

Creation of Synthetic X-Rays to Train a Neural Network to Detect Lung Cancer

Abhishek Moturu and Alex Chang

Department of Computer Science, University of Toronto, Toronto, Ontario, Canada

5 October 2018

1 Introduction

Mutated cells that proliferate uncontrollably, known as tumors (or nodules when $\leq 3\text{cm}$ in diameter), are often detected at very late stages when occurring in the lungs. With one of the lowest 5 year net survival rates amongst cancers, lung cancer is one of the deadliest cancers, making early and accurate diagnosis especially important.

X-ray radiographs are a relatively inexpensive and accessible diagnostic test which measure the intensity loss of X-rays passing through the body in a 2D image. In recent studies, neural networks trained with real X-ray radiographs have shown reasonably accurate disease detection rates, with up to 74-91% sensitivity and 75-91% specificity depending on the lung abnormality [1]. However, due to confidentiality regulations, real-patient datasets are scarce, which suggests the need for synthetic X-rays (See **Figure 1**).

Using X-ray Computed Tomography (CT) scans as 3D models of a body, nodules of various shapes, sizes, and radiodensities can be inserted into the CT scan. By inserting nodules at various lung positions, we can generate a large set of diverse synthetic chest X-rays to train a neural network to detect early stage lung cancer (See **Figure 2**).




Figure 1: A real X-ray (left) and a synthetic X-ray (right).

2 Approach

(See full reports [2], [3] and video [4])

A CT scan can be visualized as a 3D array of cuboid-shaped units called volumetric pixels (voxels), which are the smallest part of a 3D object. The value of each voxel of the visualization represents the radiodensity of the medium. Radiodensity, measured in Hounsfield Units (HU), is the relative amount of energy absorbed by a medium.

To synthesize an X-ray image, a ray connecting the point source to each pixel of the resulting X-ray image is simulated. For each ray, the intersecting distance of each voxel along its path is calculated. The parallel-ray approach was initially used as an approximation of the more realistic point source approach (with the point source infinitely far away from the CT scan) for its faster runtime. The development of an efficient ray-tracing voxel traversal algorithm allowed for the use of the more flexible point source approach. Beer's law can then relate the respective intersecting distances and radiodensities to the ray's remaining intensity, which can be translated to a grayscale value to fill in each pixel of the X-ray image.

Dual-energy soft tissue X-rays allow for an unobstructed view of the lungs by subtracting the bones from the image to avoid nodules from being concealed by the ribs. To mimic their effect, all voxels with bone-like radiodensities are replaced with the radiodensity of water prior to tracing the rays.

Growing the nodules to obtain realistic shapes took several tries. The dispersed model, which we explored initially, simulated lung infections rather than lung cancer. Now, the more realistic lobulated model

consists of a large central sphere and iteratively added hemispheres of random sizes at random positions on the surface of the existing shape. Once the nodule shape is generated, it is inserted into the voxel array by modifying the radiodensity values in a certain region of the CT array.

Positions for nodule growth within the CT array are randomly selected from the lung voxels, which can be segmented using a breadth-first search derived algorithm. Two starting positions within each lung are manually selected and neighbouring voxels are explored if their radiodensity values fall within an upper and lower lung radiodensity threshold.

To normalize the grayscale image produced by the ray tracing voxel traversal algorithm, a series of image processing techniques are employed. The images are first rotated and flipped to match the common chest X-ray orientation. Using interpolation methods, images are then resized to the neural network's input size of 256 x 256 pixels. Finally, gamma correction and histogram equalization help spread the cluster of high intensity pixels, adjusting the image contrast.

Different aspects of the project have been programmed in the most suitable languages. Python's NumPy was used to handle large matrices, C++ was used for its flexibility in memory allocation, and MatLab was used for its Image Processing Toolbox.

3 Analysis

After training a neural network with a batch of 2000 synthetic X-rays produced from 6 CT scans using the parallel-ray approach, poor sensitivity was observed. The hypothesis was that a low number of CT scans forced the neural network to recognize patient bodies over nodule presence. The next batch of X-rays will utilize more CT scans, segment lungs for more randomized nodule placement, and use the randomized point source approach for added variance. It will also have more clear bone removal and more accurate contrast. After obtaining more CT scans, we hope to inspect and remove any corrupted or suboptimal CT scans to create approximately 20000 X-rays of training data. This can help us understand if we are making progress on the path towards generating effective synthetic training data.

Reducing space and time complexity was essential to produce thousands of X-rays efficiently. We implemented the ray tracing algorithm and made the empty X-rays from empty CT scans first. Then, we only traced around inserted nodules in the CT scan to update certain "chunks" of the X-ray images to vastly improve the space and time complexity.

4 Future Work

Implementing a Generative Adversarial Network (GAN) may generate more realistic synthetic X-rays [5]. Previous research with GANs on frontal chest X-rays has shown promising results for lung abnormalities.

Additionally, considering ways of creating a balanced training dataset based on prevalence of various aspects such as size, shape, location, or radiodensity may help the neural network identify nodules more naturally. Furthermore, along with the frontal X-rays, it may be useful to train the neural network with lateral X-rays because it may reduce the possibility of nodules blending in with surrounding tissues. With better results, we can develop different methods of nodule generation for more variation in shape and explore other lung abnormalities in the future.

References

- [1] M. Cicero, A. Bilbily, E. Colak, T. Dowdell, B. Gray, K. Perampaladas, and J. Barfett, "Training and validating a deep convolutional neural network for computer-aided detection and classification of abnormalities on frontal chest radiographs," *Investigative radiology*, vol. 52, no. 5, pp. 281–287, 2017.
- [2] A. Moturu and A. Chang, "Creation of synthetic x-rays to train a neural network to detect lung cancer," 2018. Available at http://www.cs.toronto.edu/pub/reports/na/Project_Report_Moturu_Chang_1.pdf.
- [3] A. Chang and A. Moturu, "Detecting early stage lung cancer using a neural network trained with patches from synthetically generated x-rays," 2019. Available at http://www.cs.toronto.edu/pub/reports/na/Project_Report_Moturu_Chang_2.pdf.
- [4] "Training a deep convolutional neural network to detect early stage lung cancer - youtube." <https://www.youtube.com/watch?v=QZNF89cIje4>, 2019.
- [5] H. Salehinejad, S. Valaee, T. Dowdell, E. Colak, and J. Barfett, "Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks," *arXiv preprint arXiv:1712.01636*, 2017.

For more information, contact: Abhishek (abhishek.moturu@mail.utoronto.ca), Alex (le.chang@mail.utoronto.ca).




Figure 2: Examples of an assortment of synthetic X-rays, with and without nodules, generated from the 36 CT scans that we currently possess. Nodules are circled in red. Note that some of the nodules are hard to see by eye as they might be hiding behind the heart, might be nodules on the borders of the lungs (snowball lesions), might have low radiodensity, or might be blending into the surrounding tissues.