The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies

Sylvain Calinon, Irene Sardellitti and Darwin G. Caldwell

Abstract— **We propose a control strategy for a robotic manipulator operating in an unstructured environment while interacting with a human operator. The proposed system takes into account the important characteristics of the task and the redundancy of the robot to determine a controller that is safe for the user. The constraints of the task are first extracted using several examples of the skill demonstrated to the robot through kinesthetic teaching. An active control strategy based on task-space control with variable stiffness is proposed, and combined with a safety strategy for tasks requiring humans to move in the vicinity of robots. A risk indicator for human-robot collision is defined, which modulates a repulsive force distorting the spatial and temporal characteristics of the movement according to the task constraints. We illustrate the approach with two human-robot interaction experiments, where the user teaches the robot first how to move a tray, and then shows it how to iron a napkin.**

I. INTRODUCTION

Robotic applications are bringing robots into unstructured environments populated by humans. These robots are expected to achieve a large range of skills that cannot be pre-programmed. Controllers capable of handling several types of external perturbations are required to let the robot generalize the skill to new situations. Flexible learning mechanisms are also required to let non-expert users teach new skills in a user-friendly manner. Providing robots with learning by imitation capabilities is an approach to reduce the search space of the possible actions that the robot can take, while still allowing the robot to further refine its model of the demonstration through reinforcement learning [1], [2].

Research in safety is carried out in two main directions, *passive* and *active* safety. The former is mostly implemented during the design of the robot to reduce the collision forces in the case of an unexpected impact. Variable stiffness actuators have been proposed as a safe approach for driving robots that interact with humans [3], [4]. These actuators allow a robot to both absorb the energy of an impact through a compliant mechanism, and to achieve precise joint positioning through variation of the stiffness gains. The active safety approach, instead, attempts to prevent collision at the controller level [5]. Moreover, strategies for detecting the collision together with an appropriate reaction behavior have also been proposed [6].

The authors are with the Advanced Robotics Department, Italian Institute of Technology (IIT), 16163 Genova, Italy. {sylvain.calinon,irene.sardellitti, darwin.caldwell}@iit.it.

This work was supported in part by the AMARSi European project under contract FP7-ICT-248311.

The concept of risk assessment, which is mostly used for industrial applications, is contained in the standards which prescribe that safety is guaranteed by defining an area where the robot stops when a human intrusion is detected [7], [8]. These well-established standards remain valid for a broad range of stiff robots. They, however, rarely meet the requirements of close human-robot interaction in applications such as collaborative tasks, or in tasks where the robots move in the vicinity of users. With the development of torque-controlled robots such as the *Barrett WAM* arm used in this paper, novel flexible and adaptive control strategies can be explored for human-robot interaction. Consequently, novel risk management policies and associated control scheme also need to be examined.

In this paper, we refer to the kinematic redundancy of the robot when the robot possesses an infinite number of generalized inverse control strategies, see e.g. [9], [10]. We refer to task redundancy when the task can be achieved through an infinite number of solutions, see e.g. [11]. We take the perspective that both the robot and task redundancies can be exploited to regulate the dynamics of the movement and the stiffness of the robot during reproduction. After having observed several demonstrations of a similar task, the robot creates a compact model of the skill, by taking into account the variations and correlations observed along the movement. If a part of the movement was consistent across the different trials, this part of the task should probably be reproduced in this specific manner. On the other hand, if a large variability was observed during the different demonstrations, reproducing a specific reference trajectory is not required to fulfil the task requirements.

During reproduction, the robot is using this information to set an adaptive stiffness matrix compatible with the task requirements. High compliance will allow here the simultaneous consideration of other constraints. We consider two situations where the interaction can benefit from the variability and correlations of the task: (i) to let the user physically move the robot while reproducing the task; (ii) to let the robot modify the generalized trajectory to adopt gestures that are safer for a user who is close to the robot.

Instead of setting in advance a pre-determined path to follow, the robot thus makes use of the task redundancies and its kinematics redundancies to fulfill constraints related to safety and obstacle avoidance. Through this approach, the robot can still follow a specific path if the task strictly requires it to do so (i.e., the skill showed a strong invariance across the multiple demonstrations). Otherwise, the robot will loosely generalize the skill by adapting its movement

to the user's proximity and attention. The proposed strategy offers the possibility to cope with unpredicted events in a safe way, as the robot remains compliant in the parts of the task that do not require to track precisely a reference trajectory.

II. CONTROL STRATEGY

To control the robot, we exploit the torque-feedback properties of the manipulator, where the robot remains actively compliant for the degrees of freedom that are not relevant for the task. We control the n degrees of freedom (DOFs) robot through inverse dynamics solved with recursive Newton Euler algorithm [12]. The joint forces f_i at each joint $i \in \{1, \dots, n\}$ are therefore calculated as

$$f_i = f_i^a - f_i^e + \sum_{j \in c(i)} f_j,$$

where f_i^a is the net force acting on link i, f_j with $j \in c(i)$ are the forces transmitted by the child $c(i)$ of link i, and f_i^e are the external forces defined as

$$f_i^e = F_T + F_O + F_G.$$

In the above equation, $F_O = [f_O, 0]^\top \in \mathbb{R}^6$ is a repulsive force applied to the point x_O in the kinematic chain (when x_O is on the link i) and then projected at the center of gravity of the link; $F_T = [f_T, M_T]^\top \in \mathbb{R}^6$ is the vector of force and momentum requested to accomplish the task (only applied at the end-effector, i.e. when $i = n$), and $F_G = [f_G, 0]^\top \in \mathbb{R}^6$ is the gravity compensation force. Tracking of a desired path in Cartesian space is insured by the force command $f_T = m_T \hat{\ddot{x}}$, where m_T is a virtual mass and $\hat{\ddot{x}}$ is a desired acceleration command (described in next subsection).

A. Learning the task constraints

M examples of a skill are demonstrated to the robot in slightly different situations. Each demonstration $m \in \{1, \dots, M\}$ consists of a set of T_m positions x, velocities $\dot{x}$ and accelerations $\ddot{x}$ of the end-effector in Cartesian space, where each position x has $D = 3$ dimensions. A dataset is formed by concatenating the $N = \sum_{m=1}^{M} T_m$ datapoints $\{\{x_j, \dot{x}_j, \ddot{x}_j\}_{j=1}^{T_m}\}_{m=1}^{M}$. By considering flexibility and compactness issues, we propose to use a controller based on a mixture of K proportional-derivative systems

$$\hat{\ddot{x}} = \sum_{i=1}^{K} h_i(t) \Big[K_i^{\mathcal{P}} (\mu_i^{\mathcal{X}} - x) - \kappa^{\mathcal{V}} \dot{x} \Big]. \tag{1}$$

The above formulation shares similarities with the *Dynamic Movement Primitives* (DMP) framework originally proposed by Ijspeert *et al* [13], and further extended in [14], [15] (see [16] for a discussion on the similarities of the proposed controller with DMP). The principal difference is that we consider a full matrix $K_i^{\mathcal{P}}$ associated with each of the K primitives (or states) instead of a fixed $\kappa^{\mathcal{P}}$ gain. This allows us to take into consideration variability and correlation information along the movement for learning and reproduction.[1]

[1]Note that this process can generically be applied to other movement representations based on a superposition of affine linear systems, see e.g. [17], [18].

The superposition of basis force fields is determined in (1) by an implicit time dependency, but other approaches using spatial and/or sequential information could also be used [17], [18]. Similarly to DMP, a decay term defined by a canonical system $\dot{s} = -\alpha s$ is used to create an implicit time dependency $t = -\frac{ln(s)}{\alpha}$, where s is initialized with $s = 1$ and converges to zero. We define a set of Gaussians $\mathcal{N}(\mu_i^{\mathcal{T}}, \Sigma_i^{\mathcal{T}})$ in time space τ, with centers $\mu_i^{\mathcal{T}}$ equally distributed in time, and variance parameters $\Sigma_i^{\mathcal{T}}$ set to a constant value inversely proportional to the number of states. α is initially fixed depending on the duration of the demonstrations. The weights $h_i(t)$ are defined by

$$h_i(t) = \frac{\mathcal{N}(t;\ \mu_i^{\mathcal{T}}, \Sigma_i^{\mathcal{T}})}{\sum_{k=1}^{K} \mathcal{N}(t;\ \mu_k^{\mathcal{T}}, \Sigma_k^{\mathcal{T}})}. \tag{2}$$

By determining the weights through the decay term s, the system will sequentially converge to the set of attractors in Cartesian space defined by $\mu_i^{\mathcal{X}}$. The centers $\mu_i^{\mathcal{X}}$ in task space and stiffness matrices $K_i^{\mathcal{P}}$ are learned from the observed data, either incrementally or in a batch mode (through least-squares regression). For example, parts of the movement where the variations between the demonstrations are high indicate that the reference trajectory does not need to be tracked precisely. By using this information, the controller can focus on the other constraints of the task such as moving away from the user. On the other hand, parts of the movement exhibiting strong invariance among the demonstrations should be tracked precisely, i.e., the stiffness used to track the position errors needs in this case to be high.

In a batch mode, by concatenating the training examples in a matrix $Y = [\ddot{x} \frac{1}{\kappa^{\mathcal{P}}} + \dot{x} \frac{\kappa^{\mathcal{V}}}{\kappa^{\mathcal{P}}} + x] \in \mathbb{R}^{N \times D}$, and by concatenating the corresponding weights computed with (2) in a matrix $H \in \mathbb{R}^{N \times K}$, we can write the linear equation $Y = H\mu^{\mathcal{X}}$, with $\mu^{\mathcal{X}} \in \mathbb{R}^{K \times D}$ representing the concatenated attractor centers $\mu_i^{\mathcal{X}}$. The least-squares solution to estimate the attractor centers is then given by $\mu^{\mathcal{X}} = H^\dagger Y$, where $H^\dagger$ is the pseudoinverse of H.

To take into account variability and correlation along the movement and among the different demonstrations, we compute for each state $i \in \{1, \dots, K\}$ the residual errors of the least-squares estimation, in the form of covariance matrices

$$\Sigma_i^{\mathcal{X}} = \frac{1}{N} \sum_{j=1}^{N} (Y'_{j,i} - \bar{Y}'_i)(Y'_{j,i} - \bar{Y}'_i)^\top \quad \forall i \in \{1, \dots, K\},$$
$$\text{where } Y'_{j,i} = H_{j,i}(Y_j - \mu_i^{\mathcal{X}}). \tag{3}$$

$\mathcal{N}(\mu_i^{\mathcal{X}}, \Sigma_i^{\mathcal{X}})$ thus describes a Gaussian in Cartesian space $\mathcal{X}$. The set of K Gaussians defines the sequence of virtual attractor points in Cartesian space that the system will try to reach, where each attractor encapsulates variability and correlation information. The residuals terms of the regression process are then used to estimate the stiffness matrices $K_i^{\mathcal{P}}$

in Eq. (1) through eigencomponents decomposition

$$K_i^{\mathcal{P}} = V_i D_i V_i^{-1},$$
$$\text{with } D_i = \kappa_{\min}^{\mathcal{P}} + (\kappa_{\max}^{\mathcal{P}} - \kappa_{\min}^{\mathcal{P}}) \frac{\lambda_i - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}}. \quad (4)$$

In the above equation, λ_i and V_i are the concatenated eigenvalues and eigenvectors of the inverse covariance matrix $(\Sigma_i^{\mathcal{X}})^{-1}$. The basic idea is to determine a stiffness matrix proportional to the inverse of the observed covariance. For example, if high variability is observed, stiffness will become low as the tracking does not need to be precise. If D_i in (4) is set to λ_i, the eigencomponents decomposition gives $K_i^{\mathcal{P}} = (\Sigma_i^{\mathcal{X}})^{-1}$. We rescale D_i to obtain stiffnesses in the desired range $[\kappa_{\min}^{\mathcal{P}}, \kappa_{\max}^{\mathcal{P}}]$ (determined by the user and hardware's limitation) based on the initial range of eigenvalues $[\lambda_{\min}, \lambda_{\max}]$ (determined by the variability within the motion and among several demonstrations). Source-code for the algorithms in this section are available at `http://programming-by-demonstration.org/SylvainCalinon/`.

B. Risk indicator

When the robot and the user are closely interacting in an unstructured environment, the intentions of the human operator are largely unpredictable. The experiment, that we propose as a case study, considers the position and orientation of the human's head with respect to the moving robot arm. This scenario has been presented in literature as one of the most important danger events that must be addressed, see e.g. [3], [19]. In [20], we proposed an attention mechanism based on the area covered by a vision cone intersecting with a table on which a set of objects were placed. The proposed approach was restricted to predetermined 2D planes, where the level of attention was used to modify one existing task, instead of considering several constraints simultaneously. We propose in this paper a more generic mechanism in 3D Cartesian space based on the user's head pose.

A motion capture system is used to track the pose of the head, described by its position x_U and orientation matrix R_U. The closest position x_O on the robot's kinematics chain to the position x_U of the user is first computed. The vector between the user and the robot's closest point is defined by $v = x_U - x_O$, with associated norm $d = |v|$. The user's gaze direction is approximated by his/her head direction vector $u = R_U e_1$, with $e_1 = [1\ 0\ 0]^\top$ being a unit vector. The angle between the user's gaze vector and the vector directed towards the robot's closest point is determined by $\omega = \arccos(u^\top v)$. Distance d and angle ω are both considered as parameters for the determination of a risk factor related to the user's proximity and attention awareness.

When humans move their head, the indicator thus scales the risk of human-robot collision on the basis of a function $r = f(d, \omega, \sigma_d, \sigma_\omega) \in \mathbb{R}_{[0,1]}$, that depends on the angle ω in-between the user's gaze direction and the robot's closest point, and distance d of the human's head from the robot. We define the risk indicator as

$$r = \frac{\mathcal{N}(d;\ 0, \sigma_d)\, \mathcal{N}(\omega;\ \pi, \sigma_\omega)}{\mathcal{N}(0;\ 0, \sigma_d)\, \mathcal{N}(\pi;\ \pi, \sigma_\omega)},$$

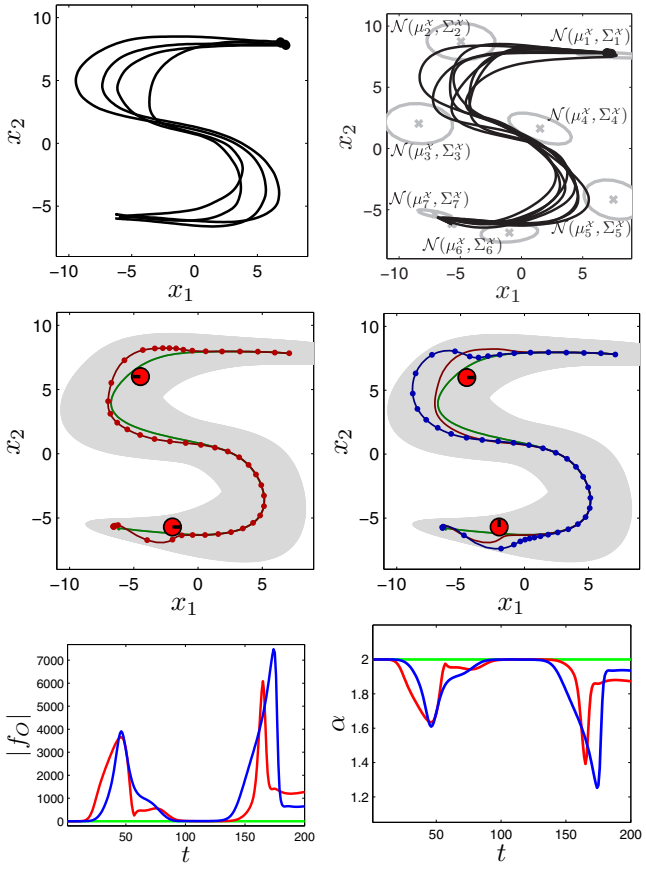

Fig. 1. Illustration of the learning and retrieval processes. *Top-left:* Four examples of the task provided as demonstrations. *Top-right:* Learned model and multiple reproduction attempts by retrieving for each reproduction a set of attractors from the Gaussian distributions $\mathcal{N}(\mu_i^{\mathcal{X}}, \Sigma_i^{\mathcal{X}})$. *Center:* Consideration of the user in the reproduction of the task. The user's head is represented by a red point, with the black line showing its orientation. The trajectory in green lines shows a generalized reproduction attempt without consideration of the user. The trajectories in red and blue lines show the retrieved path when the user is close to the robot at the beginning and at the end of the movement, for two different head directions. Those in red lines correspond to the situation where the user is on the path but remains attentive to the robot's movement (*left*). Those in blue lines correspond to the situation where the user looks away from the robot (*right*). The dots show positions at constant time intervals. *Bottom:* Norm of the repulsive force f_O and value of the decay parameter α along the movement for the different reproduction scenarios.

where σ_d and σ_ω are variance parameters determined by the experimenter. Fig. 6 *left* shows the risk function used for the experiments. The highest level is considered when the user is close to the robot and facing away. The lowest and safest level is when the human operator is out of the robot's workspace and when the user is looking in the direction of the robot's movement.

By taking this into consideration, we define a repulsive force $f_O = r f_{\max} \frac{v}{|v|}$, where $f_{\max}$ is a maximum force value defined by the experimenter. Similarly, we define the decay parameter that lets the system converge sequentially to the set of attractors modeling the task (see Sec. II-A) as $\alpha = (1-r)\alpha_{\max}$, where $\alpha_{\max}$ is determined by the experimenter.

Fig. 1 illustrates the approach using a 2-dimensional example. In the top-right graph, we see that the trajectories repro-

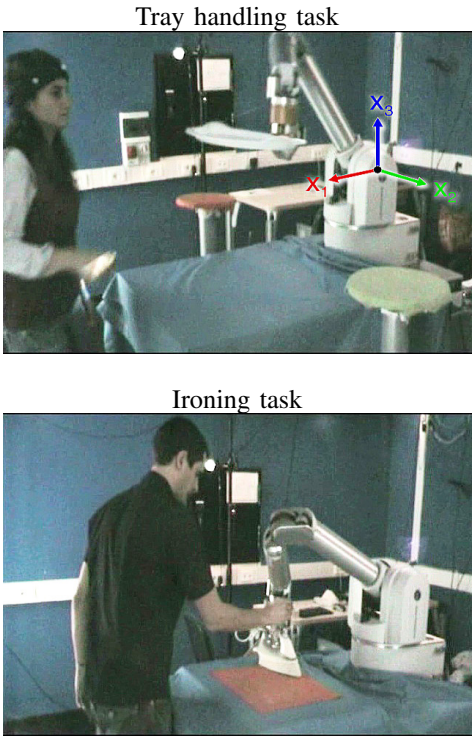

Fig. 2. Experimental setup.

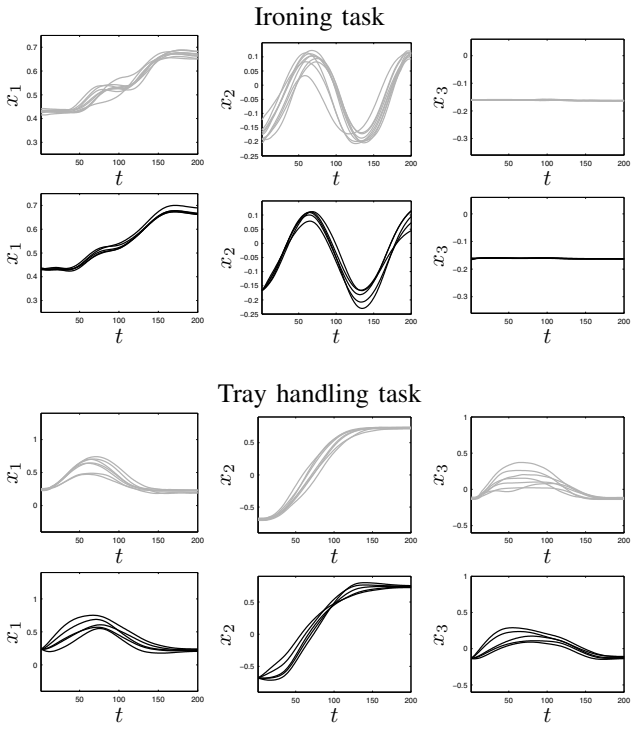

Fig. 3. Recordings (*top row*, in grey line) and stochastic reproduction from the learned model (*bottom row*, in black line). We see that the reproduced results present a variability similar to the one observed during the demonstration trials.

duced stochastically from the learned model show different levels of variability along the task. This variability shows similar characteristics to the one in the training set (*top-left* graph). In the second row, the behavior of the robot is different in the two reproduction scenarios. In the first, where collision avoidance is at the beginning of the movement, the robot goes round the user by deviating significantly from the generalized trajectory, as the constraints of the task are not very high in this area (i.e. the robot still fulfils the task constraints correctly). In contrast, in the second collision avoidance situation, near the end of the motion, the robot only moves slightly from the generalized trajectory, since the demonstrations were showing a higher level of consistency in this part of the movement (i.e. the robot can only slightly depart from the generalized trajectory to reproduce the task correctly). We see that, when the user is in the proximity of the robot's task path, the task motion is slowed down to smoothly avoid the user, with a natural behavior that takes into account the task requirements. Finally, when the user is not looking in the direction of the movement, the robot goes round the user with a larger amplitude.

III. EXPERIMENTS

A. Experimental setup

The experiment is conducted with a torque-controlled *Barrett WAM* 7 DOFs robotic arm. The position and orientation of the user's head are tracked with a marker-based *NaturalPoint OptiTrack* motion capture system. 12 cameras are used to track the position x_U and orientation of the user's head (R_U in direction cosine matrix representation), at a rate of 30 frames per second.

Two experiments are proposed. The first consists of holding a tray at the left-hand side of the robot, and moving it to a table at the right-hand side of the robot (while keeping the tray horizontal). The second consists of ironing a square napkin. These two skills have non-uniform constraints in position. For the tray handling task, the position of the tray is more constrained at the end of the movement than during its transportation, in order to bring it to the desired position. For the ironing task, the trajectory that the iron should follow is more constrained in the vertical axis than in the horizontal plane. To fulfil the task requirements, it is indeed more important to have the iron in contact with the table than to follow a very specific path on the table. These constraints are reflected in the collected data. During demonstration, gravity compensation is used to allow the user to move the robot effortless. Through this kinesthetic teaching process, 7 demonstrations are provided by recording the position and orientation of the end-effector. Each task is encoded with the proposed model, by fixing the number of states (or primitives) with respect to the length of the demonstrations. 6 and 8 states are respectively used to encode the tray handling and ironing movements.

Fig. 2 presents the experimental setup and the static frame of reference that we consider for the two tasks.

B. Experimental results

Figs 3-6 present the experimental results. Fig. 3 depicts stochastic reproduction results showing that the residuals can be used to generate reproduction attempts with similar

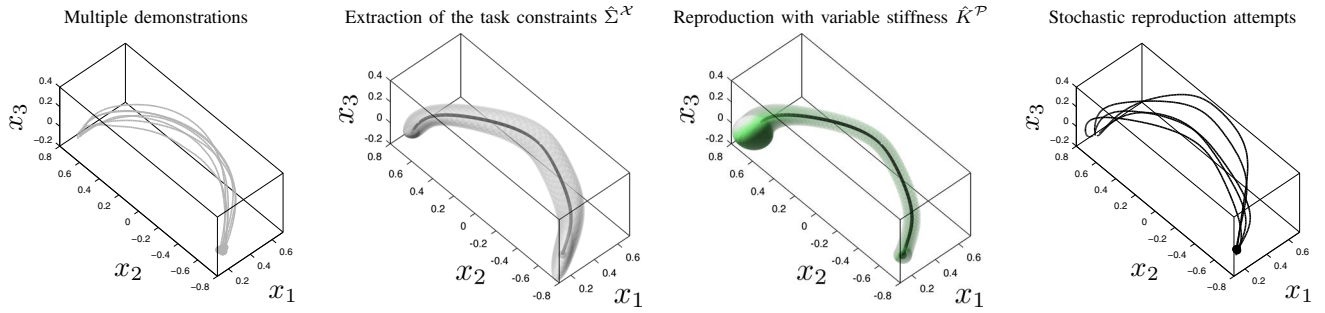

Fig. 4. Experimental results for the tray holding task. *From left to right:* Demonstrations. Extraction of the task constraints through the residuals of the regression process. Adaptive stiffness gain matrix computed from the residuals information. Stochastic reproductions of the movement.

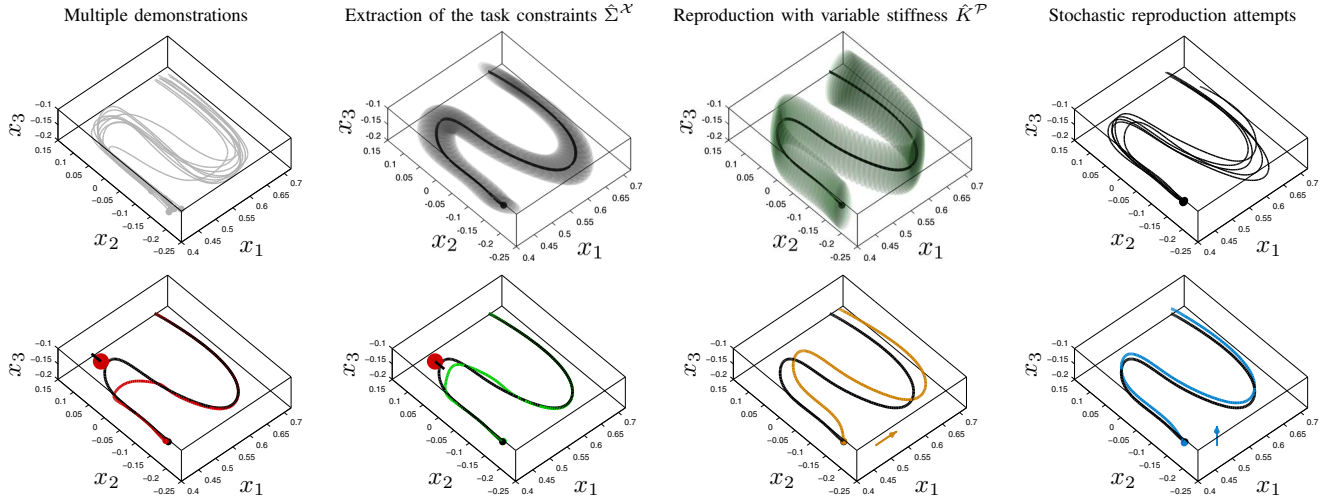

Fig. 5. Experimental results for the ironing task. *Top:* Demonstrations. Extraction of the task constraints through the residuals of the regression process. Adaptive stiffness gain matrix computed from the residuals information. Stochastic reproductions of the movement. *Bottom:* The second row shows several perturbed reproduction attempts. The user's head is represented by red dots, with the black lines showing its orientation. *From left to right:* Situation where the user is on the path of the robot and is not looking at the movement. Situation where the user is facing the robot. Reproduction by artificially applying a constant force parallel to the table at the end-effector of the robot. Reproduction by applying a constant force with the same amplitude but vertically to the table.

variabilities to those observed among the demonstration trials.

For the tray handling task in Fig. 4, the robot becomes stiffer at the end of the movement (bigger stiffness ellipsoids at the end of the movement), in order to correctly position the tray on the table (the covariance matrices $\hat{\Sigma}^{\mathcal{X}} = \sum_{i=1}^{K} h_i(t)\Sigma_i^{\mathcal{X}}$ and stiffness gain matrices $\hat{K}^{\mathcal{P}} = \sum_{i=1}^{K} h_i(t)K_i^{\mathcal{P}}$ are respectively represented with grey and green ellipsoids). During the course of the movement, the task constraints remain low, as there is no obstacle in the robot's workspace. The stiffness gains reflect this characteristic by remaining low until the robot approaches the table. We see in the second column that the model correctly encapsulates the variability of the demonstrations through the set of covariance matrices $\hat{\Sigma}^{\mathcal{X}}$ estimated from the residuals of the regression process. A variable stiffness gain matrix $\hat{K}^{\mathcal{P}}$ is automatically set in consequence in order to fulfil the learned task constraints (third column).

For the ironing task in Fig. 5, the learned model shows that it is more important to track the movement in the vertical direction than in the other two directions of the horizontal plane (nearly flat ellipsoids in the second graph). In the third graph, the stiffness matrices have consequently an elongated shape. We also see that by applying a virtual force to the end-effector during reproduction, the deformation is stronger if the force is parallel to the table than if the force is vertical (last two graphs).

Fig. 6 *right* presents snapshots of the robot reproducing the ironing task, while the user enters the workspace and grasps an object in the vicinity of the robot. A video of this experiment also accompanies this submission.

IV. DISCUSSION

We used in this paper a simple repulsive force fields policy as a first approach to avoid collisions with the user. Solutions based on superposition of additional force fields help prevent collisions [21], but do not necessarily guarantee that no collision will occur.

Such a policy is fast and computationally efficient, but it remains adequate only for a limited subset of the possible scenarios that one might expect in human-robot interaction. Considering multiple constraints within the same level of control can indeed be problematic in cases where competing

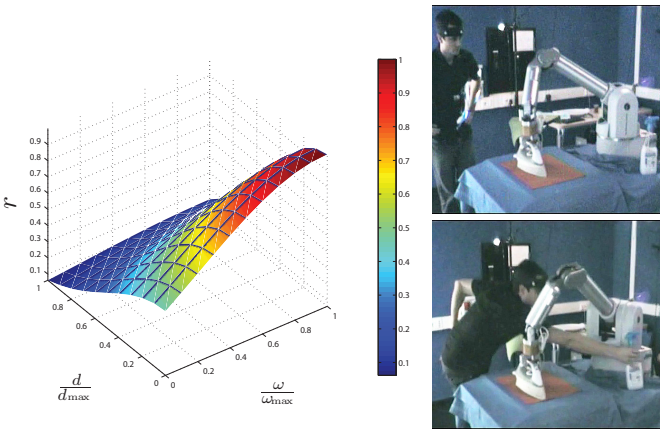

Fig. 6. *Left:* Risk indicator $r = f(d, \omega, \sigma_d, \sigma_\omega)$. *Right:* Snapshots of a reproduction attempt for the ironing task.

forces of high amplitudes sudden change in direction or application points, which can produce unpredictable behaviors if one does not limit the maximum force allowed. One possible alternative is to consider a hierarchical decomposition of the task constraints [9], [10], or to adopt a prioritized optimization strategy [22].

This first set of experiments opens the road for various further investigations. We will explore in which manner risk indicators can be integrated in an overall safety strategy, by taking into consideration that these indicators strongly depend on the user (e.g., the risk indicators can differ depending on the age or motor capabilities of the person interacting with the robot), as well as on other habituation factors. To do this, we will concentrate on learning the parameters that are relevant for the evaluation of the risk indicator. We will also explore in which manner the dynamics of the user's body can be considered in the estimation of the risk function (instead of a static pose). To do so, we plan to develop prediction strategies that take into account: (i) the context in which pre-collision occurs; and (ii) the robustness of the sensory information available to track the user's movement.

V. CONCLUSION

We proposed an active control strategy based on task-space inverse dynamics control with variable stiffness. Learning of the task is insured by the user providing multiple demonstrations of the skill. After extraction of the task constraints, the robot replicates the task by automatically selecting a variable level of compliance to reproduce the essential characteristics of the skill. The redundancy of the task and the redundancy of the robot are exploited to determine a safety control strategy through the estimation of the user's head pose. We demonstrated the feasibility of the approach with a tray handling task and an ironing task.

REFERENCES

[1] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Secaucus, NJ, USA: Springer, 2008, pp. 1371–1394.

[2] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robot. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, 2009.

[3] M. Laffranchi, N. G. Tsagarakis, and D. G. Caldwell, "Safe human robot interaction via energy regulation control," in *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, 2009, pp. 35–41.

[4] A. Bicchi and G. Tonietti, "Fast and soft arm tactics: Dealing with the safety- performance tradeoff in robots arm design and control," *IEEE Robotics and Automation Magazine*, vol. 11, no. 2, pp. 22–33, 2004.

[5] R. Schiavi and A. Bicchi, "Integration of active and passive compliance control for safe human-robot coexistence," *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pp. 1339–1345, 2009.

[6] A. De Luca, A. Albu-Schäffer, S. Haddadin, and G. Hirzinger, "Collision detection and safe reaction with the DLR-III lightweight manipulator arm," in *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, 2006, pp. 1623–1630.

[7] S. P. Gaskill and S. R. G. Went, "Safety issues in modern applications of robots," *Reliability Engineering and System Safety*, vol. 53, no. 3, pp. 301–307, 1996.

[8] ISO 10218, "Manipulating industrial robots - safety," Geneva, Switzerland, 1992.

[9] A. De Luca and L. Ferrajoli, "Exploiting robot redundancy in collision detection and reaction," in *IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, Nice, France, September 2008.

[10] J.-O. Kim, M. Wayne, and P. K. Khosla, "Exploiting redundancy to reduce impact force," *Journal of Intelligent and Robotic Systems*, vol. 9, no. 3, pp. 273–290, 1994.

[11] M. Howard, S. Klanke, M. Gienger, C. Goerick, and S. Vijayakumar, "Methods for learning control policies from variable-constraint demonstrations," in *From Motor Learning to Interaction Learning in Robots*, O. Sigaud and J. Peters, Eds. Springer Berlin / Heidelberg, 2010, pp. 253–291.

[12] R. Featherstone and D. E. Orin, "Dynamics," in *Handbook of Robotics*, B. Siciliano and O. O. Khatib, Eds. Secaucus, NJ, USA: Springer, 2008, pp. 35–65.

[13] A. J. Ijspeert, J. Nakanishi, and S. Schaal, "Trajectory formation for imitation with nonlinear dynamical systems," in *Proc. IEEE Intl Conf. on Intelligent Robots and Systems (IROS)*, 2001, pp. 752–757.

[14] S. Schaal, P. Mohajerian, and A. J. Ijspeert, "Dynamics systems vs. optimal control a unifying view," *Progress in Brain Research*, vol. 165, pp. 425–445, 2007.

[15] H. Hoffmann, P. Pastor, D. H. Park, and S. Schaal, "Biologically-inspired dynamical systems for movement generation: automatic real-time goal adaptation and obstacle avoidance," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2009, pp. 2587–2592.

[16] S. Calinon, F. D'halluin, D. G. Caldwell, and A. G. Billard, "Handling of multiple constraints and motion alternatives in a robot programming by demonstration framework," in *Proc. IEEE-RAS Intl Conf. on Humanoid Robots (Humanoids)*, Paris, France, December 2009, pp. 582–588.

[17] S. Calinon, F. D'halluin, E. L. Sauser, D. G. Caldwell, and A. G. Billard, "Learning and reproduction of gestures by imitation: An approach based on hidden Markov model and Gaussian mixture regression," *IEEE Robotics and Automation Magazine*, vol. 17, no. 2, pp. 44–54, 2010.

[18] M. Khansari and A. G. Billard, "BM: An iterative method to learn stable non-linear dynamical systems with Gaussian mixture models," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, Anchorage, Alaska, USA, May 2010, pp. 2381–2388.

[19] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *Intl Journal of Robotics Research*, vol. 28, no. 11-12, pp. 1507–1527, 2009.

[20] S. Calinon and A. Billard, "A framework integrating statistical and social cues to teach a humanoid robot new skills," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA), Workshop on Social Interaction with Intelligent Indoor Robots*, Pasadena, CA, USA, May 2008.

[21] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.

[22] M. de Lasa and A. Hertzmann, "Prioritized optimization for task-space control," in *IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, St. Louis, MO, USA, 2009, pp. 5755–5762.