Visual Computing MAGAZiNE

Towards a New Data Representation: GANs for Medical Images Segmentation.

Kenza Aoucher, Ahmed Sif Benmessaoud, Khellaf-Haned H. Faiza, Dahmane A., USTHB

Abstract

Lung cancer remains the leading cause of cancer mortality worldwide. While medical imaging enables earlier diagnosis, precise segmentation of pulmonary nodules from CT scans is critical yet challenging. This work investigates an innovative approach of harnessing generative adversarial networks (GANs) for few-shot learning of lung tumor segmentation models. We hypothesize that GANs encode semantic information when generating realistic medical images. By inverting target CTs through the GAN generator, we extract spatial activation maps to provide localization cues for segmentation with limited examples. Experiments demonstrate promising performance - a model trained on just 1 CT segmented lungs scored 10\% of traditional full supervision requiring thousands of images. Training on 8 examples matched full supervision. This methodology highlights the potential of leveraging GANs for representation learning and few-shot segmentation in medical imaging.

1. Introduction

Lung cancer has the highest mortality rate of all cancers worldwide, resulting in over 1.7 million deaths per year. Early detection is critical, and computed tomography (CT) screening enables it by identifying malignant pulmonary nodules. However, precise segmentation of lungs and nodules from CTs is challenging and critical for accurate diagnosis and treatment planning. In recent years, artificial intelligence (AI) and computer-aided diagnosis (CAD) systems have emerged to automate lung cancer screening from CT scans using deep learning. Convolutional neural networks can be trained to segment lung regions and detect suspicious nodules in CTs (Figure 1).

Figure 1. Example of a lung CT scan

2. Method

Our work comprise of two main contributions: the few-shot segmentation pipeline and the computation of highly accurate semantic tumor masks.

2.1 Accurate tumor masks

Experts' annotations in medical imaging datasets exhibit high variability, thus directly combining them reduces segmentation performance. In this work we developed a method to consolidate variable expert masks from the LIDC [1] dataset into precise lung tumor ground truth (Figure 2).

Towards a New Data Representation: GANs for Medical Images Segmentation.

Kenza Aoucher, Ahmed Sif Benmessaoud, Khellaf-Haned H. Faiza, Dahmane A., USTHB

The resulting masks allow for more effective training and evaluation of segmentation models. Tumor center points are expanded into full masks using adaptive region growing and morphology. Multiple expert masks alongside the region growing ones are consolidated into consensus ground truth via voting to mark pixels labeled tumor by over 50% of experts. This captures common tumor regions while avoiding excessive false positives or negatives.

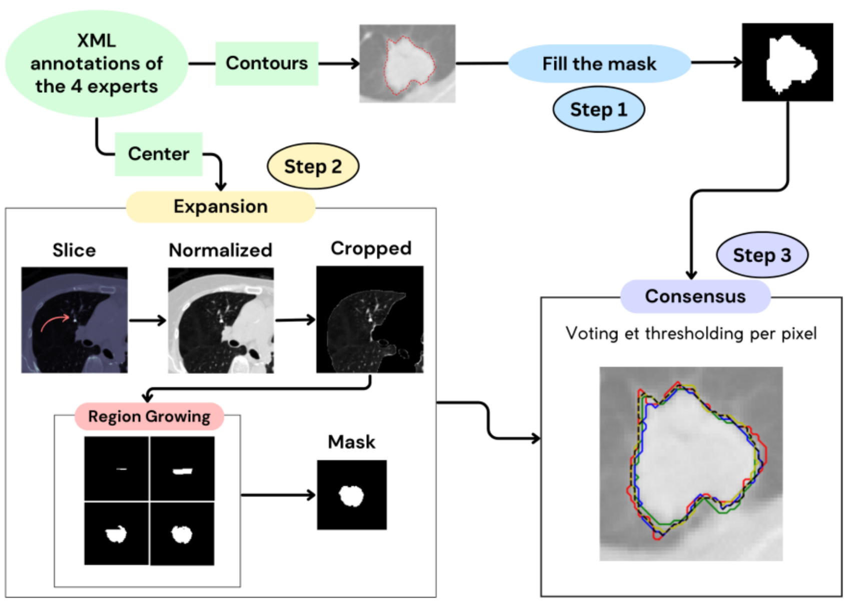

Figure 2. Tumor masks computation methodology

2.2 Few-shot segmentation

Deep learning has driven advances in medical image segmentation, but most methods are supervised thus require very large datasets of pixel-level annotations. Such data is expensive and time-consuming to obtain, especially in the medical field. We propose an innovative approach to overcome these limitations by exploiting generative adversarial networks (GANs) [2] for few-shot semantic segmentation.

GANs are deep neural networks that contain a generator and discriminator that are trained in an adversarial manner (one against the other). Our key hypothesis, as first proposed by Nontawat et al.[3], is that GANs learn to encode the semantic information when synthesizing realistic images. In other words, the generator needs to understand what makes an image an image to be able to generate it. The GAN's feature maps are then used as a pixel-wise representation of the images to train a segmentation network. This allows learning from very few examples while maintaining performance similar to traditional networks trained on massive datasets. Our methodology is exposed in Figure 3.

Visual Computing MAGAZiNE

Towards a New Data Representation: GANs for Medical Images Segmentation.

Kenza Aoucher, Ahmed Sif Benmessaoud, Khellaf-Haned H. Faiza, Dahmane A., USTHB

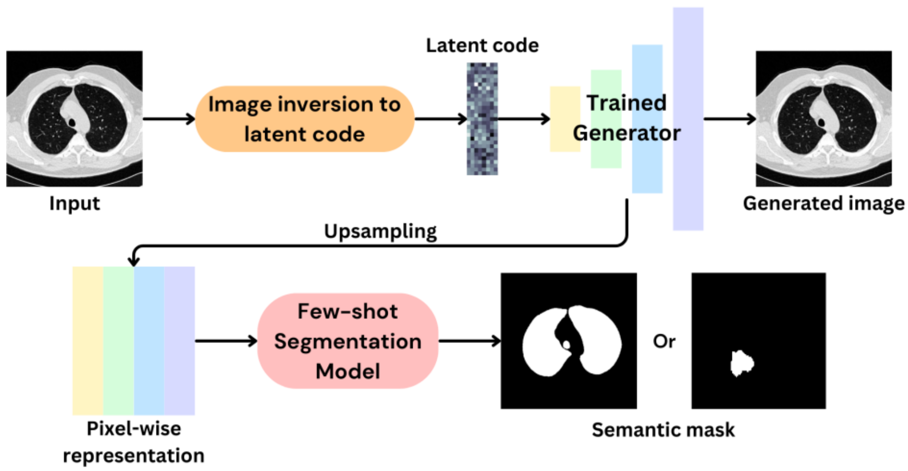

Figure 3. Few-shot segmentation methodology

After careful data preprocessing and creation of reliable semantic masks, advanced GAN architectures are leveraged along with inversion to obtain customizable feature representations of target images. These representations capture details at different levels, from simple textures to complex components. By inverting images to latent vectors, feeding them into the GAN, and concatenating the resulting multi-resolution feature maps, distinctive segmentations are produced with little data. Our work combines the generalization capabilities and component understanding of GANs with the precision of specialized segmentation networks. Specifically, we trained a StyleGAN V3 [4] on over 1 million iterations of CT scans from the Lung Nodule Analysis dataset [5] using 2 NVIDIA GT1070 GPUs over several weeks. Quantitative assessment demonstrated StyleGAN's superior image quality and resolution compared to other models. To invert images, we adapted several optimization algorithms from literature including the W, W+ and PTI methods. Feature maps were then extracted from StyleGAN generator's convolutional layers to serve as pixel-wise representations for few-shot segmentation.

3. Results

Finally, we trained a segmentation model on the GAN-derived representations in a K-shot learning approach and evaluated against full supervision baselines. Experiments show that the 1-shot achieved 90\% of the full-shot performance (Table 1), and the 8-shot demonstrated very promising segmentation of lungs, competing with the traditional fully supervised methods (Figure 4). Tumor segmentation proved more challenging, with few-shot models reaching within 10-15% of full supervision with only 1 training example. Our methodology illuminates the potential of leveraging GANs for representation learning and few-shot segmentation in medical imaging and represents a step toward reducing dependence on large labeled datasets in medical imaging.

Towards a New Data Representation: GANs for Medical Images Segmentation.

Kenza Aoucher, Ahmed Sif Benmessaoud, Khellaf-Haned H. Faiza, Dahmane A., USTHB

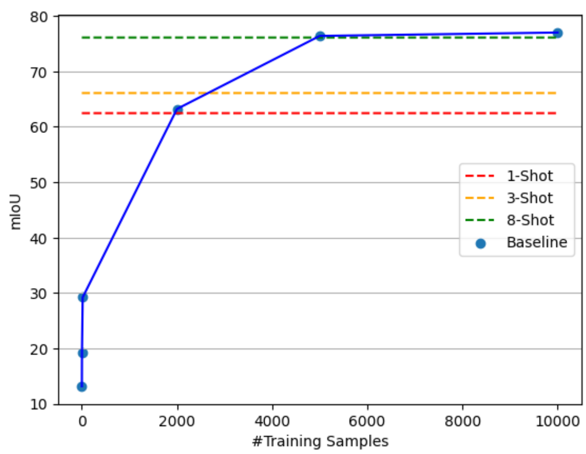

Figure 4. Performance (mIoU) of our lung segmentation model, comparing the full-shot (blue curve) with the few-shots (dotted lines)

Approach	Dice ↑	IoU ↑
Full-shot	**84.17**	**76.39**
1-Shot	73.73	62.45
3-Shot	77.51	66.08
8-Shot	**84.02**	**76.10**

Table 1. Few-shot lung segmentation results

References

[1] Armato III, Samuel G. and McLennan et al. The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference Database of Lung Nodules on CT Scans, 2004, National Cancer Institute, url = {wiki.cancerimagingarchive.net/display/Public/LIDC-IDRI.

[2] Ian J. Goodfellow and Jean Pouget-Abadie et al. Generative Adversarial Networks. 2014, arXiv. 1406.2661.

[3] Nontawat Tritrong and Pitchaporn Rewatbowornwong and Supasorn Suwajanakorn . Repurposing GANs for One-shot Semantic Part Segmentation. 2021. arXiv. 2103.04379.

[4] Tero Karras and Miika Aittala and Samuli Laine and Erik Härkönen and Janne Hellsten and Jaakko Lehtinen and Timo Aila. Alias-Free Generative Adversarial Networks. 2021, arXiv. 2106.12423.

[5] Setio, Arnaud Arindra and Ciompi, Francesco et al. Pulmonary Nodule Detection in CT Images: False Positive Reduction Using Multi-View Convolutional Networks. Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2017.