Navigating Parameter Space: Mitigating Catastrophic Forgetting in Continual Learning

Huang, Henry (School: Palos Verdes Peninsula High School)

The most sought-after goal in machine learning is achieving artificial general intelligence, which requires stable memory and flexibility in learning. Existing models are dependent on all data being available during training and do not learn during inference. If the model is required to learn new samples continuously, performance on previous data will decrease sharply. This phenomenon, coined “Catastrophic Forgetting”, is the main challenge of Continual Learning. The most effective option to prevent Catastrophic Forgetting is retraining a model from scratch with new samples. However, retraining is extremely expensive and time consuming, especially for large models. This process is also impractical in real world scenarios, because the model cannot adapt. This work focuses on improving Parameter Regularization approaches for class incremental learning, the most challenging continual learning scenario. Elastic Weight Consolidation (EWC) is thoroughly studied through hyperparameter optimization and varying model architectures. The results reveal that the structure of the last Fully Connected (FC) Layer significantly affects accuracy, with Dynamic FC performing much better than Static FC. Furthermore, optimizing the Lambda hyperparameter improves Dynamic FC but not Static FC. To improve upon EWC, a novel Range Matrix method is introduced to further constrain all parameters through a custom-built optimizer. Static FC with Range Matrix allows for an improvement over EWC Static FC through hyperparameter tuning. In conclusion, this work shows that the FC layer with hyperparameter optimization is crucial for the success of Parameter Regularization methods, enabling improvement over previous Parameter Regularization works in CIL by almost doubling accuracy.

Awards Won: