


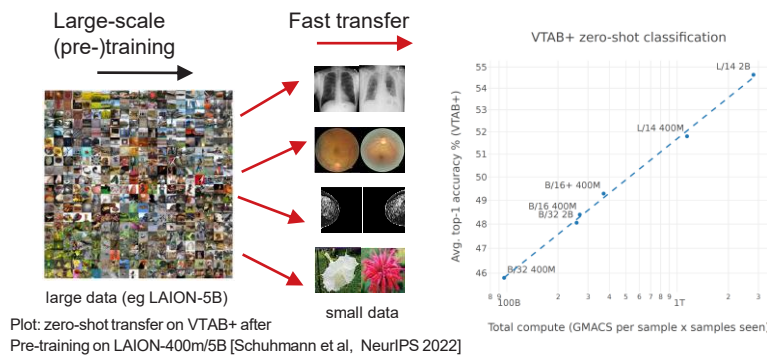
SCALABLE ML & MULTI-PURPOSE AI LAB

LARGE-SCALE LEARNING FOR EFFICIENT TRANSFER



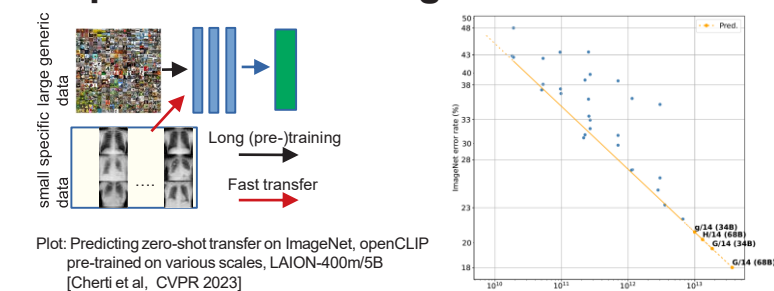
- Pre-training large generalist foundation models for efficient transfer across multiple domains and tasks
- Language-vision learning and large image-text datasets: scaling laws studies for transfer
- Scalable learning: systematic search for scalable architectures via auto-ML
- Active, continual, open-end learning
- Learning with local losses for highly energy efficient in-memory computing neuromorphic hardware
- High performance distributed deep learning on thousands of accelerators (GPUs, etc)
- Democratizing large-scale AI research by open-sourcing outcomes and workflows (part of LAION community)

Transferable language-vision learning (openCLIP, openFlamingo)




- Pre-training large models for few-shot data-efficient transfer
- Studying transfer across different domains and tasks
- Methods for efficient model fine-tuning

Reproducible scaling laws for transfer



- Predicting model properties and performance on larger scale
- Reproducible scaling laws by open-sourcing via LAION
- Transfer for specialized domains, eg medical imaging


Multi-Purpose AI



Various generative models for different tasks and domains. (a) Super-resolution for turbulence flow [Bode et al, 2021] (b) High resolution face generation and (c) Generation of solar measurements [Cherti et al, 2020, 2023]

- Studying various representation and generative learning
- Generative models transferable across domains
- Model compression for low cost deployment

Large-scale Distributed Training



- Data parallel distributed training on JUWELS, openCLIP ViT L/14 (a) Image/sec throughput (b) Scaling performance up to 1024 GPUs [Cherti et al, CVPR, 2023]
- High performance distributed deep learning on thousands GPUs
- Scalable learning based on local losses and local updates for next gen architectures