Learning Without Forgetting Approaches for Lifelong Robotic Vision

Zhengwei Wang[1], Eoin Brophy[2] and Tomás E. Ward[2]

***Abstract*— Recent advances in deep learning have achieved exciting results in the ares such as object detection, image recognition and object localization. However, robotic vision poses new challenges for applying visual algorithms due to varying distribution of images from real world and it requires that the model is able to learn knowledge continuously. This competition is about developing lifelong learning algorithms which can be applied to the robotic vision system. This work describes the approach that we submit to this open competition.**

I. INTRODUCTION

Humans have remarkable abilities to learn knowledge continuously from the real world. One of ultimate goals of a robotic vision system is to build an artificial intelligent agent which is capable of understanding the real world based on their current scenes and their previous knowledge. Object recognition in the computer vision area has achieved exciting results [1], [2], [3], [4], where some deep neural networks (DNNs) even outperform human annotators. However, these approaches still have limitations when they are applied to a robotic vision system [5], [6]. First, distributions of image datasets may vary across categories and tasks. For example, illumination can vary significantly across time (day time and night differences). A well-developed approach should have ability to recognize the object correctly with different illumination levels. Second, it is not feasible to train a model by using all images across tasks for a robotic vision system because it requires more and more computational complexity when tasks are accumulated i.e., the model needs to be retrained for all tasks when a new task comes in. This lifelong robotic vision competition mainly addresses these two issues that need to be overcame by utilizing the advanced machine learning algorithms.

II. METHODOLOGY

Lifelong learning represents a long-standing challenge for machine learning and neural networks due to the learned model catastrophically forgets existing knowledge when learning from novel observations [7], [8]. Currently, lifelong learning contains three main categories: (1) Regularization approaches, which can alleviate catastrophic forgetting by imposing constrains on the update of the neural weights [9], [10]; (2) Dynamic architectures, where this method is introduced to change architecture properties in response to new information [11]; (3) Complementary learning systems and memory replay, which provides the basis for a computational framework modeling memory consolidation and retrieval in which the complementary tasks of memorization and generalization are mediated by the interplay of the mammalian hippocampus and neocortex [12], [13].

In this work, we utilized the learning without forgetting (LwF) [7] strategy to this competition. We chose LwF because it has following advantages: (1) High computational efficiency, where the LwF only requires image datasets of the current task and only retains the trained parameters of the previous one task; (2) Simplicity in deployment. Once a task is learned, the training data does not need to be retained or reapplied to preserve performance in the adapting network.

Figure 1 illustrates the training strategy used in our approach. We deployed a pretrained MobileNet V2 [14], in which the weights up to the bottleneck are retained as θ_p (θ_p here is fine tuned during training) and we trained the bottleneck weights from scratch. In order to introduce the LwF to the architecture, we retain the θ_{old} that is trained by previous tasks in order to construct the regularization term for training new weights θ_{new}. It should be noted that there is no replay of previous task images in this structure and only the update θ_{new} is retained after training which is going to be used during the testing session. During the experiment, we find that initializing θ_p by using the pretrained weights for each task performs better than using θ_p continuously trained by all tasks. Thus we load the initial pretrained weights θ_p when processing a new task and θ_p is going to be fine tuned during the training. Details of training scheme are included in Algorithm 1.

Algorithm 1 Training details

Inputs:

Training images $\mathbf{X}$, labels $\mathbf{Y}$ of the new task and the pretrained parameters θ_p

Initialize:

1: $\mathbf{Y}_{old} \leftarrow \mathcal{M}_{\hat{\theta}_p,\theta_{old}}(\mathbf{X})$ // *Output labels using model trained by previous tasks. Both $\hat{\theta}_p$ and θ_{old} are updated by using previous tasks*

2: $\theta_{new} \leftarrow$ Xavier-init(θ_{new}) // *Use Xavier initialization for the bottleneck weights*

3: Load the pretrained weights θ_p to the new model

Train:

4: $\theta_p^*, \theta_{new}^* \leftarrow \underset{\hat{\theta}_p,\hat{\theta}_{new}}{\operatorname{argmin}}(\lambda\mathcal{L}_{old}(\mathbf{Y},\mathbf{Y}_{old})+\mathcal{L}_{new}(\mathbf{Y},\mathbf{Y}_{new}))$

5: $\theta_{old} \leftarrow \theta_{new}$ // *Cache the updated weights which are going to be used as old weights for the next task.*

*This work is supported by Science Foundation Ireland (SFI) under the Grant Number 15/RP/2776 and 17/RC-PhD/3482

[1]Zhengwei Wang is with V-SENSE, School of Computer Sciense and Statistics, Trinity College Dublin, Dublin 2, Irleand `Zhengwei.wang@tcd.ie`

[2]Eoin Brophy and Tomás E. Ward are with the Inisht Centre for Data Analytics, School of Computing, Dublin City University, Dublin 9, Ireland `eoin.brophy7@mail.dcu.ie, tomas.ward@dcu.ie`

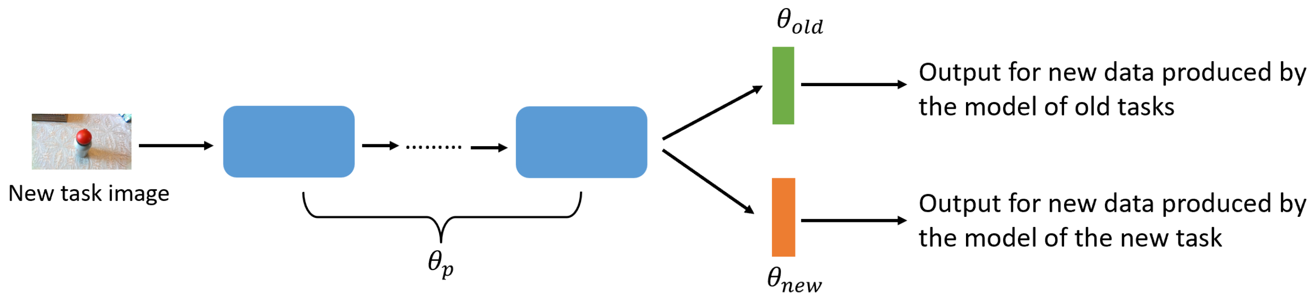

Fig. 1. LwF training strategy used in this work.

It should be noted that there are some difference between our approach and the original approach. First, we did not retain old weights θ_{old} for each task. This might deteriorate the performance for the model but this is closer to the real-world situation because we do not know which task is going to be tested. Practical situation requires a unified model which is able to learn tasks continuously. Our method is also more computationally efficient during the training compared to the original method especially when the number of tasks is huge. Second, instead of fine tuning θ_p continuously for each task, we load the pretrained weights to θ_p for each new task and then fine tune it. We find this strategy will improve the performance.

III. RESULTS

We compared our model ($\mathrm{M_{LwF}}$) to model 1 ($\mathrm{M_1}$) which does not deploy any lifelong learning strategy and model 2 ($\mathrm{M_2}$) which fine tunes θ_p continuously as we mentioned before. Performance of three models is included in Table I. It can be seen that $\mathrm{M_{LwF}}$ achieves the highest accuracy and this model is our final solution submitted to the competition.

Model	Accuracy
$\mathrm{M_{LwF}}$	76.7%
$\mathrm{M_1}$	60.9%
$\mathrm{M_2}$	74.3%

TABLE I
VALIDATION ACCURACY OF THREE MODELS.

IV. CONCLUSIONS

In this work, we described the approach we used for solving the lifelong robotic vision challenge. The core backend of our approach is LwF which was proposed to overcome the catastrophically forgetting issue arisen from the lifelong learning. We modified the original approach in order to be more suitable to deal with this challenge.

ACKNOWLEDGMENT

The authors thank organizers for organizing such a greate challenge which addresses the importance of lifelong learning in our real world and providing the extraordinary datasets. We also acknowledge the financial support by Science Foundation Ireland (SFI) under the Grant Number 15/RP/2776 and 17/RC-PhD/3482.

REFERENCES

[1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[3] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.

[4] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.

[5] Q. She, F. Feng, X. Hao, Q. Yang, C. Lan, V. Lomonaco, X. Shi, Z. Wang, Y. Guo, Y. Zhang, F. Qiao, and R. H. M. Chan, "Openloris-object: A robotic vision dataset and benchmark for lifelong deep learning," 2019.

[6] F. Feng, R. H. M. Chan, X. Shi, Y. Zhang, and Q. She, "Challenges in task incremental learning for assistive robotics," *IEEE Access*, vol. 8, pp. 3434–3441, 2020.

[7] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 12, pp. 2935–2947, 2017.

[8] S. Thrun and T. M. Mitchell, "Lifelong robot learning," *Robotics and autonomous systems*, vol. 15, no. 1-2, pp. 25–46, 1995.

[9] M. K. Benna and S. Fusi, "Computational principles of synaptic memory consolidation," *Nature neuroscience*, vol. 19, no. 12, p. 1697, 2016.

[10] S. Fusi, P. J. Drew, and L. F. Abbott, "Cascade models of synaptically stored memories," *Neuron*, vol. 45, no. 4, pp. 599–611, 2005.

[11] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, "Progressive neural networks," *arXiv preprint arXiv:1606.04671*, 2016.

[12] D. Kumaran, D. Hassabis, and J. L. McClelland, "What learning systems do intelligent agents need? complementary learning systems theory updated," *Trends in cognitive sciences*, vol. 20, no. 7, pp. 512–534, 2016.

[13] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory." *Psychological review*, vol. 102, no. 3, p. 419, 1995.

[14] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.