Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

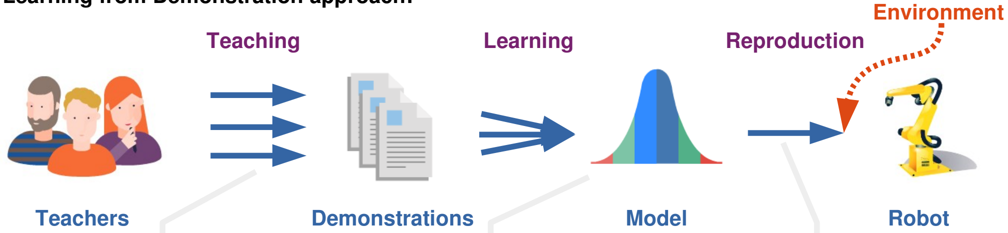

Kinesthetic Teaching

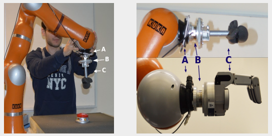

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

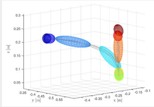

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

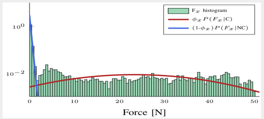

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

"The proposed LfD approach allows robots equipped with impedance control to learn from demonstration the positional and the force aspect of in-contact tasks and choose which one to prioritize during the reproduction by changing the controller stiffness"

Authors and Contacts

Mattia Racca
mattia.racca@aalto.fi
Aalto University

Prof. Ville Kyrki
ville.kyrki@aalto.fi

Joni Pajarinen
pajarinen@ias.tu-darmstadt.de
TU Darmstadt

Alberto Montebelli
alberto.montebelli@his.se
University of Skövde

Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

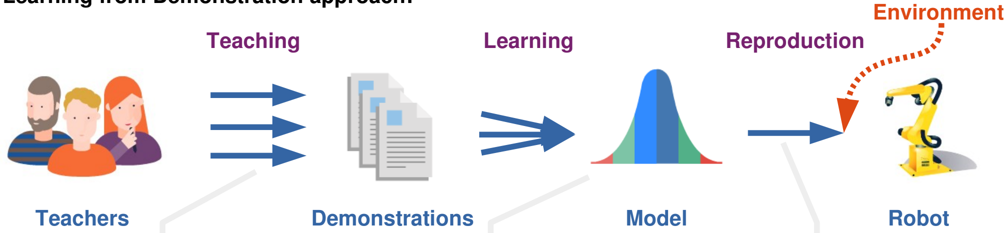

Kinesthetic Teaching

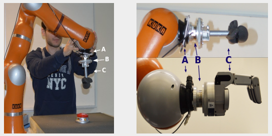

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

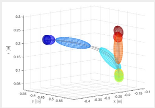

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

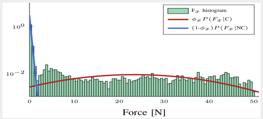

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

Simultaneous Kinesthetic Teaching of Trajectory and Exerted Forces

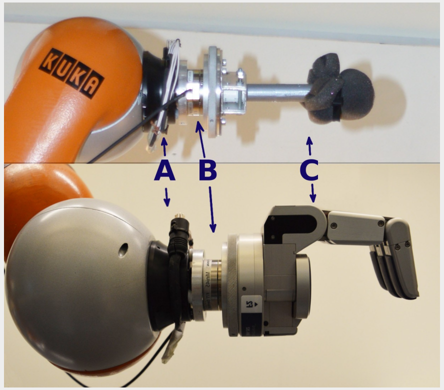

(A) robot flange (B) FT sensor (C) tool

- Avoid separate teaching sessions
- No use of external force recording devices
- The user operates the familiar tool via the robot

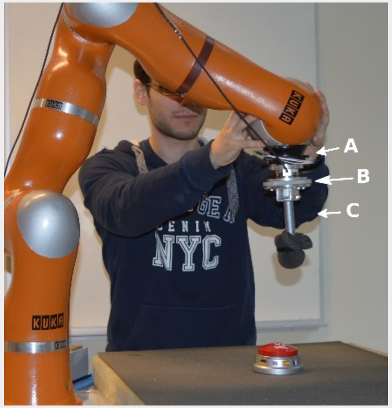

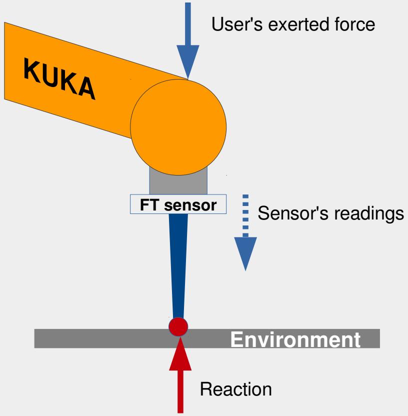

Learning In-Contact Control Strategies from Demonstration

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

Aalto University

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

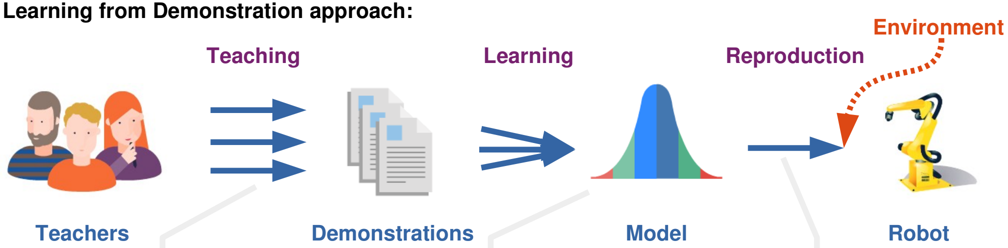

Kinesthetic Teaching

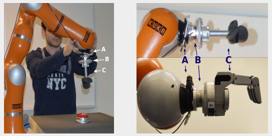

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

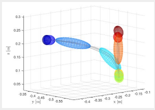

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

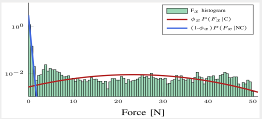

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

HSMM encoding the trajectory and the exerted forces (learned from demonstration)

A *N* state HSMM describes the task:

$$\lambda = (\pi, A, \mu, \Sigma, \mu^D, \sigma^D)$$

Duration probabilities $p_i^D(t) = \mathcal{N}(t; \mu_i^D, \sigma_i^D)$

Observation probabilities $p_i(\boldsymbol{z}_t) = \mathcal{N}(\boldsymbol{z}_t; \boldsymbol{\mu}_i, \Sigma_i)$

The model is trained with the Baum-Welch algorithm from *D* demonstrations.

Why H(S)MM?
- Captures correlation between the pose and the exerted forces during the taught task
- Easy to add new information to the learning
- Robustness against not aligned demonstrations

Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

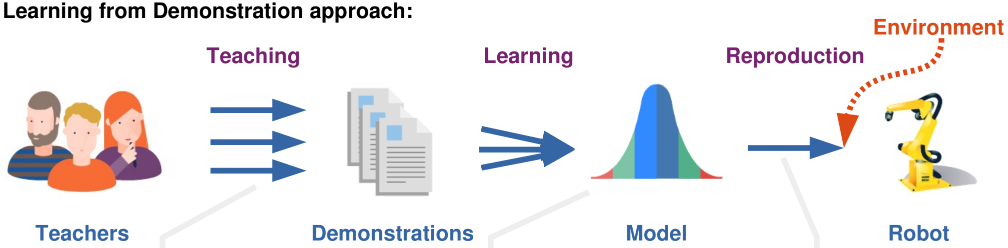

Kinesthetic Teaching

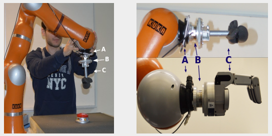

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

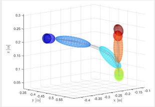

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

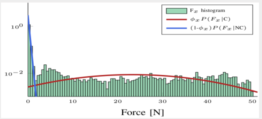

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

From the model to the reproduction...

Forward variable

$$\alpha_{i,t} = \sum_{j=1}^{N} \sum_{d=1}^{\min(t_{max},t-1)} \alpha_{j,t-d}\, a_{ji}\, p_i^D(d) \prod_{s=t-d+1}^{t} \mathcal{Q}_i\left(\begin{bmatrix} \boldsymbol{x}_t \\ \boldsymbol{q}_t \end{bmatrix}\right).$$

Gaussian Mixture Regression (commands)

$$\dot{\boldsymbol{x}}_t^* = \sum_{i=1}^{N} h_{i,t} \left[\boldsymbol{\mu}_i^{\dot{x}} + \Sigma_i^{\dot{x}x}(\Sigma_i^{xx})^{-1}(\boldsymbol{x}_t - \boldsymbol{\mu}_i^{x})\right] ,$$

$$\dot{\boldsymbol{q}}_t^* = \sum_{i=1}^{N} h_{i,t} \left[\boldsymbol{\mu}_i^{\dot{q}} + \Sigma_i^{\dot{q}q}(\Sigma_i^{qq})^{-1}(\boldsymbol{q}_t - \boldsymbol{\mu}_i^{q})\right] ,$$

$$\boldsymbol{f}_t^* = \sum_{i=1}^{N} h_{i,t}[\boldsymbol{\mu}_i^{f} + \Sigma_i^{fq}(\Sigma_i^{qq})^{-1}(\boldsymbol{q}_t - \boldsymbol{\mu}_i^{q}) + \Sigma_i^{fx}(\Sigma_i^{xx})^{-1}(\boldsymbol{x}_t - \boldsymbol{\mu}_i^{x})] .$$

... via Cartesian Impedance control

$$\boldsymbol{\tau}_{\mathrm{cmd}} = J^T(\mathrm{diag}(\boldsymbol{k}_{\mathrm{c}})(\boldsymbol{x}_{\mathrm{cmd}} - \boldsymbol{x}_{\mathrm{msr}}) + D(\boldsymbol{d}_{\mathrm{c}}) + \boldsymbol{F}_{\mathrm{cmd}}) + \boldsymbol{\tau}_{\mathrm{dyn}}(q, \dot{q}, \ddot{q}) ,$$

Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

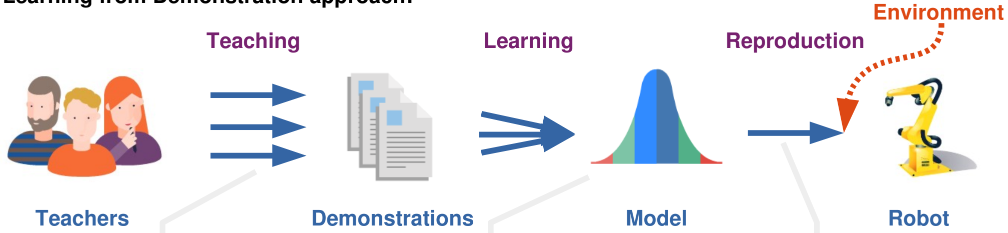

Kinesthetic Teaching

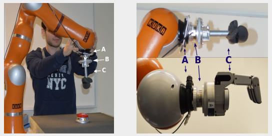

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

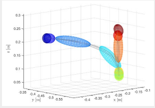

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

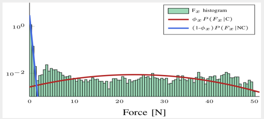

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

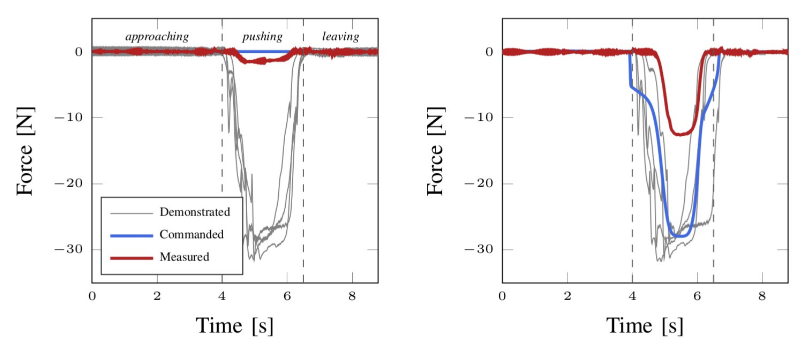

(a) No force learning

(b) Force learning and constant k_c

Constant stiffness reproduction is not enough!

$$\boldsymbol{\tau}_{\text{cmd}} = J^T(\text{diag}(\boldsymbol{k}_{\text{c}})(\boldsymbol{x}_{\text{cmd}} - \boldsymbol{x}_{\text{msr}}) + D(\boldsymbol{d}_{\text{c}}) + \boldsymbol{F}_{\text{cmd}}) + \boldsymbol{\tau}_{\text{dyn}}(q, \dot{q}, \ddot{q}) \ ,$$

The spring component of the controller is acting against the force component.
We need to know when during the reproduction the exertion of the force is more important than the trajectory following (i.e. what part of the task to prioritize)

Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

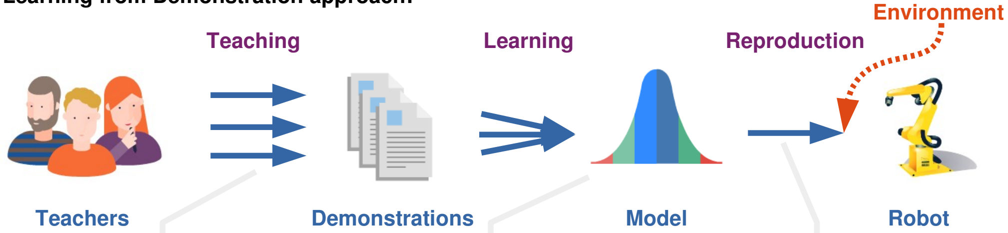

Kinesthetic Teaching

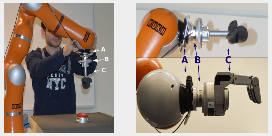

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

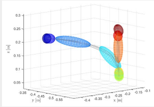

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

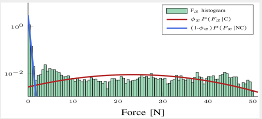

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

Modeling the in-contact and not in-contact phases

Probability of contact given exerted force

$$P(\mathrm{C}|F_j) = \frac{P(F_j|\mathrm{C})\phi_j}{P(F_j|\mathrm{C})\phi_j + P(F_j|\mathrm{NC})(1-\phi_j)}$$

$$\boldsymbol{\eta}_i = P(\mathrm{C}||\boldsymbol{\mu}_i^f|)$$

Variable Stiffness

$$\boldsymbol{k}_c(t) = \sum_{i=1}^{N} (h_{i,t}\,\mathrm{diag}(\boldsymbol{\eta}_i)(\boldsymbol{k}_{max} - \boldsymbol{k}_{min}) + \boldsymbol{k}_{min})$$

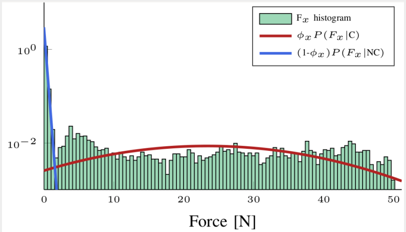

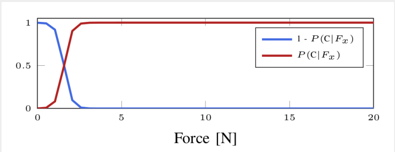

We can smoothly vary the stiffness coefficients along all Cartesian directions individually.

Aalto University

Learning In-Contact Control Strategies from Demonstration

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

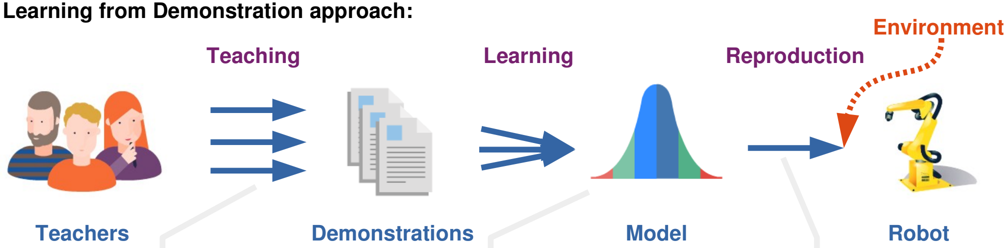

Kinesthetic Teaching

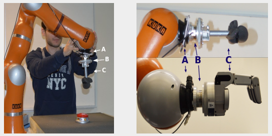

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

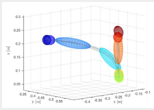

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

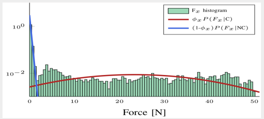

We modulate the **stiffness** during the execution of motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

Comparison: constant vs variable stiffness

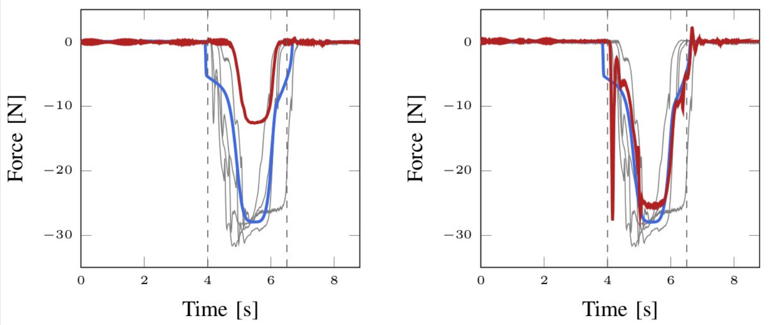

(b) Force learning and constant $\boldsymbol{k}_c$ (c) Force learning and $\boldsymbol{k}_c$ selection mech.

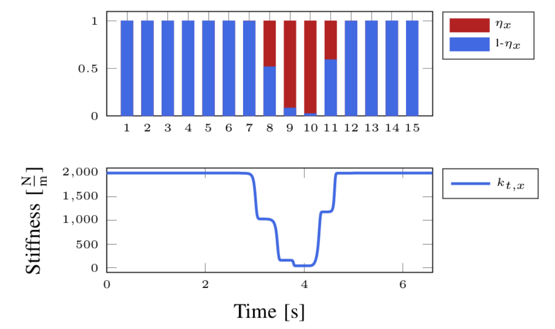

How stiffness is changed during the reproduction

Learning In-Contact Control Strategies from Demonstration

Aalto University

Mattia Racca, Joni Pajarinen, Alberto Montebelli, Ville Kyrki

ACADEMY OF FINLAND

In-contact tasks, i.e. tasks that crucially involve a distribution of ***forces*** and ***torques*** in space and time at the interface between the tool and the environment, are ***difficult*** to encode in declarative terms and, consequently, surprisingly difficult to teach to robots. We develop a ***learning from demonstration*** framework that allows a robot to learn in-contact tasks. A ***hidden semi-Markov model*** (HSMM) captures both the task's position and force constraints and provides the commands for an ***impedance controller with variable stiffness***.

Learning from Demonstration approach:

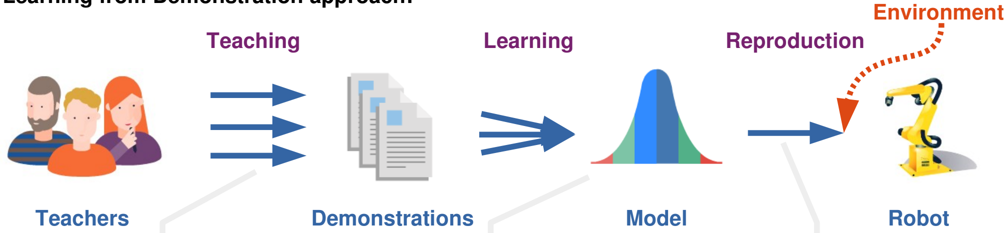

Kinesthetic Teaching

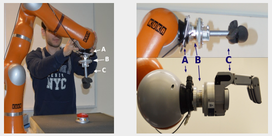

In order to collect the demonstrations, the human teacher can physically grab the robot and operate its tool.
The use of a compliant robot with a **force-torque sensor** placed between the robot flange and the tool allows for **simultaneous teaching** of the positional and force profiles of the task.

Hidden Semi-Markov Models

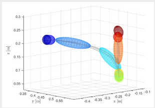

HSMMs encode the task's dynamics (**position** and **force** constraints), learned from the collected demonstrations.

Since interaction with the environment is required, we use **impedance control**. However, simply playing back the learned forces is not sufficient to successfully accomplish tasks.
A decision on which aspect of the motion (the kinematic **trajectory** or the exertion of **forces**) to **prioritize** needs to be made.

Task reproduction with impedance controller and varying stiffness

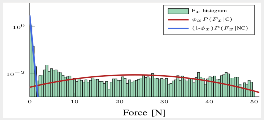

We modulate the **stiffness** during the execution of the motion, based on the force information encapsulated in each HSMM state.
We extract the two basic modes of interaction given the recorded force: **in-contact** with the environment and **not in-contact**.

2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

We proposed a LfD approach for **in-contact tasks** that uses a HSMM for encoding the dynamics of the demonstrated trajectories including the **force profiles**. We proposed a mechanism for selecting the **stiffness** parameters of a Cartesian impedance controller based on the current belief of the HSMM states and the analysis of the recorded force profiles.

In the case of tasks with complex force distributions, the proposed stiffness selection mechanism might fail in modeling the in-contact and not in-contact components of the motion. In that case, **more complex distribution** models such as mixtures of distributions could be used.

Thanks for your attention!