

Balancing Accuracy and Efficiency in Inferior Alveolar Canal Segmentation from CBCT Scans

Changkai Ji¹[0009-0007-7090-7360], Yusheng Liu¹[0009-0004-2624-9223], Yuxian Jiang¹[0009-0002-7689-5333][0009-0009-3223-0082], and Lisheng Wang¹[0000-0003-3234-7511]

School of Automation and Intelligent Sensing, Shanghai Jiao Tong University,
Shanghai 200240, People’s Republic of China
{changkaiji, lswang}@sjtu.edu.cn

Abstract. Accurate segmentation of the inferior alveolar canal (IAC) in cone-beam computed tomography (CBCT) scans is crucial for dental and maxillofacial applications, yet it remains highly challenging due to its fine-scale structure, anatomical variability, and imaging artifacts. To solve this problem, we utilized both automated and interactive approaches. Specifically, we compared an nnU-Net-based model with an nnInteractive model. While nnInteractive demonstrated promising improvements with minimal user input, our final submission was based on nnU-Net due to its favorable trade-off between accuracy and computational efficiency. To further enhance runtime performance, we incorporated inference acceleration strategies, achieving a speedup without sacrificing segmentation quality. Our method achieved a top-three ranking in the challenge test phase, highlighting its potential for accurate and efficient IAC segmentation. Our code is available at: <https://github.com/duola-wa/Toothfairy3>.

Keywords: nnU-Net · Interactive segmentation · Inferior alveolar canal · Computational efficiency

1 Introduction

The Inferior Alveolar Canal (IAC) is a crucial anatomical structure in the mandible. It plays a significant role in dental procedures and maxillofacial surgery, where precise delineation is essential for treatment planning, nerve preservation, and risk assessment [2,13,10]. The accurate segmentation of the IAC in Cone Beam Computed Tomography (CBCT) images is vital for ensuring successful outcomes in these clinical applications [16].

Despite its importance, segmenting the IAC in CBCT scans presents considerable challenges. The IAC is often small, located in intricate regions of the mandible, and exhibits significant variability across patients [1]. Its fine-grained structure, combined with artifacts and noise in the CT scans, complicates the task further [12]. Moreover, the complexity of the mandibular anatomy requires precise localization and accurate delineation, which can be difficult for fully automated models to achieve reliably [11].

In recent years, deep learning-based methods have been widely explored for automating segmentation tasks in medical imaging [17,6,14,4,5]. These approaches have shown impressive potential for segmenting complex anatomical structures. However, these models often face a trade-off between segmentation accuracy and computational efficiency. While high accuracy is crucial for clinical applicability, deep learning models can be computationally expensive, leading to long processing times or excessive memory usage, making them impractical for real-time clinical workflows [8,9].

The ToothFairy3 Task 2 in the MICCAI 2025 competition focuses on the segmentation of complex IAC structures in CBCT scans. This task aims to advance the state of the art in IAC segmentation by encouraging participants to explore both automated and interactive approaches. The challenge seeks to improve segmentation accuracy while ensuring computational efficiency. To address these challenges, we implemented and compared two methods: an interactive segmentation approach termed nnInteractive [7], and an nnU-Net-based model. The nnInteractive method incorporates user inputs, allowing clinicians to guide the segmentation process with minimal effort, enhancing precision. After careful evaluation of both methods, considering both precision and computational efficiency, we selected nnU-Net-based model as our submission version. Our contributions can be summarized as follows:

- We explored both interactive and automated segmentation approaches, comparing their performance in IAC segmentation tasks.
- We balanced segmentation accuracy with computational efficiency, showing the potential for further enhancement in clinical applicability.
- Our method secured a top-three ranking in the ToothFairy3 Task 2 test phase, showcasing its effectiveness in IAC segmentation.

2 Proposed Method

2.1 Framework Overview

As shown in Fig. 1, we utilize two methods for the segmentation of the IAC: an nnU-Net-based approach and an interactive segmentation approach nnInteractive. After obtaining the initial segmentation results from the automated nnU-Net-based approach, we apply post-processing techniques to refine the predictions and obtain the final segmentation output. The second method, nnInteractive, introduces interactive point-based refinement. In this approach, we iteratively utilize 1 to 5 interaction points as point prompts, which guide the model’s segmentation process. After each refinement step, the model updates the segmentation based on the inputs. Similar to Option 1, post-processing is applied to the results, ensuring that the final segmentation is both accurate and clinically feasible.




Fig. 1. Overview of the two segmentation approaches. The first method involves nnU-Net-based. The second method nnInteractive incorporates interaction points for iterative refinement of the segmentation.

2.2 Data Preprocessing

Before training the model, the dataset underwent several preprocessing steps to ensure proper input for the segmentation task. First, the dataset integrity was verified to ensure the quality and completeness of the images and annotations. This step involves checking for any inconsistencies, missing data, or potential errors in the dataset, ensuring that all the images are properly aligned with their corresponding ground truth annotations.

Next, the data was processed to fit the input requirements of the segmentation model. This involved standardizing the image size, orientation, and intensity values to ensure consistency across all images. The images were also resampled to match the target resolution, ensuring that the spatial dimensions were uniform.

These preprocessing steps helped to minimize variability across the dataset and ensured that the model could learn from high-quality, consistent input data.

2.3 Model Training Strategy

For the Option 1 model training, we utilized the nnU-Net framework, which automatically adapts to the dataset’s specific characteristics. The training was conducted using a 3D full-resolution architecture, focusing on the highest resolution for optimal segmentation performance. Notably, we disabled the mirror augmentation in the data preprocessing step. This was done because the left and right inferior alveolar nerves are anatomically similar, and applying mirror augmentation could increase the difficulty in accurately assigning categories during segmentation.

For Option 2, we leveraged the pre-trained weights provided by the official repository of nnInteractive, as the training process for nnInteractive was not yet publicly available. These pre-trained weights enabled us to perform a comparative analysis of the interactive segmentation method without the need for additional training from scratch.

nnInteractive adopts a classic U-Net architecture for image segmentation, but with a distinct approach compared to other segmentation models like SAM. Unlike models that use separate image and prompt encoders, nnInteractive employs an "early prompt strategy." In this method, the user-provided prompts and the original image are concatenated along the channel dimension before being fed into the segmentation network. The architecture utilizes eight channels, including the original image, existing segmentation masks, and six additional channels representing different forms of interaction: foreground points, background points, foreground scribbles, background scribbles, lasso, and bounding boxes. To handle large targets and high-resolution 3D medical images, nnInteractive incorporates an AutoZoom mechanism. This dynamic adjustment of the region of interest (ROI) and multi-stage optimization addresses the problem of target truncation and loss of details caused by fixed input block sizes. By combining flexible user interaction with advanced segmentation techniques, nnInteractive significantly enhances the model’s ability to handle anatomical structures with high precision, making it highly suitable for clinical applications where user guidance can be seamlessly integrated into the workflow.

2.4 Post-processing

The post-processing step refines the segmentation results by applying connected component analysis to identify distinct regions in the predicted mask. Label shape statistics are then used to assess the size of these regions, and a minimum area threshold is applied to remove small, irrelevant components. This step aims to ensure that only the clinically relevant structures, thereby enhancing the accuracy and reliability of the results.

3 Experiments and Results

3.1 Dataset and Assessment Metrics

The dataset consists of a total of 532 volumes. Each volume contains two classes corresponding to the left and right branches of the IAC [3,2,15].

To assess the performance of the segmentation methods, several metrics are computed for each IAC branch. These include the Dice Similarity Coefficient (DSC) and Hausdorff Distance at the 95th percentile (HD95) for the final segmentation, as well as the area under the curve (AUC) for both DSC and HD95 over the five interaction steps. Additionally, the average inference time and maximum memory usage are recorded for each case. These combined metrics provide a comprehensive assessment of both the accuracy and efficiency of the segmentation models.

3.2 Implementation details

Environments and Requirements. The nnU-Net model was trained for 1000 epochs, with the computational environment and dependencies detailed in Table 1.

Table 1. System Configuration

Ubuntu version	Ubuntu 24.04 LTS
CPU	Intel(R) Xeon(R) Platinum 8352S CPU @ 2.20GHz
RAM	503 GB
GPU	1 NVIDIA GeForce RTX 4090 (24G)
CUDA version	12.4
Programming language	Python 3.12.11
Deep learning framework	PyTorch (torch 2.6.0, torchvision 0.21.0)
Code will available at	https://github.com/duola-wa/Toothfairy3

Inference Acceleration. To improve inference efficiency, we implemented several strategies due to the importance of runtime in the challenge evaluation. First, we disabled augmentation during the inference phase of nnU-Net, which significantly reduced computational load while maintaining accuracy. Additionally, we enhanced the multi-class prediction process by refining the interpolation step. Instead of using traditional integer-based resampling methods, we utilized PyTorch’s interpolate function on floating-point tensors, which preserved numerical precision and improved throughput for large volumetric segmentation tasks. These optimizations enabled faster nnU-Net inference across the entire test set.

3.3 Results and Analysis

nnU-Net-based Model. The results of the debug and test phases are summarized in Table 2. In the debug phase, the performance of Option 1 across several

metrics was relatively strong, as demonstrated in the visualization shown in Figure 2, where the debug data results are compared to the ground truth. This comparison highlights the model’s segmentation accuracy on the debug data. However, the test phase results reveal the challenges of generalizing the model to unseen data, with some test volumes demonstrating lower performance compared to the debug phase. Additionally, due to platform limitations, the test phase results do not include runtime and AUC metrics.

Table 2. Segmentation performance on Debug and test sets.

Metric	Debug				Test			
	Mean	Std	Min	Max	Mean	Std	Min	Max
Dice_AUC	4.49	0.10	4.38	4.58	–	–	–	–
HD95_AUC	6.38	1.20	5.00	7.07	–	–	–	–
Dice_Final	0.899	0.020	0.877	0.917	0.768	0.190	0.175	0.937
HD95_Final	1.276	0.239	1.000	1.414	32.232	85.896	1.000	376.323




Fig. 2. Visualization of the debug phase results. The segmentation results are compared to the ground truth to demonstrate the accuracy of Option 1 in the debug phase.

nnInteractive. As shown in Fig. 3, we provide a visual comparison of the segmentation results for two options. For nnInteractive, we show results with the introduction of 3 and 5 interaction points, respectively. This visualization high-

lights the impact of increasing the number of user interactions on the segmentation accuracy. Due to the limited number of submissions in the competition, we did not include metric-based results in this analysis, focusing instead on the visual comparison of the segmentation outputs.




Fig. 3. The nnInteractive method is evaluated with 3 and 5 interaction points, highlighting the effect of prompt refinement on segmentation accuracy.

Upon analyzing the results, despite nnInteractive not being pre-trained on the competition dataset, still produces promising outcomes in most cases. The method demonstrates strong potential for interactive segmentation, achieving competitive results with a relatively small number of interaction points. However, considering the trade-off between accuracy and runtime, we chose nnU-Net for our submission.

Inference Acceleration. We implemented inference acceleration for the nnU-Net model to enhance runtime performance without compromising segmentation accuracy. Due to the submission limit, we did not include specific metrics. However, as shown in Fig. 4, the segmentation results before and after acceleration are nearly identical. Additionally, testing on external CBCT data shows that, after applying the acceleration, the model’s inference speed increased by a factor of four on 40 external CBCT volumes. This highlights the effectiveness of our acceleration approach, which successfully improves inference speed while maintaining stable accuracy.




Fig. 4. Comparison of segmentation results before and after acceleration. The results are compared to the ground truth, demonstrating that the segmentation performance remains nearly identical after applying the acceleration.

4 Conclusion

In this work, we addressed the challenging task of IAC segmentation in CBCT scans as part of the ToothFairy3 Task 2 challenge. To address this problem, we utilized both an automated nnU-Net-based approach and an interactive segmentation approach (nnInteractive). The automated nnU-Net model demonstrated competitive performance and was selected as our final submission due to its trade-off between segmentation accuracy and computational efficiency. However, we believe that as the training capabilities of nnInteractive become available in the future, its performance will further improve, offering even greater potential for interactive segmentation tasks. In particular, inference acceleration strategies further improved runtime efficiency while maintaining stable performance, making the method potential for applications. These enhancements not only made the model faster but also demonstrated its potential for applications requiring both efficiency and high accuracy in medical imaging tasks.

5 Acknowledgement

We would like to thank the organizers of the competition for their meticulous data organization and the opportunity to participate. We also greatly appreciate their continuous and timely support in addressing the questions and concerns that arose during the competition.

References

1. Anıl, A., Peker, T., Turgut, H., Gülekon, I., Liman, F.: Variations in the anatomy of the inferior alveolar nerve. *British Journal of Oral and Maxillofacial Surgery* **41**(4), 236–239 (2003)
2. Bolelli, F., Lumetti, L., Vinayahalingam, S., Di Bartolomeo, M., Pellacani, A., Marchesini, K., Van Nistelrooij, N., Van Lierop, P., Xi, T., Liu, Y., et al.: Segmenting the inferior alveolar canal in cbcts volumes: the toothfairy challenge. *IEEE Transactions on Medical Imaging* (2024)
3. Bolelli, F., Marchesini, K., van Nistelrooij, N., Lumetti, L., Pipoli, V., Ficarra, E., Vinayahalingam, S., Grana, C.: Segmenting maxillofacial structures in cbct volumes. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 5238–5248 (2025)
4. Cipriano, M., Allegretti, S., Bolelli, F., Pollastri, F., Grana, C.: Improving segmentation of the inferior alveolar nerve through deep label propagation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 21137–21146 (2022)
5. Di Bartolomeo, M., Pellacani, A., Bolelli, F., Cipriano, M., Lumetti, L., Negrello, S., Allegretti, S., Minafra, P., Pollastri, F., Nocini, R., et al.: Inferior alveolar canal automatic detection with deep learning cnns on cbcts: development of a novel model and release of open-source dataset and algorithm. *Applied Sciences* **13**(5), 3271 (2023)
6. Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., et al.: nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486* (2018)
7. Isensee, F., Rokuss, M., Krämer, L., Dinkelacker, S., Ravindran, A., Stritzke, F., Hamm, B., Wald, T., Langenberg, M., Ulrich, C., et al.: nninteractive: Redefining 3d promptable segmentation. *arXiv preprint arXiv:2503.08373* (2025)
8. Ji, C., Liu, Y., He, L., Jiang, Y., Huang, C., Wang, L.: Two-stage semi-supervised nnu-net framework for tooth segmentation in cbct images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 100–109. Springer (2024)
9. Ji, C., Liu, Y., He, L., Jiang, Y., Huang, C., Wang, L.: A two-stage semi-supervised nnu-net model for automated tooth segmentation in panoramic x-ray images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 91–99. Springer (2024)
10. Jiang, Y., Liu, Y., Ji, C., Wang, L.: Enhanced multi-structure segmentation in cbct images with adaptive structure optimization. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 30–40. Springer (2024)
11. Khalil, H.: A basic review on the inferior alveolar nerve block techniques. *Anesthesia Essays and Researches* **8**(1), 3–8 (2014)
12. Li, Y., Ling, Z., Zhang, H., Xie, H., Zhang, P., Jiang, H., Fu, Y.: Association of the inferior alveolar nerve position and nerve injury: a systematic review and meta-analysis. In: *Healthcare*. vol. 10, p. 1782. MDPI (2022)
13. Liu, Y., Xin, R., Yang, T., Wang, L.: Inferior alveolar nerve segmentation in cbct images using connectivity-based selective re-training. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 3–12. Springer (2024)

14. Liu, Y., Zhang, S., Wu, X., Yang, T., Pei, Y., Guo, H., Jiang, Y., Feng, Z., Xiao, W., Wang, Y.P., et al.: Individual graph representation learning for pediatric tooth segmentation from dental cbct. *IEEE Transactions on Medical Imaging* (2024)
15. Lumetti, L., Pipoli, V., Bolelli, F., Ficarra, E., Grana, C.: Enhancing patch-based learning for the segmentation of the mandibular canal. *IEEE Access* **12**, 79014–79024 (2024)
16. Pogrel, M.A.: Permanent nerve damage from inferior alveolar nerve blocks—an update to include articaine. *Journal of the California Dental Association* **35**(4), 271–273 (2007)
17. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)