Generative Data Augmentation for Pituitary Adenoma MRI Segmentation in Data-Scarce Environments

Author: Bc. Richard Szarka
Supervisor: Prof. Ing. Vanda Benešová, PhD.

Introduction

Pituitary adenomas affect individuals aged 30-35 and can cause serious symptoms including vision loss, headaches, and hormonal imbalances. Deep learning shows promise for automated diagnosis, but limited training data remains a critical challenge in specialized medical domains. Our work addresses this through a **novel generative augmentation approach** that synthesizes both MRI scans and corresponding segmentation masks. By generating anatomically consistent synthetic data with paired annotations, we **enhance segmentation network performance** with minimal real patient data, offering a solution for resource-constrained medical environments.

Our key contributions include:

• A comprehensive **synthetic medical data generation framework** designed for extremely small datasets

• A **conditional diffusion-projected GAN architecture** optimized for volumetric consistency and anatomical plausibility

• Consistent **improvement in CNN segmentation performance** when using our synthetic data, particularly in low-resource scenarios

Dataset

Our dataset contains **31 pituitary adenoma MRI** volumes with expert-annotated ground truth masks. The limited sample size represents a common challenge in specialized medical applications.

We allocated **25 volumes** for training and **6 volumes** for validation to enable comparative evaluation of our approach.

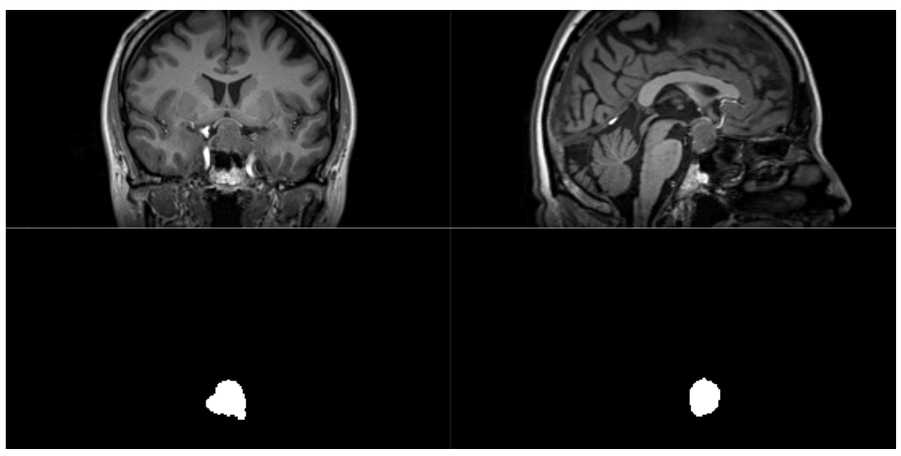

Figure 1: Representative **sagittal and coronal slices** from our dataset showing pituitary adenoma with corresponding ground truth segmentation masks highlighted in white.

Conditional Diffusion Projected (CDP) GAN

Our generative model combines the advantages of Projected GAN with **diffusion-based augmentation to ensure stable training** despite limited data. The architecture employs:

• **FastGAN** generator with **slice-index** conditioning

• **Diffusion augmentation** to balance generator/discriminator dynamics

• Dual-channel output producing both **MRI slices and segmentation masks simultaneously**

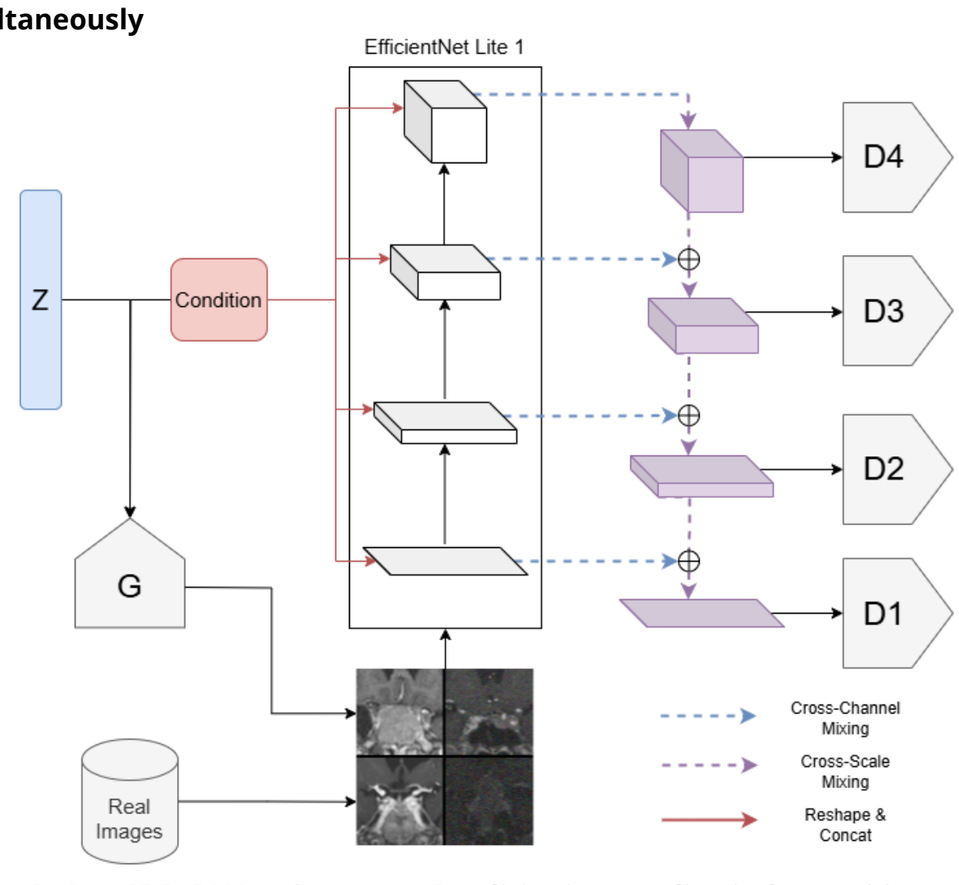

Figure 2: Our CDP-GAN architecture. **Conditioning on slice index** enables coherent 3D volume generation with matching segmentation masks.

Our Volume Generation Pipeline

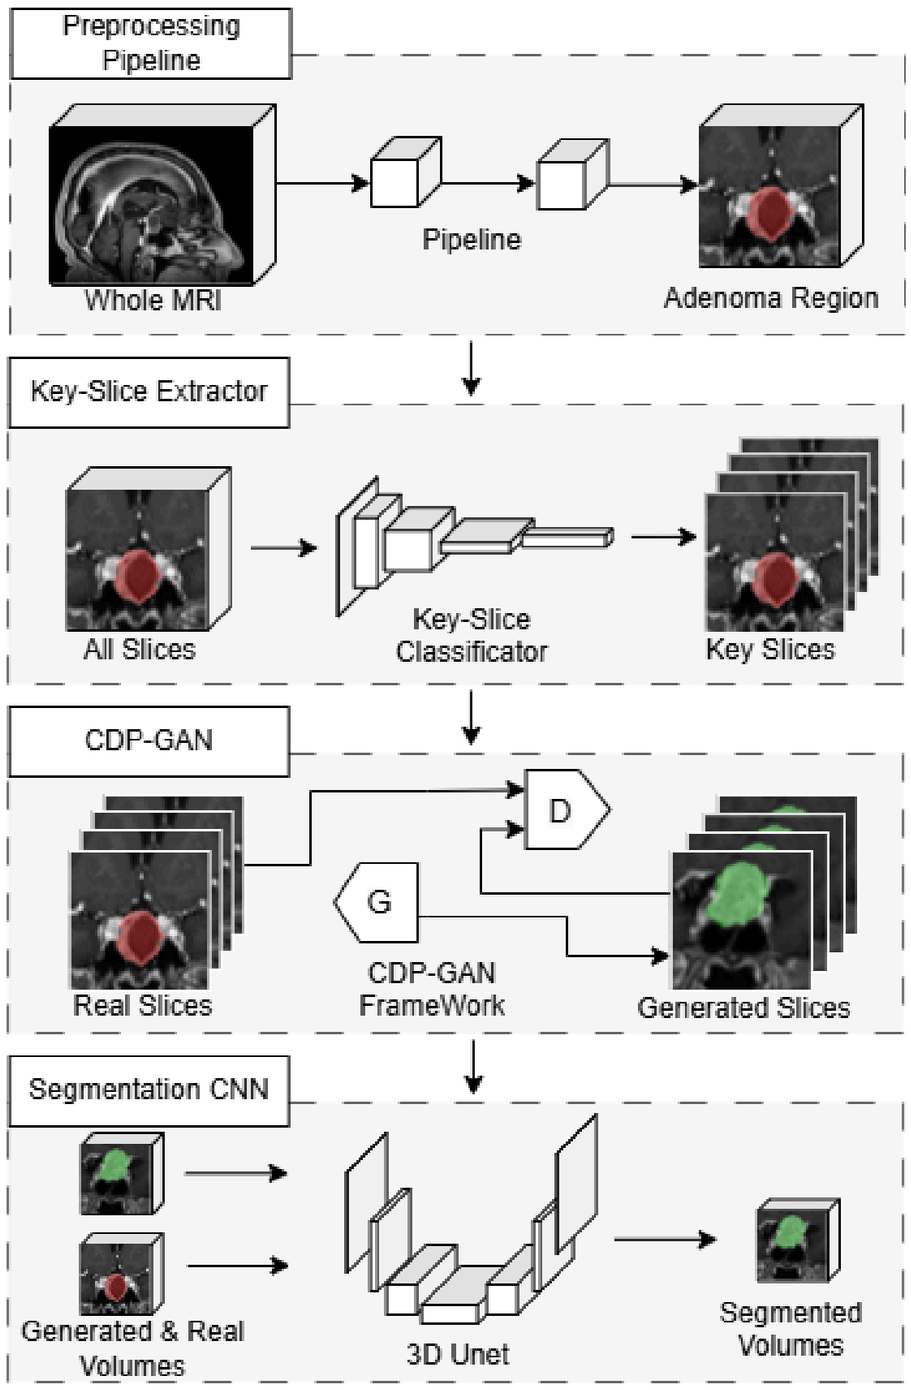

Figure 3: Our pipeline from preprocessing through segmentation. **Each module addresses a challenge in medical volume synthesis.**

Results

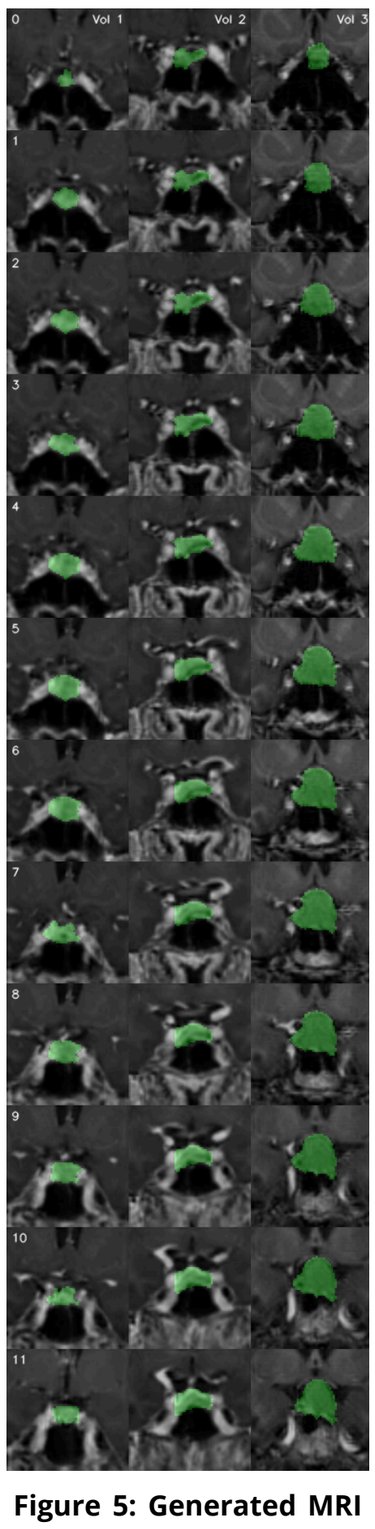

Figure 5: Generated MRI slices with corresponding segmentation masks

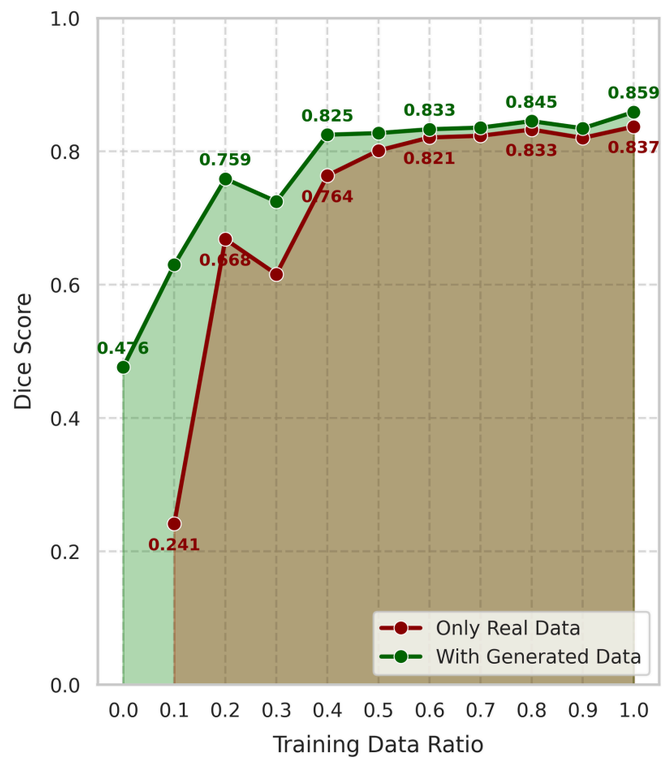

Figure 4: Dice similarity scores across varying percentages of real training data, showing **consistent improvement** when augmenting with synthetic data.

• **Key-Slice extractor successfully identified continuous sequences of relevant slices** without anatomical gaps (94% accuracy)

• **Generation speed:** 0.53 seconds per volume (0.025s per slice)

• **Blind evaluation:** medical expert correctly identified synthetic images **only 30%** of the time

• **Synthetic data consistently improved segmentation performance across all scenarios**

• Most significant improvements observed in **extremely data-limited settings** (2-5 real volumes)