Weighted sampling based reward distribution preserving corset compression for experience replay buffer in lifelong deep reinforcement learning for medical imaging

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

Deep lifelong reinforcement learning is a new and interesting area that has only been explored recently. One method to achieve lifelong learning under the deep lifelong learning framework is to introduce an Experience Replay Buffer (ERB) to encapsulate knowledge from previous rounds and to be used in future training. However, the ERB size can vary in size and can be very large. To that end, we propose a coreset based ERB compression method to compress the ERBs while preserving previous knowledge. We tested the coreset ERB based lifelong deep reinforcement learning on the brain tumor segmentation (BRATS) dataset for the task of ventricle localization. The conventional and coreset lifelong learning models (at compression ratios of 10X, 20X, 30x, and 40X) were evaluated to iteratively train on a sequence of 10 different imaging environments. The coreset lifelong learning models with compression ratios of 10X, 20X, 30x, and 40X localized the ventricle with a mean pixel distance of 12.93, 13.46, 17.75, and 18.55, respectively compared to the baseline performance of 10.87 for the conventional lifelong learning model. Our results demonstrate that the proposed coreset based ERB compression method can achieve up to 20X compression without a significant drop in performance.

Keywords: Deep reinforcement learning, lifelong learning, continual learning, medical imaging, coresets, clustering

© 2023 Anonymous.

1. Introduction

2. Related Works

3. Methods

3.1. Deep reinforcement learning

3.1.1. Experience Replay Buffer

4. Experiment and Data

4.1. Experimental Setup

4.1.1. Clinical Data

4.1.2. Training Protocol

4.1.3. Performance Evaluation

5. Results

6. Conclusion

Acknowledgments

Acknowledgments withheld.

References