

Department of Electrical, Computer, and Software Engineering

Part IV Research Project

Project report

Project Number 14:
Deep learning the geometry of the inverse problem of
electrocardiography

Author: Brandon Wood

Project Partner: Peter Briesch

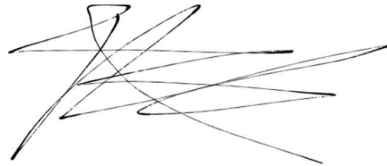
Supervisor: Dr. Avinash Malik

Co-Supervisors: Dr. Mark Trew, Tommy Peng

15th October 2021

Declaration of Originality

This report is my own unaided work and was not copied from nor written in collaboration with any other person.

A handwritten signature in black ink, consisting of several overlapping, sweeping strokes that form a stylized, abstract shape.

Name: Brandon Wood

Contents

1. Abstract.....	2
2. Introduction	2
3. Related work / Literature review	3
3.1. Background to the inverse problem.....	3
3.1.1. The Geometry of the inverse problem	4
3.2. Data selection	4
3.2.1. Data Synthesis.....	5
3.3. Neural networks	5
3.3.1. Time Dependent Neural Networks	6
3.3.2. Decoder/Encoder based Neural Networks	6
3.3.3. Convolutional Neural Networks	6
4. Methods	7
4.1. 2D domain.....	8
4.1.1. Data synthesis and forward model	8
4.1.2. Validation of HSP-BSP pair uniqueness.....	9
4.1.3. Network architecture and training.....	10
4.2. 3D domain.....	10
4.2.1. Data synthesis	11
4.2.2. Network architecture and training.....	12
4.2.3. Experimental predictions on real-world data	13
5. Results	15
5.1. Prediction of Gaussian HSPs in 2D.....	15
5.2. Prediction of gaussian HSPs in 3D.....	15
6. Discussion.....	16
7. Conclusion	17
8. References	18
9. Appendix	19
9.1. 2D prediction plots.....	22
9.2. 3D prediction plots.....	24

1. Abstract

This report proposes a neural network-based solution for the inverse problem of electrocardiography. Given the electrical potential recordings of a mesh of electrodes from a patient's chest, we predict the underlying heart surface potential (HSP) that propagated through the torso to the chest body surface potential (BSP)). Due to the lack of real-world data for training, a dataset is synthesized from a forward model based on experimental data. Both a dense fully connected architecture and deep convolutional neural network (CNN) architecture are tested and evaluated. The CNN outperforms the other architecture and is able to accurately model the relationship between the HSP-BSP pairs. Our final neural network achieves a 2.2% RMSE against the unseen test split of the synthetic dataset. These predicted HSPs can be used to non-invasively diagnose cardiac electrical dysfunctions, a significant step forward for clinical applications.

2. Introduction

The inverse problem of electrocardiography defines the determination of the electrical function on the heart surface from a non-invasive recording, such as a series of electrodes on a patient's chest[1]. A successful solution mapping the heart surface potentials (HSPs) to the body surface potentials (BSPs) would provide a significant addition to current clinical methods. These include but are not limited to source localization for cardiac procedures, pinpointing regions of tissue damage, and diagnostics for arrhythmia[2]. Despite substantial progress and research in this field over the last decade, there has yet to be a clinically adopted solution with acceptable accuracy due to the challenging nature of the problem.

The problem can be summarized by equation (1):

$$\Phi_B(t) = A\Phi_H(t) \quad (1)$$

where Φ_B and Φ_H represent the body and the heart signals respectively and A is a geometric relationship between them for any given t . The inverse problem of estimating A^{-1} is well known to be ill-posed and difficult to solve[1], [3]. Any mathematical method designed to solve for A^{-1} is difficult as small changes in the input can cause unbounded errors in the output solution. Current approaches to the inverse problem estimate HSP from electrocardiographic imaging, using a combination of body surface electrodes and CT scans to reconstruct the activation sequence[4]. It is common for modern solutions to introduce regularization techniques to mitigate the effect of input noise and ensure dependence of the solution on the data. The challenge lies in the estimation of regularization parameters, as there is no best solution for different patient geometries and signal-to-noise ratios [5].

It is proposed that application of a neural network could provide an accurate solution to the inverse problem. State-of-the-art research in this field has begun to apply AI to this problem, but results are still preliminary. We propose the generation of underlying HSPs from BSPs using a neural network trained on Gaussian basis functions, with consideration to the patient's geometry (namely rotation and position of the heart relative to the torso). By mapping recorded ECG signals to a 2d matrix of points, we can train our network on a standard dataset. Our research contribution is an evaluation of different network architectures, specifically success with a modification of U-Net [6] which had worked well with inverse problems in imaging for similar biomedical problems. We demonstrate that in the bounds of our simulation modelling and synthetic dataset, we can accurately predict HSPs from the resulting BSP. Additionally, we prove the ability of a neural network to make accurate predictions despite changes in heart geometry.

3. Related work / Literature review

To develop a broad comprehension about the current state of knowledge in this field and evaluate foundations of the project, a literature review was performed of the problem space and related work.

3.1. Background to the inverse problem

The inverse problem does not possess a complete solution that is unique. That is to say, it is impossible to quantify cardiac sources uniquely without taking measurements on the heart surface, due to the distortion caused by the propagation of the electric field through the body. Early methods for solving the inverse problem concerned epicardial potentials, which presented a simplified view of the problem and neglected more complex electrophysiological processes [7]. There has been significant development in numerical algorithms for solving the inverse problem, which has resulted in an increasing number of clinical ECGI applications. These include the Bidomain model (and its simplification, the monodomain model), Tikhonov regularization and reproducing kernel Hilbert space mapping. These methods provide more realistic solutions to the inverse problem, but still require regularization and are highly sensitive to noise. We can also define the inverse problem as a regression problem[8], approximating A^{-1} for observed BSP-HSP pairs. Our project will apply a data-driven model, training a neural network to recover HSP signals for a given BSP. The creation of this AI is composed of two problems, firstly the dataset used to train the neural network, and secondly the architecture/hyperparameters that define the network itself.

3.1.1. The Geometry of the inverse problem

This project extends the findings of existing research at UoA, *Impulse data models for the inverse problem of electrocardiography* [9]. By training a neural network on electrode measured gaussian basis functions, we can predict real world HSP for a given BSP. During robustness testing of the network BSPs were synthesized for altered heart geometry, and a quantifiable margin of error was identified as shown in figures 1 and 2. It can be proven that these changes in geometry cause significant changes in the recorded signals, and can cause false-positive diagnosis [10]. Clearly this represents a large obstacle for the clinical solution, as an acceptable network would be resilient to changes in patient geometry. For any given BSP recording, it would be necessary for either 1. the neural network to extract/learn patterns of the geometrical features of the heart, or 2. a CT scan to be performed and information about patient geometry to be fed into the network. While the former is favored as CT scans imply a certain level of radiation and are not always an option for some patients, it is largely unexplored research space and may prove impractical. It is a key outcome of the project to prove the ability of a neural network to learn these geometrical parameters, or if it will be unable to define a relationship determined from ECG signals alone.




Figure 1. HSP predictions with varying heart shift [9]

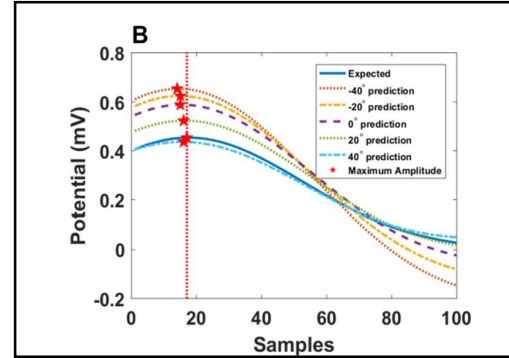


Figure 2. HSP predictions with varying heart rotation [9]

3.2. Data selection

Machine learning is no more than the data it was trained on – without careful scrutiny and an adequate dataset, a project will never produce desired results. Any neural network will only be able to make predictions based on patterns present in the training dataset. Having a sufficiently rich dataset is a common limitation of AI implementation, and our training set will need to be large. If we want our network to be able to predict HSPs regardless of patient geometry, our training set must provide plenty of examples of both different HSP-BSP pairs but additionally across multiple examples of geometry. This is a large challenge for all neural network applications to the inverse problem [11], as it is tricky to record real signal

pairs of the heart and body simultaneously, and no such physical dataset exists currently. As a result of this, it is necessary to synthesize data for the network to train on.

3.2.1. Data Synthesis

For the project, HSP-BSP pairs must be generated to train the network. Heart signals are recorded directly from a sock of recording electrodes and projected onto a 2d map of the body surface. The dataset for this project was generated by an experimental setup with a big heart suspended in a human-shaped torso tank, such that heart and body readings can be recorded simultaneously[12]. Based on this, a forward model has been developed for the synthesis of BSPs given a real HSP such that the reverse or inverse problem can be learned by the network [7], [10]. The synthesis and training pipeline can be observed in figure 3. While the experimental heart is not human, a successful modeling of the real HSPs could transfer to and motivate the need for a human dataset. Similar research applying synthesized data based on a model has proven effective for both data acquisition and network accuracy[13], [14].




Figure 3. Mapping technique for data synthesis, training, and prediction [10]

3.3. Neural networks

Neural networks have proven effective at solving data-driven regression problems[15], and are able to learn both linear and non-linear relationships. Additionally, it is mathematically proven that a simple neural network structure with only one dense, fully connected, hidden layer can approximate any smooth continuous function, given enough neurons and appropriate activation functions[16]. Therefore, after processing our dataset with the appropriate 2D mapping, a neural network should be an obvious method to the solution of the inverse problem. Multiple state-of-the-art studies have attempted to apply machine learning to the inverse problem of electrocardiography, with varying levels of success. However, a common theme among current research is the lack of consideration to changes in heart geometry (rotation, position, size). The following reviews the different networks and approaches that have been taken in recent years.

3.3.1. *Time Dependent Neural Networks*

TDNN architecture incorporates previous timesteps while training, with a window of size n previous values when processing. This tends to be ideal for electrocardiography, as we can obviously observe a high dependency between BSP/HSP values at t and previous values at $t-1$, $t-2$, etc. Mapping intracardiac and surface ECG signals (for each patient) has shown to provide improved accuracy and performance to numerical and simple artificial neural networks[17]. Another state-of-the-art application of TDNN to the inverse problem incorporates the coupling of a TDNN with a spatial operator. As our signals are recorded as an array of electrodes mapped to a 2d space, we could observe a correlation between adjacent electrodes. An adaptation of the TDNN is developed, with the first layer following a typical temporal correlation hidden layer, and the second layer an adaption of this for spatial correlation. As the temporal layer evaluates previous time values, the spatial layer evaluates adjacent electrode position values. This further assists the network to recognize patterns for HSP to BSP mapping. This neural network was applied in an experimental study mapping singular geometry synthesized BSPs to HSPs, with a high level of accuracy [18]. Despite the success of this approach, the time-invariance of our dataset makes application of this network architecture infeasible.

3.3.2. *Decoder/Encoder based Neural Networks*

These types of neural network architectures are designed to compress input data into latent or simplified vectors to extract critical features, then pass the vector to the decoder network to predict a HSP signal. As the encoder and decoder networks are completely separate, we can validate the predicted signals and frame a more generalized approach to the inverse problem. This could allow an algorithm to estimate HSPs without in-depth consideration to patient geometry. This type of network was applied experimentally to the inverse problem with a dataset of dog HSP and BSP simultaneous measurements [19]. However, it only covers a proof-of-concept application of the method due to a very limited dataset and is difficult to draw conclusions from. At the very least, this architecture is a valid method of analyzing the latent features and pattern detection the neural network can learn from a given dataset.

3.3.3. *Convolutional Neural Networks*

CNNs are a powerful tool for biomedical processing and imaging, as the nature of the convolutional layer significantly reduces the number of parameters to be learned [20]. While mainly suited to imaging problems, CNNs are great at extracting features despite spatial/spectral/temporal variation, which could make it suited to the inverse problem with geometry consideration. The network architecture utilizes convolution kernels and pooling layers to extract high-level

features from the input data. As standard approaches to the inverse problem use iterative algorithms with regularization, similarities between this and a CNN method (filtering then solving point by point), these networks seem to be an ideal solution [20].




Figure 4. U-Net Architecture [6]

Recent development in neural network architecture has enabled complex AI to be trained on increasingly smaller datasets with better accuracy than simple networks. One such architecture is U-Net, a fully convolutional network optimized for biomedical imaging problems [6]. Figure 4 illustrates the network architecture, using “up-convolution” layers with feature maps from the image before the corresponding convolution layer in the contracting path (left side). Features like data augmentation with elastic deformation can be designed to specifically teach the network variation in certain parameters, for example pixel displacement in images. State-of-the-art research shows CNNs could provide a promising solution to the inverse problem. A CNN trained to evaluate single-lead ECG signals was able to classify normal and atrial fibrillation beats with a high level of accuracy [21]. While the study worked with a spectrogram and imaging methods as opposed to electrocardiogram signals, it lays down the groundwork for ECG analysis and noted the potential for application to multi-lead ECG data for future work.

4. Methods

This report details the combined data synthesis, neural network architecture and training/testing procedures used to address research objectives. The problem is first explored a 2D abstraction of the heart in torso system, mapped as a concentric circle model. Here, we evaluate the performance of a neural network on the inverse problem, its ability to

predict when geometry is introduced, and if the relationship between HSP-BSP pairs is discrete. Following this, the 3D space is analyzed as an array of electrodes on the body surface mapped to a sock of electrodes on the heart. New data is synthesized through a forward model with variation in heart geometry and used to train a neural network. Finally, network performance is analyzed against real-world recordings, and future direction considered.

4.1. 2D domain

A simplified version of the problem is evaluated in 2D space, with the exterior circle representing the torso and the interior circle the heart in MATLAB. This is shown in figure 5 and the electrical potential of the heart activation is mapped through a volume conductor model to its corresponding potential on the body surface. Through this simulation, we can specify a gaussian activation on the heart surface, defining width and amplitude as shown by the blue plot on figure 2. The corresponding body surface potential is generated through the conductor model. This allows us to visualize and synthesize training data of HSP-BSP pairs. Geometry is also included, as the heart can be shifted in the X and Y dimensions from the origin. When the heart is mapped on a 2D surface like this, we can equate the rotation of the heart to a different angle of activation on the heart surface, or the location of the gaussian peak.

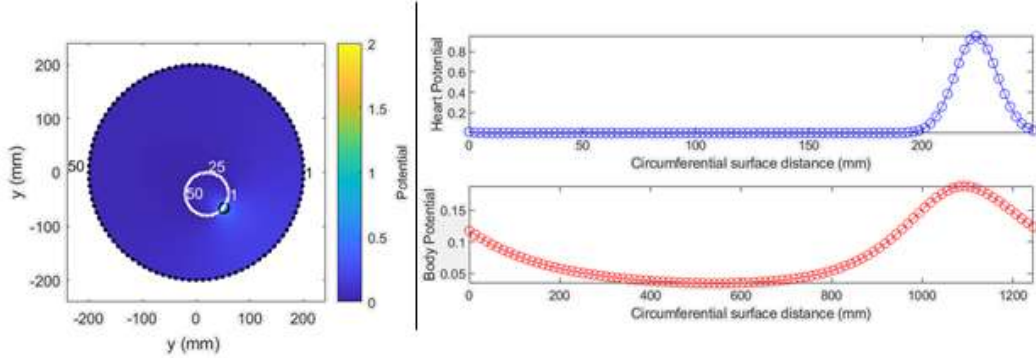


Figure 5: MATLAB 2D simulator, credit Dr. Mark Trew

As explained in section 2, the largest obstacle to applying a neural network is the dataset. At time of writing, there is no significant dataset that contains paired body and heart surface recordings. Therefore, it is necessary to synthesize data for the neural network to learn from. By using a mathematical model based upon an in-vitro model of a pig heart suspended in a torso, we can create a forward model that generates these pairs for training.

4.1.1. Data synthesis and forward model

These BSP and HSP waveforms are discretized into 100 points and used to train the model. A dataset of 73,790 training pairs was synthesized for training, testing and validation of the neural network. This is composed of 100 discrete points around the circumference of the heart with 5 variations of gaussian amplitude and width for the heart potential for each

point, repeating this iteration for x and y positional shifts of 10mm between 40 and -40mm. Each row of the dataset stores 100 discretized heart surface potentials, the corresponding 100 body surface potentials, the X/Y positional shift of the heart and the rotation of the heart (the gaussian peak with respect to 0 degrees).

For the process of data synthesis, underlying HSP gaussians are generated, then corresponding BSPs solved through simulation, and geometrical data normalized for use in the neural network. Gaussians are generated as per supporting documentation, using random uniform distribution. Bounds for gaussian amplitude and width are listed in table 3, in the appendix. Finally, the X and Y position is scaled from 40 to -40mm to a normal distribution between -1 and 1, and gaussian angle from 0 to 100 to between 0 and 2π . Normalization is common practice for neural networks and was found to greatly increase prediction accuracy.

4.1.2. Validation of HSP-BSP pair uniqueness

It is statistically significant to determine how well the model can differentiate between heart gaussian signals. Additionally, we need to verify that two or more similar HSPs corresponding to different BSPs are different enough for the neural network to identify as unique. Cases are present where the heart waveforms end up being significantly similar due to overlap in different geometries, allowing the same point in space to be replicated. These are examined, with the generated points exaggerated to increase similarity beyond what is present in the dataset, then tested by having the model predict the corresponding HSP. Figure 3 shows two different but highly similar BSPs on the left, with the blue occurring at 20,40 and the red occurring at 0,40. Despite the location difference, the location of the heart gaussian on the simulated circumference causes these two similar waveforms. The graph on the right shows the HSP prediction of both blue and red waveforms: the 'o's map true values while the 'x's map predicted values. Despite the similar BSPs the neural network is clearly able to predict the slight change in the gaussian peak – there is some error in the amplitude, but the central values are clearly matched.

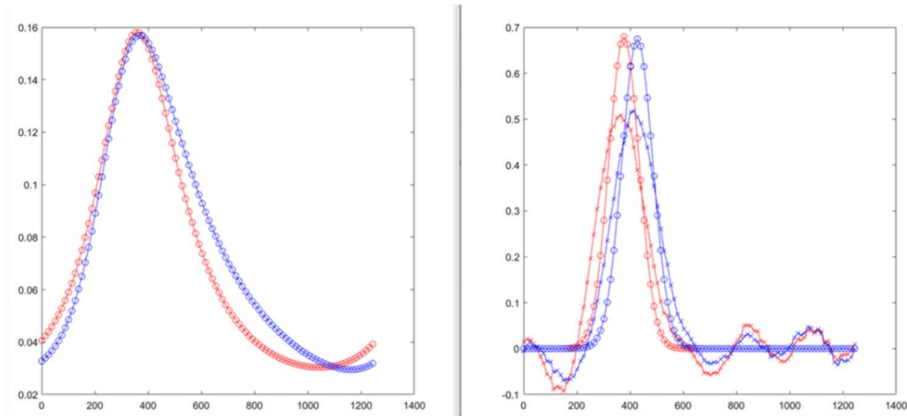


Figure 6: MATLAB visualization of test predictions, true BSPs on left, HSPs true/predicted on right

4.1.3. Network architecture and training

In the 2D scope, the neural network predicts the 100 values of the HSP vector, given an input of the 100-point BSP vector. The aim is to predict the underlying HSP with a high level of accuracy, with consideration to the complexity introduced by the changing geometry of the heart. After success with HSP prediction, the network was adapted to additionally predict underlying heart X/Y shift and rotation from an input BSP. The developed architecture uses a dense, fully connected structure with multiple layers as shown in figure 7. This type of structure is mathematically proven to have the capability of predicting any continuous function, as long as the number of layers and neurons is high enough [16]. The branching structure splits the model, allowing separate predictions to be made for HSP, x and y position, and gaussian angle. The model is developed in Python, using TensorFlow. Hyperparameters and activation functions are tuned through use of GridCV.

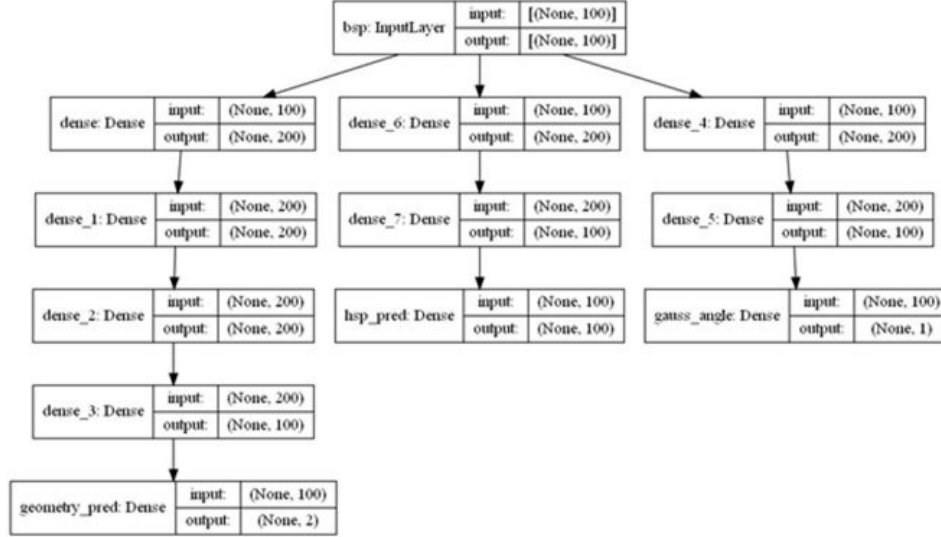


Figure 7. Structure of 2D case neural network

4.2. 3D domain

After analysis in the 2D scope, we move ahead to projecting similar HSP-BSP mappings in a 3D space. The heart surface is mapped by a 9 by 12 set of electrodes, projected onto a sock mesh that is wrapped around the heart as shown in figure 3, and the body surface is a 16 by 16 grid of electrodes placed on the chest. This is a similar mapping as taken from experimental recordings in [12]. A similar approach to the 2D problem is followed with data synthesis, evaluation, and network training. Figure 8 shows the high-level process of this process applied in 3D – the X and Y axis are the grid of electrodes, and the Z axis shows electrical potential.

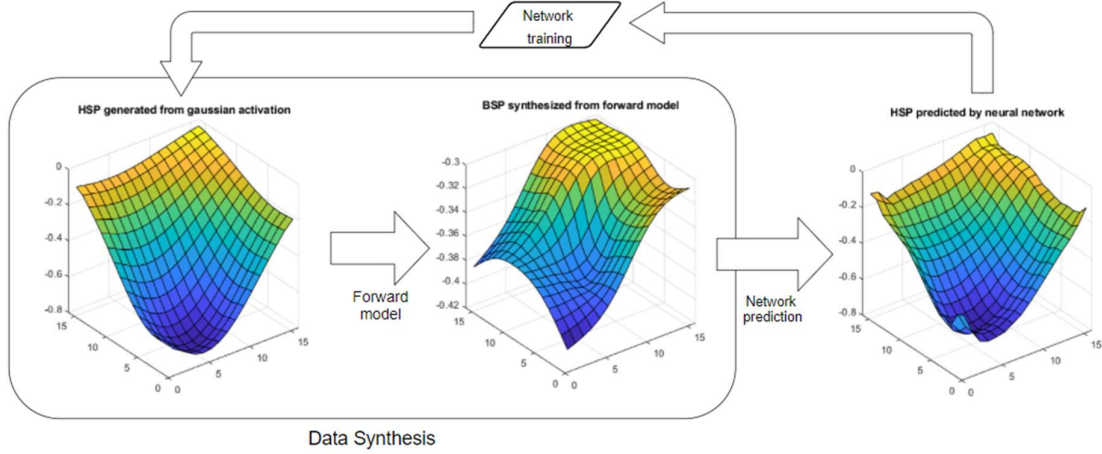


Figure 8. Synthesis and training for 3D domain

4.2.1. Data synthesis

As in the 2D case, a significant dataset for training does not exist so must be synthesized. Using MATLAB, a series of HSP potential activations on the electrode mesh are generated. The mapping of this mesh to a 3D heart is shown in figure 9. These activations are randomly generated with uniform distribution, and gaussian parameter ranges can be found in table 3 in the appendix. Some examples of generated HSPs are shown in figure 13 from the appendix. A forward model is used to map an artificial BSP to pair with the initial given HSP, as shown in figure 8. The electrode positions and HSP are combined with the real world in-vitro pig heart model through a Laplacian bi-conjugate gradient solver to synthesize this BSP. Due to the high computational load required, this task required offloading to a high-performance cluster. Using the forward model to generate BSPs for 25,111 MATLAB-synthesized HSPs took approximately two weeks, and the complexity of data synthesis for this problem remains an ongoing issue in scope.




Figure 9. Wrapping of 2D electrode mesh around 3D heart 'sock'

Geometry is also incorporated in this representation. The forward model uses a projection of the heart electrodes, mapped to the closest position in 3D space, to propagate the heart activations onto the torso through a volume conductor model. During data synthesis, we simulate different rotations of the heart by shifting the mapping of electrodes in 3D space. An example of this is shown in figure 10. This allows us to shift the heart relative to axis of the spine and generate corresponding BSPs with different rotations applied. Our final 3D dataset has rotational variation for HSPs between a -40- and +40-degree rotation, with discrete steps of 10 degrees. However, it is difficult for the model to predict the underlying rotation due to the limitations of the simulator. While the heart is rotated relative to the torso, the potential of the HSP is always recorded across the sock of electrodes on its surface. That is to say for a single gaussian activation generated, the HSP will look exactly the same despite differences in rotation of the heart. This means there is no 1:1 discrete mapping for BSP – HSP pairs, such that one HSP maps to multiple BSPs. This relationship in the dataset causes significant error in neural net predictions for rotational values, so we did not pursue this element of prediction. However, rotation is still present in the final dataset, and the neural network is able to learn BSP-HSP pairs despite the noise it generates.

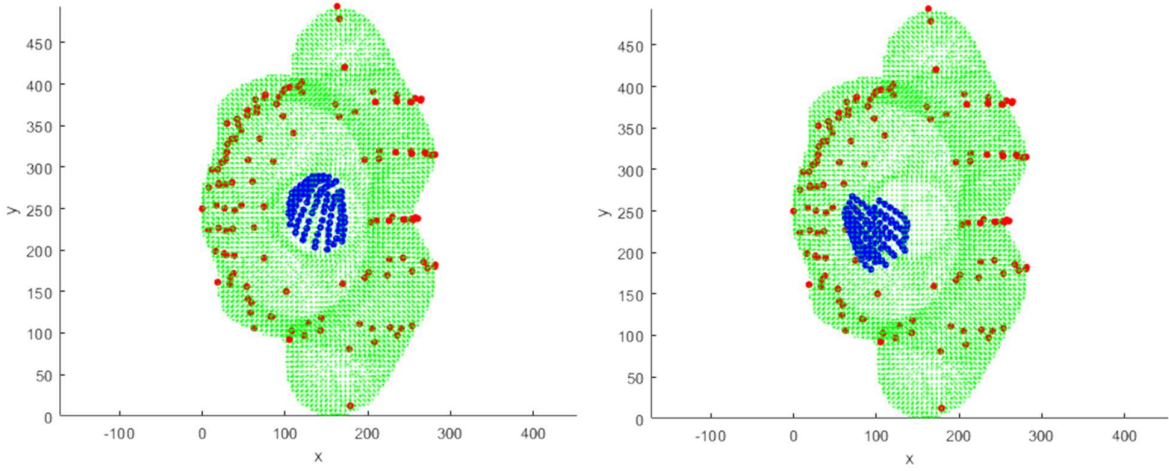


Figure 10. Visualization of rotation of electrode position mapping in 3D space, credit Dr. Mark Trew

4.2.2. Network architecture and training

After the success of the linear, dense network architecture for the 2D case, we began with a similar architecture. The extension of our initial network to 3D is shown in the appendix figure 11 (network A). Note that geometry prediction omitted from the 3D case due to limitations in prediction accuracy and the domain of the 3D simulation. While this

architecture is able to predict with decent accuracy, it was far from desirable. A more complex approach was required to fully learn the patterns present in the dataset and boost prediction accuracy. Based on existing research, a slight adaptation of U-Net was created, as it has previously proven success with inverse problems in imaging. To use CNN based architecture and convolutional layers, our HSPs are scaled from 9 by 12 to 16 by 16 through use of MATLAB's *imresize* function. This chosen CNN architecture is shown in figure 14 in the appendix (network B).

These two architectures are implemented in Python using TensorFlow. Both architectures take a 16-by-16 BSP and predict a 16-by-16 HSP, which is a resized prediction of the original 9-by-12 HSP by interpolation. From the synthesized dataset of 25,111 BSP-HSP pairs, a 70-30 training-testing split is utilized for model training. Dropout layers are utilized to reduce overfitting, and neuron bias is disabled as we would expect an activation of zero for a BSP should translate to zero activation of the HSP. Hyperparameters and activation functions are tuned using GridSearchCV from the scikit-learn library in Python. The Adam optimizer is used, combining the adaptive gradient algorithm with the root mean square propagation algorithm. Training is performed using the mean squared error loss function over 2000 epochs.

4.2.3. Experimental predictions on real-world data

As an extension to the above, we tested the abilities of the neural network based on the in-vitro pig heart experimental dataset, with left ventricular pacing in a saline torso tank, as outlined in [12]. Similar to our synthesized data, 9-by-12 HSPs are mapped to 16-by-16, paired to a BSP recorded by a 16-by-16 mesh of electrodes. A comparison between synthesized BSP-HSP pairs (left) and real-world pairs (right) is shown in figure 11.

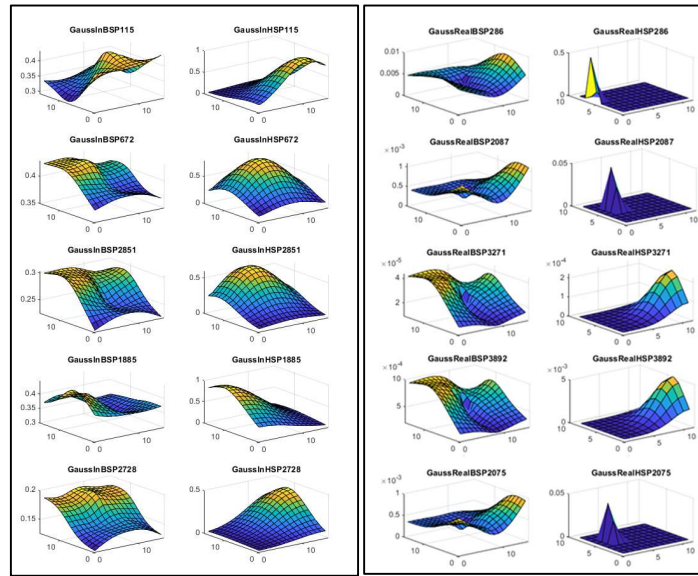


Figure 11. Comparison between synthetic (left) and real-world data (right). BSP and HSPs are labelled respectively.

There is a significant difference in amplitude size, with real world values usually in the millivolt range and synthetic values more normally distributed between -1 and 1. As part of prediction, our real-world BSP signals are normalized to this same range before prediction to ensure a standard environment for testing. Additionally, HSP activations from real-world data are a lot less smooth and sharper in character, and additional noise is propagated to the BSP. Both neural networks trained on the synthetic data struggle to accurately predict either peak location or amplitude values, due to differences between the synthetic training dataset and real-world representations. However, predictions from the feed-forward neural network show potential and are shown in figure 12. The CNN architecture is too deep to generalize towards the differences in the real-world dataset, but there is a reasonable correlation between predictions and true values for the feed-forward model shown in the figure. While outside the full scope of the project, future work needs to consider correlation and mapping between synthetic and real-world data, which remains difficult due to the absence of the latter.

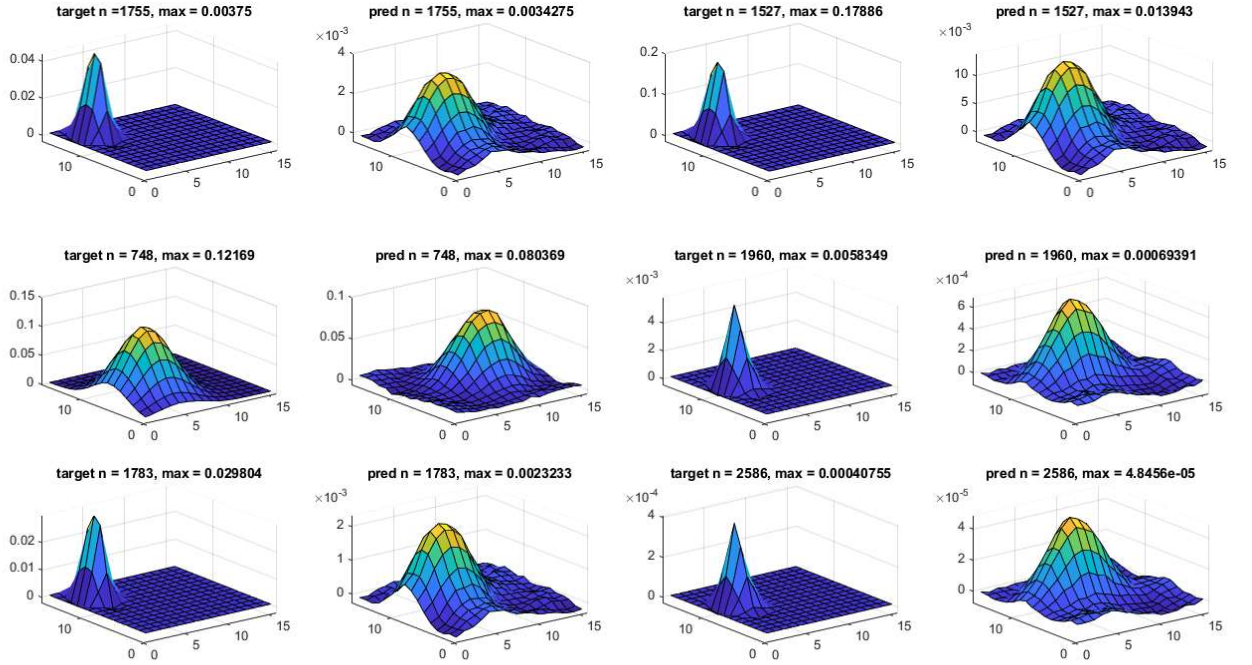


Figure 12. HSP predicted vs target values for feed-forward neural net from real-world data

5. Results

5.1. Prediction of Gaussian HSPs in 2D

Through adaptation of our neural network and data preprocessing, we are able to predict underlying HSPs reliably and accurately from BSPs. A high level of accuracy is also present for prediction of geometry, including both X, Y shift and rotation of the heart. Prediction accuracy evaluated by Root Mean Square Error (RMSE) is evaluated in Table 1, and the formula given in equation (2).

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (Predicted - Actual)^2} \quad (2)$$

The neural network architecture is defined in section 4.1.3 for this 2D case. Additional plots of the model predictions are included in section 8.1 in the appendix. Figure 16 plots some sample true HSPs in orange and the neural network prediction in blue, with a high correlation in peak location and amplitude observed. Figure 17 plots difference between true values of geometry on the x-axis and predicted values on the y-axis, for x-shift, y-shift and gaussian angle respectively. At $x = 2\pi$ on the angle plot, significant error can be observed on the right side of the plot, but this is due to 0 and 2π being identical values in the polar coordinate system. Figure 18 plots some example heart geometry predictions, with neural network predictions given in blue, true heart location mapped as a red circle, and the gaussian angle plotted as a green dot on that circle. It is tricky to see the predictions as it almost fully overlaps with the ground truth values. Finally, figure 19 evaluates the ability of the network to predict X/Y shift with respect to the X/Y values, with the highest margin of error in prediction at $(-40, 0)$ and $(0, -30)$, but a high degree of accuracy across all positions.

Table 1. Prediction accuracy for 2D abstraction

	HSP prediction	X, Y shift prediction	Gaussian angle prediction
Root Mean Squared Error	2.7%	2.2%	6.6%

5.2. Prediction of gaussian HSPs in 3D

Both described neural network architectures are trained on the synthesized 3D dataset and evaluated based on predictions from the 30% test split data. Accuracy is evaluated and calculated using RMSE, similarly using equations (1) and (2). During RMSE calculation, each point on the predicted 16-by-16 HSP is evaluated against the truth HSP. Prediction accuracy is listed in table 2. Additional plots of the neural network predictions are included in section 8.3 from the appendix. Figure 20 shows examples of HSP predictions for the feed forward neural network. This trained network can predict peak location accurately but struggles with amplitude values. Additionally, recreation of the smooth activation

function becomes rougher and more jagged. Figure 21 shows HSP predictions for the CNN based network architecture. Accuracy of prediction for amplitude is improved for the CNN, and accuracy of peak location is maintained. Although the two networks have similar RMSE, this is not fully representative of amplitude prediction accuracy due to the large number of near-zero values in the mesh, and the minimal effect on RMSE that slight differences have due to being averaged over the 16-by-16 grid. Despite the changes in heart rotation present in the dataset, we are able to accurately predict HSPs regardless of noise due to the shift in geometry.

Table 2. Prediction accuracy for 3D domain

	HSP prediction
Accuracy (Feed-forward network)	3.3%
Accuracy (CNN)	2.2%

6. Discussion

In this report, we present a neural network approach to the inverse problem of electrocardiography. Initially, a 2D simplification of the problem space is explored. We have shown that the relationship between HSP and BSP gaussian activation functions can be modelled by a neural network. Additionally, we introduce shifts in geometry in the form of x and y positional shift and rotation, introducing a more complex mapping between HSP-BSP pairs. By introducing this to our synthetic training dataset, our network proves it can maintain prediction accuracy despite changes in geometry. After modifying the network to a branch-like structure, these geometrical variables can be predicted by the network with a high level of precision. Our final trained network is able to capture HSP, x-y shift and rotation of the heart from one input BSP vector. The predictions of peak locations are almost perfect in this case, where there are only 100 discretized locations for rotation. It is interesting to note the network still makes occasional accuracy mistakes in peak amplitude, usually undershooting the value by ~20-30%, however this is acceptable considering the accuracy margins of the network. With consideration to the simplification of 2D space, it is to be expected that this relationship could be correctly modelled by the neural network. The network shows itself to be highly resilient to changes in geometry, and BSPs shown to be structurally very similar can be successfully mapped to their underlying HSPs.

In the 3D space, we demonstrate that the relationship between synthesized HSP-BSP pairs can be similarly modelled by a neural network. The network is robust to the introduction of rotational shift in the heart and its effect on this mapping. A fully connected architecture is able to satisfy this requirement, but a deeper CNN is much more accurately able to model the problem in a 3D space. CNNs have shown significant aptitude in imaging problems and is successfully applied to the

inverse problem. The model is currently unable to predict the rotational geometry of the heart, as the frame of reference measured on the heart surface cause the HSPs to look identical, and they map to multiple BSPs for each experimental rotation. Moving forward, adaptations to the simulation to measure HSPs with a different reference such as the axis of the spine could be used to generate a model that would be easier to map. Performance of the neural network is otherwise encouraging, with a high level of accuracy against the unseen test split of the synthetic data.

Lack of data and data synthesis remain one of the most significant issues in this problem space. As there is no existing significant HSP-BSP real-world dataset, and logistical and ethical concerns present a large obstacle to recording such a dataset, it seems unlikely this will change for the short-term. Data synthesis becomes the only way to approach the problem, requiring complex modelling and a high computational resource demand to generate. We have success during this project implementing, training, and testing on this, but clinical applications of neural networks to the inverse problem of electrocardiography will require more and more detailed and realistic approximations of real data. While the complexity increases exponentially from 2D to 3D, we still prove the fundamental abilities of a neural network to model the underlying relationship and effect of geometry changes. Our results in this space are highly promising and proof-of-concept for further research. Moving to 3D opens up the problem space by a large margin, when considering both mapping of synthetic and real data. Network performance from the 3D space is encouraging, and it seems feasible that a clinical level of accuracy can eventually be achieved.

The project objectives originally outlined in this report are largely met, with a successful proof of concept neural network. Although further research is required for the shift towards prediction on real-world data, implementation of the research problem is mostly successful and meets our project research intent.

7. Conclusion

This report explores a neural network approach to the inverse problem of electrocardiography, with a focus on data synthesis and the effects of change in heart geometry relative to the torso. We show in a 2D abstraction a neural network can successfully and accurately model the relationship between HSP and BSP pairs, trained on synthetic data. Additionally, the network is resilient to changes in geometry and can accurately predict geometric variables (x-y shift, rotation of heart) from the BSP. When moved into the 3D domain, a more complex CNN is able to model the relationship between HSP and BSP pairs from synthetic data. Again, this neural network is resilient to changes in geometry introduced by rotation of the heart. Future work would continue to adapt this methodology alongside existing experimental real-world data and evaluate network performance against a conclusive real-world dataset when it becomes available.

8. References

- [1] A. J. Pullan, L. K. Cheng, M. P. Nash, A. Ghodrati, R. MacLeod, and D. H. Brooks, "The Inverse Problem of Electrocardiography," in *Comprehensive Electrocardiology*, P. W. Macfarlane, A. van Oosterom, O. Pahlm, P. Kligfield, M. Janse, and J. Camm, Eds. London: Springer London, 2010, pp. 299–344. doi: 10.1007/978-1-84882-046-3_9.
- [2] Per Christopher Hansen, "The l-curve and its use in the numerical treatment of inverse problems," in *Computational Inverse problems in Electrocardiology, Advances in Computational Bioengineering*, WIT Press, 2000, pp. 119–142.
- [3] B. J. Messinger-Rapport and Y. Rudy, "Regularization of the inverse problem in electrocardiography: A model study," *Math. Biosci.*, vol. 89, no. 1, pp. 79–118, 1988, doi: [https://doi.org/10.1016/0025-5564\(88\)90113-7](https://doi.org/10.1016/0025-5564(88)90113-7).
- [4] A. Intini *et al.*, "Electrocardiographic imaging (ECGI), a novel diagnostic modality used for mapping of focal left ventricular tachycardia in a young athlete," *Heart Rhythm*, vol. 2, no. 11, pp. 1250–1252, Nov. 2005, doi: 10.1016/j.hrthm.2005.08.019.
- [5] A. J. Pullan, L. K. Cheng, M. P. Nash, C. P. Bradley, and D. J. Paterson, "Noninvasive Electrical Imaging of the Heart: Theory and Model Development," *Ann. Biomed. Eng.*, vol. 29, no. 10, pp. 817–836, Oct. 2001, doi: 10.1114/1.1408921.
- [6] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, vol. 9351, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241. doi: 10.1007/978-3-319-24574-4_28.
- [7] P. C. Franzone, B. Taccardi, and C. Viganotti, "An Approach to Inverse Calculation of Epicardial Potentials from Body Surface Maps1," in *Advances in Cardiology*, 1978, pp. 50–54. doi: 10.1159/000400421.
- [8] N. Zemzemi, S. Labarthe, R. Dubois, and Y. Coudière, "From Body Surface Potential to Activation Maps on the Atria: a Machine Learning Technique," Sep. 2012, vol. 39, p. 125. doi: 10/document.
- [9] T. Peng, A. Malik, L. Bear, and M. L. Trew, "Impulse data models for the inverse problem of electrocardiography," p. 16.
- [10] M. G. Adams and B. J. Drew, "Body Position Effects on the ECG Implication for Ischemia Monitoring," p. 7.
- [11] N. Zemzemi, R. Dubois, Y. Coudiere, O. Bernus, and M. Haissaguerre, "A Machine Learning Regularization of the Inverse Problem in Electrocardiography Imaging," p. 4.
- [12] L. Bear, P. Huntjens, R. Walton, O. Bernus, R. Coronel, and R. Dubois, "Cardiac Electrical Dyssynchrony is Accurately Detected by Noninvasive Electrocardiographic Imaging," *Heart Rhythm*, vol. 15, 2018, doi: 10.1016/j.hrthm.2018.02.024.
- [13] J. Tremblay *et al.*, "Training Deep Networks with Synthetic Data: Bridging the Reality Gap by Domain Randomization," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Salt Lake City, UT, Jun. 2018, pp. 1082–10828. doi: 10.1109/CVPRW.2018.00143.
- [14] K. W. Dunn *et al.*, "DeepSynth: Three-dimensional nuclear segmentation of biological images using neural networks trained with synthetic data," *Sci. Rep.*, vol. 9, no. 1, p. 18295, Dec. 2019, doi: 10.1038/s41598-019-54244-5.
- [15] Tin-Yau Kwok and Dit-Yan Yeung, "Constructive algorithms for structure learning in feedforward neural networks for regression problems," *IEEE Trans. Neural Netw.*, vol. 8, no. 3, pp. 630–645, May 1997, doi: 10.1109/72.572102.
- [16] K.-I. Funahashi, "On the approximate realization of continuous mappings by neural networks," *Neural Netw.*, vol. 2, no. 3, pp. 183–192, 1989, doi: [https://doi.org/10.1016/0893-6080\(89\)90003-8](https://doi.org/10.1016/0893-6080(89)90003-8).
- [17] F. Porée, A. Kachenoura, G. Carrault, R. D. Molin, P. Mabo, and A. I. Hernandez, "Surface Electrocardiogram Reconstruction From Intracardiac Electrograms Using a Dynamic Time Delay Artificial Neural Network," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 1, pp. 106–114, Jan. 2013, doi: 10.1109/TBME.2012.2225428.
- [18] A. Karoui, M. Bendahmane, and N. Zemzemi, "A Spatial Adaptation of the Time Delay Neural Network for Solving ECGI Inverse Problem," in *Functional Imaging and Modeling of the Heart*, vol. 11504, Y. Coudière, V. Ozenne, E. Vigmond, and N. Zemzemi, Eds. Cham: Springer International Publishing, 2019, pp. 94–102. doi: 10.1007/978-3-030-21949-9_11.
- [19] K. Bujnarowski, P. Bonizzi, M. Cluitmans, R. Peeters, and J. Karel, "CT-Scan Free Neural Network-Based Reconstruction of Heart Surface Potentials From ECG Recordings," in *2020 Computing in Cardiology*, Sep. 2020, pp. 1–4. doi: 10.22489/CinC.2020.200.
- [20] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep Convolutional Neural Network for Inverse Problems in Imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017, doi: 10.1109/TIP.2017.2713099.
- [21] M. Zihlmann, D. Perekrstenko, and M. Tschannen, "Convolutional Recurrent Neural Networks for Electrocardiogram Classification," *ArXiv171006122 Cs*, Apr. 2018, Accessed: Apr. 19, 2021. [Online]. Available: <http://arxiv.org/abs/1710.06122>

9. Appendix

Table 3. Parameters for HSP synthesis

	2D HSP generation	3D HSP generation
Gaussian amplitude upper limit	1mV	1V
Gaussian amplitude lower limit	0mV	-1V
Gaussian width upper limit	2mm	0.01mm
Gaussian width upper limit	5mm	0.5mm

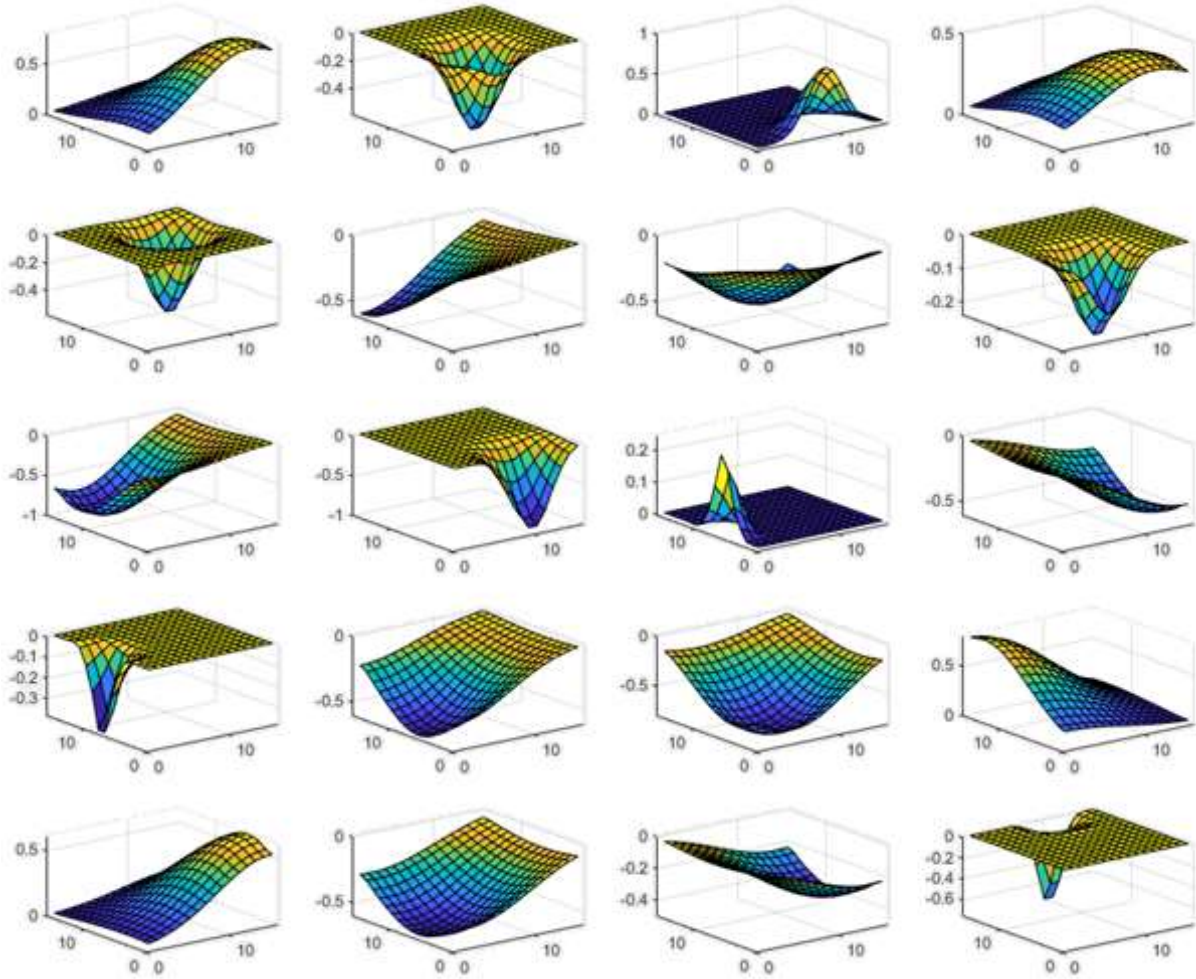


Figure 13. Example of synthesized HSPs

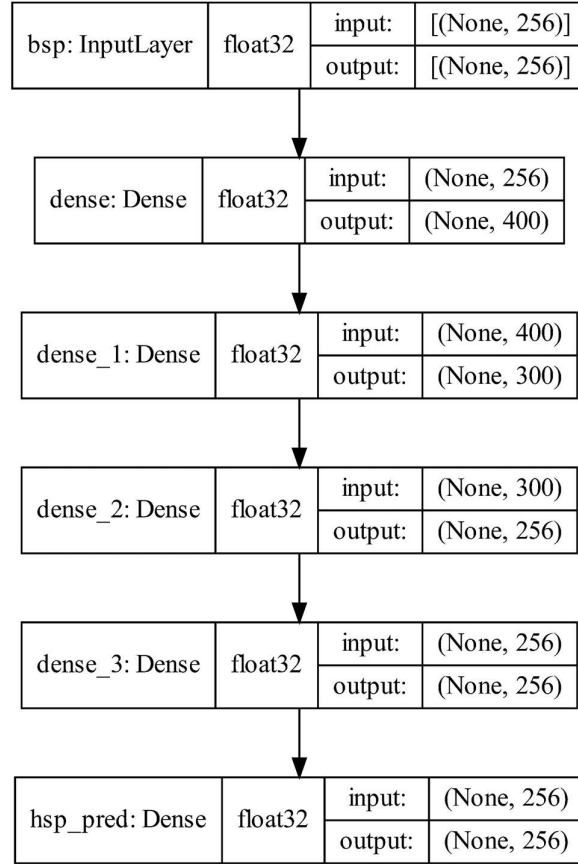


Figure 14: Fully-connected architecture adapted for 3D case

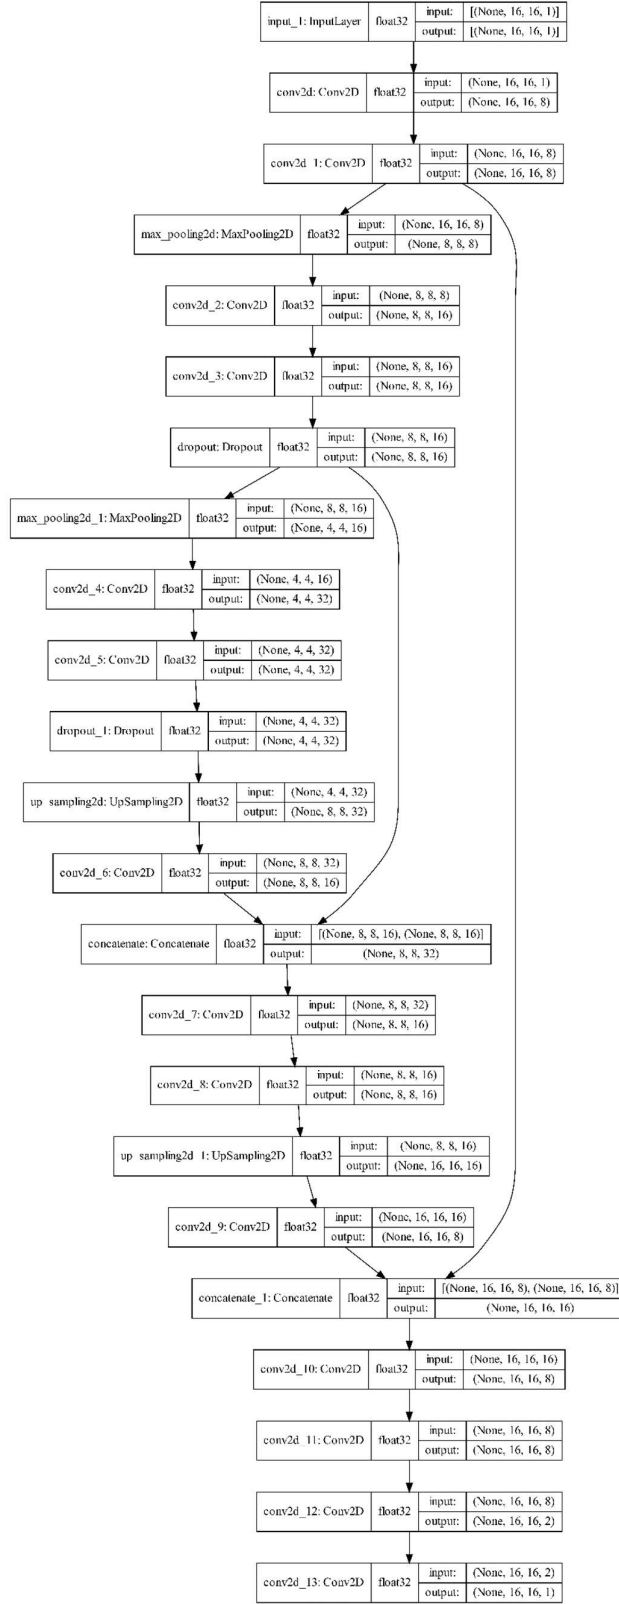


Figure 15: U-Net based CNN architecture for 3D case

9.1. 2D prediction plots

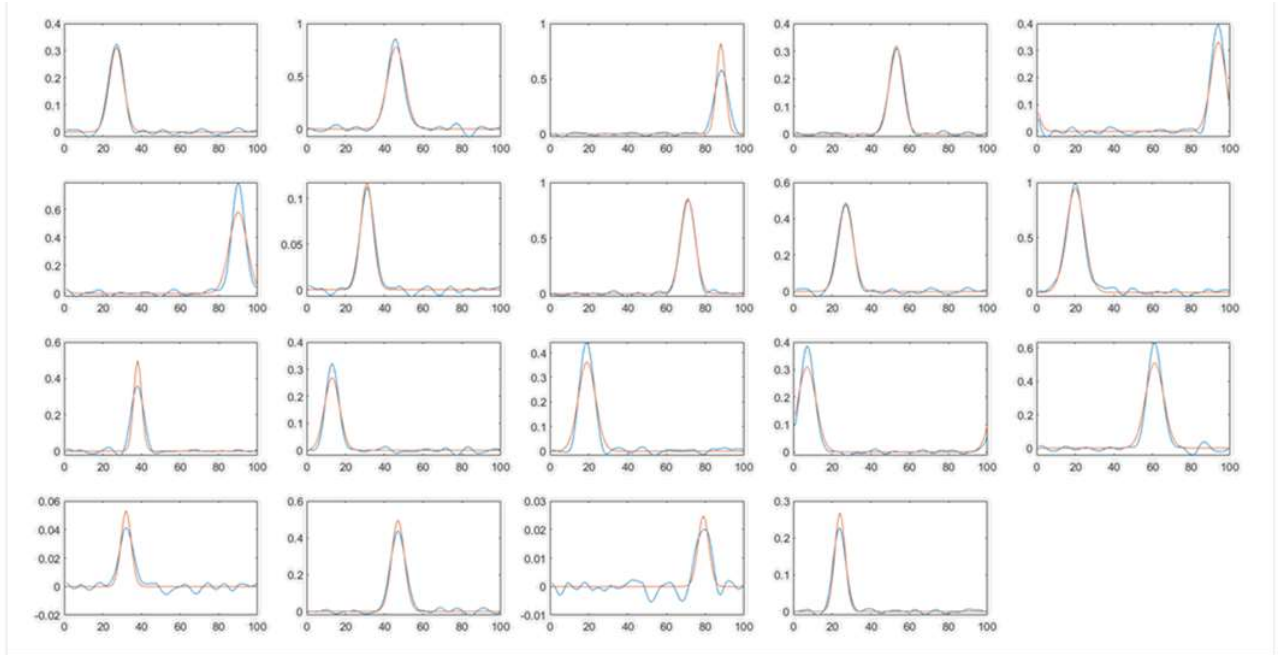


Figure 16: Neural network predictions (blue) against true values (orange) for 2D case

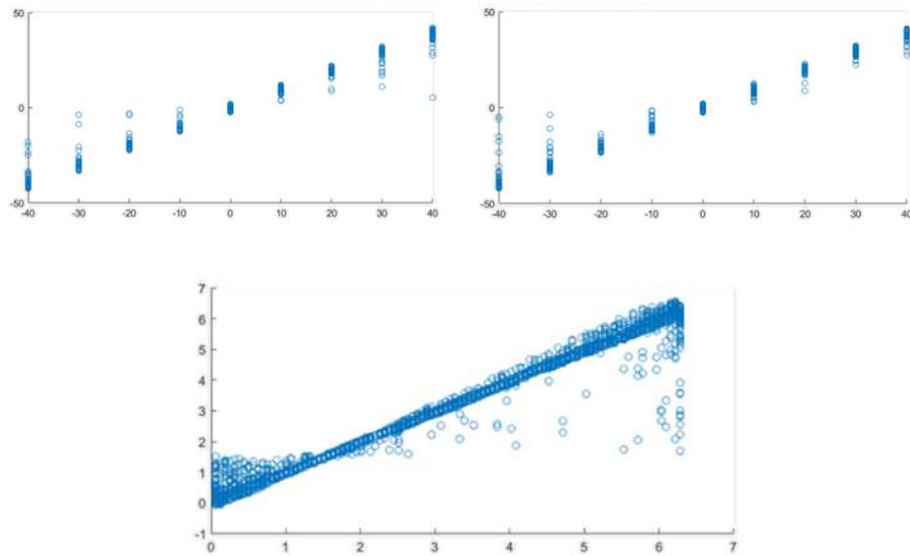


Figure 17: Neural network accuracy for geometry, predictions on the Y axis, actual values on the X axis. X-shift in mm top left, y-shift in mm top right, gaussian angle bottom middle

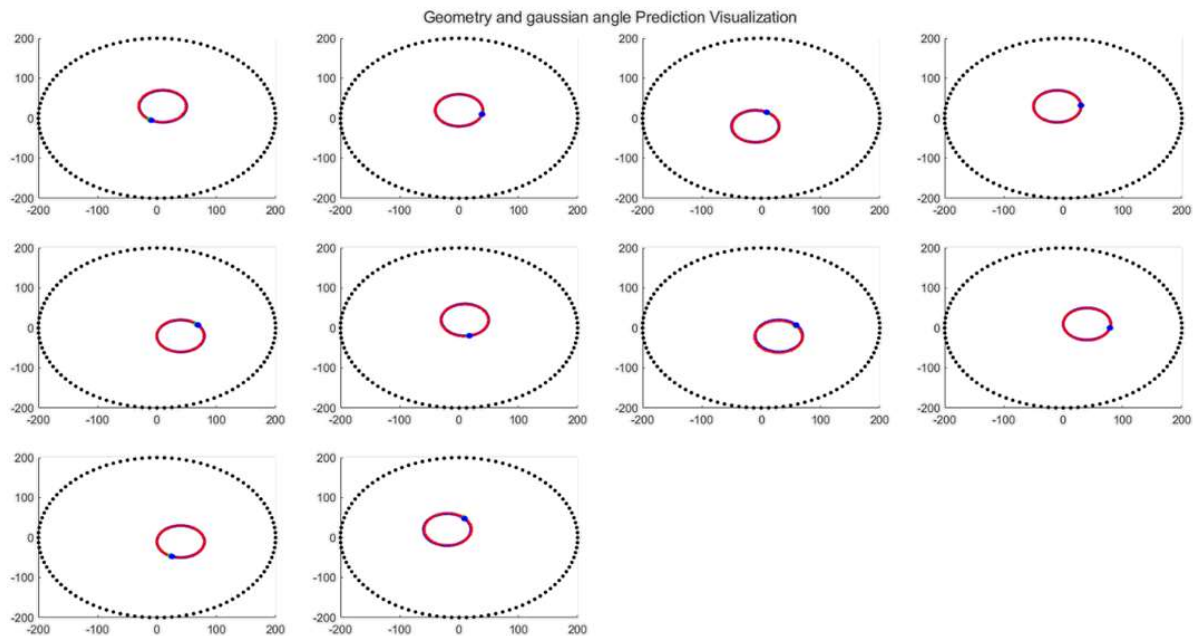


Figure 18: Visualization of network x-y shift prediction, predictions shown in blue, actual heart location mapped as a red circle and gaussian angle as a green dot. Large overlap makes it difficult to differentiate.

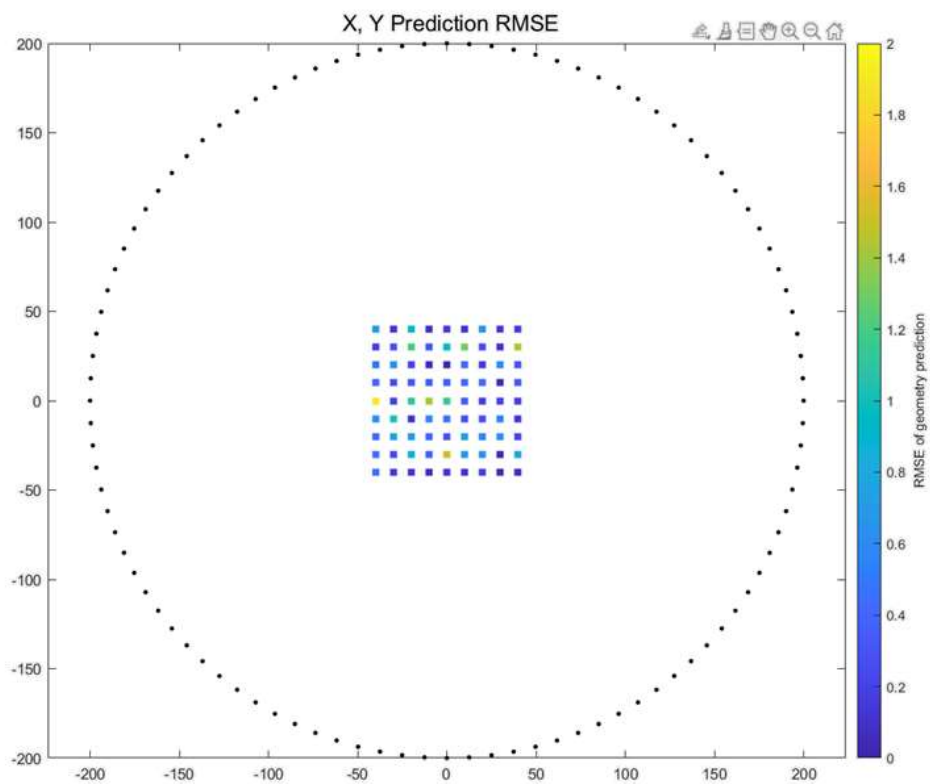


Figure 19: RMSE prediction accuracy of x-y shift relative to actual heart x-y coordinates.

9.2. 3D prediction plots

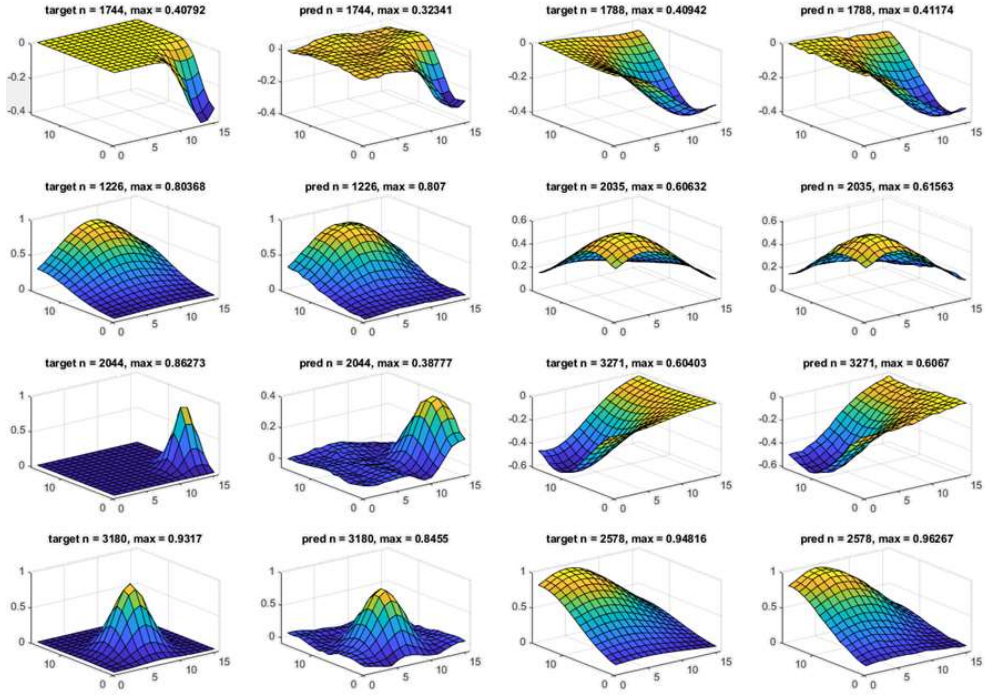


Figure 20: Feed forward neural network predictions of HSP in 3D case

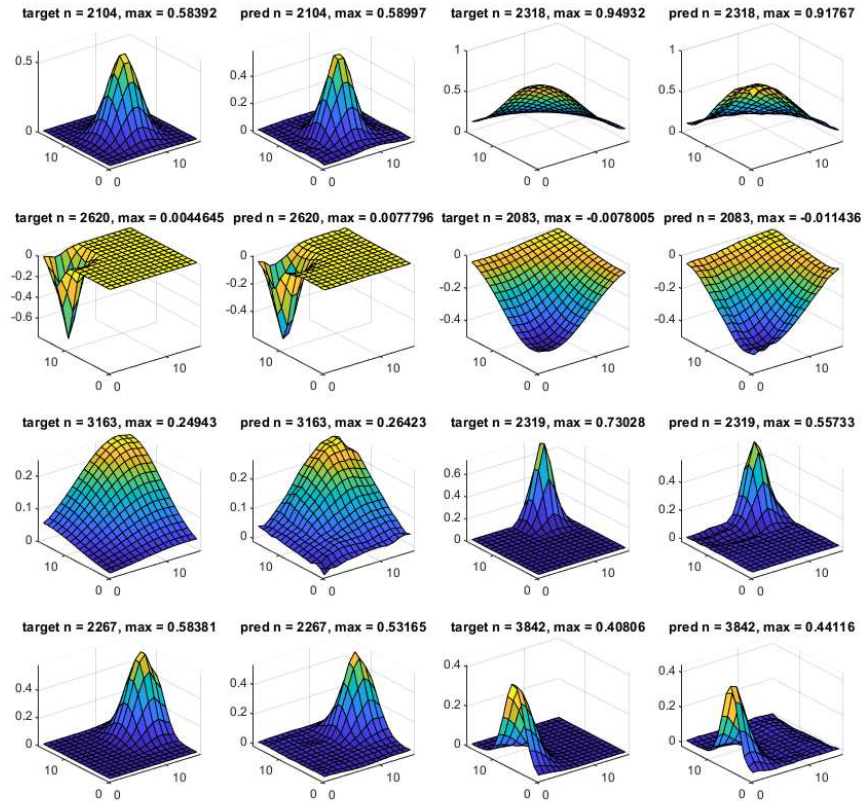


Figure 21: U-net based CNN network predictions of HSP in 3D case