



Laplace Approximations for Bayesian Deep Learning

The background is a solid teal color. On the right side, there are several decorative elements: a large, semi-transparent pie chart with four segments, and several smaller, semi-transparent pie charts of varying sizes. At the bottom right, there is a bar chart with four vertical bars of increasing height from left to right, each with a semi-transparent teal top section.



The Jacobian and GGN matrices

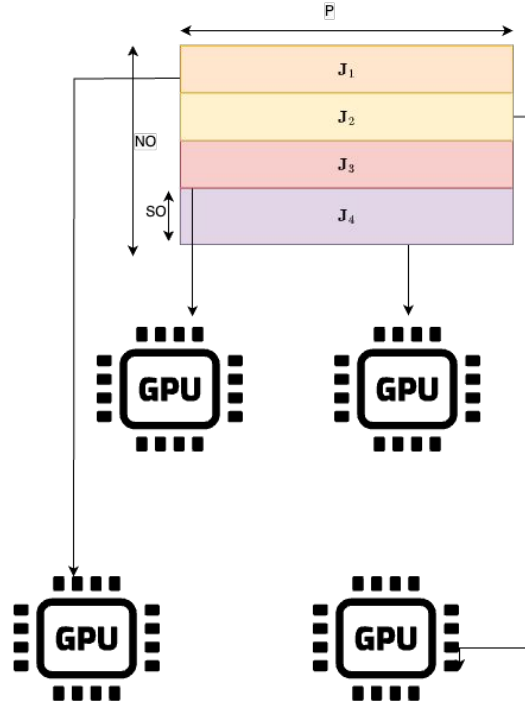

Jacobian matrix can't fit in memory



GGN-vector product is feasible thanks to autodifferentiation

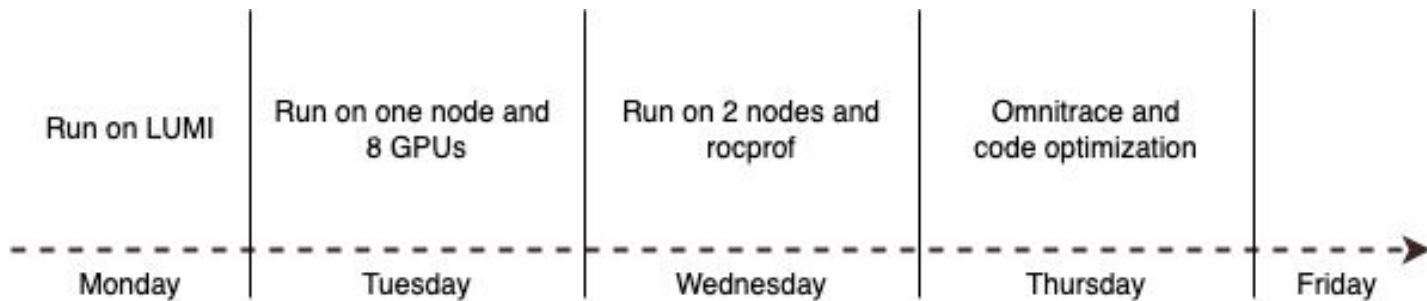


Data-Parallelism for GGN vector products





Hackathon progression





Status and next steps

Where we are now:

- Considerable speedup due to GPU parallelization (from 80s to 35s per epoch)
- Speedup due to data loading optimization (torch loaders vs tfds and pre-fetching) 15% faster

Future goals:

- Bigger models
- Bigger datasets

Thank You!

