

2018 Research Interest/Project Ideas

Tiark Rompf

tiarkrompf.github.io

TensorFlare: Accelerating SQL + Deep Learning Pipelines

Modern analytics needs to combine multiple systems such as Apache Spark and TensorFlow. In recent work, we have developed Flare, an accelerator back-end for Spark that delivers 10x-100x speedups for SQL workloads

(<https://arxiv.org/abs/1703.08219>). In this project, we plan to extend this approach to pipelines that combine SQL query plans from Spark with deep learning computation graphs from TensorFlow, providing native code generation and cross-system optimization end-to-end for entire data pipelines.