Lung Nodule Classification and Localization

Suman Gunnala, Anil Luthra, Tony Reina, Kyle M Shannon

Advisors: Dr. Mehrdad Yazdani, Dr. Bradley Voytek, Dr. Gulaka

Abstract

Although chest CT screenings are highly effective at reducing lung cancer mortality, today in 2018, radiologists are overburdened with an ever-growing volume of scans to read. We offer a solution of software-assisted lung nodule detection to aid in faster, more reliable diagnosis.

We accomplish this by training a 3D convolutional neural network (3D CNN) that provides high sensitivity and low false positive predictions. Our models are trained using data from a publicly-available dataset called LUNA16. We created a scalable solution and have tested our model on predictions of whole chest CT scans. Our optimal model achieved a dice coefficient of .814.

I. Data Processing & Storage

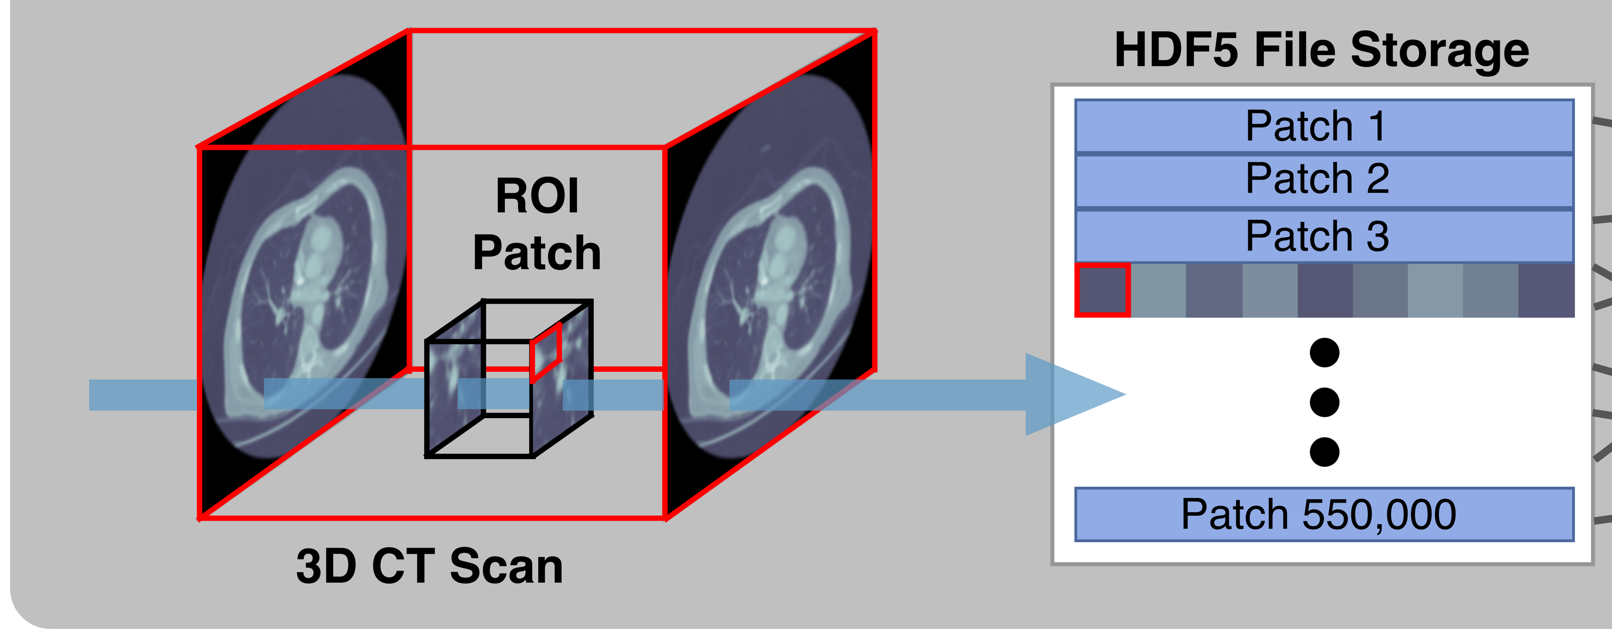

II. Extract *N* Training Data & Generate Masks

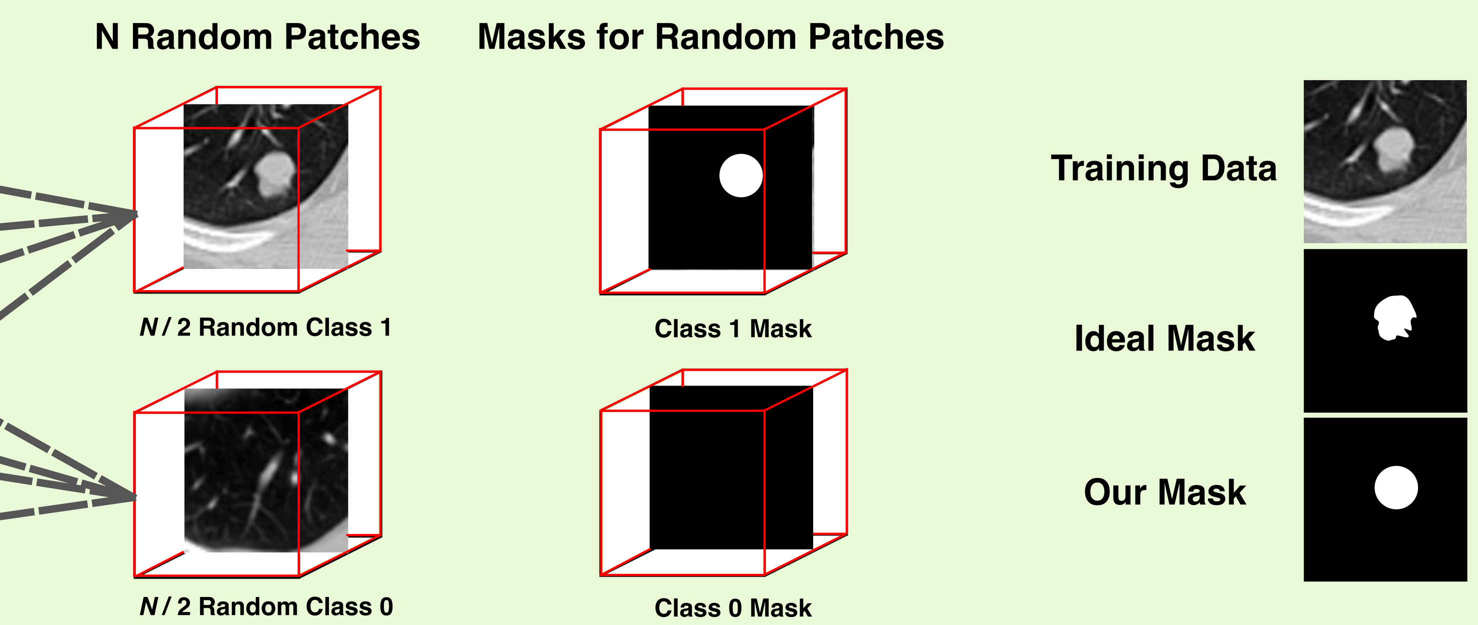

III. Data Transformations

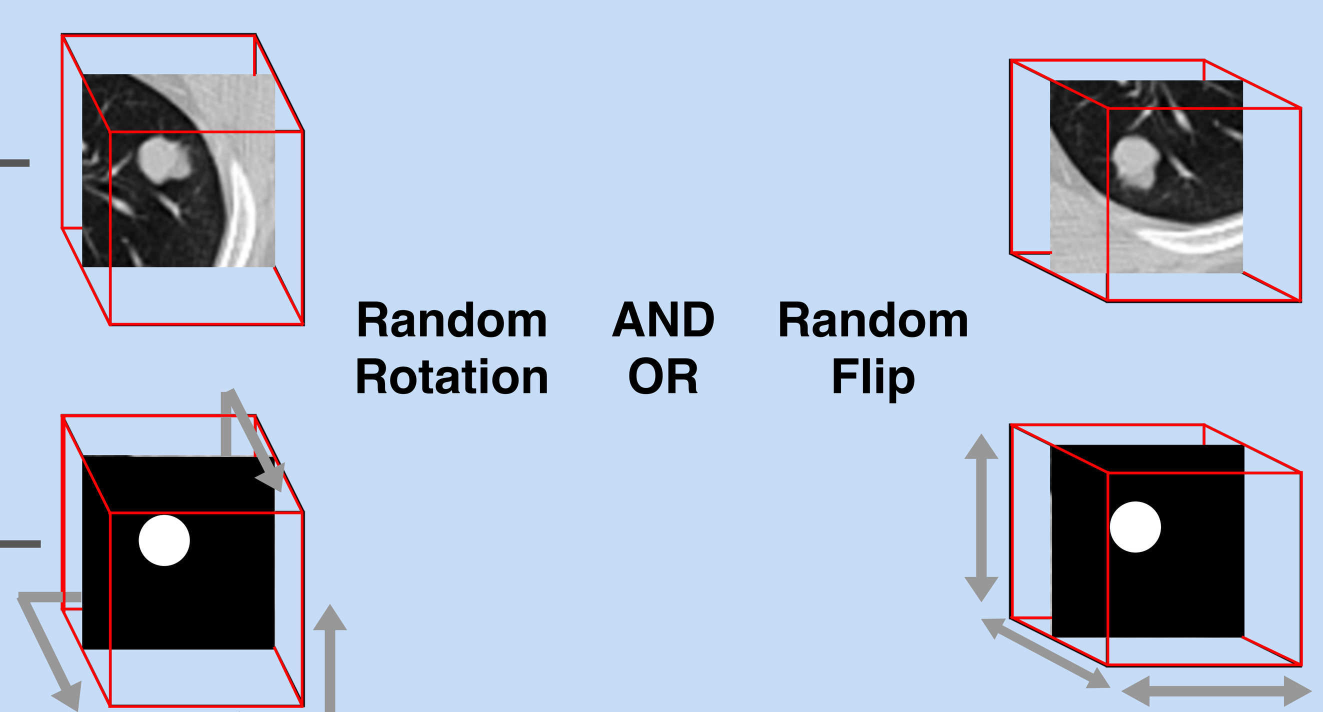

Apply Same Transformation to Both Patch And Its Mask

IV. Patch Cropping Augmentation

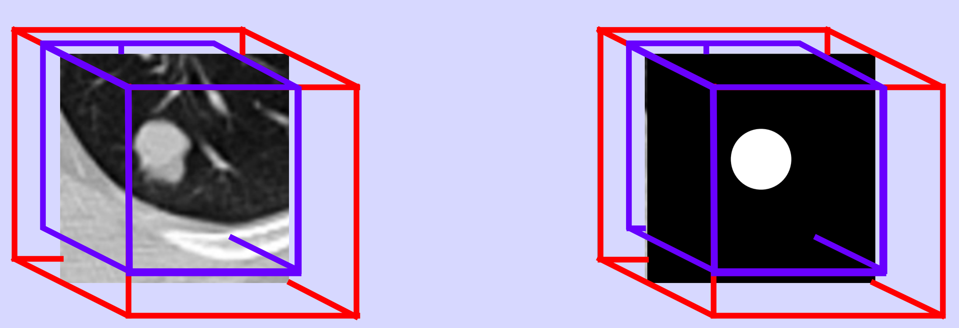

Crop ROI out of Patch Resulting in Center Shifted Smaller Patch

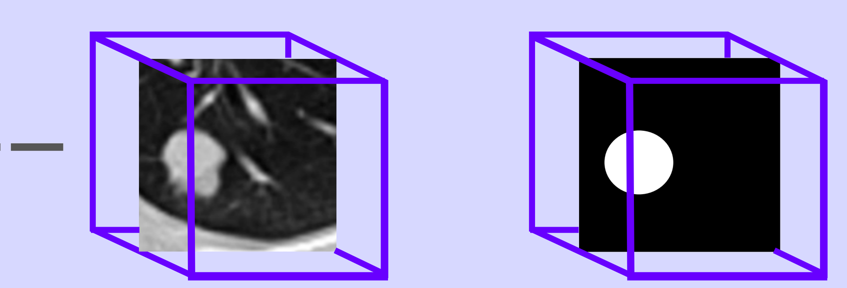

Send Cropped Patch & Mask to Model

V. Training 3D U-Net Model

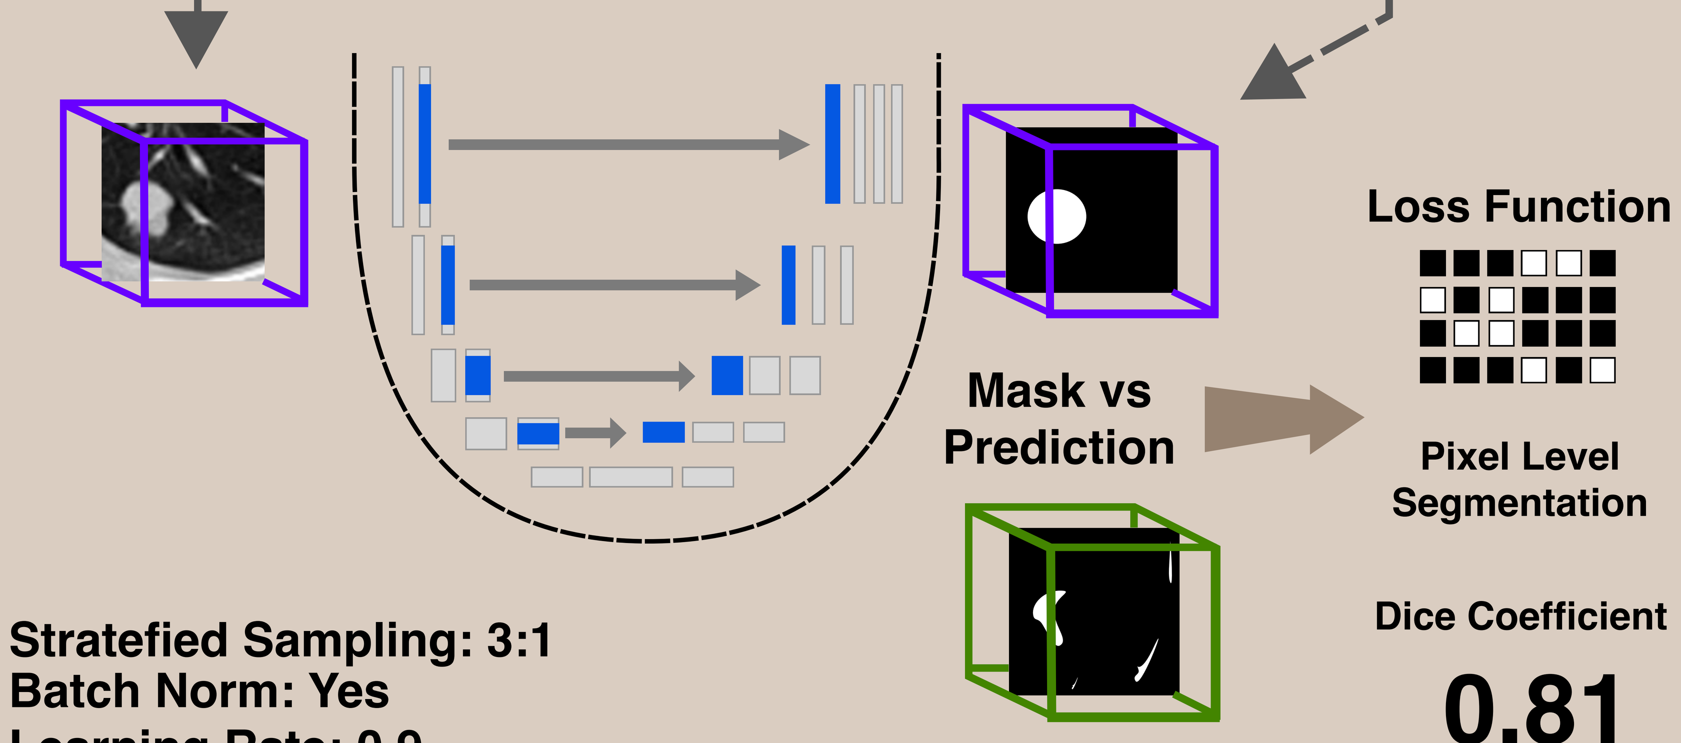

Stratefied Sampling: 3:1
Batch Norm: Yes
Learning Rate: 0.9
Resolution Layers: 4
Total Parameters: ~16 million

Methods

We created a scalable, reproducible processing pipeline that can be rapidly deployed. First, we wrote a patch extractor script that goes through each CT scan and saves regions of interest (ROI), along with associated metadata, into an HDF5 storage object. We faced several challenges in building this patch extractor, including edge clipping issues (where ROIs are too close to an edge); class overlap and centroid collisions (where part of the patch includes both class 0 & 1); and, significant class imbalance (700:1).

Due to the class imbalance, we used mini batch sampling, along with data augmentation techniques (flip, random off-center crops, and rotations) to generate a larger variance of training data. Although expert radiologists marked nodule centers and diameters, they did not segment out the nodules in 2D or 3D space. In order to train a 3D U-Net model for segmentation we generated our own 3D masks using the centroid and diameter measurements to create spherical masks. We trained 3D RESNET models for classification and 3D UNET models for localization tasks. The model architecture and code can be found on out github repository.

Results

Model scoring and validation are based on a free ROC script created by the LUNA16 challenge curators. We trained and tested our U-Net model with 10-fold cross validation. We observed improved average precision score per holdout data set as shown in the figure (below). The average precision score range was 0.563 to 0.756 for 64x64x64 patch. Our 3D segmentation model (right figure) has been shown to achieve good results with dice coefficient of 0.814.

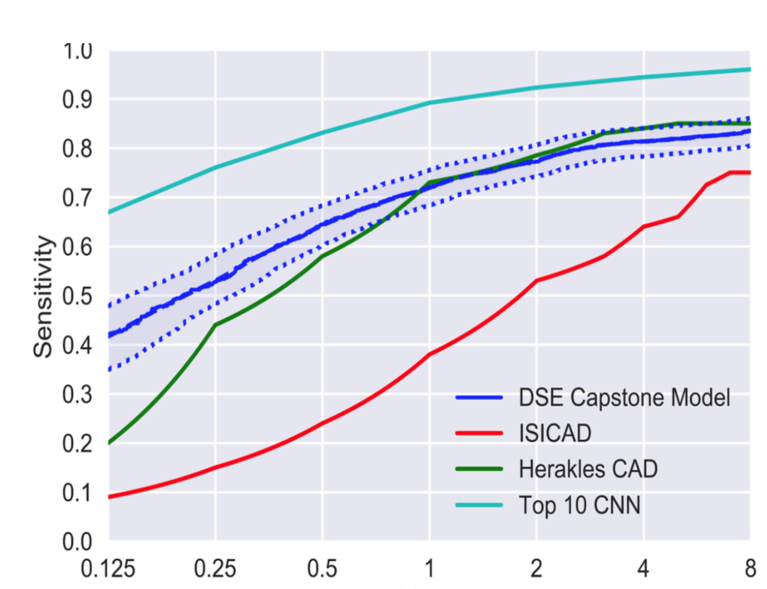

Our models are predicting nodules of varying diameters, off centers, and multiple nodules within the same slice. Inferences are made on various test image sizes which achieves the true image size agnostic segmentation. The feature maps immediately prior to the final prediction indicate that nodule contour and texture are two features the model is using to make its predictions.

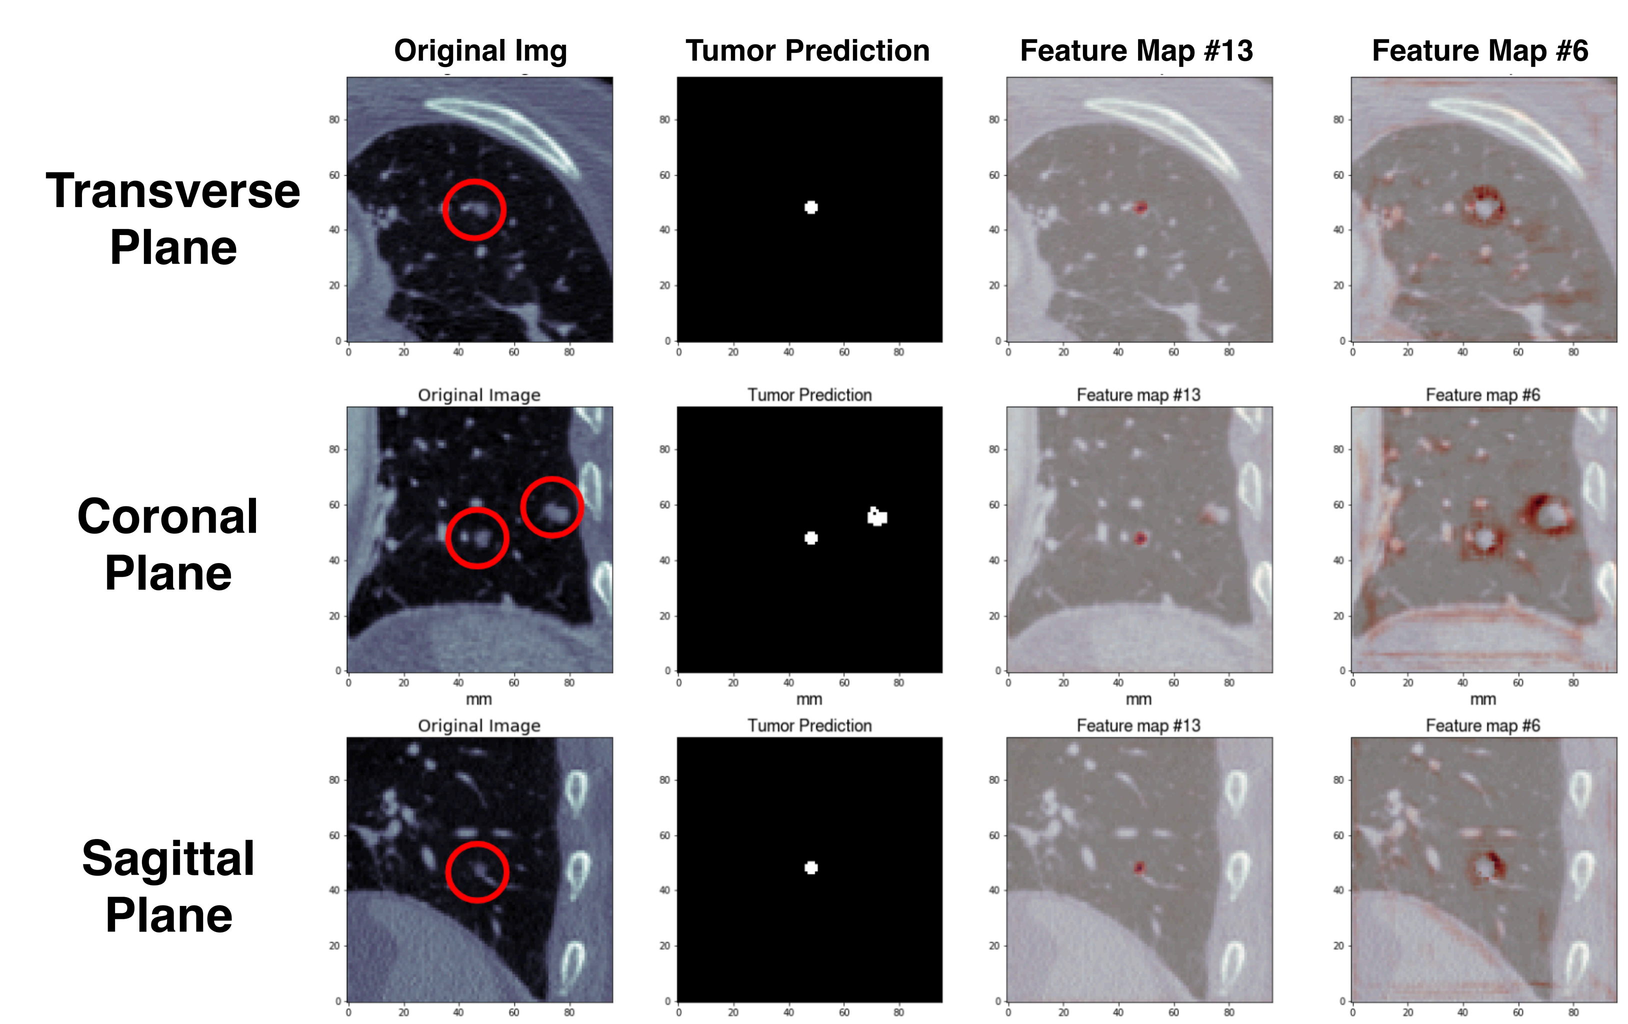

Code.
References.
More Info.
http://bit.do/github-dse

UC San Diego JACOBS SCHOOL OF ENGINEERING

UCSD CSE Computer Science and Engineering

UC San Diego Cognitive Science