

Forschungspraxis

Distributed learning on serverless compute services

Applying machine learning on large data became part of daily life. Scalability of distributed machine learning algorithms requires tolerance of slower computing nodes referred to as stragglers. To mitigate the effect of stragglers, one can add redundancy to the distributed data. However, stochastic gradient descent (SGD) is known to perform well even when computing on a part of the data. It is shown that for convex loss functions, a small redundancy or no redundancy at all is enough to guarantee good performance of SGD.

In this project, we evaluate the existent coding techniques requiring small redundancy on serverless compute functions offered by Amazon Web Services, or equivalently cloud functions offered by Google Computing Platform. Our goal is to devise a coding technique that goes beyond SGD and can be used for non-convex loss functions. The main targeted application is a full implementation of a neural network on serverless functions with straggler tolerance.

Prerequisites

Knowledge of machine learning algorithms, e.g., linear regression, logistic regression, neural network

Good programming skills

Self-motivation and dedication

Plus: knowledge of probability, statistical pattern recognition, coding theory and convex optimization

Contact

rawad.bitar@tum.de

<https://sites.google.com/site/rawadbitar1>

Advisors

Rawad Bitar