Segment Anything Model on Prostate Cancer Bone Metastases

Krish Gangal, Joseph Rich BS, Steven Cen PhD, Vinay Duddalwar MD FRCR, Assad Oberai PhD

Radiomics Lab, Department of Radiology, Keck School of Medicine, University of Southern California, Los Angeles, CA , USA
Viterbi School of Engineering, University of Southern California, Los Angeles, CA, USA
Bridge Institute, University of Southern California, Los Angeles, CA , USA

Bridge UnderGrad Science (BUGS) ***Summer Research Program***

USC University of Southern California

Abstract

Radiologists review hundreds of images in a single patient's clinical imaging study. In developing an artificial-intelligence based algorithm or platform to analyze this study, manual segmentation of lesions is necessary to develop the ground truth. To do this for thousands of images is time intensive, not practical, and subject to inter and intraobserver variability.

Deep learning, specifically a neural network, is a powerful and promising tool that has shown great success in tumor segmentation, including in prostate cancer bone metastases. The current popular artificial-intelligence-based medical segmentation tools are MONAI and nnU-Net, but Meta AI's Segment Anything Model (SAM) has shown exciting potential.

SAM is a neural network trained on a dataset containing 1 billion images of millions of different images, including tumors. In this project, I assessed the utilization of SAM to segment prostate cancer bone metastases in CT scans.

Objectives

The objective of this project is to assess the differences between SAM and current segmentation tools. We hypothesize that SAM can segment prostate cancer bone metastases with a Dice score > 0.7, the standard success threshold, and compare to other popular segmentation platforms.
The model needs to achieve a Dice (accuracy) score higher than MONAI and nnU-Net.

Segment Everything

Segment Anything works by first segmenting out each feature of the input image, then focusing in on one feature through user feature selection. This is presented as either a bounding box or a point selection (see Figs 4 and %). This process allows for the selection of multiple tumors if an image contains multiple metastases. Additionally, by highlighting multiple features, SAM can remove extraneous features to increase focus.

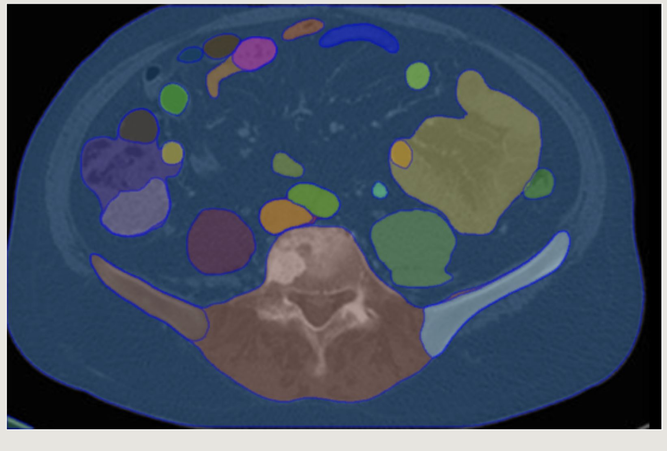

How Segment Anything Works

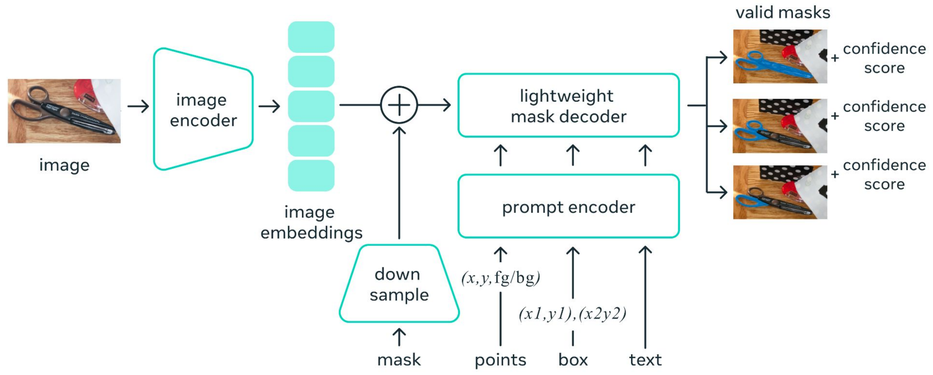

Segment Anything creates prediction masks of input images based on weights from its pretraining process.

Mask Visualization

Fig 1: Input image (first layer)

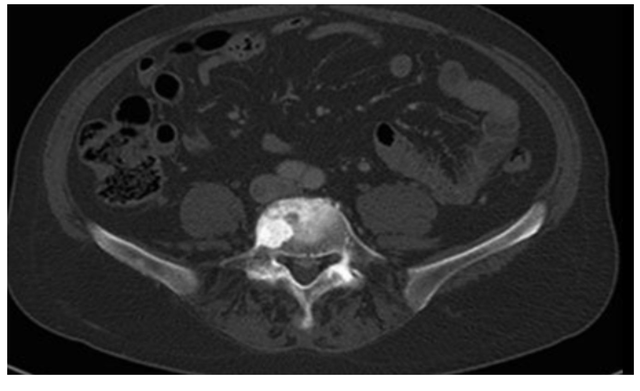

Fig 2: Surrounding features (second layer)

Fig 3: Tumor extraction (final layer)

Fig 4: Point Selection

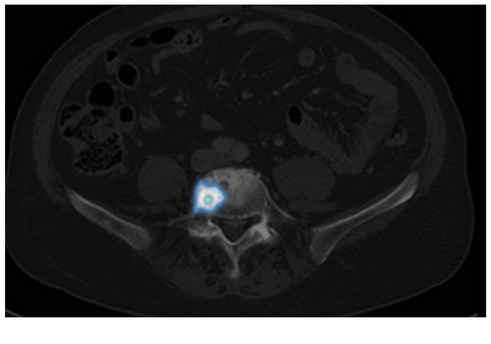

Fig 5: Bounding Box Selection

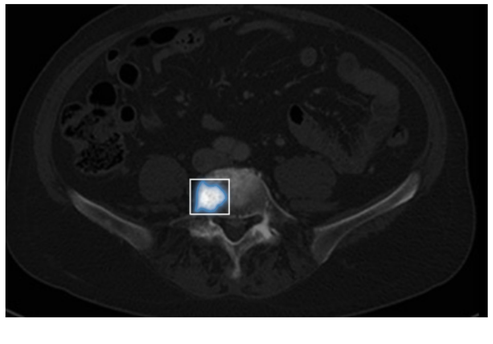

Lesion Editing

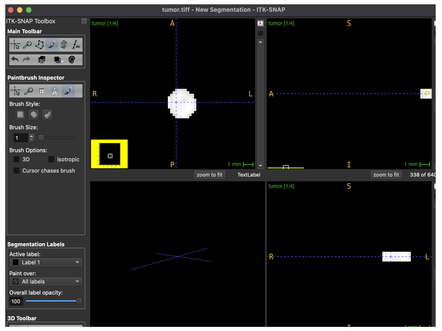

Tools such as Synapse 3D and ITK-Snap, allow for cleaning and extracting just the tumor region from the SAM prediction. This allows for clean viewing and comparisons of tumors.

Bounding Boxes

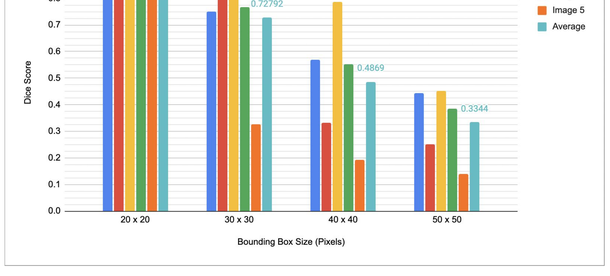

The size of the bounding box around the tumor impacts SAM's accuracy. A larger bounding box is associated with smaller Dice scores.

SAM has an average DSC of 0.728, which is above the 0.7 success threshold and above the 0.4 Dice score of both MONAI and nnU-Net.

Conclusions

- **SAM achieved expected segmentation accuracy in segmenting prostate cancer bone metastases**
- **In this dataset, SAM outperformed MONAI and nnU-Net**
- **Bounding box around the tumor needs to be carefully selected to achieve optimal performance**

All data used was private USC data. The images in this poster are from publicly available datasets.

Funding from:
RSNA and NIH/NCI (grant no. 1R01CA257610-01)

CONTACT US

bridge.usc.edu/bugs
sites.usc.edu/uscradiomicslab/