Super resolution for atrial fibrillation mapping

Stanford University

Francisco Sahli Costabal, fsc@stanford.edu

Predicting

Atrial fibrillation is the most common rhythm disorder of the heart. It is associated with chaotic electrical waves that lead to rapid and irregular beating of the upper chambers. A novel procedure to diagnose and treat this disease is to insert a catheter with 64 electrodes into the chamber to identify electrical activation patterns [1]. However, this data is sparse and difficult to interpret due to its low resolution. Increasing the quality of these activation maps will help physicians to deliver more effective treatments. This problem is particularly suitable for deep learning, super resolution techniques. Here, we take simulated low resolution images of atrial fibrillation activation times and enhance them with a convolutional neural network.

Data & Features

Because is nearly impossible and unethical to collect simultaneously low resolution and high resolution data from human hearts, we use computational models that simulate the electrical activity of the heart [2]. We generate 400 images of regular activations and more than 800 images of fibrillation of 330x330, representing the activation times. Since super resolution approaches focus only on parts of the images, we obtain 100 sub-images per example of 33x33 pixels. The labels are, at the most, 21x21, depending on the specific filter sizes. Due to the nature of the electrical waves in the heart, rotated and flipped images also represent valid examples. We can augment our dataset 6 times, considering three 90 degree rotations and two flipping operations. The final dataset contains 718,200 images. We separate 10,000 images for development. As features, we use the greyscale image values directly into the model and we also use this as the output.

Models

Among the many super-resolution algorithms available, we choose the approach described in [3]. Briefly, this method uses 3 convolutional layers to predict the high-resolution counterpart of a low resolution image, upscaled using bi-cubic interpolation. The loss is defined as mean squared error between the pixels of the predicted high-resolution image and the training data. Figure 1 shows a diagram of the model. We use Relu activations and He initialization for the parameters. We need 5 hyper-parameters to define the model: 3 filter sizes f_1, f_2 and f_3, and the number of filters for the intermediate layers n_1 and n_2.

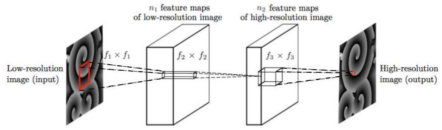

Figure 1 - Convolutional neural network for super-resolution. The model is defined by layers with different number of filters and filter sizes. Adapted from [3]

Following the recommendations outlined in [3], we keep the number of layers to 3, as it has been shown that changing this parameter does not improve the performance significantly. Instead, we focus on the number of filters and the filter size of the second layer. We set $f_1 = 9$ and $f_3 = 5$.

Results

We tested our models with two levels of scaling: 3x and 9x. Our results show that the best models are the ones with more filters ($n_1 = 128$ and $n_2 = 64$) and that increasing the filter size f_2 above 3 either decreases the accuracy or marginally improves it. For both cases, our model performs better than bi-cubic interpolation.

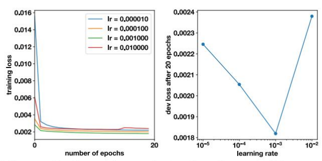

Figure 2 - Learning rate tuning for $n_1 = 64$, $n_2 = 32$ and $f_2 = 1$. The optimal rate is 0.001

scaling	n_1	n_2	f_2	training loss (N = 708200)	dev loss (N = 10000)
3	64	32	1	1.80E-03	1.82E-03
3	64	32	3	1.68E-03	1.76E-03
3	128	64	1	1.70E-03	1.75E-03
3	**128**	**64**	**3**	**1.46E-03**	**1.47E-03**
3	128	64	5	1.79E-03	1.82E-03
3		bi-cubic		3.70E-03	
9	128	64	1	5.98E-03	6.22E-03
9	**128**	**64**	**3**	**5.69E-03**	**5.95E-03**
9	128	64	5	5.67E-03	5.89E-03
9		bi-cubic		1.21E-02	

Examples

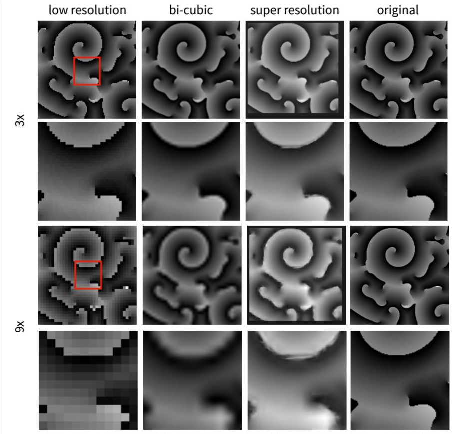

Figure 3 - Examples taken from the test set for 3x and 9x scaling

Discussion & Future

We have used a convolutional neural network to enhance atrial fibrillation maps. Although the mean squared error was reduced with respect to a simple bi-cubic interpolation, the result are not visually pleasant. This could be caused by the mean squared error loss that we used. Despite this, the model is able to better interpret the sharp edges in the image. For the future, we would like to explore different loss functions that could improve our predictions. Additionally, we would also modify all filter sizes to find the optimal architecture.

References

[1] Narayan SM, Krummen DE, Shivkumar K, Clopton P, Rappel WJ, Miller JM. Treatment of atrial fibrillation by the ablation of localized sources: CONFIRM (Conventional Ablation for Atrial Fibrillation With or Without Focal Impulse and Rotor Modulation) trial. Journal of the American College of Cardiology. 2012 Aug 14;60(7):628-36.

[2] Sahli Costabal F, Zaman JAB, Kuhl E, Narayan SM. Interpreting activation mapping of atrial fibrillation: a hybrid computational/physiological study. Ann Biomed Eng. 2018; 46:257-269.

[3] Dong C, Loy CC, He K, Tang X. Image super-resolution using deep convolutional networks. IEEE transactions on pattern analysis and machine intelligence. 2016 Feb 1;38(2):295-307.