An Efficient Approach to Scaling the Panoramic X-ray Tooth Instance Segmentation Dataset

MICCAI STS 2024 Challenge

Problems

- Object Detection Problem
- Semantic Segmentation Problem
- Active Learning Problem
- Instance Segmentation Problem

Baseline Model - Instance Segmentation

The baseline model solves the problem in an end-to-end manner, as per the emphasis on performance stressed by the organizers, we decided that we want a model which is fast and reliable.

Ultralytics YOLO v8

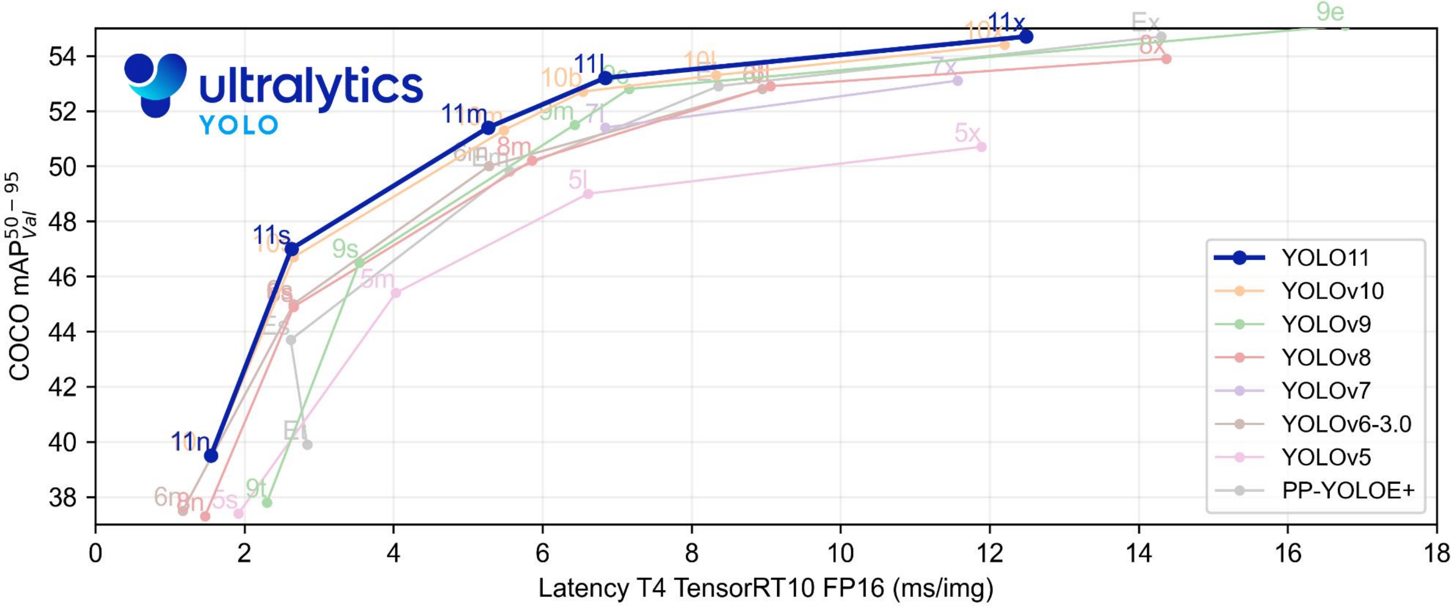

Segmentation Accuracy Problem

Several changes needed to be done for our specific task. The ultralytics implementation computes segmentation loss with a ⅛ downsampled feature map, this would not satisfy our requirement for accuracy.

Feature Pyramid Network

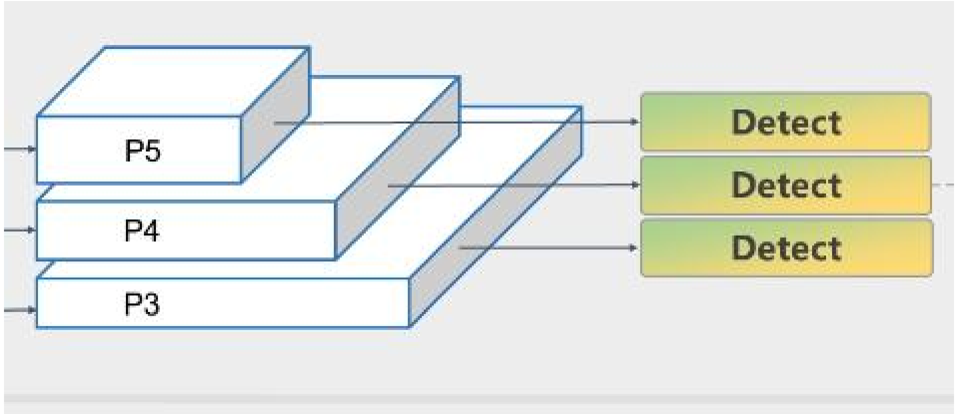

Improved FPN Layer

We modified the ultralytics implementation with additional upsampling operations, which results in more accurate segmentations on the instances.

Bounding Box Detection Problems

When the bounding box detection fails, sometimes it crops out part of the tooth. We do not want that as we would need to train the model to correctly segment the whole tooth from its background, disregarding how much of background there is.

Bounding Box Detection Problems

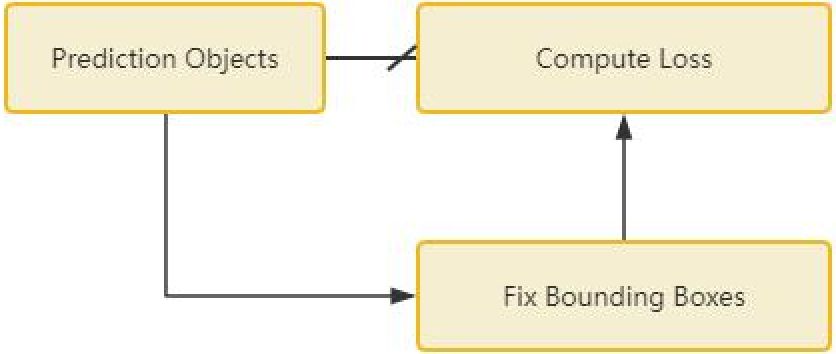

Improved Loss

As described in the previous slides, we have modified the original loss function to fit better the task at hand.

STS 2024 Dataset

- Training set: 2400 panoramic X-ray images, including 20 cases with labels and 2380 unlabeled cases
- Validation set: 20 panoramic X-ray images
- Testing set: Undisclosed to participants

Semi-Supervision Challenge

The dataset in question presents a semi-supervision challenge, as we have noted, there is a large repository of unlabelled data within the training set. Accommodations must be made for this fact.

Automated Data Procurement

Since it is always better to opt for fully supervised training if possible, we formulate this problem into an automated way to assign labels to a large unlabelled repository of data from limited data we have.

Segment Anything Model

To acquire segmentation, we harness the power of SAM.

Segment Anything Model

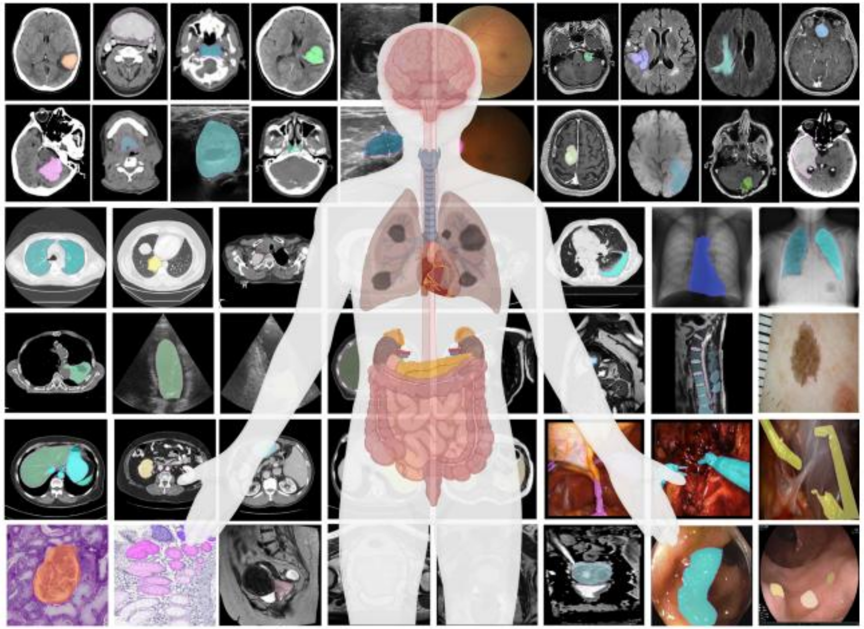

Trivial Object Detection Problem

Given the limited training data, it is possible to produce a simple object detection model to produce bounding boxes that we will need as prompts to SAM. Of course the problem of tooth number assignment must be resolved.

Tooth Number Assignment Problem

Tooth number is a rather unique problem to this challenge, many a times hard to distinguish. Tooth number is a property that is unique to each instance, and often both spatial information and structural information play a part in its determination. This is difficult to resolve alone in a object detection problem.

Maximum Weight Matching

We instead formulate the problem based on the output of our model, disregarding the possibility of duplicate classes in the model itself. We then resolve conflicts based on the confidence score the model has on each of the predictions. Which we then solve with a trivial hungarian algorithm.

Task-Specific Finetune

Some segmentation details are task specific, for instance, SAM will always segment some unerupted teeth as a circular shape, but the provided GT labels will only cover part of the crown. We collect information such as this and finetune the model to offset the differences.

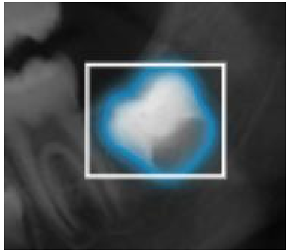

Iterative Optimization of Data

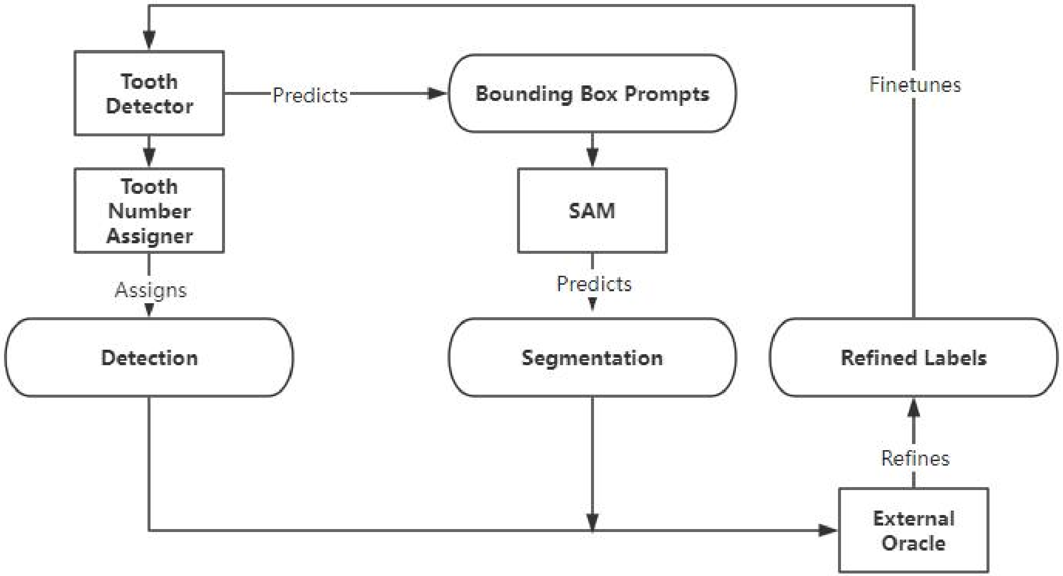

Fully Supervised Training

Finally, we produce the final model with fully supervised training when we collected all the labelled data.

References

Ma, J., He, Y., Li, F., Han, L., You, C., & Wang, B. (2024). Segment anything in Medical Images. *Nature Communications*, *15*(1). https://doi.org/10.1038/s41467-024-44824-z

Jocher, G., Qiu, J., & Chaurasia, A. (2023). Ultralytics YOLO (Version 8.0.0) [Computer software]. https://github.com/ultralytics/ultralytics