

Master's Thesis Topic: Container Image Lazy-Loading for Distributed ML Model Serving

We're looking for a motivated MSc student interested in ML orchestration and model weights distribution optimizations for container runtimes.

With growth in complexity and scale of distributed ML systems, model distribution becomes critical to ensure reduced deployment times and quick model updates and propagation. Technologies like 2DFS¹ can ensure optimized container build time for model weights and selective model partitions downloads. On the other hand, container snapshotters like stargz² can ensure lazy loading of files to optimize container startup. Combining these two technologies is still an untapped research direction that might lead to reduced startup and migration time for ML workloads, efficient model download options, and efficient propagation of model weight updates.

¹ 2DFS On-Demand Container Partitioning for Distributed ML." USENIX ATC 25.

² stargz-snapshotter <https://github.com/containerd/stargz-snapshotter>

Your Responsibilities

- Literature review for containerized model weights lazy-loading and distribution approaches.
- Assess the performance of lazy-loading snapshotters in containerd runtime for model weights download with different model sizes and model partitions configurations.
- Integrate lazy-loading with OCI+2DFS image format and compare performance results.

Requirements

Your background must match all the mandatory (★) requirements and you must have experience with at least 1 optional (♣) technologies/topics.

- Close to completion of a Master's Degree program at TUM (Technical University of Munich) ★
- Experience with Git ★
- Good knowledge of Open Container Initiative standards ★
- Knowledge of Go programming languages ★
- Split-computing ♣
- PyTorch and model training ♣
- Kubernetes/Oakestra ♣
- ML Model serving ♣



Contacts

Send **CV** and **transcript of records** to Giovanni Bartolomeo (giovanni.bartolomeo@tum.de) and Navidreza Asadi (navidreza.asadi@tum.de)

Please add a few lines about yourself and the status of your studies at TUM.