What Is It?

I applied an algorithm called *Difference Perceptual Hashing* to the "Voxel-Man Segmented Inner-Organs" dataset. This dataset contains axial CT images which have been labeled according to their visible anatomical structures (see link here: Voxelman). Perceptual Hashing allows me to create a database which can automatically label all the anatomy within a new axial CT image. I envision this technique being used in anatomy education, 3D image reconstruction, and radiotherapy planning.

How It Works

Perceptual Hashing underlies websites like TinEye and Google Reverse Image search. Essentially, it converts images into unique 8-byte numbers, which are stored in a special table for later retrieval. Google's reverse image search allows me to find an image's original location on the web even if that image has been resized, rotated, or watermarked. I've applied this technique to the domain of CT images. This means I can automatically identify a CT image's vertebral/anatomical level. From there, I can automatically label all the anatomy within that image. Starting with a new CT image, the system goes through a few steps:

1. Preprocessing - I need to ensure my image is normalized and in anatomical position. This step requires more thought on my part to achieve better end matching results.
2. Perceptual Hashing - This step converts the image to a unique binary number.
3. Matching - I find the image in the VoxelMan dataset that best matches my new image. If a perfect match is not found, I choose the image with the least Hamming distance.
4. Coregistration - this step overlays my new image with the matched database image, allowing the new CT's anatomy to be automatically labeled.

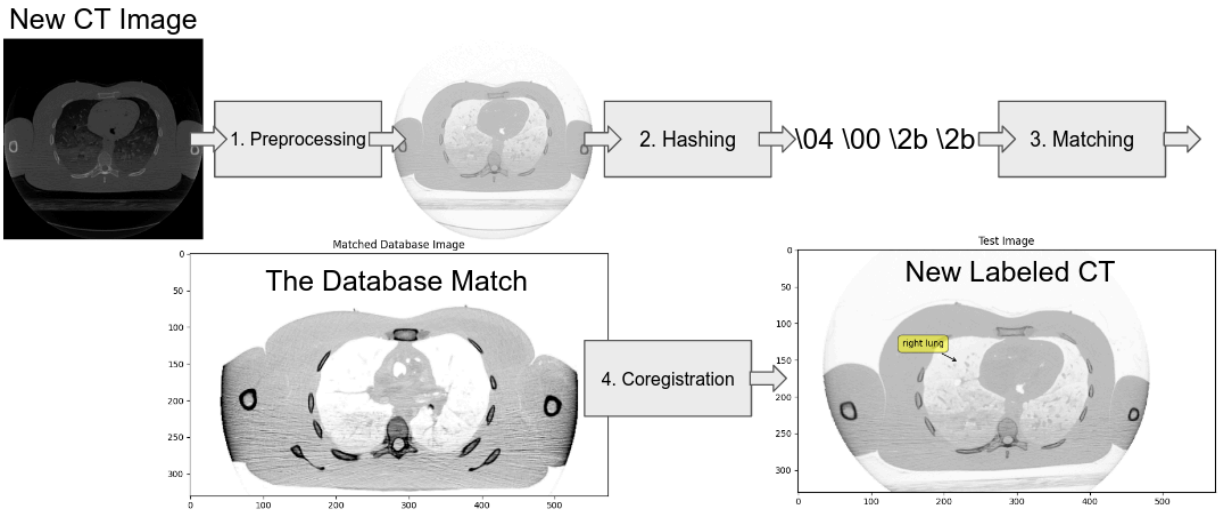

The System in Action

Here is an example of my system in action. The algorithm is by no means perfect, but this is a naive implementation. It can work far better with improved preprocessing methods.

This test image is from the thorax:

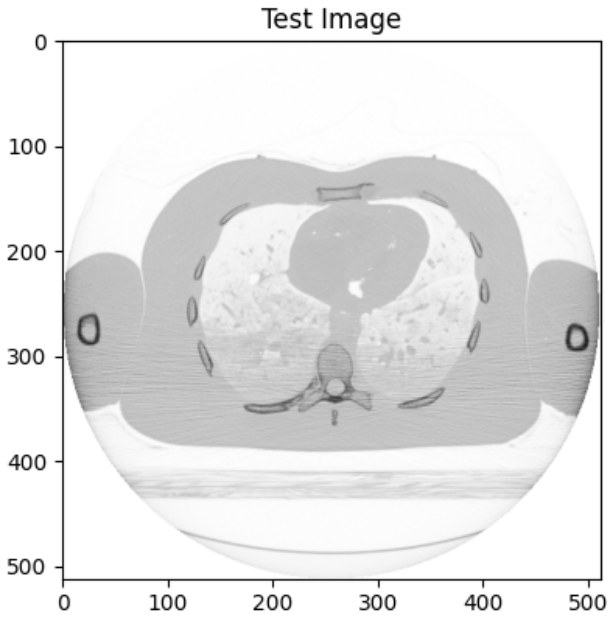

Test image from the thorax

The system then identifies a matching image within the thorax. The match is not perfect (notice the scapula is present in the match image but not the test image).

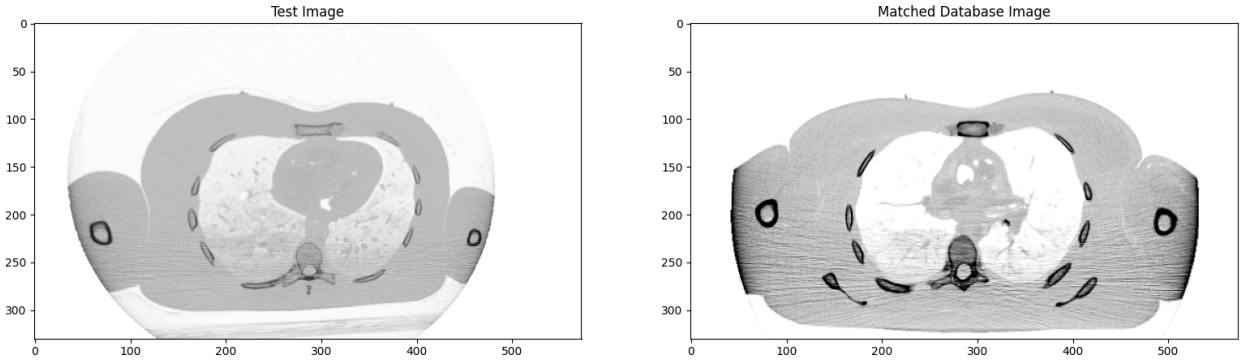

The test image side-by-side with the maged database image. Both are thorax images, but are of slightly different vertebral levels.

With a found match, the images are coregistered to automatically label the test image!

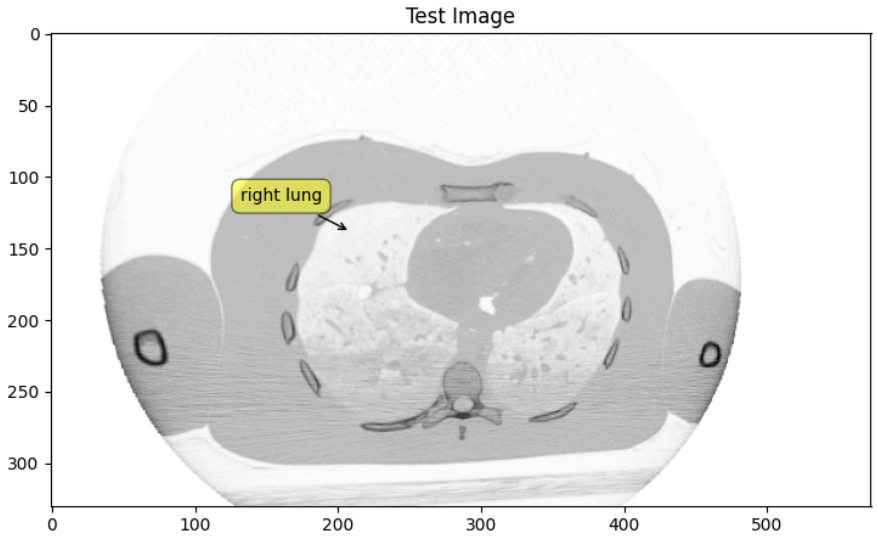

The test image now has a correctly labeled right lung.

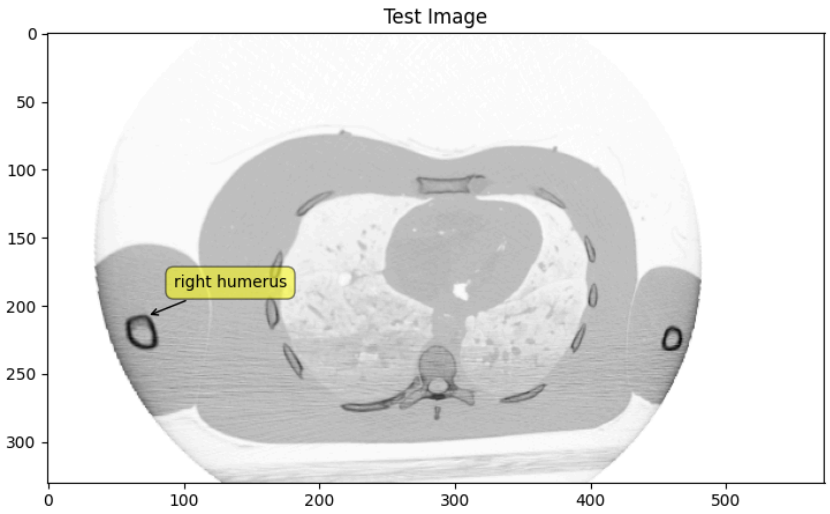

A correctly labeled right humerus.

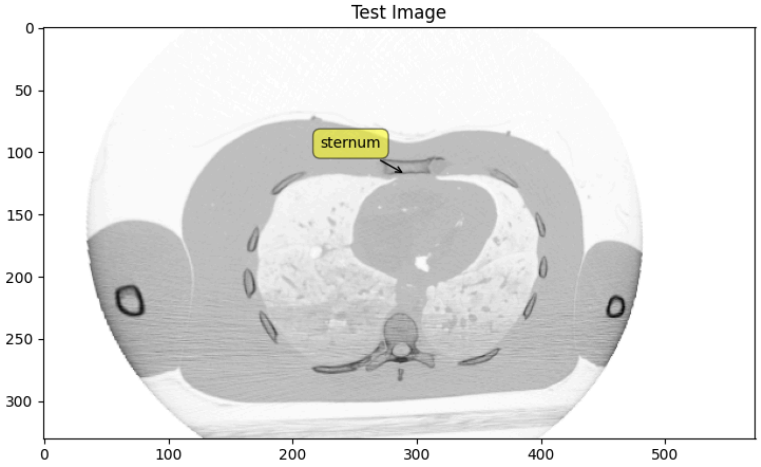

A correctly labeled sternum.

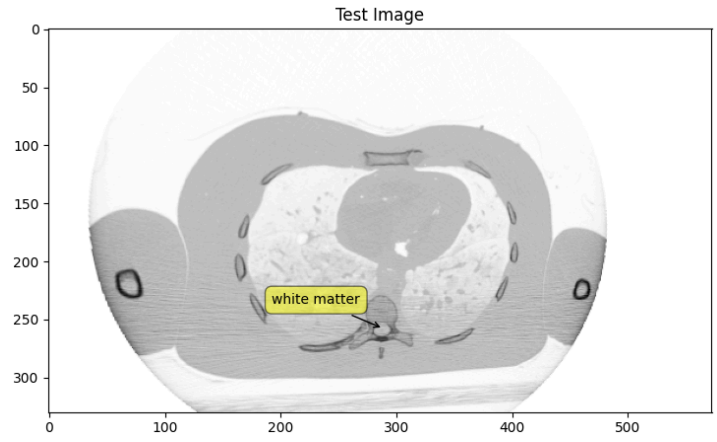

Correctly labeled spinal cord.

Advantages to My System

1. Resistance to Noise - Perceptual hashing is *relatively* invariant to noise, size changes, etc. This means my new CT image *can be correctly matched* even if it's **a different size or contains some degree of pathology**.
2. Not a Deep Learning Technique - I also haven't done too much research on similar work in this field. I think the current paradigm is to use AI/Deep learning to solve these types of problems. My technique is not an AI/Deep Learning technique. This means I don't need a large dataset for training nor large amounts of computational power (could be ran on mobile devices)

Downstream Potential

I have a few ideas of how this idea would be useful to others:

1. Teaching Anatomy to Medical Students - I had trouble learning to read CT images in MGA. I wanted to practice with lots of new images, but I could never trust that I read them correctly without them being labeled. With this technique, I can theoretically take any CT image off of Google Images and have it automatically labeled.
2. Image Segmentation and Reconstruction - In the image reconstruction workshop, we isolated the humerus using pixel intensity. We could use my technique to automatically segment based on the labeled anatomy of the database match.

IP Considerations

To get access to the Voxel-Man dataset, I had to sign an agreement. My algorithm stands alone from this particular dataset, so if this is an issue we should be able to get around it using a different dataset or by creating our own.

I found one article that proposed using Perceptual Hashing on medical images (link). However, this project attempted to label the image only at a gross level (ie shoulder joint, pelvis, GI system, etc.) and not for specific structures within that image (supraspinatus m., L4 vertebra, ascending colon, etc.). I'm not sure what effect this has on the novelty of my idea.

My Code

All my code is located on GitHub: https://github.com/thesalmonification/CTHash. This is a private repository, so if you want access I'll need to add you as a developer.