Predicting full-arm grasping motions from anticipated tactile responses

Vedant Dave* and Elmar Rueckert*

***Abstract*— Tactile sensing provides significant information about the state of the environment for performing manipulation tasks. Depending on the physical properties of the object, manipulation tasks can exhibit large variation in their movements. For a grasping task, the movement of the arm and of the end effector varies depending on different points of contact on the object, especially if the object is non-homogeneous in hardness and/or has an uneven geometry.**

In this paper, we propose Tactile Probabilistic Movement Primitives (TacProMPs), to learn a highly non-linear relationship between the desired tactile responses and the full-arm movement. We solely condition on the tactile responses to infer the complex manipulation skills. We formulate a joint trajectory of full-arm joints with tactile data, leverage the model to condition on the desired tactile response from the non-homogeneous object and infer the full-arm (7-dof panda arm and 19-dof gripper hand) motion. We use a Gaussian Mixture Model of primitives to address the multimodality in demonstrations. We also show that the measurement noise adjustment must be taken into account due to multiple systems working in collaboration. We validate and show the robustness of the approach through two experiments. First, we consider an object with non-uniform hardness. Grasping different parts of an object require different motion, and results into different tactile responses. Second, we grasp multiple objects at different locations. Our result shows that TacProMPs can successfully model complex multimodal skills and generalise to new situations.

I. INTRODUCTION AND RELATED WORK

Imagine finding an object in a box without seeing anything inside. The only stimulus that you can rely on is the sense of touch. Tactile sensation is an essential tool for robots to interact with the environment. Based on the contact configuration, touch can provide a wide range of information regarding the physical properties and dynamics of the object through a set of diverse signals. This vital information can be exploited for performing various tasks including grasping, dexterous manipulation etc.

Humans can identify physical properties (hardness, roughness, texture) of objects solely by touching [1] and also show a strong reliance on tactile responses for grasping and manipulation tasks. Monzée et al. [2] showed that the relationship between the grip and load force was disrupted if the tactile sensation in humans was eliminated. Removing the tactile feedback through anesthesia takes a heavy toll on motor skills and even the simplest tasks like grasping becomes clumsy and slow [3]. In a recent study, Chinn et al. [4] investigated the strategies in infants to reach the tactile targets on their face. Although it seems easy and obvious for humans to reach any location on the body, it requires a coordinated set of motor skills. We investigate on a similar aspect where such motions can be described through tactile targets.

Fig. 1: Illustration of different grasps/touches that the robot performs on different locations on the object. The robot must predict the action that needs to be executed solely based on the desired tactile response and reach that tactile response.

In the recent years, researchers have found various ways to extract the tactile information and utilise it for robotic tasks [5, 6]. Torres-Jara et al. [7] developed heuristics to reach and grasp the object but had no notion of behaviour inference strategy. Qiang li et al. [8] proposed a framework to realise tactile tracking through feature extraction of tactile images. However, this method requires a selection of task-specific matrix and does not generalise to complex grippers. Similarly, control actions for manipulation are learnt via Reinforcement Learning [9] and Gaussian Processes [10] to reach the desired future sensations. Wang et al. [11] explores tactile space to infer object dynamics and obtain optimal control parameters for desired swing-up angle. Calandra et al. [12] predicted joint torques from raw tactile data and force/torque sensors embedded in the robot joints through Gaussian Processes. Hogan et al. [13] addressed the problem of tactile tracking by decomposing the manipulation plans into sequences of manipulation primitives with planners. However, the motion described in the above methods is usually simple and the information about the manipulator trajectories and gripper configuration is often neglected.

In order to manipulate in an unstructured environment, the robot needs to execute simple motions and learn skills. Learning from Demonstration is a way to learn motion skills via human demonstrations [14]. Probabilistic Movement Primitives (ProMPs) [15] is one such framework that represents the trajectories in a probabilistic manner. Maeda et al. proposed Interaction ProMPs (IProMPs) [16], an extension to ProMPs to model human-robot interaction tasks. It

*The authors are associated with the Cyber-Physical-System Department at Montanuniversität Leoben, Austria.
{vedant.dave,rueckert}@unileoben.ac.at

deals with the scenario where the controlled agent's motion must be adapted based on a set of partial observations of the target agent. Marco et al. [17] improved IProMPs by learning a mixture of IProMPs to learn the multimodality in the unlabelled data. We leverage these ideas to learn the relationship between full-arm trajectories and tactile responses in a probabilistic manner.

Although there exists much work on tactile manipulation [18], there is still room for research that learns complex manipulation strategies dependent on target tactile responses. Similar to real-world cases, we focus on scenarios where large part of the trajectories contains no tactile information as the object is not in contact with the hand. Given these challenges, the keys question arises: Can we predict the robot arm and gripper manipulation, based solely on the desired future tactile response?

Concretely, our contributions are as follows. We present an approach that capture the variance in full-arm movement and exploit it to reach desired targets when conditioned only on the tactile responses. We formulate the whole trajectory with full-arm states and tactile states, which is necessary for constructing a ProMP model. We leverage the idea of interaction primitive so that full-arm trajectories can be produced through conditioning on the desired tactile sensations. We show that a single TacProMP is incapable of generalising to multimodal demonstrations and is outperformed by a mixture of TacProMPs that captures multimodality of the data and provides precise results.

To test the model, we designed two experimental setups. In the first setup, we have an object of non-homogeneous hardness. As shown in the Figure 1, grasping the same object at different parts require widely varying motion. In the second setup, we consider multiple objects at different locations. Grasping every object require distinct manipulation and grasping patterns. During testing (in both the cases), a new tactile response is provided and the full-arm motion is generalised to that response.

The experiments prove that the model is robust enough to deal with the high-dimensional data and non-linear relationships, and at the same time, flexible enough to generalise to new tactile information. We believe the idea of tactile goal-directed manipulation is a significant aspect in robotics, and the simplicity, flexibility and robustness of our approach will bring us a step closer to that goal.

In the next section, we discuss mathematics and framework of the method. In Section III, we introduce our experimental setup and implementations of the method in more detail. Finally, we discuss the work with possible future extensions and present our conclusions.

II. Tactile Probabilistic Movement Primitives

A. Problem Statement

Lets assume $\vec{\mathcal{S}}(t) \in \mathbb{R}^n$ and $\vec{\mathcal{Q}}(t) \in \mathbb{R}^m$ denote the tactile sensor observations and the full-arm trajectory data at time instance t respectively. The full-arm motion comprises of robot arm and hand/gripper trajectory data i.e. $\vec{\mathcal{Q}}(t) = \{\vec{\mathcal{Q}}_h(t), \vec{\mathcal{Q}}_a(t)\}$; $\vec{\mathcal{Q}}_h(t) \in \mathbb{R}^{m_1}, \vec{\mathcal{Q}}_a(t) \in \mathbb{R}^{m_2}, m_1 + m_2 = m$. For a grasping task, the point of contact of the hand with the object typically defines the movement of the arm. The information about the point of contact can be obtained via tactile observations. This assumption is validated in this paper.

The aim is to find the non-linear relationship between this tactile response and the full-arm trajectory:

$$\vec{\mathcal{S}}(t) = f(\vec{\mathcal{Q}}_a(t), \vec{\mathcal{Q}}_h(t)). \tag{1}$$

We formulate this problem with ProMPs approach by considering the tactile response as a trajectory which is jointly distributed with the full-arm trajectory.

B. Probabilistic Movement Primitives

Probabilistic Movement Primitives (ProMPs) [15] was proposed for representing movements based on distribution of the demonstrated trajectories and to compute feedback control laws. It also allows movement modulation, primitive combination, blending and robot control by exploiting variance in trajectories. A single demonstration trajectory is defined by $\boldsymbol{\tau} = \{y_t\}_{t=1}^{T}$, where y_t is a d-dimensional state vector that represents joint angles or cartesian position at time step t. The phase variables $z \in [0, 1]$ decouple the trajectories from time instances. For simplicity, we write $z(t) = t$. The state vector at every time step can be represented through a linear combination of basis functions,

$$y_t = \boldsymbol{\Psi_t} w + \epsilon, \tag{2}$$

where $\boldsymbol{\Psi_t} \in \mathbb{R}^{d \times dK}$ is a block diagonal matrix that contains K Gaussian basis functions ϕ_t for each degree of freedom, $w \in \mathbb{R}^{dK}$ is the weight vector and $\epsilon \sim \mathcal{N}(0, \boldsymbol{\Sigma_y})$ is the Gaussian noise with zero mean and $\boldsymbol{\Sigma_y}$ uncertainty.

Each trajectory can be represented through a weight vector w, characterised by parameters θ. The distribution of the weight vectors over multiple trajectories is assumed to be Gaussian i.e. $p(w; \theta) = \mathcal{N}(w|\mu_w, \boldsymbol{\Sigma_w})$. The distribution of the state vector can be defined as:

$$\begin{aligned} p(y_t, \theta) &= \int \mathcal{N}(y_t|\boldsymbol{\Psi_t} w, \boldsymbol{\Sigma_y})\, \mathcal{N}(w|\mu_w, \boldsymbol{\Sigma_w}) dw, \\ &= \mathcal{N}(y_t|\boldsymbol{\Psi_t}\mu_w, \boldsymbol{\Psi_t}\boldsymbol{\Sigma_w}\boldsymbol{\Psi_t}^T + \boldsymbol{\Sigma_y}). \end{aligned} \tag{3}$$

The weight vector w_i for i-th trajectory is estimated by Regularised Least Squares, $w_i = (\boldsymbol{\Psi}^T\boldsymbol{\Psi} + \lambda \boldsymbol{I})^{-1}\boldsymbol{\Psi}^T\boldsymbol{Y_i}$, where $\boldsymbol{Y_i}$ consists of all the points from i-th trajectory.

C. Tactile ProMPs

We leverage Interaction ProMPs to adapt the motion of the full arm according to the tactile data. From the definitions stated in Section II-A, let's suppose that we observed a sequence of observations that consists of tactile information $s(t) \in \vec{\mathcal{S}}(t)$, arm trajectory $q_a(t) \in \vec{\mathcal{Q}}_a(t)$ and hand trajectory $q_h(t) \in \vec{\mathcal{Q}}_h(t)$. For simplicity, we drop the vector notation. Note that instead of considering only the final tactile state, we consider all the tactile states, beginning from the time the motion is captured. This will result in large part of the tactile states $\mathcal{S}$ to be zero.

In order to capture the correlation between the tactile states and the full-arm movements, we assume that they have a joint distribution. The trajectory at time step t can be represented by $y_t = [s(t)^T, q_h(t)^T, q_a(t)^T]^T$. For every i-th demonstration, the weight vector can be written as

$$\bar{w}_i = \{[w_1^T, ..., w_n^T]^s, [w_1^T, ..., w_m^T]^q\}, \tag{4}$$

We then calculate the weight vectors $\bar{w} \in \mathbb{R}^{(m+n)K}$($K$ basis functions) for every trajectory as defined in Section II-B. Afterwards, the mean vector $\mu_{\bar{w}} \in \mathbb{R}^{(m+n)K}$ and covariance matrix $\Sigma_{\bar{w}} \in \mathbb{R}^{(m+n)K\times(m+n)K}$ are computed.

During inference, the observation from the desired tactile states are designated and used to infer the movement of the full-arm. Consequently, the observation $\boldsymbol{H_t}$ (at time instance t), will consist of two partitions, for observed and inferred data

$$\boldsymbol{H_t} = \left[\begin{array}{ccc|ccc} \boldsymbol{\Psi_t^{(s,1)}} & & 0 & & & \\ & \ddots & & & \mathbf{0} & \\ 0 & & \boldsymbol{\Psi_t^{(s,n)}} & & & \\ \hline & & & \mathbf{0}^{(q,1)} & & 0 \\ & \mathbf{0} & & & \ddots & \\ & & & 0 & & \mathbf{0}^{(q,m)} \end{array}\right] \tag{5}$$

The unobserved data to be inferred is denoted through zeros. Given the tactile state observations, the corresponding trajectories can be obtained by integrating out the weight vector

$$p(y_{1:T}|y^*_{t:t'}) = \int p(y_{1:T}|\bar{w})p(\bar{w}|y^*_{t:t'})d\bar{w}, \tag{6}$$

where $y^*_{t:t'}$ is the desired state vector, in our case, tactile states. The posterior distribution over weights can be computed both, for offline and online cases, in closed-form, i.e.,

$$\begin{aligned} \mu^{new}_{\bar{w}} &= \mu_{\bar{w}} + \boldsymbol{L}\big(y^*_{t:t'} - \boldsymbol{H}_{t:t'}\mu_{\bar{w}}\big), \\ \boldsymbol{\Sigma}^{new}_{\bar{w}} &= \boldsymbol{\Sigma}_{\bar{w}} - \boldsymbol{L}\big(\boldsymbol{H}_{t:t'}\boldsymbol{\Sigma}_{\bar{w}}\big), \\ \boldsymbol{L} &= \boldsymbol{\Sigma}_{\bar{w}}\boldsymbol{H}_{t:t'}\big(\boldsymbol{\Sigma}^*_y + \boldsymbol{H}_{t:t'}\boldsymbol{\Sigma}_{\bar{w}}\boldsymbol{H}^T_{t:t'}\big)^{-1}, \end{aligned} \tag{7}$$

where $\boldsymbol{\Sigma}^*_y$ is the measurement noise and $\boldsymbol{H}_{t:t'}$ is a concatenated $\boldsymbol{H_t}$ matrix from eq. (5) at their corresponding time instances. For implementing it online, the conditioning can be done in a recursive fashion. We condition one observation every time, compute the parameters $\theta^{new} = \{\mu^{new}_{\bar{w}}, \boldsymbol{\Sigma}^{new}_{\bar{w}}\}$ and set them as prior for the new observations.

For measurement noise, we usually use a constant across all the dimensions as they usually originate from the same source [15]. In our case, these sources have largely varying noise. The uncertainty of the estimate in tactile response states is much higher as compared to the estimate in the robot states. We compensate for the different noise levels

$$\Sigma^*_y = \begin{bmatrix} \Sigma^*_s & 0 & 0 \\ 0 & \Sigma^*_a & 0 \\ 0 & 0 & \Sigma^*_h \end{bmatrix},$$

where $\Sigma^*_s, \Sigma^*_a, \Sigma^*_h$ are diagonal matrices indicating uncertainty in tactile response, robot arm and robot hand respectively.

D. Mixture of Tactile ProMPs

For a multimodal task, a single weight vector is not enough to capture the distribution of the demonstrations. For several interaction patterns, it is essential to learn multiple weight vectors s.t. each weight vector corresponds to a different interaction pattern. In order to generalise to multiple interaction patterns, we learn a mixture of primitives from unlabelled data. This is achieved by learning Gaussian Mixture Model (GMM) in the weight space using Expectation-Maximization algorithm (EM) [19].

1) Components in the Mixture: We maintain a mixture of D Gaussians in the GMM, where each component d corresponds to a probability distribution $p(\bar{w}; \alpha_d, \theta_d)$, where $\alpha_d = p(d)$ is the prior probability and $\theta_d = \{\mu_d, \boldsymbol{\Sigma}_d\}$ are the parameters (mean and covariance matrix) of the d-th mixture. The EM algorithm iterates over Expectation step and Maximization step until convergence over the probability distribution of the weights.

Expectation Step: Probability of cluster d given weight vector (responsibility)

$$r_{id} = p(d|\bar{w}_i) = \frac{\alpha_d\, \mathcal{N}(\bar{w}_i|\mu_d, \boldsymbol{\Sigma}_d)}{\sum_{j=1}^{K} \alpha_j\, \mathcal{N}(\bar{w}_i|\mu_j, \boldsymbol{\Sigma}_j)}. \tag{8}$$

Maximization Step: Update the parameters α_d, μ_d and $\boldsymbol{\Sigma}_d$

$$\begin{aligned} n_d &= \textstyle\sum_{j=1}^{n} r_{jd}, \qquad \alpha_d = {n_d}/{n}, \\ \mu_d &= \tfrac{1}{n_d}\Big(\textstyle\sum_{j=1}^{n} r_{jd}\bar{w}_j\Big), \\ \boldsymbol{\Sigma}_d &= \tfrac{1}{n_d}\Big(\textstyle\sum_{j=1}^{n} r_{jd}(\bar{w}_i - \mu_d)(\bar{w}_i - \mu_d)^T\Big). \end{aligned} \tag{9}$$

2) Most Probable Cluster: In order to find the posterior distribution of the weight vector based on the desired observations, it is necessary to find the most probable cluster to which these observations belong. The observations can be written as $y^*_{t:t'} = \{s^*_{t:t'}, q^a_{t:t'}\}$, where $s^*_{t:t'}$ are the desired tactile observations and $q_{t:t'}$ are the full-arm point. Given this observation, the likelihood in the posterior $p(d|y^*_{t:t'}) \propto p(y^*_{t:t'}|d)p(d)$ over the clusters d can be written as

$$p(y^*_{t:t'}|d) = \mathcal{N}(\boldsymbol{H}_{t:t'}\mu_{\bar{w}}, \boldsymbol{\Sigma}^*_y + \boldsymbol{H}_{t:t'}\boldsymbol{\Sigma}_{\bar{w}}\boldsymbol{H}^T_{t:t'}). \tag{10}$$

The most probable cluster d^* can be recognised through $\underset{x}{\operatorname{argmax}}\ p(d|y^*_{t:t'})$. The cluster d^* is later conditioned over the weight vectors to find the posterior distribution over trajectories. As this conditioning is done over the tactile domain $s^* \in y^*_{t:t'}$, the inferred trajectories for the hand and arm will correspond to those tactile state observations. In other words, we can obtain full arm trajectories by just conditioning some final states (desired tactile states) from the desired hand fingers. More about these states are discussed in the Section III.

E. Error Measures

We propose two measures for error measurement.

i) **Anticipated Tactile Response Error (ATRE):** Let the mapping from anticipated tactile data to generated full-arm movements be denoted by $f : \mathbb{R}^n \mapsto \mathbb{R}^m; \{s^*\} \mapsto$

$\{q_h, q_a\}$. Consider that we implement the acquired states $\{q_h, q_a\}$ on the real hardware and extract the resulting tactile response. Let that mapping be denoted by $g : \mathbb{R}^m \mapsto \mathbb{R}^n; \{q_h^{real}, q_a^{real}\} \mapsto \{s^{real}\}$. The reproduced tactile error of the model is defined as

$$\Delta_s = \sum_{t}^{t'} \|s_t^* - s_t^{real}\|^2, \tag{11}$$

where $\|.\|$ denotes the Euclidean Norm and s_t denotes the tactile response observed at time instance t and t' is the last time instance of the selected points.

ii) **Robot RMSE:** This measure denotes the error between the robot states in test demonstrations and the robot states in actual robot when the inferred trajectories are fed into it. The robot consists of two parts: arm (position and orientation) and hand (joint angles). As these three components have separate units, we measure the error separately for all the components:

$$\begin{aligned}
\Delta_a^p &= \sqrt{\frac{1}{N}\sum_{t=t}^{t'} \|q_{(a,p)_t}^* - q_{(a,p)_t}^{real}\|^2}, \\
\Delta_a^o &= \sqrt{\frac{1}{N}\sum_{t=t}^{t'} \|q_{(a,o)_t}^* - q_{(a,o)_t}^{real}\|^2}, \\
\Delta_h &= \sqrt{\frac{1}{N}\sum_{t=t}^{t'} \|q_{h t}^* - q_{h t}^{real}\|^2},
\end{aligned} \tag{12}$$

where Δ_a^p, Δ_a^o, Δ_h are the total error in arm's position, orientation and hand's joint angles respectively, $\{q_{(a,p)_t}^*, q_{(a,o)_t}^*, q_{h t}^*\}$ is the desired robot trajectory at time instance t and $\{q_{(a,p)_t}^{real}, q_{(a,o)_t}^{real}, q_{h t}^{real}\}$ is the learned trajectory that is fed into the real robot at time instance t and N is the number of total points.

F. Feature Scaling of Tactile Responses

It is possible that the responses from the tactile sensor s have significantly different range of values than its counterparts in the model. This will result in allocating very large weights to the tactile responses compared to the other elements ($\{q_h, q_a\}$). This tactile responses are rescaled by min-max normalization and the range is scaled in $[-1, 1]$

$$s' = \frac{s - min(s)}{max(s) - min(s)}, \tag{13}$$

where s' is the normalised tactile response. For simplicity, we will denote s' as s in the next sections.

III. EXPERIMENTS

In this section, we discuss the results from two different experiments performed using the 7-dof Franka Emika Panda with a 19-dof (underactuated) Seed Robotics RH8D Hand [20] with five actuators. Each finger is equipped with tactile sensors that provide 3D tactile data. Consequently, a high-dimensional observation space is generated when these elements are combined.

The first experiment is performed on the object of non-homogeneous hardness to acquire dissimilar tactile responses when different parts of the object are exposed to contact. For the second experiment, multiple objects used in day-to-day life were selected and multiple grasps were performed on each of them. The demonstrations in both the cases have multimodal distribution. At the end, we will discuss the accuracy of our model by comparing the anticipated tactile responses to the tactile responses obtained from the real system.

A. Reaching dissimilar anticipated tactile responses

In this experiment, we have recorded the interactions of the robot with a toy object of irregular hardness at different parts. Basically, the robot encounters two types of tactile responses: high forces when grabbing the bottom part of the object and low forces when grabbing the upper part of the object. As shown in Figure 2, the robot experiences three interactions:

i) Move arm from the left and grasp the hard part of the object with all five fingers.
ii) Move arm from the right and grasp the hard part of the object with all five fingers.
iii) Move arm from the top and grasp the soft part of the object with three fingers (index, middle, thumb).

Note that the resulting hand and the robot arm movement will be different in all the cases as these regions are located apart in the object (Figure 3).

1) Experimental Settings: For each interaction, 20 demonstrations were performed to capture the variance in the movements, resulting in 60 unlabelled demonstrations (interactions are not assigned to any groups). The robot arm was moved to a point within grasping range in gravity compensation mode. Every demonstration was initialised from a random position and orientation (Section C). In order to grasp, we used the robot glove [21] that maps the movement of the human hand to the robot hand, making the process of grasping more human-like and avoid pre-coded trajectories like [17]. The details about experimental data are explained in Section A.

2) Optimal Number of Clusters: In order to decide the number of clusters in the mixture, we measure the Root Mean Square Error (RMSE) of the predicted trajectory of the hand and the arm with its corresponding ground truth using the leave-one-out cross-validation (LOOCV) over the whole dataset. We take robot's pose and hand's joints into consideration. These components are considered vectors and thus the error is the distance in Euclidean space. Figure 5 shows that the RMSE decreases as the number of cluster increases, as expected. The RMSE after 14^{th} cluster is stable, thus we select 14 clusters in the mixture. The mean error and standard deviation of (0.203 ± 0.064) units is achieved. It can be seen from Figure 6(a) that a model with single TacProMP is incapable of generalising to multiple tactile interaction patterns whereas when a mixture of TacProMPs

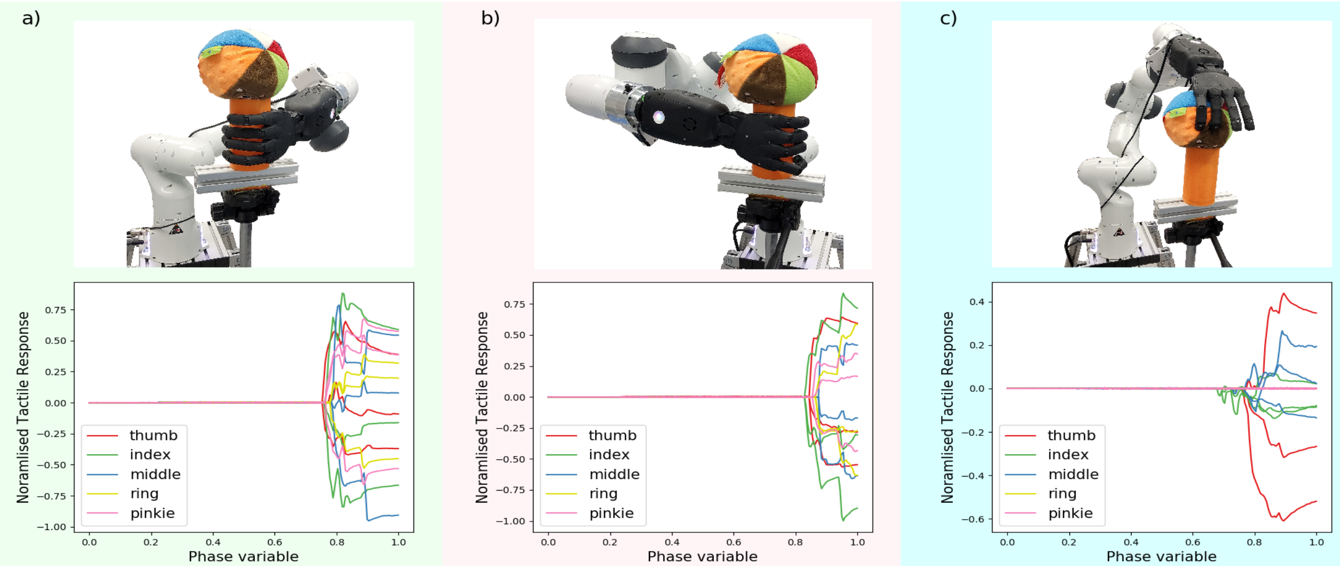

Fig. 2: The figure shows the demonstration of the three interactions of the robot with the object along with the tactile response generated as a result. The bottom row shows the normalised tactile response obtained from 5 fingers along x,y and z axes. In (a) and (b), the magnitude of the tactile response is larger as compared to (c). In (c), the tactile response was zero for the ring (yellow) and the pinkie (pink) finger. Note that the axis scaling in (c) was increased for visualization purpose.

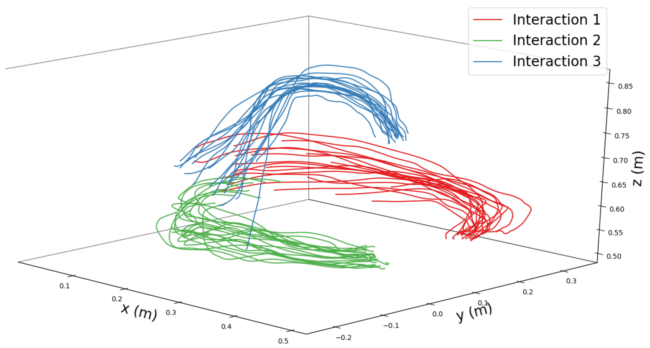

Fig. 3: The Interaction patterns are shown in this figure. The Interaction 1 and 2 are grasping the hard part of the object whereas the interaction 3 grasps the softer part.

are used, a much better performance can be observed in terms of predicting and generalising the interactions (Figure 6(b)).

3) Inference: As described in the Section II, we implement the model on the unlabelled demonstration data and determine their weights. During inference phase, the user provides the anticipated tactile response, which is done simply by establishing the contact between of the fingers with the desired point on the object. We then search the most probable primitive and find the posterior distribution through conditioning (Section II-D.2). The results can be seen in Figure 4. The demonstrated trajectories along some of the dimensions are shown in the upper row of the figure. The conditioning is performed only on the final points of the tactile data, shown by red dots. The black line is the mean and yellow is the variance of the posterior obtained after conditioning.

Feedback Controller. Finally, the mean of arm's posterior distribution is fed to a standard inverse dynamics controller. For the hand, its mean is fed to a joint position controller. Note that the resulting posterior mean for orientation will not obey the unit quaternion constraint, thus it needs to be normalised.

4) Results: The precision of the model is evaluated on the basis of error measures described in the Section II-E. For testing purposes, we capture ten completely new movements along with their tactile responses from these three interactions. For inference, conditioning is performed on the tactile responses of these testing samples and the resulting posterior mean is fed into real robot. The results are discussed below and shown in Table I.

Interaction Recognition. Provided the tactile response, the robot was able to recognise the interaction pattern 100% of the time.

Error Measures. We calculate ATRE (eq. 11) and Robot RMSE (eq. 12) as described in Section II-E. As seen from the Table I, the error in the robot pose and hand's joint positions is significantly smaller compared to tactile error, which validates our claim that tactile information can predict arm and hand movements with high precision. The reproduction error of the tactile observation for all the fingers is in the range of 100 mN, which is not very large if compared to the range of the sensors (0-50N) and the high noise from the sensors. This proves that through our model, movements can be predicted precisely even though the sensor noise is very high.

Note that the error for orientation is calculated as a Euclidean distance, as in our model, we consider it to be in Euclidean Space. More sophisticated approach for error function would be to consider Geodesic distance between two quaternions.

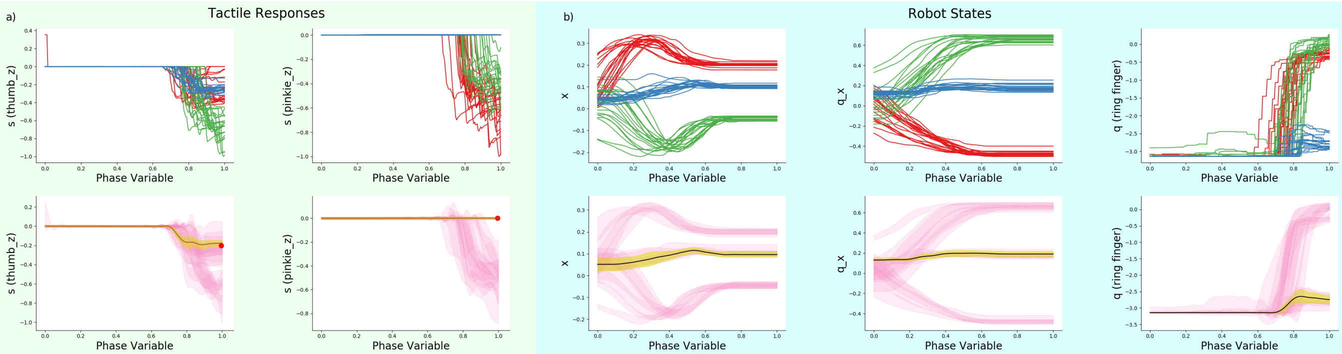

Fig. 4: The upper row shows the data of three interactions collected through demonstrations. The bottom row shows the initial and posterior distribution of the demonstrated data in pink and yellow respectively. a) It shows tactile data obtained from z-direction in thumb and pinkie finger. The red dots in the bottom row indicates the desired tactile observation on which the distribution is conditioned. b) It shows the position (x-direction) and orientation (q_x) of robot arm along with joint angles of the ring finger of the gripper. The inferred mean and variance of the posterior after conditioning is shown by black and yellow colours respectively.

TABLE I: Error measures for Experiment 1

Measure	**ProMPs***	**TacProMPs**
Δ_a^p(cm)	7.38 $\pm$ 4.69	2.97 $\pm$ 1.46
Δ_a^o	0.27 $\pm$ 0.19	0.13 $\pm$ 0.09
Δ_h (radians)	1.44 $\pm$ 0.53	1.35 $\pm$ 0.58
Δ_s (N)	1.622 $\pm$ 1.669	0.806 $\pm$ 0.386

The error measurement is calculated as defined in Section II-E. The error is the resulting mean error and standard deviation ($\mu \pm 2\sigma$) from the 10 test samples. *The result is generated from 6 samples as 4 resulted into failure.

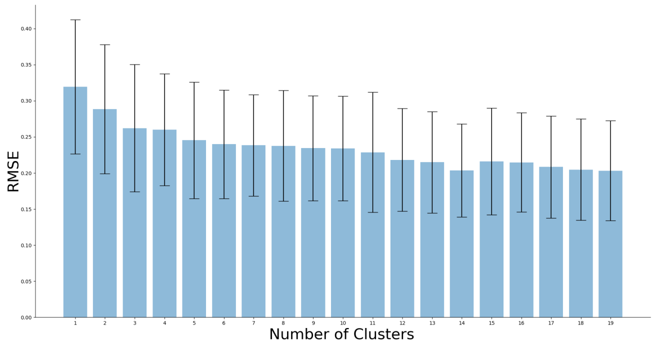

Fig. 5: The RMSE of the trajectories with its ground truth using LOOCV, averaged over all trajectories w.r.t the number of clusters. After 14 clusters, there is no significant improvement in error. Such high number of clusters is the result of high dimensional multimodal data.

B. Grasping multiple objects at different locations

For this experiment, we recorded tactile information from four different objects of different shapes and located at distinct locations (Figure 7). As they are located at different locations, the trajectory of the arm is varying and the grasping is different due to unique shapes of the objects. A total of eight grasps were performed. This increase in number of objects and grasping patterns makes the problem more challenging.

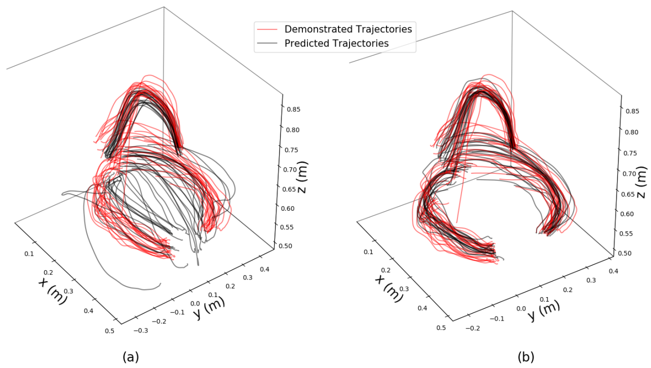

Fig. 6: The prediction of the robot arm's position in Cartesian space generated by LOOCV over the whole dataset. Both subplots show the demonstrated trajectories v/s the trajectories generated by TacProMPs. (a) Prediction using a single Gaussian. (b) Prediction using the mixture of 14 Gaussians.

We capture 20 demonstrations for each of the interactions in a similar way as described in Section III-A. For conditioning on the desired tactile response, we simply establish contact of the robot hand at randomly selected points on the object with different grasping configurations.

Figure 8 shows the RMSE of the predicted trajectory with its corresponding ground truth w.r.t. the number of clusters in the mixture model (eq. 8). Our model consists of 17 clusters in the mixture. The mean error and standard deviation was (0.0729 ± 0.0478) units. Similar to previous experiment, we take ten new samples randomly from these interactions for testing. However, in this case, we experienced some failures during experimentation. During training, the 7 out of 80 (8.75%) trajectories were not recognised in the cluster. For testing, we took ten random samples on four objects, out of which one failed, and other nine succeeded. The error measures are shown in the Table II.

Fig. 7: Four objects shown in the figure are grasped with various configurations at different positions.

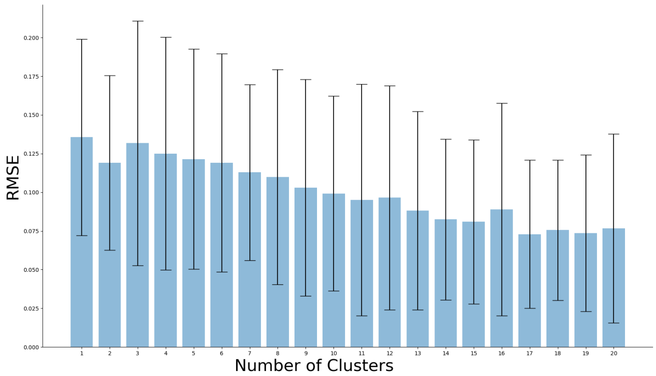

Fig. 8: The RMSE of the trajectories with its ground truth using LOOCV, averaged over all trajectories w.r.t the number of clusters.

TABLE II: Error measures for Experiment 2

Measure	ProMPs*	TacProMPs**
Δ_a^p(cm)	32.27 $\pm$ 15.59	4.90 $\pm$ 3.40
Δ_a^o	0.31 $\pm$ 0.12	0.10 $\pm$ 0.03
Δ_h (radians)	2.33 $\pm$ 1.69	1.3 $\pm$ 1.15
Δ_s (N)	NA	1.295 $\pm$ 0.784

The error is the resulting mean error and standard deviation ($\mu \pm 2\sigma$) from the 10 test samples. *The result is generated from all the samples that accounted for failure. **The result is generated from 9 out of 10 samples as one resulted in failure.

IV. Discussion

It is to be noted that object geometry plays a significant role in our experiments. The sensors receives widely varying signals based on different surfaces and shapes during contact.

The motion of the arm relies heavily on the position of the object and the specific part of its contact. Thus, even for small changes in object position (1-5 cm), the grasping operation often fails. In order to overcome this problem, we plan to track the object's 6D pose via OptiTrack Motion Capture Systems [22] to improve the robustness of our approach.

Like most of previous implementations of ProMPs, the orientation is considered in Euclidean Space for simplicity. However, this model can be extended to Riemannian Manifolds [23].

V. Conclusions

In this paper, we presented an approach to learn the mapping between the tactile responses and the full-arm motion. The problem was formulated as a problem of two agents, tactile responses and full-arm movements, where the partial observations from the tactile responses are provided and the corresponding full-arm movements are required to be inferred. Consequently, a joint trajectory of tactile responses and full-arm joints is generated. The goal was achieved through Interaction Probabilistic Movement Primitives that learns the model of non-linearly correlated trajectories. A Gaussian Mixture Model of IProMPs is incorporated to address the unlabelled multimodal demonstrations. In fact, classical IProMPs fails to solve the task (See Table I and II). Inference is achieved through conditioning the anticipated tactile responses on the most probable GMM component. To compensate for different noise levels from different sources, we change the noise measurement matrix.

In order to capture the wide range of tactile responses, an object of non-homogeneous hardness was selected and different contact points were exploited. To show the robustness of the model, we demonstrated it on multiple objects at distinct locations that mimics real-life scenario. We also incorporated grasping with varying number of fingers that leaves anticipated tactile response to zeros across many dimensions in the tactile space. We showed that only through the anticipated tactile response (even zero), the motion of a complex hand and the arm can be correctly categorised and are generalisable to unseen tactile responses. The results show that high precision manipulation skills can be inferred through our TacProMP model, even with a low cost tactile sensor that is imprecise and produces tactile responses with very high noise.

Acknowledgment

This project has received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - No #430054590 (TRAIN).

Appendix

A. Experimental Data

Here we present the details regarding the data in the experiments.

1) Tactile Response Data: The tactile responses are acquired in form of forces in the (x, y, z) directions for each five fingers, adding upto 15 dimensions in total. These tactile responses are in millinewton (mN). Before implying the algorithm, the tactile responses are normalised as described in Section II-F.

2) Robot Data: The robot arm movement is described by a 7-dimensional vector composed of position and orientation (quaternions) in Cartesian coordinates. The hand movement is described by the joint angles of the five actuators that controls these fingers. These are thumb flexion, thumb abduction, index distal flexion, middle distal flexion, ring and pinkie flexion.

B. TacProMP Parameters

TABLE III: Parameters of a single TacProMP

Parameter	**Symbol**	**Value**
Number of Basis	K	50
Length of Demonstration	T	200
Uncertainty in Tactile Measurement	Σ_s^*	$I_{15\times15} \times 10^{-2}$
Uncertainty in Robot Arm's Measurement	Σ_a^*	$I_{7\times7} \times 10^{-4}$
Uncertainty in Robot Hand's Measurement	Σ_h^*	$I_{5\times5} \times 10^{-3}$

$I_{n\times n}$ refers to $n \times n$ Identity matrix

C. Random Initial Pose

The initial pose is randomised in its working range as follows (position is in meters):
Position: $x = (0, 0.3),\ y = (0, 0.3),\ z = (0.4, 0.75)$
Orientation : $q_w = (-0.25, 0.5) \quad q_x = (-1, 0.75)$
$q_y = (0, 0.5) \quad q_z = (-1, 0.75)$

References

[1] Roberta L Klatzky, Susan J Lederman, and Victoria A Metzger. "Identifying objects by touch: An expert system". In: *Perception & psychophysics* 37.4 (1985), pp. 299–302.

[2] Joël Monzée, Yves Lamarre, and Allan M Smith. "The effects of digital anesthesia on force control using a precision grip". In: *Journal of neurophysiology* 89.2 (2003), pp. 672–683.

[3] Roland Johansson and John Flanagan. "Coding and use of tactile signals from the fingertips in object manipulation tasks". In: *Nature reviews. Neuroscience* 10 (May 2009), pp. 345–59.

[4] Lisa K Chinn et al. "Development of infant reaching strategies to tactile targets on the face". In: *Frontiers in psychology* 10 (2019), p. 9.

[5] Ravinder Dahiya et al. "Tactile Sensing—From Humans to Humanoids". In: *IEEE Transactions on Robotics* 26 (Feb. 2010), pp. 1–20.

[6] Akihiko Yamaguchi and Christopher G. Atkeson. "Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision?" In: *Advanced Robotics* 33.14 (2019), pp. 661–673.

[7] Eduardo Torres-Jara and Lorenzo Natale. "Sensitive Manipulation: Manipulation Through Tactile Feedback". In: *International Journal of Humanoid Robotics* 15.01 (2018), p. 1850012.

[8] Qiang li et al. "A Control Framework for Tactile Servoing". In: June 2013.

[9] Stephen Tian et al. "Manipulation by Feel: Touch-Based Control with Deep Predictive Models". In: *2019 International Conference on Robotics and Automation (ICRA)*. 2019, pp. 818–824.

[10] Filipe Veiga et al. "Tactile based forward modeling for contact location control". In: *RSS Workshop on Tactile Sensing for Manipulation*. 2017.

[11] Chen Wang et al. "Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation". In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2020, pp. 5633–5640.

[12] Roberto Calandra et al. "Learning Inverse Dynamics Models with Contacts". In: *Proceedings of the International Conference on Robotics and Automation (ICRA)*. Jan. 1, 2015. published.

[13] Francois R. Hogan et al. "Tactile Dexterity: Manipulation Primitives with Tactile Feedback". In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020, pp. 8863–8869.

[14] Aude Billard, Sylvain Calinon, and Rüdiger Dillmann. "Learning from Humans". In: *Springer Handbook of Robotics, 2nd Ed.* 2016.

[15] Alexandros Paraschos et al. "Probabilistic Movement Primitives". In: *Advances in Neural Information Processing Systems*. Ed. by C.J. Burges et al. Vol. 26. Curran Associates, Inc., 2013.

[16] Guilherme Maeda et al. "Learning interaction for collaborative tasks with probabilistic movement primitives". In: *2014 IEEE-RAS International Conference on Humanoid Robots*. 2014, pp. 527–534.

[17] Marco Ewerton et al. "Learning Multiple Collaborative Tasks with a Mixture of Interaction Primitives". In: vol. 2015. May 2015.

[18] Qiang li et al. "A Review of Tactile Information: Perception and Action Through Touch". In: *IEEE Transactions on Robotics* PP (July 2020), pp. 1–16.

[19] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. 1st ed. Springer, 2007. ISBN: 0387310738.

[20] *Seed Robotics RH8D Hand*. https://www.seedrobotics.com/rh8d-adult-robot-hand.

[21] Robin Denz et al. "A high-accuracy, low-budget Sensor Glove for Trajectory Model Learning". In: *2021 20th International Conference on Advanced Robotics (ICAR)*. 2021, pp. 1109–1115.

[22] *OptiTrack - Motion Capture Systems*. https://optitrack.com.

[23] Leonel Rozo and Vedant Dave. "Orientation Probabilistic Movement Primitives on Riemannian Manifolds". In: *Proceedings of the 5th Conference on Robot Learning*. Vol. 164. Proceedings of Machine Learning Research. PMLR, Nov. 2020, pp. 373–383.