Topic: Continual Learning using Memory Consolidation

Presenter: Arun George
Research Question: This paper examines how to improve the continual learning performance of deep neural networks, specifically when learning many tasks in sequence. It asks if performance can be enhanced by explicitly modeling interactions between different memory systems in the network, similar to how the human brain consolidates memories.

Challenges: A major problem in Continual Learning is catastrophic forgetting, where deep models forget previously learned knowledge when trained on new tasks, particularly over long sequences. Another key challenge is the data imbalance caused by using a small buffer of old examples alongside new task data, which can bias the model towards recent tasks.

Claims: The authors state their proposed framework, Bilateral Memory Consolidation (BiMeCo), significantly improves continual learning by separating model parameters into interacting short-term and long-term memory modules. They claim this approach effectively prevents forgetting through dynamic module interaction, is parameter-efficient, and enhances existing continual learning methods when combined. The paper also claims BiMeCo performs better than state-of-the-art methods, especially on longer task sequences and using fewer parameters.

Evidence and Analysis: The study supports its claims with empirical evidence from extensive experiments on standard image classification benchmarks (CIFAR-100, ImageNet-100, ImageNet). Performance is evaluated using metrics like Average Incremental Accuracy (AIA) and Backward Transfer (measured by average forgetting rates), with results shown in tables and graphs. The authors compare BiMeCo against various baseline and state-of-the-art methods in different settings, such as varying task sequence lengths and memory buffer sizes. Results report the mean and standard deviation across multiple runs for reliability. Additionally, detailed ablation studies analyze the impact of the framework's components, and sensitivity analysis using Gaussian noise assesses the method's robustness.