Jacek Kałużny
Generating and Validating Synthetic Image Datasets for Training AI Models in Computer Vision

Abstract

This thesis explores the generation and validation of synthetic image datasets for training models in computer vision. The core objective is to address the challenges of data scarcity and variability in training robust and accurate AI models. A multi-faceted approach was adopted, involving the development of novel techniques for synthetic data generation, the creation of realistic and diverse datasets, and the rigorous validation of these datasets through comprehensive experiments.

The initial phase of the research focuses on the reconstruction of botanical trees from single images, leveraging advanced machine learning algorithms. A method was introduced using radial bounding volumes and bi-modal growth models to accurately reconstruct 3D tree structures. This technique was validated through extensive testing against real-world datasets, demonstrating its efficacy in producing high-fidelity reconstructions.

Subsequently, the investigation was expanded to synthetic dataset creation, utilizing ControlNet integrated with Stable Diffusion to generate realistic synthetic images of various plants. This approach includes a detailed pipeline for generating annotated images, ensuring the datasets are suitable for training deep learning models for tasks.

The final phase involves the validation of the synthetic datasets. A series of experiments were conducted to compare the performance of AI models trained on synthetic data versus those trained on real data. The results indicate that models trained on these synthetic datasets perform comparably to, and in some cases exceed, those trained on traditional datasets. This highlights the potential of synthetic data to supplement or even replace real data in certain applications.

This research contributes to the field of computer vision by providing robust methods for synthetic data generation and validation, paving the way for more efficient and cost-effective training of AI models. The findings have significant implications for various applications, including agricultural automation, urban planning, and beyond.