforcement learning and we expect the approach to be highly adaptive, e.g., allowing the integration of new descriptors and the online learning of new object classes.

References

[Ale12] Alexandre L. A.: 3d descriptors for object and category recognition: a comparative evaluation. In *Workshop on Color-Depth Camera Fusion in Robotics at IROS, Vilamoura, Portugal* (2012). 1

[HPPLG11] Heider P., Pierre-Pierre A., Li R., Grimm C.: Local shape descriptors, a survey and evaluation. In *Proceedings of the 4th Eurographics conference on 3D Object Retrieval* (2011), Eurographics Association, pp. 49–56. 1

[Kei99] Keim D. A.: *Efficient geometry-based similarity search of 3D spatial databases*, vol. 28. ACM, 1999. 1

[PR99] Paquet E., Rioux M.: Nefertiti: a query by content system for three-dimensional model and image databases management. *Image and Vision Computing 17*, 2 (1999), 157–166. 1

[PRM*00] Paquet E., Rioux M., Murching A., Naveen T., Tabatabai A.: Description of shape information for 2-d and 3-d objects. *Signal Processing: Image Communication 16*, 1 (2000), 103–122. 1

[RBMB08] Rusu R. B., Blodow N., Marton Z. C., Beetz M.: Aligning point cloud views using persistent feature histograms. In *IROS 2008. IEEE/RSJ International Conference on* (2008), IEEE, pp. 3384–3391. 1, 2

[RC11] Rusu R. B., Cousins S.: 3d is here: Point cloud library (pcl). In *Robotics and Automation (ICRA), 2011 IEEE International Conference on* (2011), IEEE, pp. 1–4. 1

[SB98] Sutton R. S., Barto A. G.: *Reinforcement learning: An introduction*, vol. 1. Cambridge Univ Press, 1998. 1

[SKO00] Suzuki M. T., Kato T., Otsu N.: A similarity retrieval of 3d polygonal models using rotation invariant shape descriptors. In *Systems, Man, and Cybernetics, 2000 IEEE International Conference on* (2000), vol. 4, IEEE, pp. 2946–2952. 1

[TSDS10] Tombari F., Salti S., Di Stefano L.: Unique signatures of histograms for local surface description. In *Computer Vision–ECCV 2010*. Springer, 2010, pp. 356–369. 1, 2

[VS02] Vranic D. V., Saupe D.: Description of 3d-shape using a complex function on the sphere. In *ICME'02. Proceedings. IEEE International Conference on* (2002), vol. 1, pp. 177–180. 1