

Learning Collaborative Impedance-based Robot Behaviors*

Leonel Rozo and Sylvain Calinon and Darwin Caldwell

Department of Advanced Robotics - Istituto Italiano di Tecnologia (Via Morego, 30. Genoa, Italy)

Pablo Jiménez and Carme Torras

Institut de Robòtica i Informàtica Industrial CSIC-UPC (Carrer Llorens i Artigas, 4-6. Barcelona, Spain)

Abstract

Research in learning from demonstration has focused on transferring movements from humans to robots. However, a need is arising for robots that do not just replicate the task on their own, but that also interact with humans in a safe and natural way to accomplish tasks cooperatively. Robots with variable impedance capabilities opens the door to new challenging applications, where the learning algorithms must be extended by encapsulating force and vision information. In this paper we propose a framework to transfer impedance-based behaviors to a torque-controlled robot by kinesi-
thetic teaching. The proposed model encodes the examples as a task-parameterized statistical dynamical system, where the robot impedance is shaped by estimating virtual stiffness matrices from the set of demonstrations. A collaborative assembly task is used as testbed. The results show that the model can be used to modify the robot impedance along task execution to facilitate the collaboration, by triggering stiff and compliant behaviors in an on-line manner to adapt to the user's actions.

1 Introduction

Over the last decade Robotics is addressing new challenging problems to bring robots closer to humans, as computers have now become part of our everyday life. It is envisaged that robots should collaborate with humans to perform a large variety of tasks more easily, faster and in a safer way. To accomplish this goal, robots must be endowed with learning capabilities allowing them to acquire new knowledge from examples given by a human or through their own experience. Learning from demonstration (**LfD**) is a natural way to transfer knowledge to robots from human examples (Billard et al. 2008). Most works have focused on developing learning algorithms to encode trajectories using vision or optical systems to capture the teacher demonstrations, see e.g., (Calinon, Sardellitti, and Caldwell 2010). Nevertheless, the new variable impedance capabilities of recent robotic arms

*This work was partially supported by the STIFF-FLOP European project (FP7-ICT-287728), IntellAct European project (FP7-269959) and the Spanish project PAU+ (DPI2011-27510). L. Rozo was supported by the CSIC under a JAE-PREDOC scholarship. Copyright © 2013, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

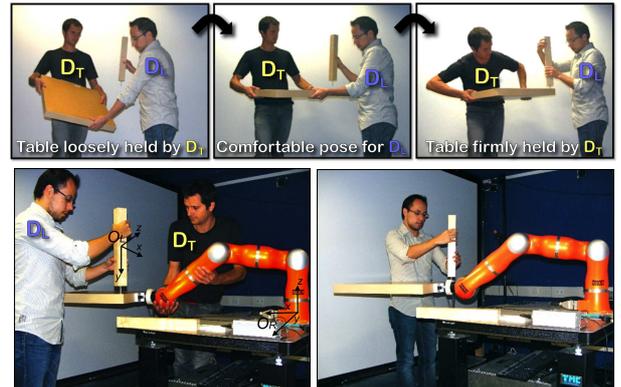


Figure 1: *Top*: Two humans assembling a wooden table. *Bottom*: demonstration (left) and reproduction (right) of the impedance-based behavior.

(Rooks 2006; Albu-Schäffer et al. 2007) demand to reformulate these methods in order to exploit their new control schemes in performing more complex tasks.

In this line, physical human-robot interaction (**pHRI**) has arisen a lot of interest recently, with the two challenging aspects of impedance control and haptic communication. On the one hand, an increasing effort has been devoted to exploiting the advantages provided by the impedance-based control of robots (Hogan 1985). On the other hand, impedance in humans has also been studied with the aim of gaining in-depth knowledge of the roles of the muscles, tendons, brain and spinal cord in modulating impedance when we interact with the environment (Burdet et al. 2001; Gomi and Kawato 1996). Also, efforts have been devoted to mimic human impedance with robots (Ganesh et al. 2010). In this context, it is desirable to transfer impedance-based behaviors from humans to robots (Kalakrishnan et al. 2011; Kronander and Billard 2012). Evrard and Kheddar (2009) tackled the problem of setting the robot's role in a collaborative lifting task by modifying the impedance controller parameters. The approach was then extended to let the robot learn this behavior autonomously from human demonstrations through probabilistic approaches (Calinon et al. 2009; Gribovskaya, Kheddar, and Billard 2011).

The second challenge in pHRI is the haptic communica-

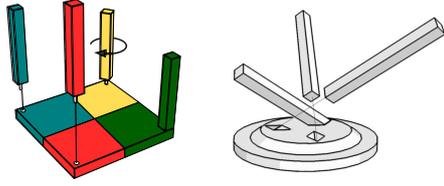


Figure 2: Each assembly task is characterized by different sequences, positions and orientations of components, with haptic and movement patterns that are specific to each item.

tion between the partners. When physical interaction takes place, forces-torques sensed by the partners constitute a very rich and valuable communication channel which is used to recognize and/or predict intentions as well as to determine the role of each participant in the task (Evrard and Kheddar 2009; Groten et al. 2013; Reed and Peshkin 2008). The challenges from the robot side are to distinguish the haptic signals related to the task from those communicating intentions, and to anticipate the human actions using information about the dynamics of the task and the force-based perceptions generated along it (Thobbi, Gu, and Sheng 2011).

In this paper, we are concerned with learning collaborative impedance-based skills¹ using visual and haptic information, where the robot behavior is conditioned by task variables as well as by its haptic perceptions. We thus require a model that can learn impedance behaviors from a set of demonstrations, where the model is modulated by haptic and vision information. The proposed approach is tested in a collaborative assembly of a wooden IKEA table. In the learning phase two humans perform the task, where one is kinesthetically² guiding the robot to demonstrate the robot’s role. The compliance behavior of the person holding the table changes to allow his/her partner to perform his/her corresponding sub-task more easily (see Figure 1). During reproduction, the robot replaces the user holding the table by automatically estimating the levels and shape of stiffness ellipsoid required during the interaction.

The remainder of the paper is organized as follows: Section 2 explains the learning algorithm and the compliance level estimation method. Section 3 describes the collaborative task and the experimental setting. Results are shown and analyzed in Section 4. Finally, conclusions and future work are presented in Section 5.

2 Proposed approach

We extend the task-parameterized movement learning approach recently proposed in Calinon et al. (2012) to a task-parameterized impedance learning problem. This approach relies on a statistical representation of dynamical systems that can be modulated with respect to task variables represented as candidate frames of reference. The model is here

¹Here, impedance-based skill refers to different stiffness levels that a robot needs to accomplish a given task.

²The term refers to the procedure where the user is holding the robot, which is gravity compensated at its links, and moves it along the trajectories that need to be followed to accomplish the task.

extended to force-based impedance behaviors requiring to adapt the stiffness of virtual springs in Cartesian space driving the robot’s behavior.

2.1 Task-parametrized Gaussian mixture model

When robots manipulate objects, their movements may largely depend on the given goals and object poses, which can be defined through reference frames. Namely, the robot motion is conditioned by a set of **task variables** representing the coordinate systems of relevant frames of reference. For generalization purposes, it is desirable to have a model enclosing different movements as a function of the these variables, instead of representing each one with a different model. The proposed approach relies on Gaussian product properties to modulate the centers and covariance matrices of a Gaussian mixture model (**GMM**). The advantages of this approach compared to other task-parameterized models such as the parametric Hidden Markov Model (**PHMM**) (Wilson and Bobick 1999) are discussed in Calinon et al. (2012).

Formally, each demonstration $m \in \{1, \dots, M\}$ contains T_m datapoints forming a dataset of N datapoints $\{\xi_n\}_{n=1}^N$ with $N = \sum_m T_m$. Each $\xi_n \in \mathbb{R}^D$ is associated with task variables $\{\mathbf{A}_{n,j}, \mathbf{b}_{n,j}\}_{j=1}^{N_P}$ representing N_P candidate frames of reference, with transformation matrices $\mathbf{A}_{n,j}$, and offset position vectors $\mathbf{b}_{n,j}$. D is the datapoint dimensionality, and the indexes n and j represent the time step and the candidate frame, respectively.

The **parameters of the model** are $\{\pi_i, \mathbf{Z}_{i,j}^\mu, \mathbf{Z}_{i,j}^\Sigma\}$, representing respectively the mixing coefficients, centers and covariances matrices for each frame j and mixture component i . With this model, for an observation of frames at iteration n , the resulting center $\boldsymbol{\mu}_{n,i}$ and covariance matrix $\boldsymbol{\Sigma}_{n,i}$ of each component i are computed as products of linearly transformed Gaussians

$$\mathcal{N}(\boldsymbol{\mu}_{n,i}, \boldsymbol{\Sigma}_{n,i}) = \prod_{j=1}^{N_P} \mathcal{N}(\mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\mu + \mathbf{b}_{n,j}, \mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\Sigma \mathbf{A}_{n,j}^\top).$$

By using the product property of normal distributions, the above equation is computed as

$$\begin{aligned} \boldsymbol{\mu}_{n,i} &= \boldsymbol{\Sigma}_{n,i} \sum_{j=1}^{N_P} (\mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\Sigma \mathbf{A}_{n,j}^\top)^{-1} (\mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\mu + \mathbf{b}_{n,j}), \\ \boldsymbol{\Sigma}_{n,i} &= \left(\sum_{j=1}^{N_P} (\mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\Sigma \mathbf{A}_{n,j}^\top)^{-1} \right)^{-1}. \end{aligned} \quad (1)$$

The parameters of the model are iteratively estimated with the following EM procedure. In the E-step, (1) are used as temporary Gaussian parameters to compute the likelihood.

E-step:

$$h_{n,i} = \frac{\pi_i \mathcal{N}(\xi_n | \boldsymbol{\mu}_{n,i}, \boldsymbol{\Sigma}_{n,i})}{\sum_k \pi_k \mathcal{N}(\xi_n | \boldsymbol{\mu}_{n,k}, \boldsymbol{\Sigma}_{n,k})}. \quad (2)$$

M-step:

$$\pi_i = \frac{\sum_n h_{n,i}}{N}, \quad \mathbf{Z}_{i,j}^\mu = \frac{\sum_n h_{n,i} \mathbf{A}_{n,j}^{-1} [\xi_n - \mathbf{b}_{n,j}]}{\sum_n h_{n,i}},$$

$$\mathbf{Z}_{i,j}^\Sigma = \frac{\sum_n h_{n,i} \mathbf{A}_{n,j}^{-1} [\boldsymbol{\xi}_n - \tilde{\boldsymbol{\mu}}_{n,i,j}] [\boldsymbol{\xi}_n - \tilde{\boldsymbol{\mu}}_{n,i,j}]^\top \mathbf{A}_{n,j}^{-\top}}{\sum_n h_{n,i}}, \quad (3)$$

with $\tilde{\boldsymbol{\mu}}_{n,i,j} = \mathbf{A}_{n,j} \mathbf{Z}_{i,j}^\mu + \mathbf{b}_{n,j}$.

Note that for force-based tasks, the datapoints, centers and covariances can be decomposed into their position, force and Cartesian torque components

$$\boldsymbol{\xi}_n = \begin{bmatrix} \boldsymbol{\xi}_n^x \\ \boldsymbol{\xi}_n^F \\ \boldsymbol{\xi}_n^T \end{bmatrix}, \quad \boldsymbol{\mu}_{n,i} = \begin{bmatrix} \boldsymbol{\mu}_{n,i}^x \\ \boldsymbol{\mu}_{n,i}^F \\ \boldsymbol{\mu}_{n,i}^T \end{bmatrix}, \quad \boldsymbol{\Sigma}_{n,i} = \begin{bmatrix} \boldsymbol{\Sigma}_{n,i}^x & \boldsymbol{\Sigma}_{n,i}^{xF} & \boldsymbol{\Sigma}_{n,i}^{xT} \\ \boldsymbol{\Sigma}_{n,i}^{F^x} & \boldsymbol{\Sigma}_{n,i}^{FF} & \boldsymbol{\Sigma}_{n,i}^{FT} \\ \boldsymbol{\Sigma}_{n,i}^{T^x} & \boldsymbol{\Sigma}_{n,i}^{TF} & \boldsymbol{\Sigma}_{n,i}^{TT} \end{bmatrix}. \quad (4)$$

The model parameters are initialized with a *k-means* procedure, modified by following a similar task-parametrized structure. Model selection is compatible with the techniques employed in standard GMM (Bayesian information criterion, Dirichlet process, etc.).

One novelty with respect to (Calinon et al. 2012) is that we augment the model with virtual stiffness matrices \mathbf{K}_i^P associated to each component i , which will be estimated as explained in Section 2.2. Thus, the complete set of parameters of the model is $\{\pi_i, \{\mathbf{Z}_{i,j}^\Sigma, \mathbf{Z}_{i,j}^\mu\}_{j=1}^{N_p}, \mathbf{K}_i^P\}_{i=1}^{N_K}$. Such extension allows us to apply the learning model to impedance-based behaviors transfer. Note that the variables of the task are obtained from the position and orientation of a set of candidate frames to learn the task. In our experimental setup, the table legs and robot frames define variables in the collaborative assembly task (described later in Section 3).

Fig. 3 illustrates the approach with a simple example. (a) shows the demonstrations where the robot behaves compliantly when another object (the green triangle) is far from its end-effector, and becomes stiff when the object approaches it with a specific orientation. (b) displays the two phases of the task, where the robot motion is driven by a set of virtual springs connected to the center of the model's Gaussians. The mean and covariance vary according to the task variables (i.e., the object and robot frames), and the influence of each model component (see Eq. (2)) determines how compliantly the robot behave.

2.2 Stiffness estimation

Several approaches have been proposed to estimate from collected data the stiffness and damping parameters to control robots. Erickson, Weber, and Sharf (2003) compared four different methods to estimate the robot impedance based on signal processing, adaptive control and recursive least squares. Flacco and Luca (2011) estimated the non-linear stiffness of robot joints with flexible transmissions by using dynamic residual signals along with least-squares and regressor-based techniques. From a different perspective, a LfD approach was proposed in Calinon, Sardellitti, and Caldwell (2010) to find a stiffness matrix using variability information extracted from training data in the form of a GMM, where the stiffness matrix is estimated from the inverse of the observed covariance in the position space. Similarly, Lee and Ott (2011) used variability encoded in the components of an HMM to define a motion refinement tube that permits a deviation from nominal trajectories for

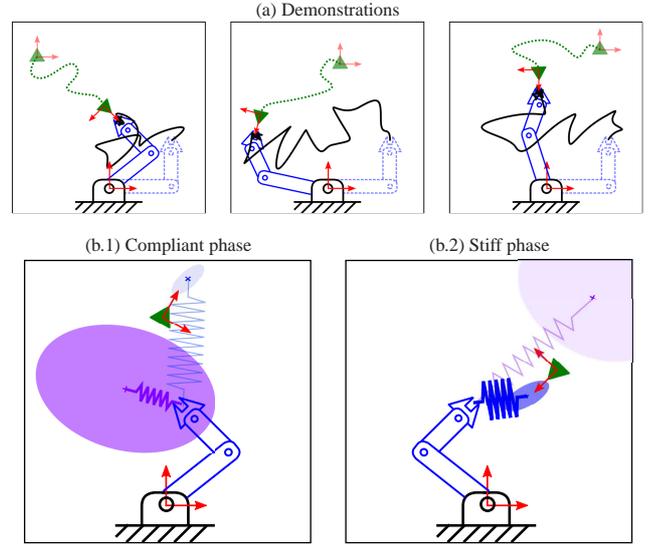


Figure 3: Simplified impedance behavior learning. (a) 3 different *demonstrations* showing compliant and stiff phases. The black line is the robot's trajectory. (b) *Reproduction* of the task, where the robot's behavior is governed by a 2-states model with virtual springs connected to the Gaussians centers. Dark ellipses and thick-line springs represent an activated Gaussian. The candidate frames are displayed in red color.

kinesthetic corrections by controlling the stiffness value at the robot joints level.

Here, we obtain an approximation through an algebraic closed-form solution to find the closest symmetric positive semi-definite stiffness matrix of a weighted least-squares (WLS) estimation. A stiffness matrix \mathbf{K}_i^P is estimated for each component i , by assuming that the robot behavior is driven by a set of virtual springs (similar to Fig. 3)

$$\mathbf{F}_n = \sum_{i=1}^{N_K} h_{n,i} \left[\mathbf{K}_i^P (\boldsymbol{\mu}_{n,i}^x - \mathbf{x}_n) \right], \quad (5)$$

where \mathbf{F}_n , $\boldsymbol{\mu}_{n,i}^x$ and \mathbf{x}_n are respectively the sensed force, the positional part of the Gaussians' centers in the model (see Eq. (4)), and the robot's end-effector position at time step n .

WLS is used to compute a first estimate $\tilde{\mathbf{K}}_i^P = [(\mathbf{X}_i^\top \mathbf{W}_i \mathbf{X}_i)^{-1} \mathbf{X}_i^\top \mathbf{W}_i \mathbf{F}]$ of the stiffness matrices by concatenating all the N datapoints in matrices $\mathbf{X}_i = [(\boldsymbol{\mu}_{1,i}^x - \mathbf{x}_1), \dots, (\boldsymbol{\mu}_{N,i}^x - \mathbf{x}_N)]^\top$ and \mathbf{F} , with a weighting matrix $\mathbf{W}_i = \text{diag}([h_{1,i}, h_{2,i}, \dots, h_{N,i}])$ (see Eq. (2)). Such estimate does not necessarily comply with the symmetric positive semi-definite constraints of a stiffness matrix. Therefore, we resort to the formulation presented in (Higham 1988), to compute \mathbf{K}_i^P as the nearest symmetric positive semi-definite (SPSD) matrix to $\tilde{\mathbf{K}}_i^P$ according to the Frobenius norm, computed as

$$\mathbf{K}_i^P = \frac{\mathbf{B} + \mathbf{H}}{2}, \quad \mathbf{B} = \frac{\tilde{\mathbf{K}}_i^P + (\tilde{\mathbf{K}}_i^P)^\top}{2}, \quad \mathbf{H} = \mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^\top. \quad (6)$$

Table 1: Learning and reproduction phases.

<p>1. Task demonstrations</p> <ul style="list-style-type: none"> - Determine N_P (number of frames or task variables) - $\forall n \in \{1, \dots, N\}$, collect ξ_n and $\{\mathbf{A}_{n,j}, \mathbf{b}_{n,j}\}_{j=1}^{N_P}$ <p>2. Model fitting</p> <ul style="list-style-type: none"> - Determine N_K (number of components of the model) - Use Eq. (3) to learn $\{\pi_i, \{\mathbf{Z}_{i,j}^\mu, \mathbf{Z}_{i,j}^\Sigma\}_{j=1}^{N_P}\}_{i=1}^{N_K}$ <p>3. Stiffness estimation</p> <ul style="list-style-type: none"> - Find \mathbf{K}_i^P for each virtual spring by using Eq. (6) <p>4. Reproduction (for each time step n)</p> <ul style="list-style-type: none"> - Collect ξ_n and $\{\mathbf{A}_{n,j}, \mathbf{b}_{n,j}\}_{j=1}^{N_P}$ - Estimate $\{\pi_i, \mu_{n,i}, \Sigma_{n,i}\}_{i=1}^{N_K}$ through Eq. (1) - Compute activation weights $h_{n,i}$ using Eq. (2) - Apply the force command computed from Eq. (5)

\mathbf{H} is the symmetric polar factor which can be found from the singular value decomposition of \mathbf{B} , namely, $\mathbf{B} = \mathbf{U}\Sigma\mathbf{V}^\top$. Table 1 summarizes the learning and estimation processes.

3 Collaborative assembly task

We consider a human-robot collaborative task where the robot’s role is to hold a wooden table while the user’s role is to screw the four legs to it. Fig. 2 presents an example of assembly instructions that can be found in “do it yourself” furniture catalog. Here, two small tables require specific sequences of force and movement to get assembled. Learning such specificities is required for an efficient collaborative assembly. Instead of manually programming those specificities for each item, we would like the robot to extract those automatically from a set of demonstrations provided by two users collaborating together to assemble the different parts of the table (see Fig. 1). After learning, the task can be reproduced by a single user, with the robot partner interacting appropriately with respect to the preferences of the user and the specificities of the item being assembled. We thus do not need to provide the robot with information about the points of assembly, the different options, orientation of table legs, etc. The robot instead learns these specificities from demonstrations.

3.1 Experimental setup

We use a KUKA lightweight 7-DoF robot (LWR) (Albu-Schäffer et al. 2007), with the *Fast-Research Interface* (Schreiber, Stemmer, and Bischoff 2010), by using a Cartesian impedance controller defined by

$$\boldsymbol{\tau}_d = \mathbf{J}^\top \mathbf{F}_d + V(\kappa_d^V) + f(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}),$$

where \mathbf{J} is the Jacobian of the robot, $\boldsymbol{\tau}_d$ is the desired torque vector, \mathbf{F}_d the desired force computed from the resulting set of virtual springs (Eq. 5), V is a damping function with desired damping values κ_d^V and $f(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})$ the dynamic model of the robot.³

³Note that we only control the Cartesian position of the robot while the rotational DOF are set to be fixed during the reproduction phase.

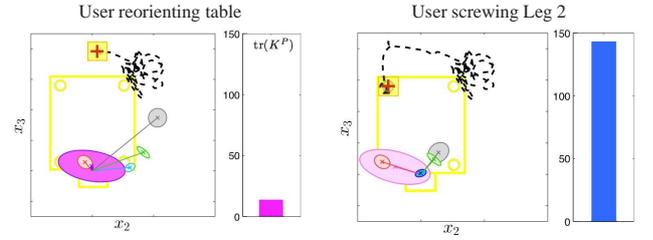


Figure 4: Reproduction results at different phases of the interaction. For each graph, the projection of the model’s Gaussians in the tool’s frame (as ellipses) is shown on the *left*, while the *right* part shows the trace of the resulting stiffness matrix. The black dotted line represents the leg’s trajectory, the table is shown in yellow with its 4 threads and the brown cross corresponds to the current position of the leg.

The position and orientation of the table legs are tracked with a marker-based NaturalPoint OptiTrack motion capture system, composed of 12 cameras working at a rate of 30 fps. A transformation matrix is computed to represent the leg configuration in the fixed robot frame \mathcal{O}_R , from which $\mathbf{b}_n^{\text{leg}}$ and $\mathbf{A}_n^{\text{leg}}$ define the Cartesian position and the orientation of the leg as a rotation matrix, respectively. During both demonstration and reproduction phases, $\{\mathbf{A}_n^{\text{leg}}, \mathbf{b}_n^{\text{leg}}\}$ are recorded at each time step n to determine the task variables. Lastly, the other candidate frame $\{\mathbf{A}_n^{\text{R}}, \mathbf{b}_n^{\text{R}}\}$ define the robot’s fixed frame of reference.⁴

The robot is equipped with a six-axis force-torque sensor (ATI Mini45) attached between its wrist and the wooden table, measuring the interaction forces generated while moving the table and screwing the legs. In order to extract the forces generated by the interaction between the partners through the table, it is necessary to remove other signals from the sensor readings. During the demonstration and reproduction stages, we consider that the sensor readings are composed of the noise, the load mass effects and the external applied forces, the latter corresponding to the interaction between the human and the table. Moreover, we assume that the interaction does not generate high linear/angular velocities or accelerations, so that the dynamical components can be neglected, similarly to Rozo, Jiménez, and Torras (2010).

3.2 Task description

Two candidate frames ($N_P = 2$) describe the task variables in the experiment: \mathcal{O}_R and the leg frame \mathcal{O}_L . We assume that one leg is used and tracked at a time. Let us define D_L and D_T as the human user and the robot, respectively (see Fig. 1). The collaborative scenario consists of screwing the legs at the four corresponding positions on the table. D_T is first compliant to allow D_L to move the table freely (compliant phase) until comfortable position and orientation are found for the work to be performed next. When D_L grasps a leg and starts inserting it into the screw thread in the table, D_T adopts a stiff posture, holding the table to facilitate D_L ’s part

⁴A 3D coordinate frame is replicated for the variables \mathbf{x} , \mathbf{F} and \mathbf{T} , the offset is only set to \mathbf{x} .

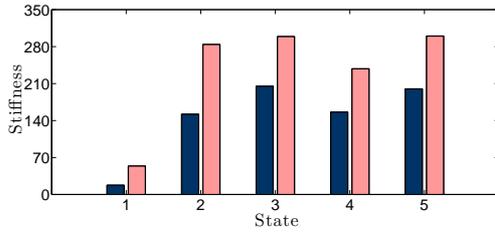


Figure 5: Stiffness level estimation for the 5-states model. The bars display the trace of the mean estimated stiffness matrices for each Gaussian. *Dark blue* bars represent to $\mathbf{K}_{\text{Frb}}^P$ (our approach) while *light red* bars show the $\mathbf{K}_{\text{Inv}}^P$ values. The first state corresponds to the compliant behavior.

of the task (stiff phase).

Note that the combination of vision and haptic information is fundamental for this task. If only vision is used, the robot cannot distinguish the phase during which the user aligns the screw with the thread. Here, D_T should instead regulate its stiffness in accordance with the sensed force pattern. If D_T 's behavior was based only on forces, the collaboration could fail because D_T could not distinguish which forces correspond to interaction with D_L and which are produced by the screwing actions. This can be problematic because these patterns might be similar in some situations. Both perception channels are thus needed to learn how the impedance behavior should be shaped.

4 Results

A model of five components ($N_K = 5$) was trained with sixteen demonstrations (i.e., each leg is assembled four times to its corresponding thread with specific vision and force patterns). The resulting model automatically discovered four stiff components corresponding to the four screwing phases, with the remaining component representing the compliant behavior. Each “stiff component” is characterized by the force-torque pattern and the relative position of the leg with respect to the robot tool frame, which are different for each leg. The “compliant component” encodes the remaining points in the data space, i.e., the interaction forces-torques as well as the varying robot end-effector and leg positions. Fig. 4 shows that the Gaussian corresponding to the compliant phase is already spatially distinguishable from the Gaussians encoding the stiff behaviors during the screwing processes (four in this case). Note that the Gaussians in the model representing the stiff phases show an elongated shape changing its orientation during the task. Such type of time-varying information encapsulated in the covariance matrices cannot be encoded by using the classic PHMM (no covariance parameterization).

Once the model is learned, the stiffness estimation is carried out as described in Section 2.2. In this experiment a stiffness matrix is locally associated with each component in the model, describing a virtual spring connected to the center of the Gaussian. During reproduction, a force command is estimated as a combination of the virtual springs (see Eq. 5).

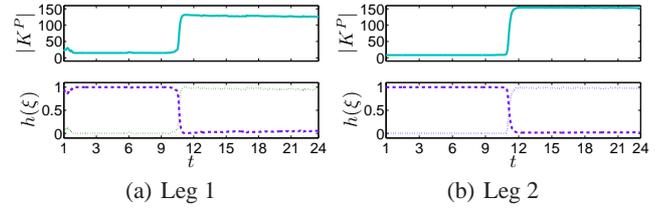


Figure 6: Resulting stiffness matrix trace described by $\mathbf{K}^P = \sum_i^{N_K} h_{n,i} \mathbf{K}_i^P$ (top), and components influence on the stiffness (bottom). Each plot displays how the weight belonging to each component changes over time – between 0 and 1 – showing its influence on the WLS-based estimation. The *activated* component at the stiff phase relies on the leg to be assembled (dashed lines), while the compliant component is *activated* regardless of the leg (dotted line). The deactivated components – whose weights are zero – are not shown in the graphs to emphasize the components participating in the current demonstration.

The proposed approach was compared to the stiffness estimation procedure based on the inverse of the observed covariance (Calinon, Sardellitti, and Caldwell 2010) (see Fig. 5), by computing the inverse of the sub-matrix $\Sigma_{n,i}^x$ for each Gaussian i at each time step n . A weighted average stiffness $\mathbf{K}_{\text{Inv}}^P$ is then calculated. This is compared to $\mathbf{K}_{\text{Frb}}^P$ obtained as described in Section 2.2. With our training set, both approaches estimate the different stiffness levels appropriately. However, the estimate of Calinon, Sardellitti, and Caldwell (2010) has the disadvantage that it takes only into account the positional information from the data, whose variability can sometimes be too weak if only a few number of demonstrations are considered. In the experiment, the users covered various portions of the workspace. In a more realistic teaching procedure, the users might not be aware of this scaffolding teaching procedure, and a smaller number of datapoints might be acquired. In such situation, variability information may not always be sufficient to estimate stiffness information. In contrast, the approach that we adopt in this paper does take into consideration the haptic inputs in the estimation process. Fig. 5 displays the trace of the estimated stiffness matrices for each Gaussian, comparing the results obtained by both approaches. The ratio between the stiff and compliant values (computed from the matrix traces) is higher using the proposed approach (10.12 as average) than those obtained from the approach based on position variability (5.19 as average), which allows a better clamping of the robot stiffness considering the obtained maximum and minimum values. This indicates that the difference between the compliant and stiff levels is more pronounced when the estimation process is based on the haptic data.

Previous experiments with force and position data have shown that it was often more challenging to exploit force recordings than positions in practice, mostly due to the lower accuracy of the sensors and disturbing forces (Roza, Jiménez, and Torras 2013). However, a small variability has been observed for force signals in the second phase of the

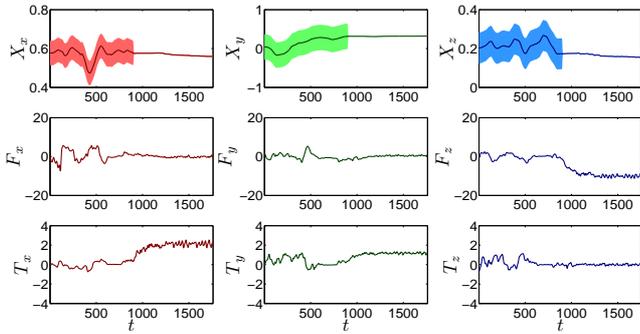
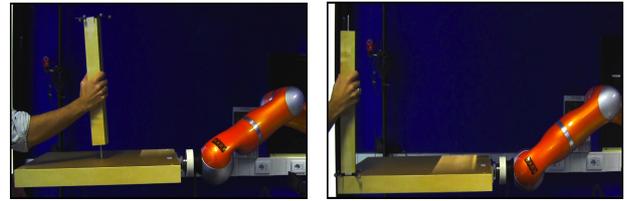


Figure 7: Estimated stiffness along the Cartesian axes of the robot (represented as an envelope surrounding the reproduced trajectory) and the corresponding force-torque profiles. The envelopes (light color area) in the first row show the robot compliance (the wider the zone is, the more compliant the robot is). When the envelope is wide, the robot can be moved easily by D_L , while narrow envelope represents the table holding phase with high stiffness.

task where the leg is screwed to the table. We attribute this low variability to two main factors: 1) the force signals pre-processing and 2) the fact that our experiment protocol facilitated the collaborative task to the users, which resulted in high consistency in the demonstrations. The two users demonstrated the collaborative assembly without talking, but by using their hands and their eyes. During reproduction, the robot however only had access to a limited subset of haptic and vision cues. For example, the users could observe the changes of posture and even the gaze of each other (Strabala et al. 2012; Mutlu, Terrell, and Huang 2013), which likely provided additional cues to jointly agree about the change of stiffness behaviors.

We tested the reproduction and generalization capabilities of the system by carrying out the assembly process for all the legs. In Fig. 6, we can observe how the compliant component (purple dotted line) is influential during the first half of the reproduction, dominating the other components. After this, the robot becomes stiff, with specific patterns depending on which leg is being screwed. This means that not all the components influence the robot impedance at the stiff phase, but mostly the Gaussian encoding the stiff behavior for the corresponding leg (as observed from the different colors representing the different stiff components), while the remaining activation weights stay close to zero.

The proposed approach does not only learn when to change the compliance in an on/off fashion, but also the manner to switch between the two behaviors. The sharp stiff/compliant difference is a characteristic of the collaborative task presented here (mostly binary, but with smooth transitions between the two compliance levels), which is correctly learned by the proposed approach. Fig. 7 shows the resulting stiffness matrix for the demonstration corresponding to leg 1, where the Cartesian robot position is shown along with the corresponding stiffness for each axis. We can see how these values vary along the different Cartesian axes,



(a) Leg away from the threads (b) Leg turned upside-down

Figure 8: Situations not shown in the demonstration phase. Both correctly result in the robot behaving compliantly.

which is useful when the robot behavior demands to be stiff in a specific direction and compliant in the others.

In order to show the relevance of combining visual and haptic data for generalizing the learned task, two situations that did not appear in the demonstrations were presented to the robot (Fig. 8). First, D_L tried to screw the leg at the center of the table, which means that the leg was placed at an incorrect position. In the second situation, D_L positioned the leg in one of the table threads but the leg was tilted, making the screwing process unfeasible. In both cases, the robot behaved compliantly as expected, because neither corresponded to a correct screwing phase. A video of the experiment and the task-parameterized GMM sourcecode are available at <http://programming-by-demonstration.org/AAAI2013/>

5 Conclusions and Future Work

We presented a learning framework to encode and reproduce impedance behaviors using a task-parameterized statistical dynamical system. Our method allows to encode behaviors that rely on task variables, yielding only one model to encode the whole task. In contrast to previous approaches where robot impedance is learned from position variability, our framework extracts the impedance behavior from the manner in which the teacher behaves during the task, relying on recorded force patterns and visual information. We use forces not only to encode the skill, but also to estimate the stiffness of virtual springs governing the collaborative behavior, thus emphasizing that interaction forces-torques vary during different phases of the task.

The proposed approach is used to learn a reactive behavior, where the model automatically provides soft clusters of the different impedance behaviors that the robot might adopt. They are estimated as impedance controllers whose parameters are learned from demonstrations and associated with each model component. We plan in future work to provide the robot with a more active role. Initially, the roles would be considered as time independent. Then, more complex task preferences between the partners would be acquired, where the robot could adopt a proactive role in the collaboration. We devise to exploit the weighting mechanism in the model to influence the switching from reactive to proactive behaviors, by anticipating the next part of the task depending on user preferences. For example, if the user holding the leg is too far from the robot, it could take the initiatives to anticipate the movement and move to the leg.

References

- Albu-Schäffer, A.; Haddadin, S.; Ott, C.; Stemmer, A.; Wimbeck, T.; and Hirzinger, G. 2007. The DLR lightweight robot - design and control concepts for robots in human environments. *Industrial Robot: An International Journal* 34(5):376–385.
- Billard, A.; Calinon, S.; Dillmann, R.; and Schaal, S. 2008. Robot programming by demonstration. In Siciliano, B., and Khatib, O., eds., *Springer Handbook of Robotics*. Secaucus, NJ, USA: Springer. 1371–1394.
- Burdet, E.; Osu, R.; Franklin, D.; Milner, T.; and Kawato, M. 2001. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414(6862):446–449.
- Calinon, S.; Evrard, P.; Gribovskaya, E.; Billard, A.; and Kheddar, A. 2009. Learning collaborative manipulation tasks by demonstration using a haptic interface. In *Intl. Conf. on Advanced Robotics*, 1–6.
- Calinon, S.; Li, Z.; Alizadeh, T.; Tsagarakis, N.; and Caldwell, D. 2012. Statistical dynamical systems for skills acquisition in humanoids. In *Intl. Conf. on Humanoid Robots*, 323–329.
- Calinon, S.; Sardellitti, I.; and Caldwell, D. 2010. Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies. In *Intl. Conf. on Intelligent Robots and Systems*, 249–254.
- Erickson, D.; Weber, M.; and Sharf, I. 2003. Contact stiffness and damping estimation for robotic systems. *Intl. Journal of Robotics Research* 22(1):41–57.
- Evrard, P., and Kheddar, A. 2009. Homotopy switching model for dyad haptic interaction in physical collaborative tasks. In *EuroHaptics*, 45–50.
- Flacco, F., and Luca, A. D. 2011. Residual-based stiffness estimation in robots with flexible transmissions. In *Intl. Conf. on Robotics and Automation*, 5541–5547.
- Ganesh, G.; Albu-Schaffer, A.; Haruno, M.; Kawato, M.; and Burdet, E. 2010. Biomimetic motor behavior for simultaneous adaptation of force, impedance and trajectory in interaction tasks. In *Intl. Conf. on Robotics and Automation*, 2705–2711.
- Gomi, H., and Kawato, M. 1996. Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272(5258):117–120.
- Gribovskaya, E.; Kheddar, A.; and Billard, A. 2011. Motion learning and adaptive impedance for robot control during physical interaction with humans. In *Intl. Conf. on Robotics and Automation*, 4326–4332.
- Groten, R.; Feth, D.; Klatzky, R.; and Peer, A. 2013. The role of haptic feedback for the integration of intentions in shared task execution. *Trans. on Haptics* 6(1):94–105.
- Higham, N. 1988. Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra and its Applications* 103:103–118.
- Hogan, N. 1985. Impedance control: An approach to manipulation: Part i, ii, iii. *Journal of Dynamic Systems, Measurement, and Control* 107(1):1–24.
- Kalakrishnan, M.; Righetti, L.; Pastor, P.; and Schaal, S. 2011. Learning force control policies for compliant manipulation. In *Intl. Conf. on Intelligent Robots and Systems*, 4639–4644.
- Kronander, K., and Billard, A. 2012. Online learning of varying stiffness through physical human-robot interaction. In *Intl. Conf. on Robotics and Automation*, 1842–1849.
- Lee, D., and Ott, C. 2011. Incremental kinesthetic teaching of motion primitives using the motion refinement tube. *Autonomous Robots* 31:115–131.
- Mutlu, B.; Terrell, A.; and Huang, C. 2013. Coordination mechanisms in human-robot collaboration. In *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI) - Workshop on Collaborative Manipulation*, 1–6.
- Reed, K., and Peshkin, M. 2008. Physical collaboration of human-human and human-robot teams. *Trans. on Haptics* 1(2):108–120.
- Rooks, B. 2006. The harmonious robot. *Industrial Robot: An International Journal* 33(2):125–130.
- Rozo, L.; Jiménez, P.; and Torras, C. 2010. Sharpening haptic inputs for teaching a manipulation skill to a robot. In *Intl. Conf. on Applied Bionics and Biomechanics*, 370–377.
- Rozo, L.; Jiménez, P.; and Torras, C. 2013. A robot learning from demonstration framework to perform force-based manipulation tasks. *Intelligent Service Robotics* 6(1):33–51.
- Schreiber, G.; Stemmer, A.; and Bischoff, R. 2010. The fast research interface for the KUKA lightweight robot. In *ICRA-Workshop on Innovative Robot Control Architectures for Demanding (Research) Applications*, 15–21.
- Strabala, K.; Lee, M.; Dragan, A.; Forlizzi, J.; and Srinivasa, S. 2012. Learning the communication of intent prior to physical collaboration. In *Intl. Symposium on Robot and Human Interactive Communication*, 968–973.
- Thobbi, A.; Gu, Y.; and Sheng, W. 2011. Using human motion estimation for human-robot cooperative manipulation. In *Intl. Conf. on Intelligent Robots and Systems*, 2873–2878.
- Wilson, A., and Bobick, A. 1999. Parametric hidden Markov models for gesture recognition. *Pattern Analysis and Machine Intelligence* 21(9):884–900.