Online Unsupervised Cumulative Learning For Life-Long Robot Operation

Gatsoulis, Y., Burbridge, C., & McGinnity, TM. (2011). Online Unsupervised Cumulative Learning For Life-Long Robot Operation. In *Unknown Host Publication* (pp. 2486-2490). IEEE.

[Link to publication record in Ulster University Research Portal](#)

Published in:
Unknown Host Publication

Publication Status:
Published (in print/issue): 07/12/2011

Document Version
Publisher's PDF, also known as Version of record

General rights
Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

ONLINE UNSUPERVISED CUMULATIVE LEARNING FOR LIFE-LONG ROBOT OPERATION

Y. GATSOU LIS

C. BURBRIDGE

T. M. MCGINNITY

Intelligent Systems Research Centre, University of Ulster, Londonderry, BT48 7JL, UK

{i.gatsoulis, cj.burbridge, tm.mcginfinity}@ulster.ac.uk

Abstract—The effective life-long operation of service robots and assistive companions depends on the robust ability of the system to learn cumulatively and in an unsupervised manner. For a cumulative learning robot there are particular characteristics that the system should have, such as being able to detect new perceptions, being able to learn online and without supervision, expand when required, etc. Bag-of-Words is a generic and compact representation of visual perceptions which has commonly and successfully been used in object recognition problems. However in its original form, it is unable to operate online and expand its vocabulary when required.

This paper describes a novel method for cumulative unsupervised learning of objects by visual inspection, using an online and expanding when required Bag-of-Words. We present a set of experiments with a real-world robot, which cumulatively learns a series of objects. The results show that the system is able to learn cumulatively and recall correctly the objects it was trained on.

I. INTRODUCTION

The effective life-long operation of service robots and assistive companions depends on the robust ability of the system to learn cumulatively and in an unsupervised manner. The aim of cumulative learning is to provide a system with developmental programs that allow it to evolve and learn through prolonged periods of observation and interaction with its environment. In order to efficiently achieve this, a mechanism that identifies observations that are new to the robot is needed. In previous work [6] we have identified particular characteristics that are important for the effective operation of a cumulative learning system. In particular, a cumulative learning system should be able to identify new perceptions, learn online and unsupervised, expand when required, cope with noise, and fuse information from different sensors.

A first phase in cumulative learning is perceptual learning, i.e. learning and being able to recall objects previously observed. There is an extensive body of existing research into object recognition, and one popular approach is using the compact and generic Bag-of-Words representation of visual perceptions (BoW) [8], [11]. Inspired by the BoW text classification algorithm, a “vocabulary” of unordered feature descriptors is constructed offline. The frequency of these features/words within an image is then used to categorise the imaged object online.

This work was supported by EU FP7-ICT-2007-3: 2.2 Cognitive Systems, Interaction and Robotics PROJECT: IM-CLeVeR (2009-2013), and by the Centre of Excellence in Intelligent Systems (CoEIS) project, funded by Northern Ireland Integrated Development Fund and InvestNI.

Applications of the BoW method in robotics have ranged from visual-SLAM [3] to tactile-sensing driven object recognition [10]. However, due to the need to train and fix the vocabulary offline, the algorithm is not directly suitable for robotics applications that require the ability to cumulatively learn online in an unsupervised way. In this work we are interested in creating a system that can autonomously and cumulatively learn to identify objects without any user interaction. In order to achieve this, a growable BoW vocabulary must be created dynamically online.

This work has been inspired by work with expandable learning structures, in particular the Grow When Required [12] network and Growing Cell Structure [13]; but, the overall architecture of the network and the way that new nodes are added are two key differences from these previous expandable networks. The closest related work is that of Filliat et al. [1], [4]. In [4], an online BoW approach is proposed for localisation and mapping that offers the ability to incrementally learn by expanding the vocabulary. However, semi-supervised learning is carried out in an interactive training process, where as the system proposed here is purely unsupervised. Also, in [1] an unsupervised BoW approach to robot visual mapping is proposed. Loop-closure is detected in a video sequence by assessing the similarity of video frames incrementally online. However, in order to assist the BoW matching process, geometric constraints are used to validate image matches. In the system proposed here such constraints do not need to be considered, and complete images do not need to be stored.

II. EXPANDABLE BAG-OF-WORDS

The architecture of the proposed expandable BoW is shown in Figure 1. It consists of a neural network, which each node keeps a vector of a feature descriptor, and a binary vector, named “ownership vector”, defining the objects that the node belongs to.

The network expands when a particular feature descriptor is not a close match to any of the nodes in the network. In this case, a new node is created with the feature descriptor vector initialised to the input descriptor. Else if the input feature is a close match, then the best-matching node is trained. In both cases (expanding or training) the ownership vector of the new or of the trained node is updated to reflect that this node belongs to the object of the input feature.

During classification, the set of input features from a perceived object are matched to the neural network, and those

Fig. 1. Expandable Bag-of-Words

Fig. 2. Sample histogram of votes

with a close match consist the set of known features/words for this particular perception. The rest of the input features populate the “unknown” category. The sum of the ownership vectors of the best matching nodes of the network produces a histogram of votes, an example is shown in Figure 2, which together with the number of unknowns features determines the classification of the perception. This histogram shows how well the perceived object matches the known objects. In general terms if the number of unknowns is significantly higher than the number of votes of the most popular category then the object is classified as a new object, otherwise the object is classified as the most popular category.

III. OPERATIONAL PROCEDURE

Our object learning system operates cumulatively. The robot inspects an object from all directions by driving around it and taking images of it, providing a 360° perception of the object. As with many BoW implementations, robust SURF [2] features are extracted from the images to provide scale, rotation and partial illumination invariance. As new images are taken, the data is dealt with immediately and then discarded.

The operational process of the system, shown in Figure 3, consists of the following phases that are continuously performed in a repetitive loop:

- *Inspection phase*, during which the robot carries out a single 360° visual inspection of the object. An inspection loop consists of the following steps:

Fig. 3. Operation loop

- 1) For every perception x_o of object o extract its SURF features S ; $x_o \rightarrow S$
- 2) For every feature descriptor s in S ; $\forall s \in S$:
  - a) Calculate its Euclidean distances from every node n with weights vectors w_n of the network N ; $D_s = \{\|s - w_n\|\}, \forall n \in N$
  - b) Find the best matching node b_1 with distance d_{b_1} ; $\{b_1\} = \text{argmin}(D_s)$, $d_{b_1} = \min(D_s)$
  - c) Append b_1 and d_{b_1} to inspection loop log I ; $I \cup \{[b_1, d_{b_1}]\}$
- *Decision phase*, during which the robot decides based on the inspection phase whether the perceived object is an already known object or a new one. In the case that the object is known, the number of unknown new features is taken into account to decide whether further training is required. In more detail,
  - 1) Using the inspection log I from the inspection loop; calculate the histogram of votes, H , by summing the ownership vectors, B , of all the inspection instances, $i \in I$, that are good matches; $H = \sum_{i \in I_g} B_i$, where $I_g \subseteq I$
  - 2) From H get the most popular category; $c = \text{argmax}(H)$, $v = \max(H)$
  - 3) Count the number of bad matches; $u = |I_b|$, where $I_b = I - I_g$
  - 4) If $u \gg v$ then the perceived object is a new unknown object
  - 5) Else the object is known and requires further training if $u \approx v$, or it is well known if $u \ll v$
- *Do nothing phase*, during which the robot stands still waiting for the object to change.
- *Training phase*, during which the robot circles around the table perceiving and learning in an online, continuous, cumulative and unsupervised manner the 360° perception of either a previously known object or a new object, depending on the outcome of the preceding decision step. Training carries on until the object is well learnt. The training phase consists of the following steps:
  - 1) For current perception x_o of object o extract its SURF features S ; $x_o \rightarrow S$
  - 2) For every feature descriptor s ; $\forall s \in S$:
    - a) Calculate its Euclidean distances from every node n with weights vectors w_n of the network N ; $D_s = \{\|s - w_n\|\}, \forall n \in N$

- b) Find the best matching node b_1 with distance d_{b_1} ; $b_1 = \text{argmin}(D_s)$, $d_{b_1} \in D_s$
  - c) If the best matching node b_1 is a good match, then train node b_1 ,
    - i) Update the weights vector w_{b_1} according to $w_{b_1} = w_{b_1} + \eta \times (s - w_{b_1})$, where, $\eta \in (0, 1]$ is the learning rate
    - ii) Update the ownership vector B to reflect that the node represents the training object
  - d) Otherwise if node b_1 is not a good match, then create a new node; $N = N \cup \{n_n\}$
    - i) Initialise the weights vector w_{n_n} of the new node n_n to the input vector s ; $w_{n_n} = s$
    - ii) Update the ownership vector B to reflect that the node represents the training object
- 3) Repeat steps 1 and 2 until the robot completes a loop

IV. EXPERIMENTS

A. Experimental Setup

The robot used in this work was a MetraLabs Scitos-G5 differential drive mobile robot equipped with a SCHUNK 7 degrees-of-freedom manipulator and a Microsoft Kinect camera attached to the end-effector. The 3D data from the Kinect were used to reliably segment the object from the background; this was achieved by using RANSAC for identifying the table and then using depth filtering and projecting the object points to the 2D image to extract the region of the object.

In order to obtain 360° perceptions of objects, the robot was programmed to autonomously drive around a table. One object was placed on the table at a time, and the manipulator was programmed to maintain the object in the camera view at all times.

The set of objects consisted of 10 real-world objects, shown in Figure 4. The surface of the table was a random sheet of gift wrap paper to make the world as realistic as possible.

The experimental setup with one of the objects is shown in Figure 5.

B. Experimental Procedure

The experimental procedure consisted of a training phase and a validation phase. In the training phase the robot inspects, decides and learns the training objects autonomously. The only time a user interferes with the system is at the end of the training phase when the robot has satisfactorily learnt the object it was trained on previously, and it is on a “do nothing” phase. At that point a user can change the object on the table.

In the validation phase the robot is asked to recall the 10 objects it has previously learnt in the training phase, by performing 5 consecutive inspection-decision loops without any training.

Fig. 5. The experimental setup with one training object on the table

V. RESULTS AND DISCUSSION

A. Training phase

Figure 6 shows the results of training the system on the 10 objects shown in Figure 4. In total, 64 training and inspection cycles were required to learn the 10 objects, and the size of the network grew to 63000 nodes. The most substantial jump in network size occurred when training on Object 2 (a toy robot). This was due to the higher number of SURF features found on this object.

Figure 6 shows that the ratio between known and unknown features quickly dropped as the training on each object proceeded. This demonstrates the learning effectiveness of the proposed system, and the stability of the SURF features used. However, extreme changes in lighting were found to effect the object learning. This is apparent when training on Object 9 (epochs 47 – 54), where a spike in the ratio of known features to unknowns can be seen at epoch 52. This occurred when bright sunlight saturated the image.

Regardless of such real-world online difficulties, the system was able to learn all 10 objects successfully in a completely unsupervised manner. As each new object was presented to the robot, it was correctly identified as unknown and a new object class was created.

B. Recall phase

The same 10 objects as had been learnt by the system were each placed in front of the robot 5 times and an inspection loop executed. Figure 7 shows the classification results of all 5 loops for each object.

Each object was correctly identified despite there being a large number of unknown features also discovered, particularly on Objects 2, 9 and 10. The unknown features are the result of erroneous features in the image background (around the object edge), illumination changes causing features to not be matched to previously seen features, and new features on the object that had not been detected while training. Since the system works with online images in real-time, unknown features will always be present. However, the system demonstrated reliability even under these conditions.

Fig. 4. The 10 objects used in the experiments

Fig. 6. Graph showing the growing size of the network during the training experiment on the 10 objects, and the learning performance ratio for each one of the objects

The overlap in features between objects is most observable between Objects 3 and 8 (two books), where although each object was identified correctly, there is evident similarity between them. On the other hand, the percentage of votes for objects other than Object 2 when presented with Object 2 were low. This is again due in part to the large number of features present on Object 2, reflecting its distinctiveness from the other objects.

VI. CONCLUSIONS

This paper described an unsupervised cumulative online learning method for continuous learning of a robot system, based on a novel growing Bag-of-Words approach, which can be used in service assistive robotics that have to adapt and learn new environments and skills. The experimental results from the training phase have shown that the robot was able to continuously and cumulatively learn the different objects that were presented to it. During the validation phase it was able to correctly recall all the objects it has previously learnt.

Despite the successful demonstration of the proposed system, there are certain areas for improvement. The first improvement is dealing with the increasing size of the network by using a pruning technique that removes the nodes that represent the background noise. Although in the experiments

presented here there was no issue with computational performance, it is expected that as the number of objects learnt increases, then the size of the network will become so large that computational performance will degrade. An additional common step dealing with computational performance issues is using some kind of hierarchical organisation of the data to avoid an exhaustive search over the whole network. Lastly, from a practical point of view, the system should be able to simultaneously segment and recognise multiple objects existing in a scene, as currently it is capable of segmenting only a single object.

REFERENCES

- [1] A. Angeli, S. Doncieux, J.A. Meyer, and D. Filliat, "Real-time visual loop-closure detection", Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on, IEEE, 2008, p. 18421847.
- [2] H. Bay, T. Tuytelaars, and L. Van Gool, Surf: Speeded up robust features, Computer Vision/ECCV 2006, May. 2006, p. 404417.
- [3] M. Cummins and P. Newman, "Probabilistic Appearance Based Navigation and Loop Closing", Proceedings of the 2007 International Conference on Robotics and Automation (ICRA 2007), IEEE, 2007, pp. 2042-2048.
- [4] D. Filliat, "A visual bag of words method for interactive qualitative localization and mapping", Proceedings of the 2007 International Conference on Robotics and Automation, Apr. 2007, pp. 3921-3926.
- [5] M.A. Fischler and R.C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and

Fig. 7. Graph showing the mean histograms of 5 inspection loops for each one of the 10 learnt objects in the recall phase

Automated Cartography”, Communications of the ACM, vol. 24, Jun. 1981, pp. 381-395.

- [6] Y. Gatsoulis, E. Kerr, J.V. Condell and T.M. McGinnity, “Novelty detection for cumulative learning”, in *Towards Autonomous Robotic Systems (TAROS'10)*, Plymouth, 2010, pp. 62-67.
- [7] T. Kohonen, The self-organizing map, Neurocomputing, vol. 21, 1998, pp. 1-6.
- [8] T. Leung and J. Malik, “Representing and Recognizing the Visual Appearance of Materials using Three-dimensional Textons”, International Journal of Computer Vision, vol. 43, 2001, pp. 29-44.
- [9] U. Nehmzow, *Mobile robotics: A practical introduction*, Springer, London; 2003.
- [10] A. Schneider, J. Sturm, C. Stachniss, M. Reiser, H. Burkhardt, and W. Burgard, “Object identification with tactile sensors using bag-of-features”, IEEE/RSJ IROS'2009, Oct. 2009, pp. 243-248.
- [11] J. Sivic and A. Zisserman, “Video Google: Efficient visual search of videos”, Toward Category-Level Object Recognition, Springer Berlin, 2006.
- [12] J. Shapiro, S. Marsland, and U. Nehmzow, “A self-organising network that grows when required”, Neural Networks, vol. 15, 2002, pp. 1041-1058.
- [13] B. Fritzke, “Growing cell structures—A self-organizing network for unsupervised and supervised learning”, Neural Networks, vol. 7, 1994, pp. 1441-1460.