Tyche: Stochastic in Context Learning Model for Medical Segmentation

Marianne Rakic, Jose Javier Gonzalez Ortiz, Hallee Wong, Beth Cimini, John V. Guttag & Adrian V. Dalca

Motivation

Challenge 1: Task Diversity

Many learning tasks -> need to retrain for every new dataset.
Degrees of variability: Anatomy, Structure in one anatomy, Modality, Machine settings…

Solution: Use examples to generalize (context).

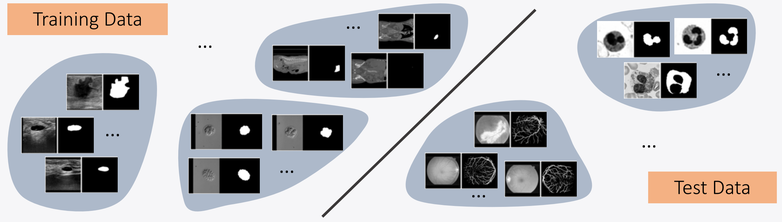

Goal: Avoid need to retrain

Challenge 2: Inherent Stochasticity

- Data inherently ambiguous.
- Labels can be segmented differently by different people.
- Existing methods output a single prediction and don't model this.

Solution: Predict multiple segmentation candidates to capture ambiguity.

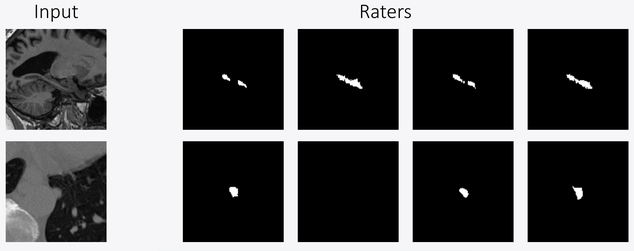

Goal: Capture stochasticity

Intuition behind our loss function

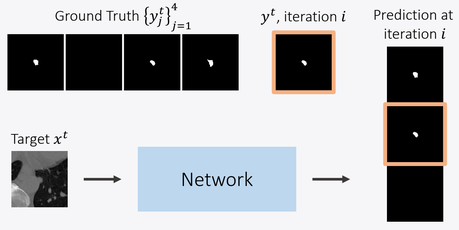

Goals
- Get the network to output diverse predictions.
- The network can hedge its bests.
- Network can predict more than the average.

Tyche produces multiple candidate segmentations for images from unseen biomedical datasets without fine-tuning

Method

Architecture

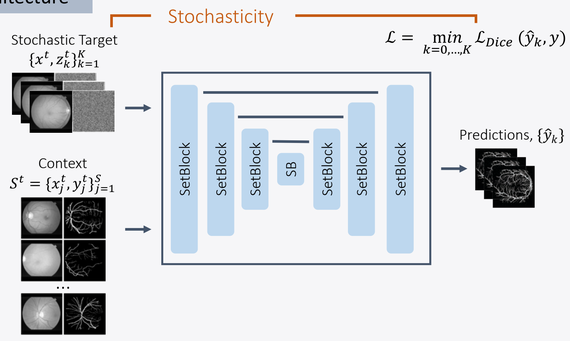

SetBlock: interaction between two sets

The candidate segmentations can interact through the SetBlock.

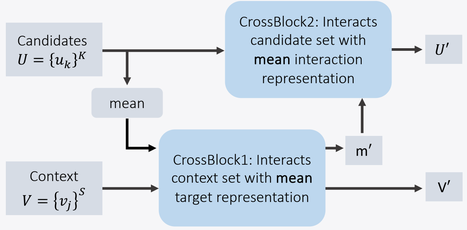

CrossBlock: interaction between single feature and set

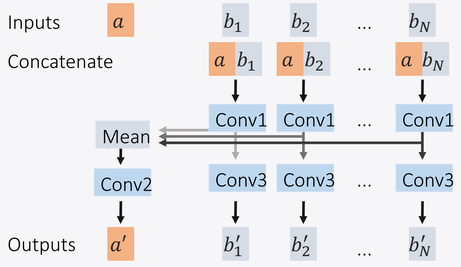

Results

Multi-Annotator Visualization

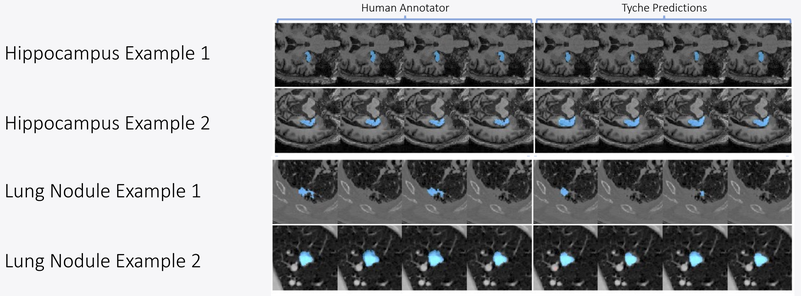

Tyche outputs diverse examples from 16 image segmentation pairs.

Single Annotator Visualization

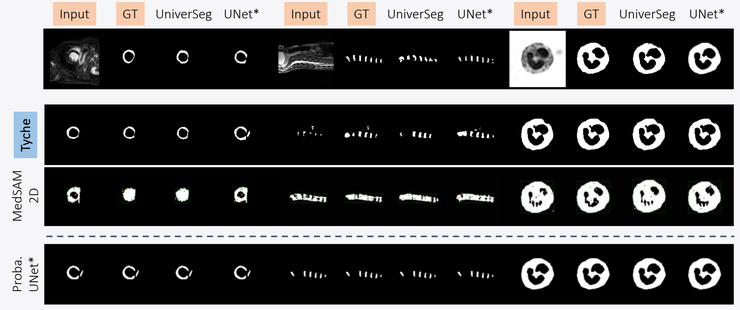

Quantitative Results

- Tyche uses a set of examples (context) to generalize to new tasks and output diverse segmentation candidates.
- Tyche outperforms baselines and is comparable to upper bound benchmarks on various metrics: best candidate Dice, Generalized Energy Distance, diversity.

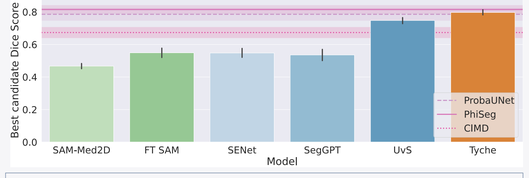

Funding
- Eric and Wendy Schmidt Center
- Quanta
- National Institutes of Health under award number R01EB033773

CVPR 2024 paper