Reg. No: 2024/53/N/ST6/03156; Principal Investigator: mgr inż. Damian Sójka

Reliable and efficient real-world test-time adaptation

Limitation of Deep Neural Networks. Over the last decade, deep neural networks have revolutionized various technical areas, finding applications in everything from depth estimation to text generation. Despite their advancements, they are not without flaws. One of the most significant drawbacks comes from poor performance on data that diverges from their training set - the test data is out-of-distribution. This issue is especially noticeable in computer vision, where the network trained on clean photos taken in daylight, performs poorly on corrupted or nighttime images.

This limitation is less problematic in controlled environments, such as robotic applications in warehouses, where training data can be carefully curated. However, in open-world applications like autonomous driving, it is nearly impossible or prohibitively expensive to anticipate every possible data variety and prepare it for model training, given the ever-changing and unpredictable environments.

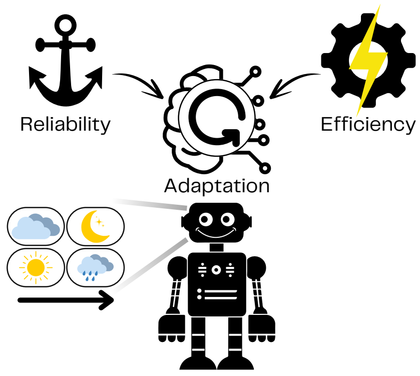

Figure 1: **Project aim**: Our goal is to tackle the neural network's problem with out-of-distribution data by advancing the field of test-time adaptation. We aim to develop reliable and efficient TTA techniques specifically designed for real-world applications.

Test-Time Adaptation. Recently, the paradigm of Test-Time Adaptation (TTA) has emerged as a rapidly growing research area. TTA aims to adapt a pre-trained neural network to the unlabeled data on the fly during test time. The goal is to prevent performance degradation by adjusting the network to continually changing data distributions without any previous assumptions about the test conditions.

Current TTA methods are not without significant drawbacks. Firstly, it was recently shown that no existing TTA approach can handle all types of distribution shifts, making them unreliable for real-world deployment. Secondly, many of these methods are developed without considering computational efficiency, which is crucial for time-sensitive applications and resource-constrained platforms. This project aims to address these limitations by developing novel TTA methods that focus on both reliability and efficiency, and apply the findings to real-world computer vision tasks.

Anchoring Reliability and Stability. Real-world applications require machine learning systems to be reliable. However, it might still be difficult to trust a fixed "black box" in safety-critical applications, let alone the neural network adapted on the fly. Therefore, we intend to focus on reliability across deployment scenarios. Our research will encompass a broad spectrum of testing conditions to identify trustworthy techniques that consistently perform well. We will draw inspiration from the fields of self-supervised and continual learning.

Enhancing Efficiency In real-world applications, e.g., related to mobile robotics, models are deployed on edge devices with limited computational resources. The limitations include computational power, memory capacity, and energy consumption. The model inference itself for an edge device can be computationally demanding. On top of the inference, TTA aims to adapt the model creating a significant computational burden. We will analyze the efficiency of current TTA methods, which will let us understand the trade-off between their efficiency and adaptation performance. Given the insights, we will develop our techniques pushing the boundary of TTA efficiency further.

Applying the Findings to Real-world Tasks. TTA methods are frequently developed using an image classification task as this is a simple setup allowing for easier development and testing. However, focusing solely on a single task limits the TTA applicability. Real-world problems faced by embodied AI agents during an open-world operation are far more sophisticated. Additionally, they do not work with random pictures, but sequences of perception data, which includes potentially useful information for adaptation, considering time-domain and geometrical constraints of surroundings. We plan to apply our findings on reliable and efficient TTA to tasks related to embodied AI agents, such as depth estimation, and leverage all available contextual clues for effective adaptation.

Expected Results and Implications. The outcomes of this project will advance the understanding of test-time adaptation and drive significant progress across a variety of open-world applications utilizing neural networks. Enabling a neural network to adapt reliably and efficiently will allow for more robust deep learning systems in many industries, ranging from robotics to voice assistants.