Deep robot path planning from demonstrations for breast cancer examination*

Marta Crivellari[1], Oluwatoyin Sanni[2], Andrea Zanchettin[1], and Amir Ghalamzan E.[2]

[1] Politecnico di Milano, Italy
[2] University of Lincoln, UK
aghalamzanesfahani@lincoln.ac.uk

Abstract. In 2020, breast cancer affected around two million people worldwide. Early cancer detection is, therefore, needed to save many lives and reduce treatment costs. Nowadays, mammography and self- palpation are the most popular monitoring methods. The high number of cases and the difficulty of correct self-diagnosis has prompted this research work to design a fully autonomous robot to perform breast palpation. Specifically, this study focuses on learning the path for a successful breast examination of a silicone model. Learning from demonstrations proved to be the most suitable approach to reproduce the desired path. We implemented a teleoperation control between two Franka Emika Panda robots with tactile and force feedback to perform palpation on both simple and complex shapes. Moreover, we created a dataset of simple palpation strategy. Finally, we developed and tested different sequential neural networks such as Recurrent Neural Network (RNN), Long short-term memory (LSTM), Gated recurrent unit (GRU) and Temporal Convolutional Network (TCN) to learn the stochastic behaviour of the acquired palpation trajectories. The results showed that TCN is capable of reproducing the desired behaviour with more accuracy and stability than the other models.

Keywords: Robotic Palpation · Learning from Demonstration · Teleoperation · Force and Tactile Feedback · Sequential Neural Networks

1 Introduction

Breast cancer affects many people in the world every year. According to the *GLOBOCAN Cancer Today* database from the International Agency for Research on Cancer, 2.261.419 new cases were identified in 2020 [4]. Early cancer detection is of utmost importance as it can allow faster, simpler and more effective treatment; hence saving many lives. Breast self-examination, expert palpation and Mammography are currently the means of detecting breast cancer. Expert and Self-examination are composed of a visual inspection of the breasts

* This work was partially supported by Cancer Research UK C24524/A30038 in ARTEMIS project.

and palpation of the breasts and lymph nodes. Nonetheless, these are subjective approaches and may result in many false negative [15]. On the other hand, in mammography, the body is exposed to radiation. Hence, early breast cancer detection is not well practised illustrating a technology gap. Robot palpation is a solution to fill this gap. For example, during robotic minimally invasive surgery [24, 22, 19], palpation is essential to identify anomalies [6]. Nowadays, the palpation action for breast cancer detection is performed by the patient, consequently of the subjects, the diagnosis is not always reliable, or by expert which is not convenient for many subjects, revealing an autonomous robot for palpation an interesting solution.

Advanced Robotic breasT ExaMination Intelligent System (ARTEMIS) project funded by Cancer Research UK aims to develop a completely autonomous robot that substitutes the palpation action in order to obtain higher accuracy in the diagnosis. There are studies focused on the estimation of mechanical characterisation ([5, 16, 17, 19]); but, to the best of our knowledge, this is the first work in literature that study data collection and palpation path planning for autonomous palpation to be used for breast cancer examination. Robotic palpation is complex because combined haptic feedback and visual information determine the interaction between the robotic end-effector and breast tissue, which is complex geometry and has varying mechanical properties across different breasts. Here we consider two phases of path/trajectory planning: (1) palpation path/trajectory planning based only on visual information provided by RGB-D sensor/s and (2) trajectory adaptation based on the tactile information– i.e. according to change of stiffness in the examined tissue by the tactile sensor the robot generates exploration trajectory for that part to gain more information necessary for later cancer detection classifier. The complexity of path planning comes from the alternation between contact and no-contact actions and the correlation between the path and RGB-D sensory information representing different geometry, i.e. the path has to change in accordance to the patient's breast.

Palpation of a breast is an intuitive test technique in which the pads of the three middle fingers are used to identify and locate diseased tissue. It is very important to examine the entire surface, hence, the doctors suggest following a predefined path that can be circles, wedges and/or lines during the palpation ([1–3]). Recent research highlights the potential of Learning from demonstration (LfD) in path planning [20]. Recent LfD approaches are reviewed in [23]. Elliot et al. [8] presented LfD to extract the tool application pattern for cleaning task. However, they use many hard-coding and preprocessing operations which is not feasible in breast palpation case due to high geometrical variability across subjects at different palpation tasks. In this work, we aim to develop a data-set and deep model learning the palpation trajectory from human demonstrations without human interference. Palpation trajectory is complicated and stochastic due to human factors. Thus, it is difficult to capture a correct path using traditional machine learning methods. Deep models instead seem to be an appropriate solution correlating visual information to proper palpation trajectory.

[11] used a deep model to learn the robot control action for inserting a needle into a soft tissue. Zhang et al. [25] used a Temporal Convolutional Network (TCN) to predict the long-term lane-changing behaviour trajectory. The results showed that the model could successfully predict the trajectory behaviour with higher performance than RNN and CNN. The contribution of this paper is threefold: (1) a data collection setup for tele-operating palpation actions, (2) a palpation data-set, (3) a series of deep recurrent models generating palpation trajectories based on visual information. These models can capture the behaviour of the palpation trajectory by sequential architectures Recurrent Neural Network (RNN), Long Short Term Memory (LSTM) [10], Gated Recurrent Units (GRU) [9] and TCN. This study demonstrates deep time series are suitable tools for learning palpation trajectory for a phantom at different position and orientation. We created a palpation data set publicly available and baseline deep time series model. This is a proof of concept study. Future works include the big data set more suitable for deep time series model, considering forceful interactions in our experimentation and the use of a breast phantom.

2 Methodology

Learning from the demonstration method achieved good where a human demonstration is not expensive while programming the robot can be time consuming and expensive. The main concept is to record demonstrations of palpation action and extrapolate from them the main features to predict a successful path. It is important to understand what to learn and how to learn before going through the details of the methodology. As such, we made demonstrations using a simple silicon phantom and moving the follower arm end-effector from right to left with a contact motion and from left to right with a no-contact motion, as shown in Fig. 1c), and repeat the motion until all the surface is examined. Then, a deep model will be trained on the data acquired. In the following, we describe the experimental setup, data-set acquisition and deep-model implementation.

2.1 Data-set acquisition

Set-up: We use two Franka Emika Panda robots where the right one (called follower-arm) mirrors the left one (called leader-arm). I.e. the left arm follows the movement of the right arm in Fig. 1a. A real-sense RGB-D camera is mounted on the wrist of The follower and a 4 by 6 Xela sensor is attached to the left finger of Panda arm as shown in Fig. 1b. Each sensing cell of the Xela sensor measures a normal force and two tangent forces during palpation actions. The leader is compliant controlled so that it feels zero mass while a human is moving the follower far from the phantom. If the follower is close to the surface of the phantom the leader will guide the user to keep the leader be normal to the phantom surface and does not allow too much penetration in the phantom (see the Control for data acquisition section!). We have two phantoms (1) breast phantom which is a model of a real breast, (2) a simple flat surface silicon

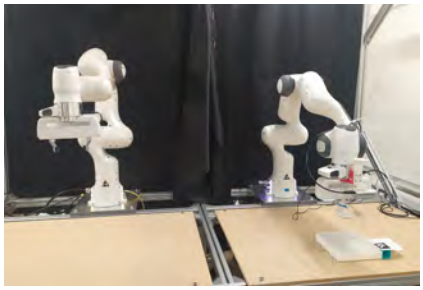

(a) Set-up

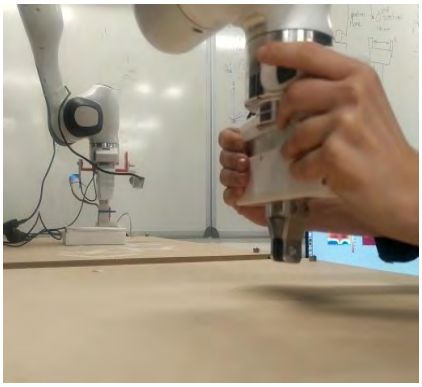

(b) Teleoperation

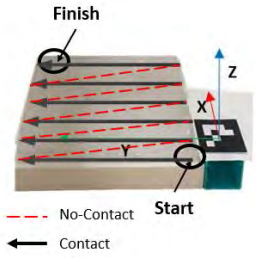

(c) Phantom motion

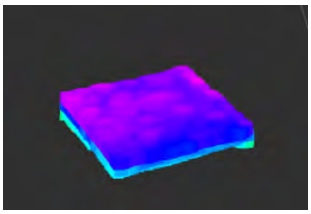

(d) Phantom point-cloud

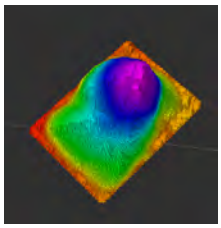

(e) Breast point-cloud

Fig. 1: The experimental setup to perform teleoperated palpation actions consisting of 2 Panda arm manufactured by Franka Emika Gmbh (a) and (b); the palpation are applied on a phantom with rectangle cross section made of silicon. The movements pattern during palpation (c); the depth map of the silicon phantom (d) and the breast phantom (e).

phantom. In our first implementation in this paper, we use the silicone piece (2A Shore scale) to develop necessary control, learning and experimental setups, shown in figure 1c. We fabricated two silicone layers of 15x15x2cm dimension and stacked one on the other where two harder objects were embedded as simulation of lumps betweeen the layers. We cover the top silicone surface with a paper to help the sensor sliding.

Control for Data Acquisition The robot has to learn a behaviour from one or several demonstrations performed by a human who is considered an expert. A teleoperation control with force feedback is implemented to give input to the Franka arm. For terminologies, robots are going to be called 'leader', the one used by operator, and 'follower', the one that computes the task on the silicone phantom. Looking at figure 1a, leader is on the left and follower on the right. The two robots have to communicate with each other. For that purpose, a PD control is implemented. The control's goal is to impose a motion to the follower arm equal to the one applied by the operator on leader side. Therefore, the torque is calculated in order to apply same in the follower's actuator T_f as follows:

$$T_f = K_p(\theta_l - \theta_f) + K_d(\dot{\theta}_l - \dot{\theta}_f) \tag{1}$$

where K_p and K_d are the proportional and derivative coefficient, θ is the joint position and $\dot{\theta}$ is the angular velocity for leader (l) and follower (f) side. Notice that the control works only for small $\Delta\theta$ and $\Delta\dot{\theta}$. To satisfy this hypothesis, the initial positions of the leader and the follower are set equal, and the frequency of the control loop has to be higher than a limit value estimated equal to 200Hz. The palpation action is performed in a perpendicular direction. In order to help the operator, who does not have a direct contact with the object to estimate the normal direction and to fill the reaction force given by the contact, a force feedback control is implemented. The torque on the leader side is modelled and was set T_l equal to:

$$T_l = J^T \cdot F_f^{EE} \tag{2}$$

where J is the Jacobian matrix and F_f^{EE} is the force apply on the follower end-effector. F_f^{EE} is a 6x1 dimensional vector defined as follow:

$$F_f^{EE} = \begin{bmatrix} F_x^{EE} \\ F_y^{EE} \\ F_z^{EE} \\ T_\alpha^{EE} \\ T_\beta^{EE} \\ T_\gamma^{EE} \end{bmatrix} = K \cdot f \left(\begin{bmatrix} x^{EE} \\ y^{EE} \\ z^{EE} \\ \alpha^{EE} \\ \beta^{EE} \\ \gamma^{EE} \end{bmatrix} - \begin{bmatrix} x^{ref} \\ y^{ref} \\ z^{ref} \\ \alpha^{ref} \\ \beta^{ref} \\ \gamma^{ref} \end{bmatrix} \right) \tag{3}$$

Where ref are the variable related to the frame attached to each point of the object's surface with the z axis equal to the normal vector in that point to the surface, while EE refers to the frame attached to the follower's end-effector. The force in x and y direction are considered null, instead the one along z is estimated as:

$$\begin{cases} F_z^{EE} = 0, & z > z_{ref} \\ F_z^{EE} = k_1 \cdot |z - z_{ref}|, & z_{ref} > z > z_{ref} - \delta z \\ F_z^{EE} = k_1 \cdot |\delta z| + k_2(e^{|z-(z_{ref}-\delta z)|} - 1), & z < z_{ref} - \delta z \end{cases} \tag{4}$$

k_1 and k_2 are coefficient chosen experimentally and δz is the area within the object is easily deformed. The three torques T^{EE} are calculated as a coefficient

times the difference between the Euler angles of the end-effector frame and reference surface frame. During the demonstrations, the object's point closer to the end-effector is used as reference. The point cloud of the object was used for the estimation of the reference frame in each points of the object (figure 1e and 1d). During the teleoperation, the distance between the end-effector and all the points is calculated. If the minimum value of the distance is lower than 2cm, then the force feedback control is activated.
The z^{ref} is estimated on-line taking the average value of z of the closest 1000 points with respect to the end-effector position. On the other hand, the reference of the orientation, the rotational matrix which corresponds to the closest point is considered. The control was tested on both the simple phantom as well as the real breast with excellent results.

2.2 Deep-model

The extraction of palpation path from the raw demonstration is a crucial task for robot motion planning. Therefore, the goal is to capture the common features of an entire execution and estimating a global pattern to follow.

Data-set 31 demonstrations were collected with the phantom fixed to the table and only one subject as operator. The data-set is composed by X, Y and Z Cartesian's coordinates of the follower's end-effector. The first and the last contact between the end-effector and the object are the initial and last point acquired. The trajectory is described in a relative frame fixed to the object in such a way that the x coordinates corresponds to the repetition direction, y to the task direction and z to the contact direction (1c). One trial lasts about $2.5 - 3$ min. In the phase of modelling our trajectories, shallow methods like Hidden Markov models or Bayesian inverse reinforcement learning lack the capacity to model data that has the long-term dependencies property, this is why RNN was chosen for the implementation of our model. They have demonstrated strength in modeling variable length sequence. Hao Wu et al adopted RNN to predict trajectories on a real world taxi GPS trajectory datasets. Similarly, the behaviour of our trajectories coordinates are periodical, and this makes it sufficient for the network to learn the behaviour of a time series. They can learn the time-dependent mechanisms underlying the expert palpation on the breast phantom. In this work, RNN and its variants like LSTM, GRU and TCN are designed to work with sequential data, they have all been tested to observe which performs best.

Model Architecture The experiments have been tested on RNN, LSTM, GRU and TCN since each of them have a unique distinction to their structure. RNN is not enough because they cannot solve the vanishing gradient problem, a common problem encountered when training artificial neural networks with gradient-based learning methods and backpropagation. During training, the gradient becomes very small in a vanishing fashion and prevents the weights of the model

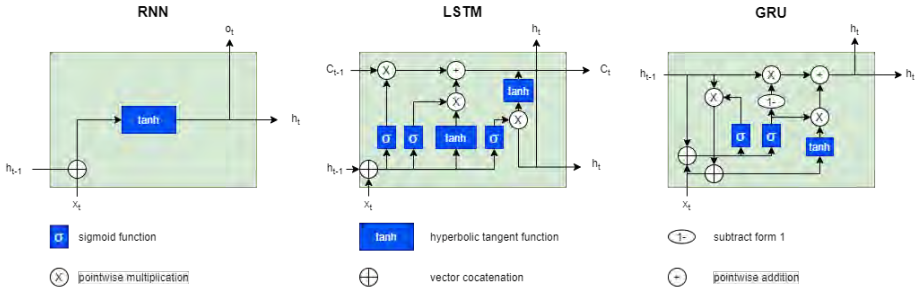

Fig. 2: Schematic comparisons of RNN, LSTM and GRU

from changing its value. LSTM 's and GRU's were created as the solution to the vanishing gradient problem, they both use a gate system that controls the memorization process. LSTM composes of a cell, an input gate, an output gate and a forget gate. The cell remembers values over arbitrary time intervals and the three gates regulate the flow of information into and out of the cell, it learns to bridge minimal time lags in discrete time steps by enforcing constant error flow through the cell. GRU on the other hand, composes of two gates, a reset gate and update gate. The network has fewer tensor operations since it has a simpler network than the LSTM, and this makes it computationally faster than LSTM.

Conversely, TCN is a recent network different from previous models described. Lea et al [14] introduced it as a modification of Convolutional Neural Networks for sequence modelling tasks, by combining RNN and CNN architectures. It's a combination of two different network with dissimilar architectures. The first network is a CNN that functions as a feature extractor by encoding spatial-temporal information, while the second network, usually a RNN, takes the low-level dimensional features output of the CNN and acts as a classifier by capturing a high-level temporal information from these features. The CNN is a Casual convolution, a type of network which ensures the model cannot violate the ordering in which we model the data. This means prediction emitted by the model at timestep t cannot depend on any of the future t+1 timesteps. This work focus is not to much focus on how the sequential models shown in Fig. 2, however, we invite readers to get more information about these architectures from the original authors, RNN [21], LSTM [13], GRU [7] and TCN [14].

Defining a sequence modelling as $f : X^{T+1} \rightarrow Y^{T+1}$ with the mapping: $\hat{y}_0, ..., \hat{y}_T = f(x_0, ..., x_T)$ where $(x_0, ..., x_T)$ is the input sequence that predict the output $y_0, ..., y_T$. In our case, T is set as 50, i.e. 50 time steps in past are considered to predict 50 time steps in the future. Even if 50 time steps in the future are predicted, the input temporal window is shifted only by one time step for the next iteration with respect to the previous one, i.e if the first input is from t_0 to t_{49}, the second input will be from t_1 to t_{50} and so on. 50-time steps in the past are considered to predict 50-time steps in the future. For this

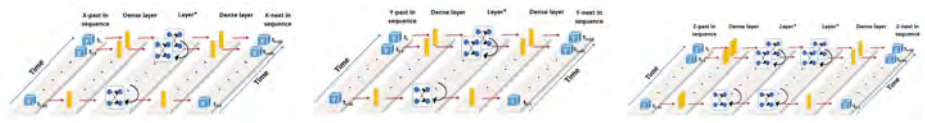

(a) $XModel$ (b) $YModel$ (c) $ZModel$

Fig. 3: Deep-model

research, 67% of the data-set are considered for training and validation while the remaining 33% were used in test phase.
The four sequential models described above were trained to predict each coordinate i.e X,Y,Z separately. The structures of the models for the 3 coordinates are very similar: they start and finish with one or more Time Distributed Dense (TDD) layers and in the middle, there are one or more sequence layers (*layer**). The X and Y model consists of: *InputLayer* - $TimeDistributedLayer_1$ - *NNLayer* - $TimeDistributedLayer_2$.
The Z model consists of: *InputLayer* - $TimeDistributedLayer_1$ - $TimeDistributedLayer_2$ - $NNLayer_1$ - $NNLayer_2$ - $TimeDistributedLayer_3$ where NNLayer is replaced with either RNN, LSTM, GRU or TCN layer.
The Time Distributed layer serves as a wrapper by applying a Dense layer to every temporal slice of the input. The addition of the extra NNLayer and TimeDistributedLayer in the Z model was done to optimize the prediction of the trajectory efficiently because it's a different type of trajectory in comparison to X and Y. The optimizer chosen is Adam and the loss function used is Mean Absolute Error between the trajectory predictions and ground truth for each coordinate. The number of epochs in accordance with the validation losses is set to 15 for all the models. The performance of models are discussed in the next section. We have used Adam optimiser, the Learning rate is 0.001, 15 epochs, and the loss function used was Mean Absolutes Error.

3 Results

The implemented models show interesting results in both test phase and real-time prediction. In test phase, the inputs sequences given to the model are taken directly from the acquisitions; in real-time instead, the sequences used as inputs are taken from the previous prediction, obviously an exception is made for the first one. The interest is not in replicating the single demonstration but only the correct trend of each coordinate. As a matter of fact, the performance of a network is compared on the capability of estimating the common trend of each demonstrations.

Fig 4 (a) represents the X prediction for all the possible networks. The trends of the curves are very similar to each other and there are not relevant differences. For the Y prediction (Fig 4 (b)), it is evident that the peaks of the ground-truth do not correspond to the predicted one. This should not be surprising

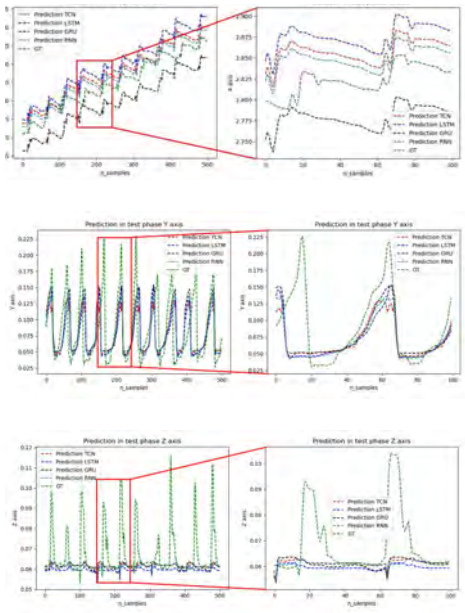

Fig. 4: Test results: (a) Prediction test results of behaviour for X; (b) Prediction test results of behaviour for Y; (c) Prediction test results of behaviour for Z.

since the periodicity is related to each execution. On the contrary, this fact is interpreted as a strength because the model can learn a medium time period to complete the task. It is significant that all the models identify the same periodicity. Another strength is that the amplitudes of the data acquired have variations in repetitions, instead the predicted trajectory replicates a constant behaviour in accordance with the fixed shape of the phantom. The peaks have a smoother shape along the positive slope and steeper shape along the negative one; the trend reflects perfectly what happens during the execution. In the ascent phase, the sensor is in contact with the surface and the velocity is slower, while in the descent phase, the end-effector is no more in contact and it can reach higher velocity.

Looking at the prediction, one can conclude that Y best estimation is given by GRU, thanks to the smoother shape. Finally, some considerations about Z are needed. It is very difficult to correctly estimate the amplitude of the peaks (no-contact motion) due to the low number of samples with respect to the contact phase since the velocity of the robot is very high. Moreover, most of the samples have same value because the phantom has a constant surface, this is why the model tends to predict a constant output. Nevertheless, a behaviour similar to the ground-truth is achieved by the TCN (Fig. 5). To conclude, unseen inputs were tested on the the trained models, and they exhibited the ability to predict

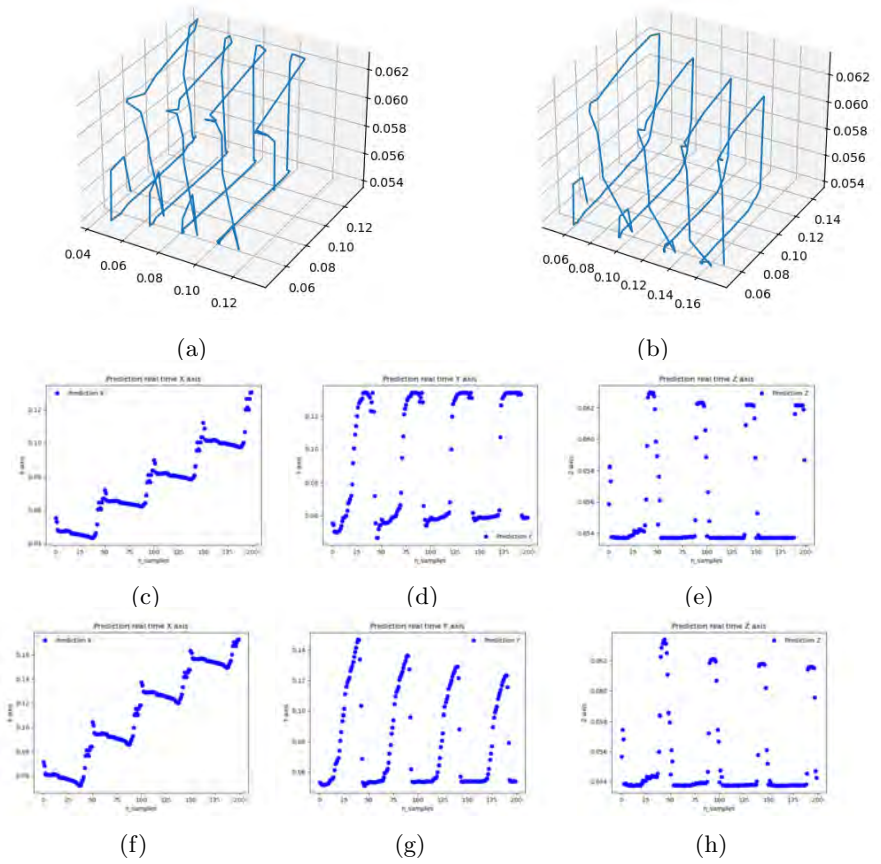

Fig. 5: Predicted 3D trajectory by TCN (a) and GRU in real-time (b); TCN predicted values for (f) x, (g) y, (h) z and GRU predicted values for (i) x, (j) y, (k) z.

a path with a correct trend. Out of all four, GRU had the best performances for Y prediction while TCN performed best for Z.

In both train and test phase, the sequences given as inputs to the network are taken from the ground-truth. In the following results, only the first sequence is considered as known and the others inputs are taken directly from the prediction. The goal is to find a unique path (Γ) for a correct palpation on the specific phantom. The first 50 points of Γ was given in advance to predict the whole curve. The function is initialised with 50 points and it uses the first element of the output vector as the first element of the input for the next iteration.
TCN and GRU result to be the best models from the testing phase, for the real-time only they are implemented. Figure 5 and 5 summarise the results of TCN and GRU model respectively. The x axis is designed to be equal to repeated direction thus an ideal shape for the X coordinates will look like an ascendant step function since the repetitions should be parallel to x axis. The X prediction in TCN is closer to this ideal trend with respect to the GRU.

	TCN	LSTM	GRU	RNN
X	16.39	39.71	14.5	14.46
Y	57.93	27.40	41.37	57.80
Z	15.89	17.50	16.94	15.75

Table 1: Dynamic Time Warping Distance

Y coordinate prediction is unstable in the GRU model (decrease in time), but in TCN, the amplitude remains constant. Considering that the length of the phantom to be examined is equal in each repetition, TCN prediction achieves better performance with respect to GRU. Similar consideration applies to the Z coordinate, the TCN was more stable in Z prediction. On the other hand, the predicted value of X is like the real one in term of dimension. The contact points of Z are comparable to the ground-truth, but the no-contact points are only 1cm higher than the phantom surface, therefore these values have to be increased for a real application. Finally, Y predictions were 0.5 less with respect to the real value.

In conclusion, TCN reflects the best trajectory for the phantom examination. The model can predict a whole trend trajectory that is very close to the expected using only the first 50 time steps, but the dimensions of the final trajectory do not fit perfectly the real values. Nevertheless, the obtained results must be considered a fundamental basis for future works in this research topic. The implementation of the Neural Networks have been made available on github: `https://github.com/imanlab/artemis_dpd`.

The the cumulative mean absolute error between the ground truth and predictions for X, Y and Z in our test set (33% of the data) is $X = 0.03412817658471579$, $Y = 0.03194648418311379$, and $Z = 0.006836229386848896$.

4 Conclusion

Mammography and self-examination are the only tools known for detecting breast cancer in the early stage. The increasing number of new cases all over the years and the importance of correct diagnosis motivated the development of a completely autonomous robot for breast cancer examination. One of the key elements of such a system is learning how to palpate different breasts from expert palpation. This work built the basis of learning from demonstration for breast cancer examination. The highly complex geometry of breasts and different palpation practice, makes it challenging to hard code the palpation actions; hence, we built the palpation trajectory directly from human demonstrations.

We implemented a teleoperation control with force and tactile feedback for data acquisition where this setup enables us to perform palpation on different environments. We collected data of a series of palpation actions on a silicon phantom. Finally, we studied the implementation and results of different deep models for trajectory planning from demonstrations including LSTM, GRU, TCN and GRU. These baseline implementations revealed the pros and cons of each deep

model. Although TCN is among the best model, considerable improvements on the trajectory generation is still needed before making autonomous palpation in future developments.

We identified the limitations of optimisation based approaches, e.g. [18], and conventional learning from demonstration approaches, e.g. DMP [12] can only generalise to new start and goal points whereas in our application the entire palpation path should be adapted to a new breast geometry. Hence, we explored the use of deep model as the mean to generalise the palpation action across different breast geometry.

References

1. Breast self-examination (it can save your life) (2017), \url{https://www.youtube.com/watch?v=LrfE6JUwIms&t=142s&ab_channel=RafflesHospital}
2. The breast exam - stanford medicine 25 (2018), \url{https://www.youtube.com/watch?v=pJ55UtP0_nA&t=410s&ab_channel=StanfordMedicine25}
3. Examination of female breasts (2018), \url{https://www.youtube.com/watch?v=LrfE6JUwIms&t=142s&ab_channel=RafflesHospital}
4. Cancer today (2021), \url{https://gco.iarc.fr/today/home}
5. Ahn, B., Kim, Y., Oh, C.K., Kim, J.: Robotic palpation and mechanical property characterization for abnormal tissue localization. Medical & biological engineering & computing **50**(9), 961–971 (2012)
6. Ayvali, E., Ansari, A., Wang, L., Simaan, N., Choset, H.: Utility-guided palpation for locating tissue abnormalities. IEEE Robotics and Automation Letters **2**(2), 864–871 (2017)
7. Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259 (2014)
8. Elliott, S., Xu, Z., Cakmak, M.: Learning generalizable surface cleaning actions from demonstration. In: 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). pp. 993–999. IEEE (2017)
9. Fu, R., Zhang, Z., Li, L.: Using lstm and gru neural network methods for traffic flow prediction. In: 2016 31st Youth Academic Annual Conference of Chinese Association of Automation (YAC). pp. 324–328. IEEE (2016)
10. Gers, F.A., Schmidhuber, J., Cummins, F.: Learning to forget: Continual prediction with lstm. Neural computation **12**(10), 2451–2471 (2000)
11. Ghalamzan, A., Nazari, K., Hashempour, H., Zhong, F.: Deep-lfd: Deep robot learning from demonstrations. Software Impacts **9**, 100087 (2021)
12. Ghalamzan E., A., Ragaglia, M.: Robot learning from demonstrations: Emulation learning in environments with moving obstacles. Robotics and autonomous systems **101**, 45–56 (2018)
13. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
14. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 156–165 (2017)
15. McDonald, S., Saslow, D., Alciati, M.H.: Performance and reporting of clinical breast examination: a review of the literature. CA: a cancer journal for clinicians **54**(6), 345–361 (2004)

16. Nichols, K.A., Okamura, A.M.: Autonomous robotic palpation: Machine learning techniques to identify hard inclusions in soft tissues. In: 2013 IEEE International Conference on Robotics and Automation. pp. 4384–4389. IEEE (2013)
17. Nichols, K.A., Okamura, A.M.: Methods to segment hard inclusions in soft tissue during autonomous robotic palpation. IEEE Transactions on Robotics **31**(2), 344–354 (2015)
18. Pardi, T., Ortenzi, V., Fairbairn, C., Pipe, T., Ghalamzan E, A., Stolkin, R.: Planning maximum-manipulability cutting paths. IEEE Robotics and Automation Letters **5**(2), 1999–2006 (2020)
19. Pastor, F., Gandarias, J.M., García-Cerezo, A.J., Gómez-de Gabriel, J.M.: Using 3d convolutional neural networks for tactile object recognition with robotic palpation. Sensors **19**(24), 5356 (2019)
20. Pérez-Higueras, N., Caballero, F., Merino, L.: Learning human-aware path planning with fully convolutional networks. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). pp. 5897–5902. IEEE (2018)
21. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning internal representations by error propagation. Tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science (1985)
22. Xiao, B., Xu, W., Guo, J., Lam, H.K., Jia, G., Hong, W., Ren, H.: Depth estimation of hard inclusions in soft tissue by autonomous robotic palpation using deep recurrent neural network. IEEE Transactions on Automation Science and Engineering **17**(4), 1791–1799 (2020)
23. Xie, Z., Zhang, Q., Jiang, Z., Liu, H.: Robot learning from demonstration for path planning: A review. Science China Technological Sciences pp. 1–10 (2020)
24. Yan, Y., Pan, J.: Fast localization and segmentation of tissue abnormalities by autonomous robotic palpation. IEEE Robotics and Automation Letters **6**(2), 1707–1714 (2021)
25. Zhang, Y., Zou, Y., Tang, J., Liang, J.: A lane-changing prediction method based on temporal convolution network. arXiv preprint arXiv:2011.01224 (2020)