Praxis der Forschung: Self-Supervised Learning for Real-World Object Detection

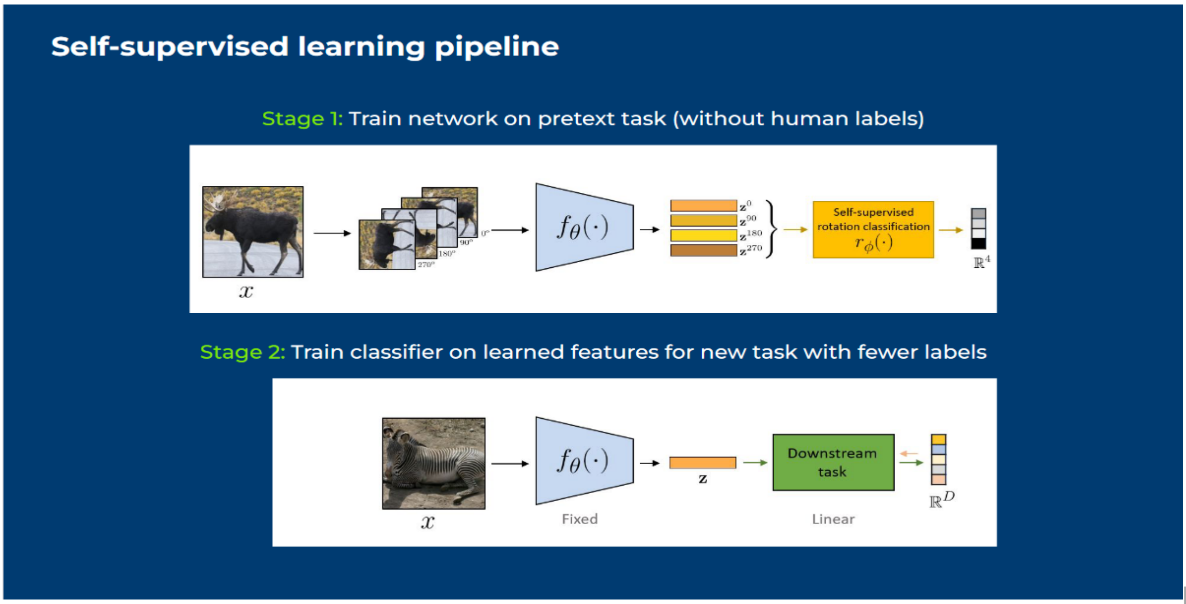

[5]

Self-supervised learning (SSL) is a promising approach for leveraging large amounts of unlabeled data. SLL thus offers potential improvements in object detection models, especially in scenarios where labeled data is limited or expensive to obtain. This project aims to apply SSL techniques to enhance object detection in real-world settings, with a focus on improving detection of small, distant, or occluded objects.

Research Agenda for Students:

1. <u>Understanding Self-Supervised Pre-training:</u> Students will start by exploring existing SSL techniques such as contrastive learning, masked image modeling, and clustering-based approaches. They will investigate how pre-trained models can leverage features learned from unlabeled data to improve downstream object detection tasks.

2. <u>Collecting Performance Data Across SSL Methods:</u> Students will gather and analyze performance data for various SSL methods applied to object detection tasks. This will involve comparing the effectiveness of different models (e.g., EfficientDet, YOLO) when pre-trained with SSL techniques such as SimCLR, MoCo, or BYOL on real-world datasets like KITTI and MS COCO.

3. <u>Analyzing Key Differences in SSL for Object Detection:</u> The project will involve a detailed comparison of top SSL methods to identify their strengths and weaknesses in object localization and classification. Students will focus on how these methods improve classification and localization for object detection in real-world environments, such as urban and indoor scenes.

4. Investigating Cropping and Augmentation Strategies: SSL techniques often benefit from advanced data augmentation strategies. Students will research how augmentations such as color jittering, geometric transformations, and context cropping can improve SSL performance. Additionally, novel loss functions that encourage consistency between augmented views of the same image will be explored.

5. Implementing Findings to Improve SSL for Object Detection: Based on the results of the earlier analyses, students will implement adjustments to SSL methods, optimizing them for real-world object detection task. This may involve integrating insights on augmentation and cropping strategies to enhance multi-instance SSL approaches for real-world object detection scenarios.

Project Goals:

1. Developing a Self-Supervised Pre-training Pipeline**:** Students will develop a pipeline that uses SSL for pre-training on large-scale, unlabeled datasets (e.g., ImageNet without labels, COCO unlabeled images). They will experiment with various SSL methods, such as SimCLR, MoCo, or BYOL, and current state-of-the-art to determine which techniques best improve the feature learning for subsequent object detection tasks.
2. Testing and Benchmarking on Real-World Datasets: The effectiveness of the SSL-based approach will be validated on real-world object detection datasets, such as KITTI (autonomous driving) and MS COCO (general object detection). Students will compare their SSL-enhanced detectors against standard baseline models to measure improvements in performance.
3. End-to-End Detection System Integration: Students will integrate their pre-trained SSL models into standard detection architectures and implement fine-tuning techniques. They will evaluate performance enhancements, particularly focusing on challenging scenarios such as detecting small, distant objects or objects in low-light conditions.

Deliverable: Code and report (that may lead to a peer-review publication or conference paper)

Prerequisites: Proficient in Python

Number of students: ideally 2 (but can be flexible)

Contact and Supervisor: Moussa Kassem Sbeyti, Postdoc in the Research Group Methods for Big Data, Prof. Dr. Nadja Klein, Scientific Computing Center, Karlsruhe

Institute of Technology. Email: moussa.sbeyti@kit.edu Website: kleinlab-statml.github.io)

References:

1. A Cookbook of Self-Supervised Learning. https://arxiv.org/abs/2304.12210
2. Self-Supervised Learning, presentation by Andrew Zisserman. https://project.inria.fr/paiss/files/2018/07/zisserman-self-supervised.pdf
3. Blog on Self-Supervised Learning. https://lilianweng.github.io/posts/2019-11-10-self-supervised/
4. Self-supervised representation learning toolbox. https://github.com/open-mmlab/mmselfsup
5. ECCV 2022 Tutorial on Self-Supervision on Wheels: Advances in Self-Supervised Learning from Autonomous Driving Data. https://gidariss.github.io/ssl-on-wheels-eccv2022/
6. ICCV 2021 Workshop: Self-supervised Learning for Next-Generation Industry-level Autonomous Driving. https://sslad2021.github.io/
7. A curated list of awesome Self-Supervised Learning resources. https://github.com/jason718/awesome-self-supervised-learning
8. Implementation of the unsupervised object discovery method LOST. https://github.com/valeoai/LOST
9. A Study on Self-Supervised Object Detection Pretraining. https://arxiv.org/pdf/2207.04186
10. Self-supervised Learning for Object Detection in Autonomous Driving. https://link.springer.com/chapter/10.1007/978-3-030-92659-5_31
11. SODA10M: A Large-Scale 2D Self/Semi-Supervised Object Detection Dataset for Autonomous Driving. https://arxiv.org/pdf/2106.11118
12. MultiSiam: Self-supervised Multi-instance Siamese Representation Learning for Autonomous Driving. https://arxiv.org/abs/2108.12178
13. Know Your Self-supervised Learning: A Survey on Image-based Generative and Discriminative Training. https://arxiv.org/pdf/2305.13689
14. Self-Supervised Multi-Object Tracking with Cross-Input Consistency. https://arxiv.org/pdf/2111.05943