



Contribution ID: 20

Type: **Project talk**

Leveraging Rehearsal Buffers to Enable Efficient Data-Parallel Continual Learning

Thursday 23 March 2023 17:00 (20 minutes)

Deep Learning (DL) emerged as a way to extract valuable information from ever-growing volumes of data. However, when trained on sequential tasks *ie. without full access to the dataset at the beginning of the training*, typical Deep Neural Networks (DNNs) suffer from catastrophic forgetting, a phenomenon causing them to reinforce new patterns at the expense of previously acquired knowledge. This limitation prevents updating models incrementally, which is problematic in many real-life scenarios where the aforementioned datasets are replaced by data streams generated over time by distributed devices. It is unfeasible to train models from scratch every time new samples are being ingested either, as this would lead to prohibitive time and/or resource constraints.

In this talk, we will present techniques based on rehearsal to achieve Continual Learning at scale. Rehearsal-based approaches leverage representative samples previously encountered during training to augment future minibatches with. The key novelty we address is how to adopt rehearsal in the context of data-parallel training, which is one of the main techniques to achieve training scalability on HPC systems. The goal is to design and implement a distributed rehearsal buffer that handles the selection of representative samples and the augmentation of minibatches asynchronously in the background. We will discuss trade-offs introduced by such a continual learning setting in terms of training time, accuracy and memory usage.

JLESC topic

Continual Learning at Scale

Primary authors: COSTAN, Alexandru (Inria); NICOLAE, Bogdan (ANL); ANTONIU, Gabriel (Inria); BOUVIER, Thomas

Presenter: BOUVIER, Thomas

Session Classification: Project Talks on AI/ML/DL

Track Classification: AI and ML/DL