

# Real-time Local GP Model Learning

Duy Nguyen-Tuong, Matthias Seeger and Jan Peters

## Abstract

For many applications in robotics, accurate dynamics models are essential. However, in some applications, e.g., in model-based tracking control, precise dynamics models cannot be obtained analytically for sufficiently complex robot systems. In such cases, machine learning offers a promising alternative for approximating the robot dynamics using measured data. However, standard regression methods such as Gaussian process regression (GPR) suffer from high computational complexity which prevents their usage for large numbers of samples or online learning to date. In this paper, we propose an approximation to the standard GPR using local Gaussian processes models inspired by [1, 2]. Due to reduced computational cost, local Gaussian processes (LGP) can be applied for larger sample-sizes and online learning. Comparisons with other nonparametric regressions, e.g., standard GPR, support vector regression (SVR) and locally weighted projection regression (LWPR), show that LGP has high approximation accuracy while being sufficiently fast for real-time online learning.

## 1 Introduction

Precise models of technical systems can be crucial in technical applications [3, 4]. In robot tracking control, only well-estimated inverse dynamics models allow both high accuracy and compliant, low-gain control. For complex robots such as humanoids or light-weight arms, it is often hard to analytically model the system sufficiently well and, thus, modern regression methods can offer a viable alterna-

---

Duy Nguyen-Tuong, Jan Peters  
Max Planck Institute for Biological Cybernetics, 72076 Tübingen, e-mail: {duy.nguyen-tuong, jan.peters}@tuebingen.mpg.de,

Matthias Seeger  
Saarland University, 66123 Saarbrücken, e-mail: mseeger@mmci.uni-saarland.de

tive [1, 5]. However, highly accurate regression methods such as Gaussian process regression (GPR) suffer from high computational cost, while fast real-time learning algorithms such as locally weighted projection regression (LWPR) are not straightforward to use, as they require manual adjustment of many data dependent parameters.

In this paper, we attempt to combine the strengths of both approaches, i.e., the high accuracy and comfortable use of GPR with the fast learning speed of LWPR. We will proceed as follows: firstly, we briefly review both model-based control as well as standard Gaussian process regression. We will discuss the necessity of estimating the inverse dynamics model for compliant, low-gain control. Subsequently, we describe our local Gaussian process models (LGP) approach.

In Section 3, the learning accuracy and performance of the presented LGP approach will be compared with several relevant regression methods, e.g., standard GPR [6],  $\nu$ -support vector regression ( $\nu$ -SVR) [7], sparse online GP (OGP) [8] and LWPR [1, 5]. The applicability of the LGP for low-gain model-based tracking control and real-time learning is demonstrated on a Barrett whole arm manipulator (WAM). We can show that its tracking performance exceeds analytical models [9] while remaining fully compliant.

### 1.1 Model-based Control

Model-based control, e.g., computed torque control [10], enables high speed and compliant robot control while achieving accurate control with small tracking errors for sufficiently precise robot models. The controller is supposed to move the robot that is governed by the system dynamics [10]

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) + \boldsymbol{\varepsilon}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}) = \mathbf{u}, \quad (1)$$

where  $\mathbf{q}$ ,  $\dot{\mathbf{q}}$ ,  $\ddot{\mathbf{q}}$  are joint angles, velocities and accelerations of the robot, respectively,  $\mathbf{u}$  denotes the applied torques,  $\mathbf{M}(\mathbf{q})$  the inertia matrix of the robot and  $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$  Coriolis and centripetal forces,  $\mathbf{G}(\mathbf{q})$  gravity forces and  $\boldsymbol{\varepsilon}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})$  represents nonlinearities of the robot which are not part of the rigid-body dynamics.

The model-based tracking control law determines the joint torques  $\mathbf{u}$  necessary for following a desired trajectory  $\mathbf{q}_d$ ,  $\dot{\mathbf{q}}_d$ ,  $\ddot{\mathbf{q}}_d$  using a dynamics model while employing feedback in order to stabilize the system. For example, the inverse dynamics model of the robot can be used as a feed-forward model that predicts the joint torques  $\mathbf{u}_{\text{FF}}$  required to perform the desired trajectory [9, 10], while a feedback term  $\mathbf{u}_{\text{FB}}$  ensures the stability of the tracking control with a resulting control law of  $\mathbf{u} = \mathbf{u}_{\text{FF}} + \mathbf{u}_{\text{FB}}$ . The feedback term can be a linear control law such as  $\mathbf{u}_{\text{FB}} = \mathbf{K}_p \mathbf{e} + \mathbf{K}_v \dot{\mathbf{e}}$ , where  $\mathbf{e} = \mathbf{q}_d - \mathbf{q}$  denotes the tracking error and  $\mathbf{K}_p$ ,  $\mathbf{K}_v$  position-gain and velocity-gain, respectively. If an accurate model in the form of Equation (1) can be obtained for the feed-forward model, e.g., for negligible unknown non-

linearities  $\varepsilon$ , the resulting feed-forward term  $\mathbf{u}_{\text{FF}}$  will largely cancel the robots nonlinearities [10].

For complex robots such as humanoids or light-weight arms, it is often hard to model the system sufficiently well using the rigid body dynamics. Unknown nonlinearities  $\varepsilon(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})$  such as flexible hydraulic tubes, complex friction, gear boxes, etc, couple several degrees of freedom together and result in highly altered dynamics. Such unknown nonlinearities can dominate the system dynamics and deteriorate the analytical model [11]. The resulting tracking error needs to be compensated using large gains [10]. However, high feedback gains prohibit compliant control and, thus, make the robot less safe for the environment while causing many practical problems such as actuator saturation, excitation of unmodeled dynamics, may result in large tracking errors in presence of noise, increase energy consumption, etc. To avoid high-gain feedback, it is essential to improve the accuracy of the dynamics model for predicting  $\mathbf{u}_{\text{FF}}$ . Since  $\mathbf{u}_{\text{FF}}$  is a function of  $\mathbf{q}_d, \dot{\mathbf{q}}_d, \ddot{\mathbf{q}}_d$ , it can be obtained with supervised learning using measured data. The resulting problem is a regression problem that can be solved by learning the mapping  $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}} \rightarrow \mathbf{u}$  on sampled data [12, 13] and, subsequently, using the resulting mapping for determining the feed-forward motor commands. As trajectories and corresponding joint torques are sampled directly from the real robot, learning the mapping will include all nonlinearities and not only the ones described in the rigid-body model.

## 1.2 Regression with Standard GPR

As any realistic inverse dynamics is a well-defined functional mapping of continuous, high-dimensional inputs to outputs of the same kind, we can view it as a regression problem. Given the input  $\mathbf{x} \in \mathbb{R}^n$  and the target  $\mathbf{y} \in \mathbb{R}^n$ , the task of regression algorithms is to learn the mapping describing the relationship from input to target using samples. A powerful method for accurate function approximation in high-dimensional space is Gaussian process regression (GPR) [6]. Given a set of  $n$  training data points  $\{\mathbf{x}_i, y_i\}_{i=1}^n$ , we would like to learn a function  $f(\mathbf{x}_i)$  transforming the input vector  $\mathbf{x}_i$  into the target value  $y_i$  given a model  $y_i = f(\mathbf{x}_i) + \varepsilon_i$ , where  $\varepsilon_i$  is Gaussian noise with zero mean and variance  $\sigma_n^2$  [6]. As a result, the observed targets can also be described by a Gaussian distribution  $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I})$ , where  $\mathbf{X}$  denotes the set containing all input points  $\mathbf{x}_i$  and  $\mathbf{K}(\mathbf{X}, \mathbf{X})$  the covariance matrix computed using a given covariance function. Gaussian kernels are probably the most frequently used covariance functions [6] and are given by

$$k(\mathbf{x}_p, \mathbf{x}_q) = \sigma_s^2 \exp\left(-\frac{1}{2}(\mathbf{x}_p - \mathbf{x}_q)^T \mathbf{W}(\mathbf{x}_p - \mathbf{x}_q)\right), \quad (2)$$

where  $\sigma_s^2$  denotes the signal variance and  $\mathbf{W}$  represents the widths of the Gaussian kernel. Other choices for possible kernels can be found in [6, 7]. The joint distribution of the observed target values and predicted value  $f(\mathbf{x}_*)$  for a query point  $\mathbf{x}_*$  is

given by

$$\begin{bmatrix} \mathbf{y} \\ f(\mathbf{x}_*) \end{bmatrix} \sim \mathcal{N} \left( \mathbf{0}, \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I} & k(\mathbf{X}, \mathbf{x}_*) \\ k(\mathbf{x}_*, \mathbf{X}) & k(\mathbf{x}_*, \mathbf{x}_*) \end{bmatrix} \right). \quad (3)$$

Conditioning the joint distribution yields the predicted mean value  $f(\mathbf{x}_*)$  with the corresponding variance  $V(\mathbf{x}_*)$

$$\begin{aligned} f(\mathbf{x}_*) &= k_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} = k_*^T \boldsymbol{\alpha}, \\ V(\mathbf{x}_*) &= k(\mathbf{x}_*, \mathbf{x}_*) - k_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} k_*, \end{aligned} \quad (4)$$

with  $k_* = k(\mathbf{X}, \mathbf{x}_*)$ ,  $\mathbf{K} = \mathbf{K}(\mathbf{X}, \mathbf{X})$  and  $\boldsymbol{\alpha}$  denotes the so-called prediction vector. The hyperparameters of a Gaussian process with Gaussian kernel are given by  $\theta = [\sigma_n^2, \sigma_f^2, \mathbf{W}]$  and remain the only open parameters. Their optimal value for a particular data set can be automatically estimated by maximizing the log marginal likelihood using standard optimization methods such as Quasi-Newton methods [6].

## 2 Local Gaussian Process Regression

Due to high computational complexity of nonlinear regression techniques, inverse dynamics models are frequently only learned offline for pre-sampled desired trajectories [13]. In order to take full advantage of a learning approach, online learning is an absolute necessity as it allows the adaption to changes in the robot dynamics, load or the actuators. Furthermore, a training data set will never suffice for most robots with a large number of degrees of freedom and, thus, fast online learning is necessary if the trajectory leads to new parts of the state-space. However, for most real-time applications online model learning poses a difficult regression problem due to three constraints, i.e., firstly, the learning and prediction process should be very fast (e.g., learning needs to take place at a speed of 20-200Hz and prediction may take place at 200Hz up to 5kHz). Secondly, the learning system needs to be capable of dealing with large amounts of data (i.e., with data arriving at 200Hz, less than ten minutes of runtime will result in more than a million sampled data points). And, thirdly, the data arrives as a continuous stream, thus, the model has to be continuously adapted to new training examples over time.

Model learning with GPR suffers from the expensive computation of the inverse matrix  $(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1}$  which yields a cost of  $\mathcal{O}(n^3)$ , see Equation (4). Inspired by locally weighted regression [1,5], we propose a method for speed-up the training and prediction process by partitioning the training data in local regions and learning an independent Gaussian process model (as given in Section 1.2) for each region. The number of data points in the local models is limited, where insertion and removal of data points can be treated in a principled manner. The prediction for a query point is performed by weighted average similar to LWPR [1]. For partitioning and weighted prediction we use a kernel as similarity measure. Thus, our algorithm consists out of three stages: (i) clustering of data, i.e., insertion of new data points into the local

models, (ii) learning of corresponding local models and (iii) prediction for a query point.

## 2.1 Partitioning of Training Data

Clustering input data can be performed efficiently using a similarity measure between the input point  $\mathbf{x}$  and the centers of the respective local models. From a machine learning point of view, the similarity or proximity of data points can be defined in terms of a kernel. Kernel functions represent the dot product between two vectors in the feature space and, hence, naturally incorporate the similarity measure between data points. The clustering step described in this section results from the basic assumption that nearby input points are likely to have similar target values. Thus, training points that belong to the same local region (represented by a center) are informative about the prediction for query points next to this local region.

A specific characteristic in this framework is that we take the kernel for learning the Gaussian process model as similarity measure  $w_k$  for the clustering process. If a Gaussian kernel is employed for learning the model, the corresponding measure will be

$$w_k(\mathbf{x}, \mathbf{c}_k) = \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{c}_k)^T \mathbf{W}(\mathbf{x} - \mathbf{c}_k)\right), \quad (5)$$

where  $\mathbf{c}_k$  denotes the center of the  $k$ -th local model and  $\mathbf{W}$  a diagonal matrix represented the kernel width. It should be emphasized that for learning the Gaussian process model any admissible kernel can be used. Thus, the similarity measure for the clustering process can be varied in many ways, and, for example, the commonly used Matern kernel [14] could be used instead of the Gaussian one. For the hyperparameters of the measure, such as  $\mathbf{W}$  for Gaussian kernel, we use the *same* training approach as introduced in Section 1.2. Since the hyperparameters of a Gaussian process model can be achieved by likelihood optimization, it is straightforward to adjust the open parameters for the similarity measure. For example, we can subsample the available training data and, subsequently, perform the standard optimization procedure.

After computing the proximity between the new data point  $\mathbf{x}_{new}$  and all available centers, the data point will be included to the *nearest* local model, i.e., the one with the maximal value of  $w_k$ . As the data arrives incrementally over time, a new model with center  $\mathbf{c}_{k+1}$  is created if all similarity measures  $w_k$  fall below a threshold  $w_{gen}$ . The new data point is then used as new center  $\mathbf{c}_{k+1}$  and, thus, the number of local models will increase if previously unknown parts of the state space are visited. When a new data point is assigned to a particular  $k$ -th model, i.e.,  $\max_k w_k(\mathbf{x}) > w_{gen}$  the center  $\mathbf{c}_k$  will be updated to the mean of corresponding local data points.

**Algorithm 1:** Partitioning the training data with incremental model learning.

---

**Input:** new data point  $\{\mathbf{x}_{\text{new}}, y_{\text{new}}\}$ .  
**for**  $k = 1$  **to** number of local models **do**  
  Compute proximity to the  $k$ -th local model:  
    $w_k = k(\mathbf{x}_{\text{new}}, \mathbf{c}_k)$   
**end for**  
Take the nearest local model:  
 $v = \max_k w_k$   
**if**  $v > w_{\text{gen}}$  **then**  
  Insert  $\{\mathbf{x}_{\text{new}}, y_{\text{new}}\}$  into the nearest local model:  
    $\mathbf{X}_{\text{new}} = [\mathbf{X}, \mathbf{x}_{\text{new}}]$ ,  $\mathbf{y}_{\text{new}} = [y, y_{\text{new}}]$   
  Update the corresponding center:  
    $\mathbf{c}_{\text{new}} = \text{mean}(\mathbf{X}_{\text{new}})$   
  Update the Cholesky matrix and the prediction vector of local model:  
   Compute  $l$  and  $l_*$   
   Compute  $\mathbf{L}_{\text{new}}$   
  If the maximum number of data points is reached  
   delete another point by permutation.  
  Compute  $\alpha_{\text{new}}$  by back-substitution  
**else**  
  Create new model:  
    $\mathbf{c}_{k+1} = \mathbf{x}_{\text{new}}$ ,  $\mathbf{X}_{k+1} = [\mathbf{x}_{\text{new}}]$ ,  $\mathbf{y}_{k+1} = [y_{\text{new}}]$   
  Initialize of new Cholesky matrix  $\mathbf{L}$  and new prediction vector  $\alpha$ .  
**end if**

---

**Algorithm 2:** Prediction for a query point.

---

**Input:** query data point  $\mathbf{x}$ ,  $M$ .  
Determine  $M$  local models closest to  $\mathbf{x}$ .  
**for**  $k = 1$  **to**  $M$  **do**  
  Compute proximity to the  $k$ -th local model:  
    $w_k = k(\mathbf{x}, \mathbf{c}_k)$   
  Compute local prediction using the  $k$ -th local model:  
    $\hat{y}_k = \mathbf{k}_k^T \alpha_k$   
**end for**  
Compute weighted prediction using  $M$  local models:  
 $\hat{y} = \sum_{k=1}^M w_k \hat{y}_k / \sum_{k=1}^M w_k$ .

---

## 2.2 Incremental Update of Local Models

During online learning, we have to deal with an endless stream of data (e.g., at a 500 Hz sampling rate we get a new data point every 2 ms and have to treat 30 000 data points per minute). In order to cope with the real-time requirements, the maximal number of training examples needs to be limited so that the local models do not end up with the same complexity as a standard GPR regression. Since the number of acquired data points increases continuously over time, we can enforce this limit by incrementally deleting old data points when newer and better ones are included. Insertion and deletion of data points can be achieved using first order principles, for

example, maximizing the information gain while staying within a budget (e.g., the budget can be a limit on the number of data points). Nevertheless, while the update of the target vector  $\mathbf{y}$  and input matrix  $\mathbf{X}$  can be done straightforwardly, the update of the covariance matrix (and implicitly the update of the prediction vector  $\alpha$ , see Equation (4)) is more complicated to derive and requires thorough analysis given here.

The prediction vector  $\alpha$  can be updated incrementally by directly adjusting the Cholesky decomposition of the Gram matrix ( $\mathbf{K} + \sigma_n^2 \mathbf{I}$ ) as suggested in [15]. For doing so, the prediction vector can be rewritten as  $\mathbf{y} = \mathbf{L}\mathbf{L}^T \alpha$ , where the lower triangular matrix  $\mathbf{L}$  is a Cholesky decomposition of the Gram matrix. Incremental insertion of a new point is achieved by adding an additional row to the matrix  $\mathbf{L}$ .

**Proposition 1.** *If  $\mathbf{L}$  is the Cholesky decomposition of the Gram matrix  $\mathbf{K}$  while  $\mathbf{L}_{new}$  and  $\mathbf{K}_{new}$  are obtained by adding additional row and column, such that*

$$\mathbf{L}_{new} = \begin{bmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{l}^T & l_* \end{bmatrix}, \mathbf{K}_{new} = \begin{bmatrix} \mathbf{K} & \mathbf{k}_{new}^T \\ \mathbf{k}_{new} & k_{new} \end{bmatrix}, \quad (6)$$

with  $\mathbf{k}_{new} = k(\mathbf{X}, \mathbf{x}_{new})$  and  $k_{new} = k(\mathbf{x}_{new}, \mathbf{x}_{new})$ , then  $\mathbf{l}$  and  $l_*$  can be computed by solving

$$\mathbf{L}\mathbf{l} = \mathbf{k}_{new} \quad (7)$$

$$l_* = \sqrt{k_{new} - \|\mathbf{l}\|^2} \quad (8)$$

*Proof.* Multiply out the equation  $\mathbf{L}_{new}\mathbf{L}_{new}^T = \mathbf{K}_{new}$  and solve for  $\mathbf{l}$  and  $l_*$ .  $\square$

Since  $\mathbf{L}$  is a triangular matrix,  $\mathbf{l}$  can be determined from Equation (7) by substituting it back in after computing the kernel vector  $\mathbf{k}_{new}$ . Subsequently,  $l_*$  and the new prediction vector  $\alpha_{new}$  can be determined from Equation (8), where  $\alpha_{new}$  can be achieved by twice back-substituting while solving  $\mathbf{y}_{new} = \mathbf{L}_{new}\mathbf{L}_{new}^T \alpha_{new}$ . If the maximal number of training examples is reached, an old data point has to be deleted every time when a new point is being included. The deletion of the  $m$ -th data point can be performed efficiently using a permutation matrix  $\mathbf{R}$  and solving  $\mathbf{y}_{new} = \mathbf{R} \mathbf{L}_{new} \mathbf{L}_{new}^T \mathbf{R} \alpha_{new}$ , where  $\mathbf{R} = \mathbf{I} - (\delta_m - \delta_n)(\delta_m - \delta_n)^T$  and  $\delta_i$  is a zero vector whose  $i$ -th element is one [15]. In practice, the new data point is inserted as a first step to the *last* row ( $n$ -th row) according to Equation (6) and, subsequently, the  $m$ -th data point is removed by adjusting  $\mathbf{R}$ . The partitioning and learning process is summarized in Algorithm 1. The incremental Cholesky update is very efficient and can be performed in a numerically stable manner as discussed in detail in [15].

Due to the Cholesky update formulation, the amount of computation for training can be limited due to the incremental insertion and deletion of data points. The main computational cost for learning the local models is dominated by the incremental update of the Cholesky matrix which yields  $\mathcal{O}(N_l^2)$ , where  $N_l$  presents the number of data points in a local model. Importantly,  $N_l$  can be set in accordance with the computational power of the available real-time computer system.

### 2.3 Prediction using Local Models

The prediction for a mean value  $\hat{y}$  is performed using weighted averaging over  $M$  local GP predictions  $\bar{y}_k$  for a query point  $\mathbf{x}$  similar to LWPR [1]. The weighted prediction  $\hat{y}$  is then given by  $\hat{y} = \mathbb{E}\{\bar{y}_k | \mathbf{x}\} = \sum_{k=1}^M \bar{y}_k p(k | \mathbf{x})$ . According to the Bayesian theorem, the probability of the model  $k$  given query point  $\mathbf{x}$  can be expressed as

$$p(k | \mathbf{x}) = \frac{p(k, \mathbf{x})}{p(\mathbf{x})} = \frac{p(k, \mathbf{x})}{\sum_{k=1}^M p(k, \mathbf{x})} = \frac{w_k}{\sum_{k=1}^M w_k}. \quad (9)$$

Hence, we have

$$\hat{y} = \frac{\sum_{k=1}^M w_k \bar{y}_k}{\sum_{k=1}^M w_k}, \quad (10)$$

Thus, each local GP prediction  $\bar{y}_k = k(\mathbf{X}_k, \mathbf{x})^T \alpha_k$  is additionally weighted by the similarity  $w_k(\mathbf{x}, \mathbf{c}_k)$  between the corresponding center  $\mathbf{c}_k$  and the query point  $\mathbf{x}$ . The search for  $M$  local models can be quickly done by evaluating the proximity between the query point  $\mathbf{x}$  and all model centers  $\mathbf{c}_k$ . The prediction procedure is summarized in Algorithm 2.

## 3 Learning Inverse Dynamics for Model-based Control



(a) SARCOS arm (b) Barrett WAM

Fig. 1: Robot arms used for data generation and experiments.

Learning models for control of high-dimensional systems in real-time is a difficult endeavor and requires extensive evaluation. For this reason, we evaluate our algorithm (LGP) using high-dimensional data taken from two real robots, e.g., the 7 degree-of-freedom (DoF) anthropomorphic SARCOS master arm and 7-DoF Barrett WAM both shown in Figure 1. Subsequently, we apply LGP for online learning of inverse dynamics models for robot tracking control. The tracking control task with model online learning is performed on the Barrett WAM in real-time. Finally, we highlight the

advantages of online-learned models versus offline approximation and analytical model in a more complex experiment for learning of character writing.

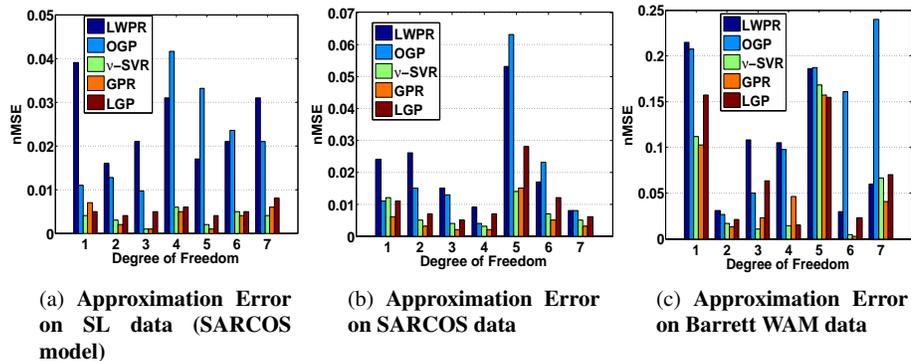


Fig. 2: The approximation error is represented by the normalized mean squared error (nMSE) for each DoF (1–7) and shown for (a) simulated data from physically realistic SL simulation, (b) real robot data from an anthropomorphic SARCOS master arm and (c) measurements from a Barrett WAM. In all cases, LGP outperforms LWPR and OGP in learning accuracy while being competitive to  $v$ -SVR and standard GPR. The small variances of the output targets in the Barrett data results in a nMSE that is a larger scale compared to SARCOS; however, this increase has no practical meaning and only depends on the training data.

### 3.1 Learning Accuracy Comparison

In this section, we compare the learning performance of LGP with the state-of-the-art in nonparametric regression, e.g., LWPR,  $v$ -SVR [7], standard GPR and online Gaussian Process Regression (OGP) [7] in the context of approximating inverse robot dynamics. For evaluating  $v$ -SVR and GPR, we have employed the libraries [16] and [17], respectively. The code for LGP contained also parts of the library [17].

For comparing the prediction accuracy of our proposed method in the setting of learning inverse dynamics, we use three data sets, (i) SL simulation data (SARCOS model) as described in [13] (14094 training points and 5560 test points), (ii) data from the SARCOS master arm (13622 training points and 5500 test points) [1] as well as (iii) a data set generated from our Barrett arm (13572 training points, 5000 test points). Given samples  $\mathbf{x} = [\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}]$  as input, where  $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}$  denote the joint angles, velocity and acceleration, respectively, and using the corresponding joint torques  $\mathbf{y} = [\mathbf{u}]$  as targets, we have a well-defined, proper regression problem. The considered seven degrees of freedom (DoF) robot arms result in 21 input dimensions (i.e., for each joint, we have an angle, a velocity and an acceleration) and seven target or output dimensions (i.e., a single torque for each joint). The robot inverse dynamics model can be estimated separately for each DoF employing LWPR,  $v$ -SVR, GPR, OGP and LGP, respectively.

The training examples for LGP can be partitioned either in the same input space where the local models are learned or in a subspace that has to be physically consistent with the approximated function. In the following, we localize the data depending on the position of the robot. Thus, the partitioning of training data is performed in a seven dimensional space (i.e., consisting of the seven joint angles). After deter-

mining the similarity metric  $w_k$  for all  $k$  local models in the partitioning space, the input point will be assigned to the *nearest* local model, i.e., the local model with the maximal value of distance measure  $w_k$ . For computing the localization, we will use the Gaussian kernel as given in Equation (2) and the corresponding hyperparameters are optimized using a subset of the training set.

Note that the choice of the limit value  $w_{\text{gen}}$  during the partitioning step is crucial for the performance of LGP and, unfortunately, is an open parameter requiring manual tuning. If  $w_{\text{gen}}$  is too small, a large number of local models will be generated with small number of training points. As these small models receive too little data for a stable GPR, they do not generalize well to unknown neighboring regions of the state space. If  $w_{\text{gen}}$  is large, the local models will include too many data points which either results in over-generalization or, if the number of admitted data points is enlarged as well, it will increase the computational complexity. Here, the training data is clustered in about 30 local regions ensuring that each local model has a sufficient amount of data points for high accuracy (in practice, roughly a hundred data points for each local model suffice) while having sufficiently few that the solution remains feasible in real-time (e.g., on the test hardware, an Intel Core Duo at 2GHz, that implies the usage of up to a 1000 data points per local model). On average, each local model includes approximately 500 training examples, i.e., some models will not fill up while others actively discard data. This small number of training data points enables a fast training for each local model using the previously described fast Cholesky matrix updates.

Figure 2 shows the normalized mean squared error (nMSE) of the evaluation on the test set for each of the three evaluated scenarios, i.e., a physically realistic simulation of the SARCOS arm in Figure 2 (a), the real anthropomorphic SARCOS master arm in Figure 2 (b) and the Barrett WAM arm in Figure 2 (c). Here, the normalized mean squared error is defined by  $\text{nMSE} = \text{Mean squared error}/\text{Variance of target}$ . During the prediction on the test set using LGP, we take the most activated local models, i.e., the ones which are next to the query point.

When observing the approximation error on the test set shown in Figure 2(a-c), it can be seen that LGP generalizes well to the test data during prediction. In all cases, LGP outperforms LWPR and OGP while being close in learning accuracy to the offline-methods GPR and  $\nu$ -SVR. The mean prediction for GPR is determined according to Equation (4) where we pre-computed the prediction vector  $\alpha$  from training data. When a query point appears, the kernel vector  $\mathbf{k}_*^T$  is evaluated for this particular point.

### 3.2 Online Learning for Model-based Control

In this section, we apply the inverse dynamics models for a model-based tracking control task [9]. Here, the model is used for predicting the feedforward torques  $\mathbf{u}_{\text{FF}}$  necessary to execute a given the desired trajectory  $[\mathbf{q}_d, \dot{\mathbf{q}}_d, \ddot{\mathbf{q}}_d]$ . First, we compare standard rigid-body dynamics (RBD) models with several models learned offline on

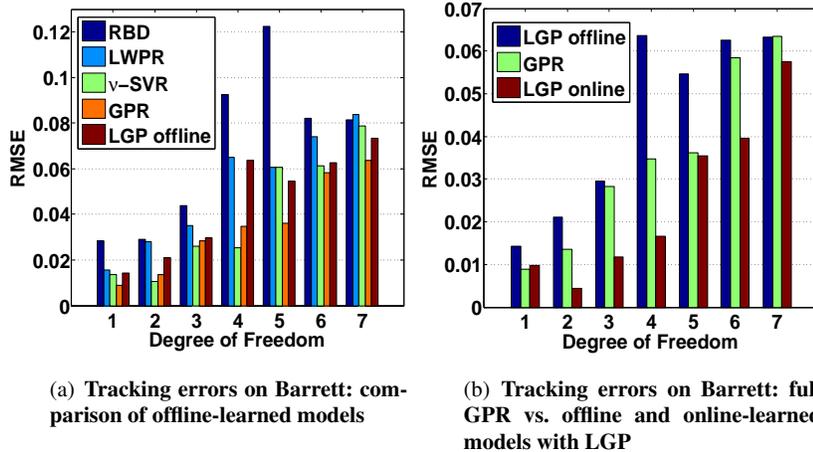


Fig. 3: (a) and (b) show the tracking errors (RMSE) on the Barrett WAM. For offline-learned models, LGP is competitive with full GPR and  $v$ -SVR while being better than LWPR and rigid-body model. When employing online-updates, LGP can largely improve the tracking results outperforming the offline-learned models using full GPR. The reported results are computed for a test trajectory executed on the robot.

training data sets. The RBD-parameters are estimated from the corresponding CAD model. During the control task, the offline-learned models are used for an online torque prediction. For this comparison, we use LWPR,  $v$ -SVR, standard GPR as well as our LGP as compared learning methods. We show that our LGP is competitive when compared with its alternatives. Second, we demonstrate that LGP is capable of online adaptation while being used for predicting the required torques. During online learning, the local GP models are updated in real-time, and the online improvement during a tracking task outperforms the fixed offline model in comparison. Our goal is to achieve compliant tracking in robots without exception handling or force sensing but purely based on using low control gains. Our control gains are three orders of magnitude smaller than the manufacturers in the experiments and we can show that using good, learned inverse dynamics models we can still achieve compliant control. Due to the low feedback gains, the accuracy of the model has a stronger effect on the tracking performance in this setting and, hence, a more precisely learned model will also result in a significantly lower tracking error.

For comparison with offline-learned models, we also compute the feedforward torque using rigid-body (RB) formulation which is a common approach in robot control [9]. The control task is performed in real-time on the Barrett WAM, as shown in Figure 1. As desired trajectory, we generate a test trajectory which is similar to the one used for learning the inverse dynamics models. Figure 3 (a) shows the tracking errors on test trajectory for 7 DoFs using offline-learned models. The error is computed as root mean square error (RMSE) which is a frequently used measure in time series prediction and tracking control. Here, LGP provides a competitive control performance compared to GPR while being superior to LWPR and the state-of-the-art rigid-body model.

Figure 3 (b) shows the tracking error after online learning with LGP in comparison with offline learned models. It can be seen that the errors are significantly reduced for LGP with online updates when compared to both standard GPR and LGP with offline learned models. During online-learning, the local GP models are adapted as new data points arrive. Since the number of training examples in each local model is limited, the update procedure is sufficiently fast for real-time application. For doing so, we employ the joint torques  $\mathbf{u}$  and the resulting robot trajectories  $[\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}]$  as samples which are added to the LGP models online as described in Section 2.2. New data points are added to the local models until these fill up and, once full, new points replace previously existing data points. The insertion of new data point is performed with information gain [18] while for the deletion we randomly take an old point from the corresponding local model. A new data point is inserted to the local model, if its information gain is larger than a given threshold value. In practice, this value is set such that the model update procedure can be maintained in real-time (the larger the information gain threshold, the more updates will be performed).

### 3.3 Performance on a Complex Test Setting

In this section, we create a more complex test case for tracking with inverse dynamics models where the trajectories are acquired by kinesthetic teach-in, i.e., we take the Barrett WAM by the end-effector and guide it along several trajectories



Fig. 4: The figure illustrates the data generation for the learning task.

which are subsequently used both in learning and control experiments. In order to make these trajectories straightforward to understand for humans, we draw all 26 characters of the alphabet in an imaginary plane in task space. An illustration for this data generation process is shown in Figure 4. During the imagined writing, the joint trajectories are sampled from the robot. Afterwards, it will attempt to reproduce that trajectory, and the reproductions can be used to generate training data. Subsequently, we used several characters as training examples (e.g., characters from **D** to **O**) and others, e.g., **A**, as test examples. This setup results in a data set with 10845 samples for training and 1599 for testing.

Similar as in Section 3.1, we learn the inverse dynamics models using joint trajectories as input and joint torques as targets. The robot arm is then controlled

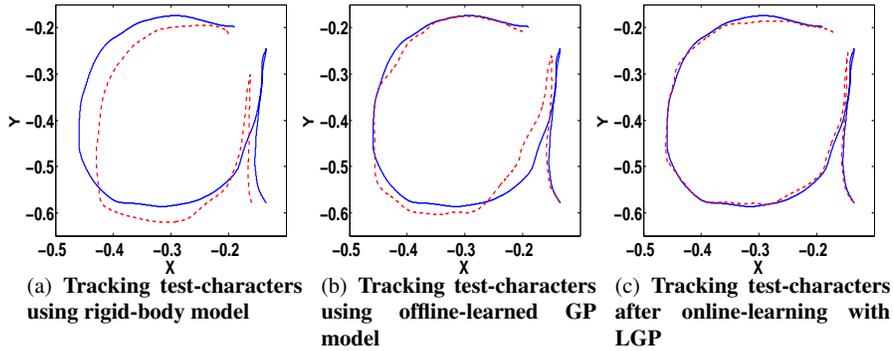


Fig. 5: Compliant tracking performance on Barrett WAM for the test character **A**, where the controlled trajectory lies in joint-space while our visualization is in task space for improved comprehensibility. We compare the corresponding rigid body model, an offline trained GP model and an online learning LGP. The thick, blue line denotes the desired trajectory, while the dashed, red line represents the robot trajectory during the compliant tracking task. The results indicate that online learning with LGP outperforms the offline-learned model using full GPR as well as the rigid-body dynamics.

to perform the joint-space trajectories corresponding to the test characters using the learned models. For LGP, we additionally show that the test characters can be learned online by updating the local models, as described in Section 3.2. The Figure 5 shows the tracking results using online-learning with LGP in comparison to the offline trained model with standard GPR and a traditional rigid body model.

It can be observed that the offline trained models (using standard GPR) can generalize well to unknown characters often having a better tracking performance than the rigid-body model. However, the results can be improved even further if the dynamics model is updated online – as done by LGP. The LGP results are shown in Figure 5 and are achieved after three trials on the test character.

## 4 Conclusion

The local Gaussian process regression LGP combines the strength of fast computation as in local regression with the potentially more accurate kernel regression methods. As a result, we obtain a real-time capable regression method which is relatively easy to tune and works well in robot application. When compared to locally linear methods such as LWPR, the LGP achieves higher learning accuracy while having less computational cost compared to state of the art kernel regression methods such as GPR and  $\nu$ -SVR. The reduced complexity allows the application of the LGP for online model learning which is necessary for realtime adaptation of model errors or changes in the system. Model-based tracking control using online learned LGP models achieves a superior control performance for low gain control in comparison to rigid body models as well as to offline learned models.

Future research will focus on several important extensions such as finding kernels which are most appropriate for clustering and prediction, and how the choice of a similarity can affect the LGP performance. Partitioning in higher dimension space is still a challenging problem, a possible solution is to perform dimensionality reduction during the partitioning step. Furthermore, alternative criteria for insertion and deletion of data points need to be examined more closely. This operation is crucial for online learning as not every new data point is informative for the current prediction task, and on the other hand deleting an old but informative data point may degrade the performance. It is also interesting to investigate further applications of the LGP in humanoid robotics with 35 or more DoFs and learning other types of the control such as operational space control.

## References

1. S. Vijayakumar, A. D'Souza, and S. Schaal, "Incremental online learning in high dimensions," *Neural Computation*, 2005.
2. E. Snelson and Z. Ghahramani, "Local and global sparse gaussian process approximations," *Artificial Intelligence and Statistics*, 2007.
3. J. W. Roberts, J. Zhang, and R. Tedrake, "Motor learning at intermediate reynolds number: Experiments with policy gradient on a heaving plate," *From Motor to Interaction Learning in Robots*, 2009.
4. M. Fumagalli, A. Gijsberts, S. Ivaldi, L. Jamone, G. Metta, L. Natale, F. Nori, and G. Sandini, "Learning how to exploit proximal force sensing: a comparison approach," *From Motor to Interaction Learning in Robots*, 2009.
5. S. Schaal, C. G. Atkeson, and S. Vijayakumar, "Scalable techniques from nonparametric statistics for real-time robot learning," *Applied Intelligence*, pp. 49–60, 2002.
6. C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Massachusetts Institute of Technology: MIT-Press, 2006.
7. B. Schölkopf and A. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond*. Cambridge, MA: MIT-Press, 2002.
8. L. Csato and M. Opper, "Sparse online gaussian processes," *Neural Computation*, 2002.
9. J. J. Craig, *Introduction to Robotics: Mechanics and Control*, 3rd ed. Prentice Hall, 2004.
10. M. W. Spong, S. Hutchinson, and M. Vidyasagar, *Robot Dynamics and Control*. New York: John Wiley and Sons, 2006.
11. J. Nakanishi, J. A. Farrell, and S. Schaal, "Composite adaptive control with locally weighted statistical learning," *Neural Networks*, 2005.
12. S. Schaal, C. G. Atkeson, and S. Vijayakumar, "Real-time robot learning with locally weighted statistical learning," *International Conference on Robotics and Automation*, 2000.
13. D. Nguyen-Tuong, J. Peters, and M. Seeger, "Computed torque control with nonparametric regression models," *Proceedings of the 2008 American Control Conference (ACC 2008)*, 2008.
14. M. Seeger, "Gaussian processes for machine learning," *International Journal of Neural Systems*, 2004.
15. M. Seeger, "Low rank update for the cholesky decomposition," University of California at Berkeley, Tech. Rep., 2007. [Online]. Available: <http://www.kyb.tuebingen.mpg.de/bs/people/seeger/>
16. C.-C. Chang and C.-J. Lin, *LIBSVM: a library for support vector machines*, 2001, <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
17. M. Seeger, *LHOTSE: Toolbox for Adaptive Statistical Model*, 2007, <http://www.kyb.tuebingen.mpg.de/bs/people/seeger/lhotse/>.

18. M.Seeger, "Bayesian gaussian process models: Pac-bayesian generalisation error bounds and sparse approximations," Ph.D. dissertation, University of Edinburgh, 2005.