Topology-Aware GAN (TopoGAN): Transforming Medical Imaging Advances

Hrishi Patel, Colleen Farrelly, Quincy Hathaway, Jennifer Z Rozenblit, Deepa Deepa, Yashbir Singh, Ashok Chaudhary, Yassine Himeur, Wathiq Mansoor and Shadi Atalls

December 7, 2023

Topology-Aware GAN (TopoGAN): Transforming Medical Imaging Advances

Hrishi Patel
North Broward Preparatory School
Coconut Creek,USA
contact2hriship@gmail.com

Colleen Farrelly
Staticlysm
LLC
Miami,USA
cfarrelly@med.miami.edu

Quincy A. Hathaway
Department of Medical Education
West Virginia University
Morgantown, USA
qahathaway@mix.wvu.edu

Jennifer Z Rozenblit
Department of Mathematics
The University of Texas
Austin, USA
jrozenblit@utexas.edu

Deepa Deepa
Institute of Genomic Health
Icahn Sschool of Medicine At Mount Sinai
NYC, USA
justdeepahbti@gmail.com

Yashbir Singh
Department of Radiology
Mayo Clinic
Rochester, USA
singh.yashbir@mayo.edu

Ashok Chaudhary
Ashok Chaudhary
Mayo Clinic
Rochester
USAashok91chaoudhary@gmail.com

Yassine Himeur
University of Dubai
Dubai, UAE
yhimeur@ud.ac.ae

Wathiq Mansoor
University of Dubai
Dubai, UAE
wmansoor@ud.ac.ae

Shadi Atalls
University of Dubai
Dubai, UAE
satalla@ud.ac.ae

***Abstract*—Generative Adversarial Networks (GANs) have gained prominence in medical imaging due to their ability to generate realistic images. Traditional GANs, however, often fail to capture intricate topological features such as holes and connectivity components in real images. This study applies TopoGAN, a recently developed model tailored for medical imaging. TopoGAN dynamically learns and incorporates topological features like connectedness and loops, addressing a real-world medical data augmentation problem. Utilizing a topological GAN loss function based on Persistent Homology and a new success metric, TopoGAN minimizes the topological discrepancy between synthetic and actual images. Experimental results, showcasing a Wasserstein distance of 0.0021 and a Dice coefficient of 0.995, highlight the model's efficacy in producing qualitatively rich synthetic images. This approach not only improves the realism of generated images but also enhances performance in downstream tasks such as image segmentation, offering a groundbreaking solution with significant implications for medical image analysis, diagnosis, and treatment planning.**

***Index Terms*—Topological data analysis, Persistent Homology, Generative Adversarial Network, Mathematics, Medical imaging.**

I. Introduction

Generative adversarial networks (GANs) [1] have emerged as a powerful technique in the realm of medical imaging, demonstrating remarkable success in generating realistic images. GANs employ dual network architecture, training a generator to synthesize images resembling real ones, while simultaneously training a discriminator to differentiate between fake and genuine images. This adversarial interplay facilitates the convergence of the generator towards generating synthetic images that align with the distribution of real images [1]–[4]. However, a crucial challenge in GAN design lies in bridging the gap not only in appearance but also in semantics between the synthetic and real image distributions. Traditional GAN approaches primarily focus on matching the first-order moments of the image distributions within a CNN-based feature space [1], [2], [5], [6]. Recent advancements have explored incorporating higher- order statistics, such as second-order statistics of image features, to better align the synthetic and real image distributions.

In the context of medical imaging, Kossaifi et al. have introduced a novel approach by integrating a statistical shape prior specifically for face images into the generator [7]. This approach aims to capture and leverage higher-order information, thereby enhancing the semantic realism of the generated images. The underlying intuition suggests that the more intricate and high-order information a generator can assimilate, the more faithful and meaningful the synthesized medical images become. By extending GAN methodologies to encompass higher-order statistical measures and incorporating domain-specific priors, the field of medical imaging stands to benefit from enhanced image synthesis, enabling the generation of images that not only possess realistic appearances but also exhibit compelling semantic characteristics relevant to various medical applications.

In this paper, we overview TopoGAN, a topology-based GAN model specifically designed for medical imaging and apply it to a real-world data augmentation problem. TopoGAN represents the first GAN to actively learn and incorporate the crucial aspect of topology from real data. By

directly capturing structural complexity, such as the presence of connected components and holes, TopoGAN offers a unique and powerful approach to understanding and generating medical images. Through this innovative model, we aim to unlock new possibilities in medical image analysis, diagnosis, and treatment planning by leveraging the intrinsic topological properties of the data. Our key technological contribution is a novel topological GAN loss that explicitly matches the topology of synthetic and real image distributions, according to the notion of persistent homology [6], [8]–[10]. This study examines the real-world performance of a revolutionary Generative Adversarial Network (GAN) model that dynamically learns and incorporates the topological features of real images, such as their connectedness and looping nature, to address the deficiencies of existing GANs in reproducing fine structural characteristics seen in actual images.

II. METHODS AND MATERIALS

To address the limitations of traditional GAN models, particularly in the topological discrepancies between synthetic and real images, we examine a relatively new GAN model specialized for medical imaging scenarios such as Computed Tomography (CT). This model centers around a unique loss function that targets topological feature spaces explicitly, mitigating discrepancies between synthetic and actual image distributions. The loss function, leveraging persistent homology, emphasizes topological features across various scales, ensuring synthesized images maintain essential structural similarities. This enhances the reliability and accuracy of the comparative analysis between synthetic and real images in multidimensional spaces.

To validate the effectiveness of our model, we employ novel GAN evaluation metrics that focus on assessing the topological realism of the generated images. Unlike conventional metrics that prioritize visual similarity, our evaluation criteria delve deeper into the structural aspects, ensuring that the synthetic images are visually convincing and structurally coherent with real-world data. This allows for quick validation of model results and assurance that data augmentation through GANs is biologically-sound.

Through this model and new evaluation metrics, we aim to set a new standard in the quality of synthetic images generated by GANs, specifically in the realm of medical imaging. The approach promises to contribute significantly to advancements in data augmentation, diagnostic accuracy, and treatment planning by producing images that more faithfully represent the intricate topological features seen in actual medical images.

TopoGAN matches synthetic and real image distributions for image and topology features using a topology-based loss term based on persistent homology, in addition to the conventional generator and discriminator. This loss term measures how close generated images are to real images in terms of topology, and hence minimization forces more topologically accurate images to be generated.

$$\text{Discriminator Loss: } \arg\max_{D} \mathbb{E}_{x\sim p_{\text{data}}}[\log D(x)] + \mathbb{E}_{z\sim p_z}[\log(1 - D(G(z)))]$$

$$\text{Generator Loss: } \arg\max_{G} \mathbb{E}_{z\sim p_z}[\log(1 - D(G(z)))] + \lambda \mathcal{L}_{\text{topo}}(p_{\text{data}}, G)$$

* λ represents the weight of the topological GAN loss.

As we generate binary images, the generator outputs a real- valued grey-scale image as the synthetic mask, and the discriminator treats the input image (real or synthetic) as a real-valued grey-scale image ranging between 0 and 1. Following mask synthesis a separately trained pix2pix network fills in textures based on the mask.

The Philips computerized tomography device was used to capture cardiac CT images for 17 individuals (both males and females) with A-fib, aged between 50 and 62. This device recorded ten sets of timed frames at the same location using various contrast agents, spanning an entire cardiac cycle. Each data set comprised 409 images, each with a resolution of 512×512 pixels [11], [12].

III. RESULTS

The efficacy of TopoGAN in generating topologically accurate synthetic images was rigorously tested through a series of experiments. One of the most compelling pieces of evidence supporting the model's effectiveness is demonstrated in Figure 1, where synthetic images were found to accurately replicate the structural topology seen in actual medical images. Quantitative metrics further substantiate this claim: the calculated Wasserstein distance between synthetic and real images was found to be a minimal 0.0021, indicating an extremely close match between the two distributions.

Additionally, the Dice coefficient, a statistical measure used to gauge the similarity between two samples, was recorded at an exceptionally high 0.995. Such high scores confirm that the generated images are visually similar and have a deep structural resemblance to real medical images. Beyond the generation of synthetic images, TopoGAN has also shown considerable promise in enhancing performance in downstream applications.

Notably, this model combined with our novel evaluation methodology significantly improved the quality of image segmentation tasks. Traditional GAN models have struggled with this application due to their inability to capture intricate topological features, which are often critical in medical imag-

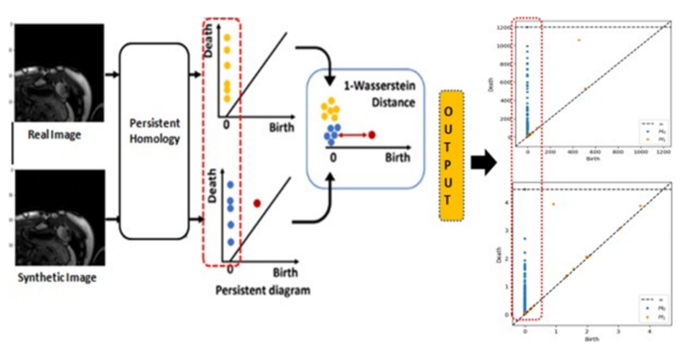

Fig. 1. A plot of a square

ing. However, the topological awareness of TopoGAN seems to overcome these limitations, leading to more accurate and reliable segmentation results.

These outcomes have profound implications for the field of medical imaging. They not only validate the effectiveness of incorporating topological features into GAN models but also signify a major step forward in enhancing the utility of synthetic images in medical applications. From data augmentation to diagnosis and treatment planning, TopoGAN's ability to generate topologically accurate synthetic images opens new horizons for research and clinical practice.

IV. Conclusion

Our research benchmarks a groundbreaking topology-based data augmentation GAN to a real-world data augmentation problem through the use of new metrics. The study demonstrates that this methodology can effectively produce synthetic images that closely mirror the intricate topological properties seen in real-world medical images.

The implications of this work are myriad. By offering a more accurate representation of real images, our model sets the stage for advancements in various facets of medical imaging, including data augmentation, machine learning algorithms, and diagnostic tools. Furthermore, the approach promises to revolutionize medical education by providing more realistic and topologically accurate images for study and analysis.

In essence, this research not only fills a critical gap in current synthetic image generation techniques but also offers a robust solution for capturing the complex topological features that are often pivotal for accurate medical image interpretation and subsequent clinical applications.

Acknowledgment

This research work has been done in the collaboration of WVU, Morgantown, Mayo Clinic, Rochester and University of Dubai.

References

[1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.

[2] S. Nowozin, B. Cseke, and R. Tomioka, "f-gan: Training generative neural samplers using variational divergence minimization," *Advances in neural information processing systems*, vol. 29, 2016.

[3] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," *Advances in neural information processing systems*, vol. 29, 2016.

[4] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*. PMLR, 2017, pp. 214–223.

[5] Y. Mroueh, T. Sercu, and V. Goel, "Mcgan: Mean and covariance feature matching gan," in *International conference on machine learning*. PMLR, 2017, pp. 2527–2535.

[6] F. Wang, H. Liu, D. Samaras, and C. Chen, "Topogan: A topology-aware generative adversarial network," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 118–136.

[7] J. Kossaifi, L. Tran, Y. Panagakis, and M. Pantic, "Gagan: Geometry-aware generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 878–887.

[8] H. Edelsbrunner and J. Harer, "Computational topology: an introduction american mathematical soc.(2010)," *MR2572029*.

[9] Y. Singh, C. M. Farrelly, Q. A. Hathaway, T. Leiner, J. Jagtap, G. E. Carlsson, and B. J. Erickson, "Topological data analysis in medical imaging: current state of the art," *Insights into Imaging*, vol. 14, no. 1, pp. 1–10, 2023.

[10] Y. Singh, C. Farrelly, Q. A. Hathaway, A. Choudhary, G. Carlsson, B. Erickson, and T. Leiner, "The role of geometry in convolutional neural networks for medical imaging," *Mayo Clinic Proceedings: Digital Health*, vol. 1, no. 4, pp. 519–526, 2023.

[11] D. Deepa, Y. Singh, M. C. Wang, and W. Hu, "An automated method for detecting atrial fat using convolutional neural network," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 235, no. 11, pp. 1329–1334, 2021.

[12] Y. Singh, D. Shakyawar, and W. Hu, "An automated method for detecting the scar tissue in the left ventricular endocardial wall using deep learning approach," *Current medical imaging*, vol. 16, no. 3, pp. 206–213, 2020.