2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

Robotic Force Estimation Using Motor Torques and Modeling of Low Velocity Friction Disturbances

Magnus Linderoth, Andreas Stolt, Anders Robertsson, Rolf Johansson

***Abstract*— For many robot operations force control is needed, but force sensors may be expensive and add mass to the system. An alternative is to use the motor torques, though friction causes large disturbances. The Coulomb friction can be quite well known when a joint is moving, but has much larger uncertainties for velocities close to zero.**

This paper presents a method for force estimation that accounts for the velocity-dependent uncertainty of the Coulomb friction and combines data from several joints to produce accurate estimates. The estimate is calculated by solving a convex optimization problem in real time. The proposed method was experimentally evaluated on a force-controlled dual-arm assembly operation and validated with data from a force sensor. The estimates were shown to improve with the number of joints used, and the method can even exploit data from an arm that is controlled not to move.

I. Introduction

The traditional way of programming industrial robots is to use position control and follow pre-defined trajectories, using the joint position sensors. Modern robot controllers are very good at this and perform these tasks both fast and with high accuracy. In tasks where the robot has to physically interact with the environment, however, this control strategy is less advantageous. The accuracy of the robot and the location and geometry of everything in the workspace have to be known with high precision, and this is usually hard to achieve. A remedy to this problem is to introduce additional sensing, e.g., a force sensor that gives the robot capabilities to handle position uncertainties by sensing the contact forces. A force sensor can thus be used to make the robot system more robust towards uncertainties. The main drawbacks with using a force sensor are that it may be expensive and add mass to the system.

An alternative to using a force sensor is to estimate the external forces applied to the robot based on sensing already available in the robot. Usually this includes position sensors in the joints and torques exerted by the motors. The main problems with estimating forces is how to handle the large disturbances that are present, e.g., originating from friction. In a previous work by the authors [17], the joint control errors were used to estimate the external force.

Magnus Linderoth, magnus.linderoth@control.lth.se, Andreas Stolt, Anders Robertsson, and Rolf Johansson are with the Department of Automatic Control, LTH, Lund University, Sweden.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 – Challenge 2 – Cognitive Systems, Interaction, Robotics – under grant agreement No 230902 - ROSETTA. This document reflects only the author's views and the European Community is not liable for any use that may be made of the information contained herein.

The authors are members of the LCCC Linnaeus Center and the eLLIIT Excellence Center at Lund University.

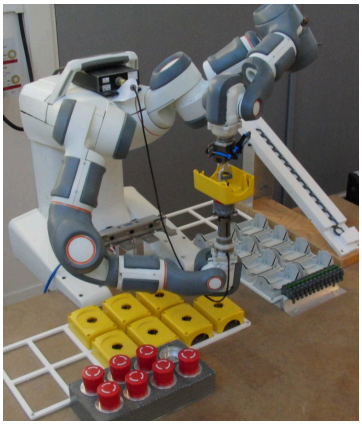

Fig. 1. The setup for the assembly task used in the experiments. The force/torque sensor, mounted at the wrist of the robot's right arm, is only used for verification and evaluation of the estimated forces.

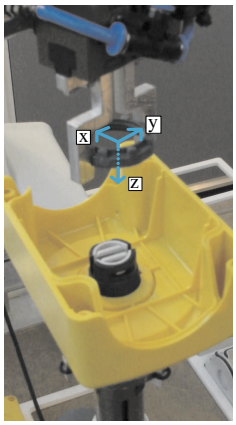

Fig. 2. Zoom in and illustration of the frame in which the contact forces are measured.

Force estimation has been considered many times before in the literature. A common approach is to use model-based disturbance observers. A dynamical model of the robot is used and forces are estimated by an observer based on the deviations from the model. Apart from a model of the robot, these methods need a model for interaction with the environment. Examples of previous work using this technique can, e.g., be found in [2], [6], and [11].

Methods for force estimation are also available in commercial industrial robot systems, e.g., both ABB [1] and Toshiba [18] provide such products. These systems are, however, designed to work well for large forces, and would therefore not be suitable for use in tasks where it is important to react already for smaller forces, which is considered in this paper.

In previous work, e.g., [9], [10], [14], [15], it was assumed that the joints were always moving, and no attention was given to the large uncertainties in the friction torques at velocities close to zero. It has also been common to only consider experiments with three or fewer joints.

In [13], [15] the performance of force estimation was improved by modeling the friction carefully and considering position-dependent torque variations, and [5], [12] include modeling of low-velocity friction phenomena. These models, however, require knowledge of many parameters that are challenging to identify and prone to change, due to, e.g., temperature and wear.

The focus of this paper is not on estimating the friction torques of the individual joints rigorously, but on modeling the velocity-dependent uncertainties in the friction torques and combining measurements from multiple joints to com-

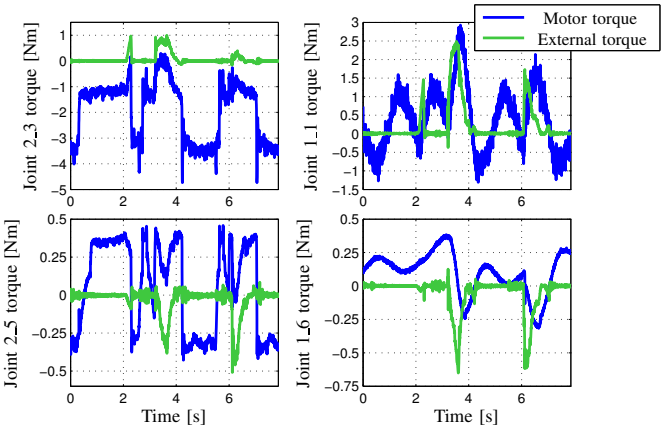

Fig. 3. Measured motor torque and applied external torque from an assembly sequence. The diagrams in the left column show data from the left arm, which was controlled to move, and the diagrams in the right column show data from the right arm, which was controlled to be still. The upper diagrams show data from base joints, and the lower diagrams show data from wrist joints. The joints chosen for display are those where the external torques were the most visible. It can be seen that the disturbances were as large or even larger than the signal of interest.

pute an accurate estimate of the contact force. In particular, we take the noise in the velocity measurement into account, and model that the Coulomb friction is quite well known when a joint is moving, but has much larger uncertainty for velocities close to zero. The force is estimated by solving a convex optimization problem, and an approximate confidence interval is also calculated.

The validity of the approach is investigated by comparing the estimated forces with measurements from a force sensor. The method is finally tested in a dual-arm screwing assembly task with the setup shown in Figs. 1 and 2

A. Motivation

The motivation behind this work is to be able to use force control in cases when no force sensor is available. Measured motor torques are often available, but they contain large disturbances. An example of measured motor torque from a dual-arm assembly execution is displayed in Fig. 3. The right robot arm was controlled to be still while the left robot arm was manipulating an object held by the right robot arm. The actual external torque, as measured by a force sensor, is also displayed in the figure. It can be seen that the disturbances are larger than the signal of interest. It can further be seen that the disturbances are different for the two arms, with a distinct Coulomb friction pattern for the arm that is controlled to move, while the arm that is controlled to be still appears to have less predictable disturbances.

II. Method

A. Modeling

The method used for force estimation in this paper is based on the measured joint motor torques. The model used is

$$\tau_m = \tau_{grav} + \tau_{dynamic} + \tau_{ext} + \tau_e \tag{1}$$

where τ_m denotes the measured motor torques, τ_{grav} denotes the torques originating from gravity, $\tau_{dynamic}$ denotes dynamic torques originating from accelerations of the robot, τ_{ext} denotes external torques, and τ_e denotes disturbances due to, e.g., friction, measurement noise, and modeling errors.

The influence from gravity, τ_{grav}, can be calculated if the mass and center of mass are known for each link of the robot. If they are not known, it is fairly easy to perform identification experiments to find these parameters. The actual calculation is, e.g., described in [16, p. 271]. The dynamic torques, $\tau_{dynamic}$, can also be calculated if the dynamic parameters of the robot are known, i.e., moment of inertia for each link of the robot. The dynamic torques will, however, be small in tasks where it is interesting to use force estimation, as the robot will be interacting with the environment and thus needs to move quite slow. It is therefore assumed that the dynamic torques can be neglected.

The external joint torques originate from external forces and torques applied to the robot. If it is assumed that all external forces are applied to the end effector of the robot, the external joint torques are given by

$$\tau_{ext} = J^T F \tag{2}$$

where $J = J(q)$ is the Jacobian of the robot, q the joint coordinates, and F denotes the force/torque applied to the end effector.

B. Disturbance torques

The disturbance, τ_e, influencing each joint mostly consists of Coulomb friction, which can be modeled to give the following contribution for joint i

$$\tau^i_{Coulomb} = \begin{cases} \tau^i_{C,max}\ , & \dot{q}_i > 0 \\ \tau^i_{C,min}\ , & \dot{q}_i < 0 \end{cases} \tag{3}$$

where $\dot{q}_i$ denotes the velocity of joint i, and $\tau^i_{C,max}$ and $\tau^i_{C,min}$ denote the constant friction levels for positive and negative velocities respectively. What happens at zero velocity is not given by the model, and the torque might be anywhere in the interval $[\tau^i_{C,min}, \tau^i_{C,max}]$. Therefore, for low velocities close to zero, the Coulomb friction contribution can be modeled by a uniform random variable.

Another type of friction is viscous friction. It can be modeled to give the following contribution for joint i

$$\tau^i_{viscous} = c_i \dot{q}_i \tag{4}$$

where c_i is a constant specific for each joint.

Another large disturbance is measurement noise, which can be modeled to have a zero-mean Gaussian distribution.

C. Disturbance model

To find out the disturbance characteristics, an identification experiment was performed for each joint. The joint was then moving back and forth with a low piecewise constant acceleration, without any external forces applied to the robot. Two versions of this experiment are displayed in Fig. 4; the upper diagram shows an experiment with low velocities, and the lower diagram an experiment with higher velocities. The raw data, sampled at 250 [Hz], were low-pass filtered

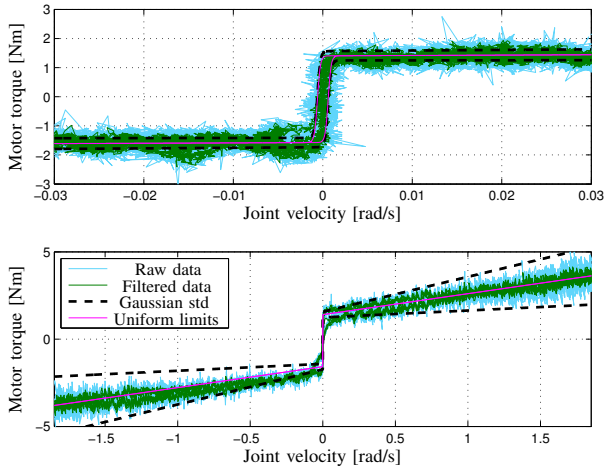

Fig. 4. Experimental data from an experiment where one joint of the robot was controlled to move back and forth with piecewise constant acceleration. The upper diagram shows an experiment with only low velocities, and the lower diagram shows an experiment with higher velocities. The disturbance characteristics are clearly visible in this experiment. Also shown is one standard deviation of the measurement noise multiplied with a velocity dependent factor, and the upper and lower limits for the uniform distribution describing the Coulomb and the viscous friction.

with the discrete-time filter (5) to remove some of the noise influence.

$$H(z) = \frac{0.4}{1 - 0.6z^{-1}} \tag{5}$$

The Coulomb friction is easy to see in both experiments. As was suggested earlier, the amount of friction for zero velocities varies between $\tau_{C,min}$ and $\tau_{C,max}$, and due to noise in the velocity measurements this is true also for measured velocities slightly different from zero. Aside from Coulomb friction, the experiment with large velocities shows viscous friction. Further, there is also noise present.

A probabilistic model of the disturbances is therefore that the Coulomb and the viscous friction are the outcome of a uniform random variable with a velocity-dependent range. This range is zero for large velocities and the range grows for low velocities. One way to describe this range is by using sigmoid functions for describing the upper and the lower limits. To incorporate also viscous friction, a linear term is added. The upper and lower limit for each joint can be described by the following functions (joint index skipped)

$$\begin{aligned} \tau_{f,max}(\dot{q}) &= \tau_{C,min} + \frac{\tau_{C,max} - \tau_{C,min}}{1 + e^{-A(\dot{q}+B)}} + c\dot{q} \\ \tau_{f,min}(\dot{q}) &= \tau_{C,min} + \frac{\tau_{C,max} - \tau_{C,min}}{1 + e^{-A(\dot{q}-B)}} + c\dot{q} \end{aligned} \tag{6}$$

The parameter A determines the slope of the sigmoid function, and the parameter B the width of the area between the curves. Parameters for such functions were manually tuned for each joint of the robot, and an example is seen as magenta curves in Fig. 4.

A Gaussian noise term is used to account for measurement noise, uncertainty in the friction limits and unmodeled disturbances. It can be seen in Fig. 4 that the variance of the noise increases when the velocity increases. The model used is therefore that the variance of the noise is velocity dependent and the standard deviation for different velocities is calculated as the standard deviation for low velocities multiplied with a factor $(1 + k|\dot{q}_i|)$. One standard deviation of the noise is displayed in Fig. 4. Data recorded during assembly operations indicated that the actual disturbances at high velocities were higher than the measured data in Fig. 4 indicate. Hence, the one-standard-deviation limit may appear overly pessimistic in this figure.

To conclude, the total disturbance torque is modeled as

$$\tau_e = \tau_f + e \tag{7}$$

where $\tau_{f,min}(\dot{q}) \leq \tau_f \leq \tau_{f,max}(\dot{q})$, and e is zero-mean Gaussian with diagonal covariance matrix $\mathbb{E}[ee^T] = R_e(\dot{q}) = \mathrm{diag}(\mathbf{1} + k|\dot{q}|)^2 R_e(0)$.

D. Force estimation

Let $\bar{\tau}$ be the motor torques compensated for gravity, calculated as

$$\bar{\tau} = \tau_m - \tau_{grav} \tag{8}$$

Using (1), (2), (7), and the assumption that the dynamic torques are negligible, this gives

$$\begin{aligned} \bar{\tau} &= \tau_{ext} + \tau_e \\ &= J^T F + \tau_f + e \end{aligned} \tag{9}$$

where $\bar{\tau}$ and J are given, and τ_f and e are random variables with uniform and Gaussian distributions respectively. The ML (Maximum Likelihood) estimate of F is then given by

$$\begin{aligned} \underset{\text{over } F,\tau_f}{\text{minimize}} \quad & \left(\bar{\tau} - J^T F - \tau_f\right)^T R_e^{-1} \left(\bar{\tau} - J^T F - \tau_f\right) \\ \text{subject to} \quad & \tau_{f,min} \leq \tau_f \leq \tau_{f,max} \end{aligned} \tag{10}$$

The estimate given by (10) can be improved by adopting a Bayesian approach and using prior knowledge of F in the particular task. The type of prior knowledge that can be used is, for instance, that the contact torques are small compared to the torque disturbances, and by reflecting this knowledge in the distribution of F it is possible to improve the quality of the estimated contact forces.

Assuming that the prior on F is Gaussian with $\mathbb{E}[F] = \bar{F}$ and $\mathbb{E}[(F - \bar{F})(F - \bar{F})^T] = R_F$, and that F and e are uncorrelated, the ML estimate of F is given by

$$\begin{aligned} \underset{\text{over } F,\tau_f}{\text{minimize}} \quad & \left(\bar{\tau} - J^T F - \tau_f\right)^T R_e^{-1} \left(\bar{\tau} - J^T F - \tau_f\right) \\ & + \left(F - \bar{F}\right)^T R_F^{-1} \left(F - \bar{F}\right) \\ \text{subject to} \quad & \tau_{f,min} \leq \tau_f \leq \tau_{f,max} \end{aligned} \tag{11}$$

The problem (11) is convex and can be solved numerically in real time, as described in Sec. II-F.

E. Confidence interval estimation

The uncertainty of the estimate given by (11) varies significantly with, e.g., the velocity of the different joints and the robot Jacobian. Hence, it is important to calculate the uncertainty of every estimate individually.

It is difficult to compute exact quantiles for the solution of (11), but this section describes a method for extracting

approximate confidence intervals that can be computed in real time. The method is first described for the case with a single robot joint without prior, and then generalized to handle multiple joints and a prior distribution on F.

1) One-dimensional case: The proposed confidence interval for the case of a single robot joint with no prior and the Jacobian $J = 1$ is illustrated in Fig. 5. The limits are calculated as

$$\begin{aligned} \tau_{conf,min} &= \tau_{f,min} - \lambda\sigma \\ \tau_{conf,max} &= \tau_{f,max} + \lambda\sigma \end{aligned} \tag{12}$$

where σ is the standard deviation of the Gaussian tails and λ is a parameter deciding the confidence level of the confidence interval.

For the special case where $\tau_{f,min} = \tau_{f,max}$ (the distribution of τ_e is Gaussian), the portion of the measurements outside the confidence interval is $2(1-\Phi(\lambda))$, where $\Phi(\cdot)$ is the cumulative distribution function of the zero-mean unit-variance Gaussian distribution.

An alternative way of finding the limits (12), is to minimize the negative log-likelihood function of τ_e and adding a gradient to push the solution toward the upper or lower limit. The log-likelihood of a zero-mean Gaussian with standard deviation σ is given by

$$\log \mathcal{L}(e) = -\frac{e^2}{2\sigma^2} + \text{const.} \tag{13}$$

$$\frac{d}{de}\left(\log \mathcal{L}(e)\right) = -\frac{e}{\sigma^2} \tag{14}$$

Hence, at the limits of the confidence interval, the derivative of the negative log-likelihood function of τ_e is

$$-\frac{d}{de}\left(\log \mathcal{L}(\pm\lambda\sigma)\right) = \pm\frac{\lambda}{\sigma} \tag{15}$$

Consequently, adding a gradient with one of the slopes (15) to the negative log-likelihood of τ_e and finding the minimum, gives one of the limits (12) as the solution. This way of calculating the limits is described because it generalizes to higher dimensions better than (12).

2) Multi-dimensional case: For the multi-joint problem (11), first assume that $\tau_{f,min} = \tau_{f,max}$, (i.e., Gaussian distribution). The standard deviation σ of the marginal distribution of F in the direction of the unit vector v is then given by

$$\sigma = \sqrt{v^T \left(JSR_e^{-1}J^T + R_f^{-1}\right)^{-1} v} \tag{16}$$

where S is the identity matrix for the Gaussian case but may have other values for the general case, as described later in this section. The limits of the confidence interval in the direction v are then given by the F solving

$$\begin{aligned} \underset{\text{over } F,\tau_f}{\text{minimize}} \quad & \left(\bar{\tau} - J^T F - \tau_f\right)^T R_e^{-1} \left(\bar{\tau} - J^T F - \tau_f\right) \\ & + \left(F - \bar{F}\right)^T R_F^{-1} \left(F - \bar{F}\right) \mp \tfrac{\lambda}{\sigma} v^T F \\ \text{subject to} \quad & \tau_{f,min} \le \tau_f \le \tau_{f,max} \end{aligned} \tag{17}$$

where the "$-$" in the "$\mp$" is for the upper limit, and the "$+$" is for the lower limit. This formulation is obtained by adding the gradient (15) to the problem (11).

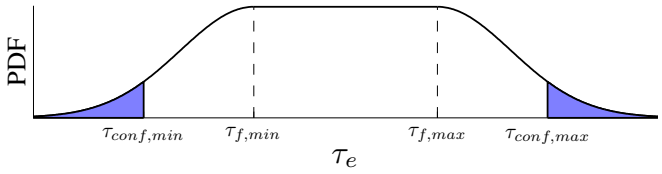

Fig. 5. Illustration of the proposed confidence interval on the probability density function (PDF) of $\tau_e = \tau_f + e$, a flat part and two Gaussian tails. The blue areas indicate the portion of the measurements expected to be outside the confidence interval.

Returning to the general case, when $\tau_{f,min} \neq \tau_{f,max}$, some of the joints may get an estimated τ_e in the range $\tau_{f,min} < \tau_e < \tau_{f,max}$. The cost function for that joint is then locally flat, cf. Fig. 5, and should not be considered when calculating (16).

To find out for which joints the estimated τ_e ends up in the Gaussian part of the distribution, we propose the following algorithm. It is assumed that n joints are used for force estimation and that R_e is diagonal.

1) Set $S = 0_{n\times n}$
2) Calculate (16)
3) Solve (17)
4) For the joints where $\tau_f = \tau_{f,min}$ or $\tau_f = \tau_{f,max}$, set the corresponding diagonal elements of S to 1
5) If S was modified in step 4), go to step 2). Else quit.

The intuition behind the above algorithm is the following. The problem (17) is first solved using a gradient based only on the prior. If the Coulomb friction for all joints is large, the resulting τ_e may all be within the flat part of the distribution and only the prior is used for determining the confidence interval. If any of the estimated τ_e reaches the Gaussian parts of the distributions, the gradient based only on the prior will not be able to push the estimate far down the Gaussian tails. The value of σ in (16) is then modified to include all joints where the τ_e estimate is in the Gaussian part, resulting in a steeper gradient, which may in turn push the estimate of additional joints to the Gaussian part of the distribution. The process (steps 2-5) is iterated until convergence.

F. Implementation

The optimization problem (11) is a convex optimization problem of fairly small size and can be solved in real time in a reliable manner. To this purpose, CVXGEN [8] was used. It is a code generator for embedded convex optimization. The generated code is library-free C code, and this code has been connected to the robot controller via an Ethernet connection.

The generated solver is run on a Linux PC and the computation time to arrive to a solution is in the order of 0.3 [ms]. The robot controller is run with a sampling time of 4 [ms], and the speed of the solver is therefore sufficient to be run in each sample. The solution with the Ethernet connection introduces a delay of one sample, as the indata to the solver is sent one sample before the solution is returned.

III. Assembly Scenario

A part of the assembly of an emergency stop button, shown in Fig. 2, was considered to illustrate the use of

the force-estimation method. The subtask considered was the attachment of the button onto the box by screwing a nut. This assembly operation was performed in a dual-arm setting, with one arm holding the button with the yellow box on, and the other arm performing the screwing (2.5–3.5 revolutions). The task was modeled and specified using the iTaSC framework [4].

The arm holding the button was *static*. By this we mean that it was controlled not to move, but forces applied to it could cause it to move slightly.

The assembly sequence consisted of first putting the nut on the thread of the button, using the contact forces to find the hole of the nut, though the exact grip of the button was uncertain. Before the nut was properly screwed onto the thread, large side forces could be caused by a bad grip of the nut. If such a force was detected during the first two revolutions, the gripper was opened and re-closed, which usually gave a more centric grip. Due to the way that the nut was gripped, large side forces were generated when the nut was tightened, which could be used to detect when the screwing operation was completed.

The robot used in the assembly scenario was the ABB FRIDA [7]. It is a dual-arm manipulator with 7 joints in each arm, developed for automation of assembly operations. The robot was controlled with the IRC5 control system, extended with an open control system [3], which made it possible to modify the references for the low-level joint control loops. The motor torques used for force estimation were calculated from the motor currents. An ATI Mini40 6-DOF force/torque sensor was mounted on the wrist on one of the arms to give ground truth data, see Fig. 1.

IV. Experimental Results

A. Calibration

An experiment where the robot was programmed to slowly move around in its workspace was performed to identify the parameters used for calculating the gravity torque, τ_{grav}. The resulting parameters resulted in a mean absolute error ranging from 0.05 [Nm] for the wrist joints to 0.3 [Nm] for the base joints.

The friction parameters were tuned by performing experiments that were described in Sec. II-C.

B. Force estimation

The force estimation method was tested in an experiment where forces were applied to a static arm (controlled not to move). The estimated forces and confidence intervals are displayed in Fig. 6, together with ground truth data from the force sensor. All confidence intervals in this paper were estimated with $\lambda = 1.96$ in (17), which would give a 95 % confidence interval for a Gaussian random variable. The parameter k was set to 5 [s/rad].

Fig. 6 shows that the estimated force tracked the measured force well, but the confidence intervals seem to be overly pessimistic. The Coulomb friction is, however, a very large disturbance for low velocities. When the robot is moving, the uncertainty in the Coulomb friction is much smaller, as can be seen on the magenta-colored curve in the upper diagram of Fig. 4. Large external forces make the robot move slightly, and this gives significantly tighter confidence intervals than when the robot is still, see, e.g., the z-force at $t = 15$ [s] and $t = 18$ [s] in Fig. 6.

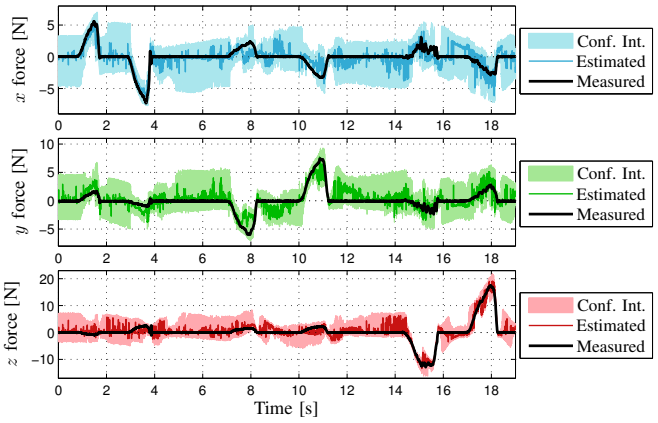

Fig. 6. Data from an experiment where forces were applied to the end effector of the robot. A wrist-mounted force sensor was used to measure a ground truth, and data from this are shown with black lines. The estimated force is displayed together with a confidence interval.

Force direction	x	y	z
With prior	0.67	0.69	0.87
Without prior	1.58	1.22	2.12

TABLE I
Mean absolute error [N] for the experiment in Fig.6.

The prior used in this experiment was $\bar{F} = 0_{6\times1}$ and $R_f = \text{diag}(10\text{[N]}, 10\text{[N]}, 10\text{[N]}, 0.1\text{[Nm]}, 0.1\text{[Nm]}, 0.1\text{[Nm]})^2$, i.e., only small contact torques were expected. The benefit of using the prior is shown in Table I, where the mean absolute estimation errors with and without the use of the prior are listed, showing significantly decreased estimation errors when the prior was used.

C. Screwing assembly task

Estimated and measured forces from an execution of the screwing assembly task are displayed in Fig. 7(a). The forces are given in the coordinate frame illustrated in Fig. 2. It can be seen that the estimated forces tracked the measured forces, at least when the measured forces were non-zero, i.e., during contact operations. When the measured force was zero, however, the estimated force was sometimes a bit wrong, e.g., in the y- and the z-directions around $t = 1.6$ [s]. This estimation error was most likely due to modeling errors, as the robot was moving quite fast in this part of the assembly, but it is known that the motion will be performed in free space when the robot is moving fast, and therefore it is not that important to get a perfect force estimate.

For the assembly task, some of the important forces to detect were the contact forces in the z-direction when the nut was put on the thread. They were correctly detected at $t = 0.2$ [s] and $t = 3$ [s], and the confidence interval was tight at these moments. The screwing was finished when a large side force was detected at $t = 12.6$ [s]; a zoom in on this part

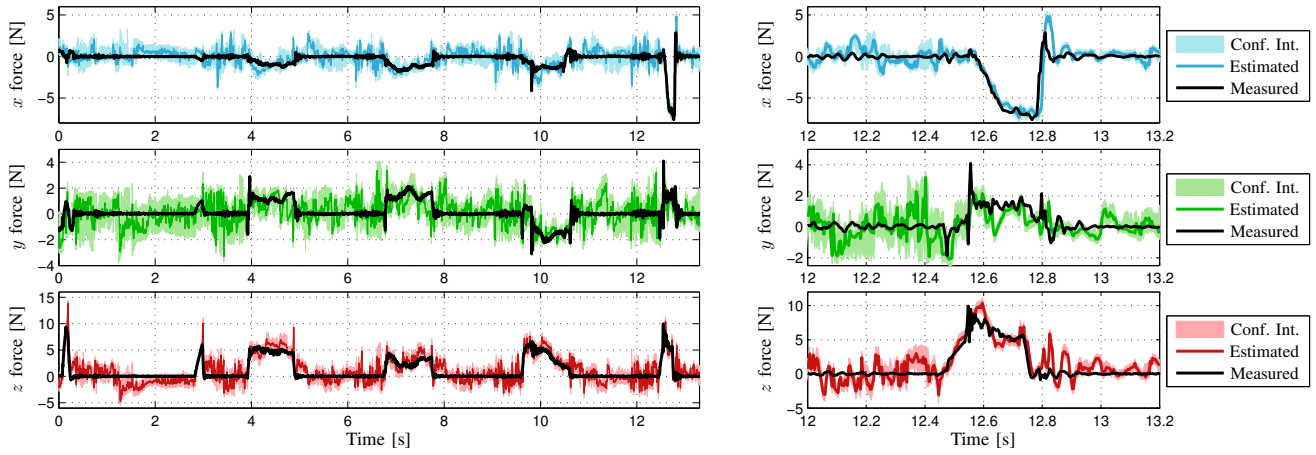

(a) The entire assembly sequence.

(b) Zoom in on the time interval 12–13.2 [s].

Fig. 7. Measured and estimated forces from an execution of the screwing assembly task. Forces are given in the coordinate frame illustrated in Fig. 2. The computed confidence intervals are also shown.

	All samples			Samples when external force was non-zero		
	Percentage of ground truth force/torque measurements within confidence interval					
Wrench component	Static arm	Moving arm	Both arms	Static arm	Moving arm	Both arms
x force	96.2 %	66.7 %	78.6 %	91.2 %	65.2 %	70.7 %
y force	95.6 %	64.4 %	82.4 %	90.0 %	64.2 %	74.8 %
z force	91.7 %	61.9 %	58.4 %	81.2 %	67.5 %	57.2 %
x torque	97.1 %	87.4 %	86.9 %	93.2 %	64.7 %	80.5 %
y torque	97.9 %	80.5 %	72.2 %	95.3 %	85.8 %	78.4 %
z torque	99.8 %	72.8 %	67.7 %	99.6 %	89.9 %	77.5 %
	Mean absolute error					
Wrench component	Static arm	Moving arm	Both arms	Static arm	Moving arm	Both arms
x force	0.60	1.08	0.60	1.05	1.23	0.62
y force	0.63	1.40	0.64	1.14	1.42	0.64
z force	0.91	1.38	1.09	1.33	1.29	1.02
x torque	0.066	0.22	0.095	0.16	0.25	0.096
y torque	0.036	0.15	0.11	0.073	0.13	0.099
z torque	0.014	0.050	0.074	0.036	0.057	0.091

TABLE II

STATISTICS FOR THE FORCE ESTIMATION IN THE SCREWING ASSEMBLY TASK. THE UPPER HALF SHOWS THE PERCENTAGE OF THE SAMPLES WHERE THE MEASURED FORCE WAS WITHIN THE COMPUTED CONFIDENCE INTERVAL. THE LOWER HALF SHOWS THE MEAN ABSOLUTE ESTIMATION ERROR (FORCES IN [N], TORQUES IN [NM]). THE LEFT PART SHOWS STATISTICS FOR WHEN THE ENTIRE ASSEMBLY SEQUENCE IS CONSIDERED, WHILE THE RIGHT PART ONLY CONSIDERS THOSE SAMPLES WHEN THE MEASURED FORCE WAS NON-ZERO, I.E., WHEN THE ARMS WERE IN CONTACT.

of the data is displayed in Fig. 7(b). It can be seen that the force estimate was both quite correct and confident when the forces occurred. Some oscillations can be seen in the force estimate, e.g., in the z-force around $t = 12.2$ [s]. This might be caused by unmodeled disturbances, like cogging torques in the motors or mechanical resonances.

The force estimates presented in Fig. 7 were calculated using data from both arms, and in Table II they are compared to estimates based on only the right (static) or left (moving) arm. It can be seen that the estimate using the static arm, gave the most measurements within the confidence interval. This data should, however only be used to evaluate the quality of the confidence intervals, not the force estimates. When the joints were not moving, the large uncertainty in the Coulomb friction resulted in wide confidence intervals, and hence many measurements were inside the confidence intervals. When the robot was moving, however, the model seemed to be a bit too confident about the estimate. Only looking at the samples when external forces were present did not change much.

The lower part of Table II shows the mean absolute estimation error. Here it can be seen, that when all samples were considered, the estimates from the static arm were the best, but only slightly better than the estimate using both arms. It may be surprising that using only one arm can give better results than using both arms, but when only the static arm was used, the confidence interval was large and usually enclosed the prior, which was $\bar{F} = 0_{6\times1}$ in our example. Hence most estimates were pulled to the prior, which was almost equal to the actual forces for most samples in this sequence. When there were external forces present, which is the situation when the force estimation is really useful, the estimates based on both arms were significantly better than those based on any single arm, as seen in the lower right part of Table II.

For most samples, the confidence intervals of the estimated

contact torques included the zero torque, since the friction torques were very large in comparison to the contact torques. Consequently, larger contact torques would be required for the estimator to be able to detect them reliably.

The prior knowledge used in the assembly scenario was that the contact torques should be quite small, but not zero. This prior knowledge gave a slight increase in estimation performance compared to not using a prior, but it was not dramatic. More is gained in scenarios with only a point contact, where it is known that the contact torques should be zero, such as the experiment in Sec. IV-B.

V. Discussion

Experiments showed that, when the arms were in contact, complementing motor torque data from a moving arm with data from a static arm, significantly improved the quality of the estimated forces. The static arm could not have been exploited with previously published force estimation methods, which do not account for the uncertainty of the Coulomb friction at low velocities.

Using a prior distribution on F with small variance on the contact torques, can be seen as putting a soft constraint on the contact torques. Instead of estimating a 3-DOF force and a 3-DOF torque, the problem is then almost reduced to estimating only a 3-DOF force. The increased redundancy in the problem gives a better force estimate.

The prior distribution on F was chosen to be Gaussian, which is a really crude approximation of the true distribution. Using a Gaussian prior, however, leads to fast calculations, and it has been shown that it also can give a significant performance increase, despite its simplicity. The variance of the prior should be chosen according to process knowledge; how large forces and torques that are expected.

The magnitudes of the contact torques used in this paper were small compared to the uncertainties in the torque estimates, both with and without the use of a prior distribution on F. This was mainly caused by the relatively large disturbances. The contact torques considered were in the order of 0.1–0.3 [Nm], which was in the same order of magnitude as the errors in the gravity torque compensation. Also errors in the Coulomb friction modeling gave disturbances in the same order of magnitude as the contact torques. Estimating forces was more advantageous, as relatively small forces could give rise to relatively large joint torques through long lever arms.

Some of the parameters used for force estimation in this paper were manually tuned, including A and B in (6), the low-pass filter, and the velocity dependence of the Gaussian noise term. This should be possible to do in an automatic fashion, i.e., make experiments of the type presented in Sec. II-C and choose the parameters by optimizing some criterion. This would further simplify the use of the method and it is considered some of the future work.

The force estimation method has so far only been applied to the FRIDA robot. The future work includes to test also on more common industrial robots, which usually are much heavier and stiffer than the light-weight structure of FRIDA.

VI. Conclusions

A method for estimating contact forces from motor torques was presented. The method explicitly models the uncertainty of the Coulomb friction for low velocities and combines data from several joints. The estimate and confidence intervals are calculated by solving convex optimization problems in real time. The method was experimentally verified in a dual-arm assembly operation and validated with data from a wrist-mounted force sensor.

References

[1] ABB Robotics. *Application manual - SoftMove*, 2011. Robot documentation M2004, rev E, RW5.14.

[2] A. Alcocer, A. Robertsson, A. Valera, and R. Johansson. Force estimation and control in robot manipulators. In *Proc. 7th IFAC Symp. Robot control (SYROCO)*, pages 31–36. Wrocław, Poland, Sep. 2003.

[3] A. Blomdell, G. Bolmsjö, T. Brogårdh, P. Cederberg, M. Isaksson, R. Johansson, M. Haage, K. Nilsson, M. Olsson, T. Olsson, A. Robertsson, and J. Wang. Extending an industrial robot controller–Implementation and applications of a fast open sensor interface. *IEEE Robotics & Automation Magazine*, 12(3):85–94, Sep. 2005.

[4] J. De Schutter, T. De Laet, J. Rutgeerts, W. Decré, R. Smits, E. Aertbeliën, K. Claes, and H. Bruyninckx. Constraint-based task specification and estimation for sensor-based robot systems in the presence of geometric uncertainty. *Int. J. Robotics Research*, 26(5):433–455, 2007.

[5] H. Du and S.S. Nair. Modeling and compensation of low-velocity friction with bounds. *IEEE Trans. Control Systems Technology*, 7(1):110–121, 1999.

[6] K.S. Eom, I.H. Suh, W.K. Chung, and S.R. Oh. Disturbance observer based force control of robot manipulator without force sensor. In *Proc. Int. Conf. Robotics and Automation (ICRA)*, pages 3012–3017, Leuven, Belgium, May 1998.

[7] S. Kock, T. Vittor, B. Matthias, H. Jerregård, M. Källman, I. Lundberg, R. Mellander, and M. Hedelind. Robot concept for scalable, flexible assembly automation: A technology study on a harmless dual-armed robot. In *Proc. IEEE Int. Symp. Assembly and Manufacturing (ISAM)*, pages 1–5, Tampere, Finland, May 2011.

[8] J. Mattingley and S. Boyd. CVXGEN: A code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, 2012.

[9] T. Murakami, F. Yu, and K. Ohnishi. Torque Sensorless Control in Multidegree-of-Freedom Manipulator. *IEEE Trans. Industrial Electronics*, 40(2):259–265, 1993.

[10] K. Ohishi. Sensorless force control using h infin; acceleration controller. In *Motion Control Proceedings, 1993., Asia-Pacific Workshop on Advances in*, pages 13–18, Jul.

[11] K. Ohishi, M. Miyazaki, and M. Fujita. Hybrid control of force and position without force sensor. In *Proc. Int. Conf. Industrial Electronics, Control, Instrumentation, and Automation, Power Electronics and Motion Control*, pages 670–675, San Diego, USA, Nov. 1992.

[12] H. Olsson, K.J. Åström, C. Canudas de Wit, M. Gäfvert, and P. Lischinsky. Friction models and friction compensation. *European Journal of Control*, 4(3):176–195, 1998.

[13] M.R. Popovic and A.A. Goldenberger. Modeling of Friction Using Spectral Analysis. *IEEE Transactions on Robotics and Automation*, 14(1):114–122, 1998.

[14] P. Rocco, G. Ferretti, and G. Magnani. Implicit Force Control for Industrial Robots in Contact with Stiff Surfaces. *Automatica*, 33(11):2041–2047, 1997.

[15] J.W.L. Simpson, C.D. Cook, and Z. Li. Sensorless Force Estimation for Robots with Friction. In *Proc. Australasian Conf. Robotics and Automation*, pages 94–99. Auckland, New Zealand, Nov. 2002.

[16] M. W. Spong, S. Hutchinson, and M. Vidyasagar. *Robot modeling and control*. John Wiley & Sons, 2006.

[17] A. Stolt, M. Linderoth, A. Robertsson, and R. Johansson. Force Controlled Robotic Assembly without a Force Sensor. In *Proc. Int. Conf. Robotics and Automation (ICRA)*, pages 1538–1543. St. Paul, Minnesota, USA, May 2012.

[18] Toshiba. Sensor-less compliance control for assembly robots. http://www.toshiba-machine.co.jp/en/technology/tech_catalog/e3.html, Sep. 2012.