

EMG-Based Teleoperation and Manipulation with the DLR LWR-III

Jörn Vogel, Claudio Castellini and Patrick van der Smagt

Abstract—In this paper we describe and practically demonstrate a robotic arm/hand system that is controlled in real-time in 6D Cartesian space through measured human muscular activity. The soft-robotics control architecture of the robotic system ensures safe physical human robot interaction as well as stable behaviour while operating in an unstructured environment. Muscular control is realised via surface electromyography, a non-invasive and simple way to gather human muscular activity from the skin. A standard supervised machine learning system is used to create a map from muscle activity to hand position, orientation and grasping force which then can be evaluated in real time—the existence of such a map is guaranteed by gravity compensation and low-speed movement. No kinematic or dynamic model of the human arm is necessary, which makes the system quickly adaptable to anyone. Numerical validation shows that the system achieves good movement precision. Live evaluation and demonstration of the system during a robotic trade fair is reported and confirms the validity of the approach, which has potential applications in muscle-disorder rehabilitation or in teleoperation where a close-range, safe master/slave interaction is required, and/or when optical/magnetic position tracking cannot be enforced.

I. INTRODUCTION

Surface electromyography (sEMG) for prosthetic hand control has since a long time been used, even in the commercial setting, for the control of prosthetic hands, albeit the control is mostly limited to one or two degrees of freedom and predefined hand postures. sEMG is a totally non-invasive and relatively cheap technique to measure muscular activity. The obtained signal is strongly and stably related to the force exerted by the measured muscle(s) thanks to the electrical activity of ensembles of motor units. The signal is noisy, but the fact that many hundreds of motor units combine together makes it is relatively insensitive to small influences, to the point that sEMG is widely demonstrated in literature [1], [2], [3], [4], [5] as an excellent means to control dexterous hand prostheses.

The use of EMG for copying arm movement however is a totally different problem: in this case, the human skeletomuscular system does not exert a certain force to hold or grasp an object; rather, muscles are activated to move an arm to or track a desired Cartesian position. In most finger movement settings, the precise finger position is secondary to the exerted force; in arm movement, it is rather the other way around. The nature of the sEMG signal may therefore indicate its inapplicability for this task; nonetheless,

This work is partially supported by the European FP7-Project THE Hand Embodied (FP7-IST-248257).

The authors are with the Institute of Robotics and Mechatronics, German Aerospace Center, D-82234 Oberpfaffenhofen, Germany {joern.vogel,claudio.castellini,smagt}@dlr.de



Fig. 1. The system during a live demonstration at the *Automatica 2010* trade fair (www.automatica-munich.com).

interesting results have already been obtained by Artemiadis and Kyriakopoulos [6], [7], [8].

In this work we follow the approach outlined in that work, at the same time lifting some of its limitations. In particular, we address the low number of controlled degrees-of-freedom (3), the low generalisability (a detailed model of the arm of the human subject is required) and the use of an industrial robotic setup, which makes the whole system unsuitable for close cooperation with human beings. In particular we introduce a method which matches a 9-dimensional sEMG signal to a 6-DoF end-effector position plus 1-DoF grasp force. The system, which has been demonstrated online, requires a short training period and can then be used for a prolonged period of time to grasp and move objects using the sEMG signal as user interface only. No precise positioning of the sEMG electrodes, nor any model of the human arm is required; the adaptivity of the machine learning approach used also automatically incorporates compensation for muscle fatigue, an issue which usually requires non-trivial handling [7].

Even though the use of the system for precise position teleoperation is limited, the system can be used for rehabilitation purposes as well as for impedance-controlled, stiffness-based teleoperation. Furthermore, the current system allows, where necessary, for increased accuracy by optimisation of the data acquisition methods.

II. APPROACH

The system, as described here, was demonstrated at the *Automatica 2010* trade fair in Munich, Germany (see www.




Fig. 2. An abstract block-diagram representation of the system.

automatica-munich.com). All data presented in this paper was collected live during the fair.

A picture of the setup in action is visible in Figure 1. The robot is mounted on a pedestal in a right-arm like posture. A table is placed within the workspace of the robot, to allow the user to pick up and put down objects; a soft ball is used to this end. Figure 2 depicts a simplified schematic overview of the core elements of the system, which consists of two independent parts: (1) an EMG-decoder which calculates Cartesian wrist position and orientation as well as grasp force from measured muscle activity, and (2) the DLR Light-Weight Robot III equipped with the DLR-HIT Hand II, which performs the decoded motion and grasps. Additionally, a visual tracking system is integrated in the setup in order to gather the human arm position for training the machine learning algorithm. As is customary in (supervised) machine learning, the system operation consists of two phases: the training phase, during which an EMG-to-(arm/hand/force) map is built; and the prediction phase, when the map is employed to predict new, previously unseen arm/hand/grasping configurations. Before entering the training phase, the subject is equipped with EMG electrodes to record the muscular activity.

To build the map, EMG data needs to be acquired, as well as the real position, orientation and grasp force of the user's hand. Therefore a tracking marker is fixed to the upper side of the subjects wrist and a rubber ball equipped with force sensitive resistor is given to the subject (see Figure 3). Once the mapping is created, the EMG signals can be used to directly control the robotic system.

A. Data acquisition

Muscular activity is gathered using nine OttoBock MyoBock 13E200 surface EMG electrodes (www.ottobock.com). The electrodes already provide an amplified, bandpass-filtered and rectified signal, eliminating the need of further processing onboard the card and/or the computer (their usefulness was already demonstrated at least in [5], [9]). They are connected to a wireless DAQ card sampling the EMG signals at 100 Hz.

There are three sets of three electrodes. Each set is tied to a velcro elastic band and roughly uniform spacing of the




Fig. 3. Schematic view of the subject's arm equipped with the EMG electrodes, the motion tracking marker and the ball with the FSR on top.

electrodes is visually enforced. The bands are placed on the subject's forearm about five centimeters below the elbow, on the upper arm midway between elbow and shoulder, and on the shoulder (see Figure 3). No precise positioning of the electrodes is enforced — this is a great simplification of the operations and has already been demonstrated effective, even on amputees [9]. The exerted grasp force is measured with an Interlink *Standard 400 FSR* force-sensing resistor (see www.interlinkelectronics.com). The standard amplification circuit connected to the FSR returns a voltage signal which is univocally (logarithmically) related to the force applied to its surface. The above wireless DAQ card is used to digitise this signal, too, making the whole setup rather easy to wear and take away. The FSR is mounted on a rigid rubber ball and the subject is instructed to press it to teach a grasping signal to the system.

Motion capture is enforced by a Vicon MX (www.vicon.com) motion tracking system. A Vicon "rigid object" consisting of 4 passive markers rigidly connected to one another is fixed to the subject's wrist, and six near-infrared cameras use it to reconstruct the object's position (x, y, z in centimeters) and orientation (α, β, γ in radians) in real time. The coordinates are relative to an inertial frame set up during the Vicon calibration phase. The Vicon has a sampling rate of 200 Hz and generates accordingly a UDP stream of data.

The resulting global data stream (from the Vicon and acquisition card) is received by a standard desktop machine equipped with Matlab and synchronously subsampled at 25 Hz, and its moving average over 400 milliseconds (10 samples) is evaluated.

B. Machine Learning

After the data acquisition is concluded in the training phase, the data set is organised into samples (each sample consisting of the 9 values of the EMG electrodes) and 7

target values for each sample (the position/orientation of the hand and the grasping force). The resulting (sample, target) pairs are used to train a Support Vector Machine (SVM, see [10]) with Gaussian kernel, with the purpose of building a function mapping the EMG values to each target. One of the most popular machine learning nowadays, SVMs build an approximation of the underlying function as the sum of a finite, and hopefully small, number of Gaussian functions, centered on a subset of the training samples called Support Vectors. Since in our case every training set would consist of several thousands of samples, we subsampled sequentially the training set in order to have always 750 and, later on, 1000 samples available. The training samples are normalised by subtracting the means and dividing by the standard deviations, dimension-wise; 10-fold cross-validation and grid-search are conducted in order to find the best SVM hyperparameters. Lastly, a regression model for each target value is created using the 750 or 1000 training samples at once. These 7 models are the map used in prediction.

C. Robotic Setup

The position and grasp decoding from the EMG signals is used to control the DLR Light-Weight Robot III (LWR III) with the DLR-HIT Hand II attached. The LWR III [11] is a seven degree of freedom (7 DoF) robotic arm weighing 14kg and able to lift payloads equal to its own weight. The kinematics of the robotic arm allow to replicate a reasonable large part of the workspace of a human arm. The LWR III is equipped with joint torque sensing in each joint, which makes it possible to realise special soft-robotic features such as Cartesian Impedance Control and collision detection and reaction [12], [13]. These features are essential when operating in rather unknown environments and physical contact with rigid objects needs to be established. The active compliance of the robot provided by the impedance controller copes with these uncertainties and thereby provides stable behaviour of the system. In addition to that, the integrated torque sensors allow detecting external forces that appear whenever the robot establishes a physical contact. Depending on the magnitude of the force, the robot can react with different strategies and thereby provide safety to the operator as well as to the robotic hardware (see [14], [15]). All these soft-robotic features are embedded in a human-friendly state-based control architecture [16].

A simplified schematic overview of the robot control architecture is depicted in Figure 4. From the EMG signals the human hand position, orientation, and force is decoded and send to the robot controller via UDP. The UDP stream is received by the high-level task execution layer of the robot controller. This layer operates at a rate of $\approx 100Hz$ and allows defining complex tasks in a state-based manner illustrated by the Hybrid task state machine in the block diagram. The low-level robot controller runs on a VxWorks realtime machine at a rate of $1kHz$. Communication between the two layers is realised via the ARDNet Interface [17]. The core component of the Robot Control kernel is the Cartesian Impedance Controller (CIC) which receives a desired Carte-

sian frame as input and calculates the desired joint torques (further details on this control scheme can be found in [12]). As Cartesian human wrist position and orientation are directly decoded at a rate of 20Hz it is necessary to employ an interpolation to fit the $1ms$ cycle of the robot controller and generate a smooth robot motion. Within the interpolation the maximum achievable translational and rotational velocity is limited to $\leq 0.25m/s$ and $\leq 0.3rad/s$, respectively. The scaling between human motion and robot motion was kept 1:1 in this sets of experiments, though it would easily be possible to up or downscale the motion when it is exerted by the robotic system. To provide additional safety to the operator as well as to the robotic hardware, known obstacles in the workspace, such as the table the user operates on, and the pedestal the robot is mounted on, are internally represented in a virtual environment. This consist of virtual planar walls which the robot is unable to pass. To achieve this, the virtual walls create a repelling Cartesian force which then is transformed via the Jacobian transposed in joint torques. These are added to the desired joint torques of the Cartesian impedance controller.

In addition to wrist position and orientation, the human grasp force is also decoded from the EMG signals. This then is commanded to the DLR-HIT Hand II. The DLR-HIT Hand II is a five-finger robotic hand with three degrees of freedom in each finger. Similar to the LWR III, the hand is equipped with joint-torque sensors in each of its 15 joints, allowing it to be used with joint impedance control. In our implementation, the hand performs a preprogrammed grasp motion whenever the decoded grasp force exceeds a certain threshold. Furthermore, the joint stiffness of each finger is increased proportional to the decoded grasp force and thereby enabling the operator to increase or decrease the force exerted on a grasped object. Dropping below the predefined force threshold causes the hand to release the grasp again.

D. Experimental protocol

Three male, healthy subjects (age 28, 38 and 44) joined the experiment and controlled the system over a time span of 4 days. (Actually, a fourth subject, namely a TV reporter, joined the demonstration, but her training set was very simple, limited to arm positioning, due to the lack of time. See the attached video and Section III-A for more details). A particularly interesting point is that of task-oriented training, which was enforced via direct teleoperation between the motion capture system and the robotic system during the training phase.

In particular, in the non-task-oriented (NTA) modality, the subject would be equipped with the electrodes, FSR-ball and Vicon object. He would then be placed within the reach of the Vicon cameras (about 8 cubic meters) and asked to relax his arm for a few seconds, during which the initial position and orientation of the rigid object on the wrist would be gathered—this enables mapping muscle activity to a body-solidal frame rather than to the Vicon’s absolute frame. The




Fig. 4. Schematic overview of the system and its core components

subject performed then random arm movements and grasps all over the reach space.

The allowed hand movement speed was willfully kept relatively low, topping a maximum of about 0.5 m/s. Actually, the whole idea of EMG-to-position mapping relies on gravity compensation: one end-effector position is characterised by one (or more) isometric/isotonic configuration(s) of the arm muscles, and this in turn corresponds to a single, quasi-stationary muscular activity (and EMG signal) pattern. Therefore low end-effector speeds are required in order for the muscular activity to be only negligibly affected by acceleration. The data acquisition usually lasted until about 5000 to 7000 samples were gathered (3.5 to 4.5 minutes). No subject reported fatigue or cramps during the acquisition phase.

On the other hand in the task-oriented (TA) modality, the subject was in sight of the robot’s workspace, including the robot itself, the table and the objects to manipulate. The hand’s position, orientation and grasp force as detected by the Vicon and FSR was directly transmitted to the robotic system, so that the subject could *teleoperate*. Meanwhile, the subject’s muscular activity was gathered by the EMG setup. In this modality the subject would actively try and grasp the ball from its pedestal, move it around (such as, e.g., stretching the arm along the coronal plane with the ball in hand) and put it back on it, or release it somewhere else or in someone else’s hands.

Figure 5 shows the 3D motion capture plots of two typical training sessions in the NTA/TA modalities. The difference in the sample distribution is apparent.

Once the map is built and validated, the subject enters the loop again and now without use of the tracking system and FSR data the EMG readings are directly mapped onto hand positions, orientations and grasp forces. Decoded position/orientation/force commands are then sent to the robotic system. The sEMG gathering setup (electrodes, velcro bands and the digital acquisition card) is light and can easily be carried around by the user, for instance in a small bag (check Figure 1 again). Since the muscular activity of the shoulder /

arm / hand is largely independent of that required by walking the subject is allowed to move freely and as the robotic system is operated in a human friendly control architecture, the user may as well interact with it while being in control, as seen in the attached video.

III. EXPERIMENTAL RESULTS

The system was practically demonstrated during the Automatica 2010 robot trade fair in Munich, Germany, over a period of 4 days (see also <http://www.automatica-munich.com>). The subjects trained and tested the system several times each day, generating a total number of 38 data sets and corresponding models. The TA phase was activated during the last day (10 datasets and models).

A. System demonstration

The video attached to this paper shows an exemplar prediction phase (NTA phase): one of the subjects controls the robot system and uses it to grasp/release a soft ball, remove it from the pedestal, carry it around and place it back. Although a noticeable delay is seen between the subject’s operation and the robot response, and the overall position of the arm does not always reflect the subject’s one, the whole qualitative appearance of the coordinated human/robotic motion is good, and after a few failed attempts the subject succeeds in grasping the ball and placing it back on the pedestal.

All subjects reported the same subjective impression of motion accuracy and ease of operation; this is due to the adaptability of the supervised machine learning approach. Additionally, a reporter from the n-tv Germany Broadcasting company (<http://www.n-tv.de>) trained the system in a quicker fashion, limited to a few simple arm movements. In the end she reported the same feeling of easiness and compliance by the robot. A short clip of the demonstration is available, at the time of writing, on the tv-station’s website itself (<http://www.n-tv.de/mediathek/videos/technik/Automatica-zeigt-Roboter-fuer-zu-Hause-article916558.html>, starting approximately at second 40 of the clip).




Fig. 5. 3D plots of non-task-oriented (left) and task-oriented (right) motion capture data sets. The densely-sampled zone in the right panel corresponds to the position of the object to be grasped in the robot’s workspace. The red/green/blue axes indicate the reference frame (x, y, z axes).

Better robot response and ease of use was reported by the subjects in the TA modality, with respect to the NTA. Especially, the lack of an input space sampling strategy was particularly frustrating in the NTA modality, since the subjects needed to wave randomly with little relation with what they would do later on. This resulted in frequent unexpected decrease in the decoding quality when getting close to the grasping zone. This drawback disappeared in the TA modality, as one would expect.

B. Numerical evaluation

In this section the performance of the models obtained during the demonstration is evaluated. The evaluation is accomplished by testing the SVM model on the samples *not* included in the training set (normalised using the statistics of the training set), and reporting the Mean-Squared-Error normalised over the testing samples variance (NMSE). It must be noted that this is an approximation of the performance actually obtained during the demonstration—the latter cannot be numerically evaluated any longer, since no ground truth is available (the subjects would step outside the Vicon detection zone during the demonstration). We expect the real performance to be numerically slightly worse than what we show here; nevertheless, notice that evaluating the real performance would also require a subjective success metric, such as, e.g., the number of successful grasp/release trials, the time needed to grasp, etc.

Figure 6 reports the results for each target value and for each model, in chronological order. Consider the Figure. For the position coordinates (upper panel) the overall NMSE values are 0.0982 ± 0.0591 , 0.0789 ± 0.0455 and 0.0674 ± 0.0510 (mean / standard deviation); for the orientation, 0.1689 ± 0.1518 , 0.0989 ± 0.0864 and 0.1249 ± 0.1208 and for the grasping force 0.0321 ± 0.0341 . Models 29 to 38, built during the TA phase, show an apparent better performance; in fact for these models only, the NMSE values are remarkably better than those obtained for NTA models (see Table I).

TABLE I
PREDICTION MEAN-SQUARE-ERROR NORMALISED W.R.T. THE TARGETS’ VARIANCE, MEAN VALUES \pm ONE STANDARD DEVIATION, FOR NTA AND TA PHASES.

	NTA	TA
x	0.1101 ± 0.0615	0.0647 ± 0.0364
y	0.0885 ± 0.0482	0.0519 ± 0.0214
z	0.0741 ± 0.0573	0.0485 ± 0.0178
α	0.1817 ± 0.1644	0.1330 ± 0.1081
β	0.1130 ± 0.0961	0.0593 ± 0.0242
γ	0.1362 ± 0.1292	0.0934 ± 0.0921
force	0.0357 ± 0.0386	0.0223 ± 0.0128

In general, these values denote a remarkably low error also in physical terms; for instance, in the TA phase, the non-normalised squared MSE for x, y, z is in turn 3.12 ± 0.54 , 3.14 ± 0.95 and 3.54 ± 0.72 cm. These values are comparable with those reported in [6].

Further on, it is interesting to consider the effect of task-orientedness on the distribution of the error in the input space. Figure 7 shows two typical cases. Model 20 (NTA, NMSE 0.2953, 0.1582, 0.0799 for x, y, z) has a uniformly distributed error, while for model 38 (TA, 0.1418, 0.0807, 0.0576) the error is mostly seen outside the zone where the object to be grasped was located, which appears as a darker, almost black cloud of well-predicted points. Lighter, and therefore less accurately predicted “clouds” are clearly seen around the path the arm followed from the resting position to the fully stretched (right-hand side) and on the left of the zone where the ball was. Figure 8 shows an example of real and predicted position, namely the x coordinate for these models, i.e., 20 and 38.

IV. CONCLUSIONS

In this paper we reported about a robotic arm/hand system controlled through sEMG-measured human muscular activity. The control has such a level of accuracy that a small object can be repeatedly grasped, carried around, and released.




Fig. 6. Prediction Mean-Square-Error normalised w.r.t. the targets' variance, for each dataset and target dimension. Models from 1 to 28 have been obtained in the NTA modality, 29 to 38 with the TA modality.

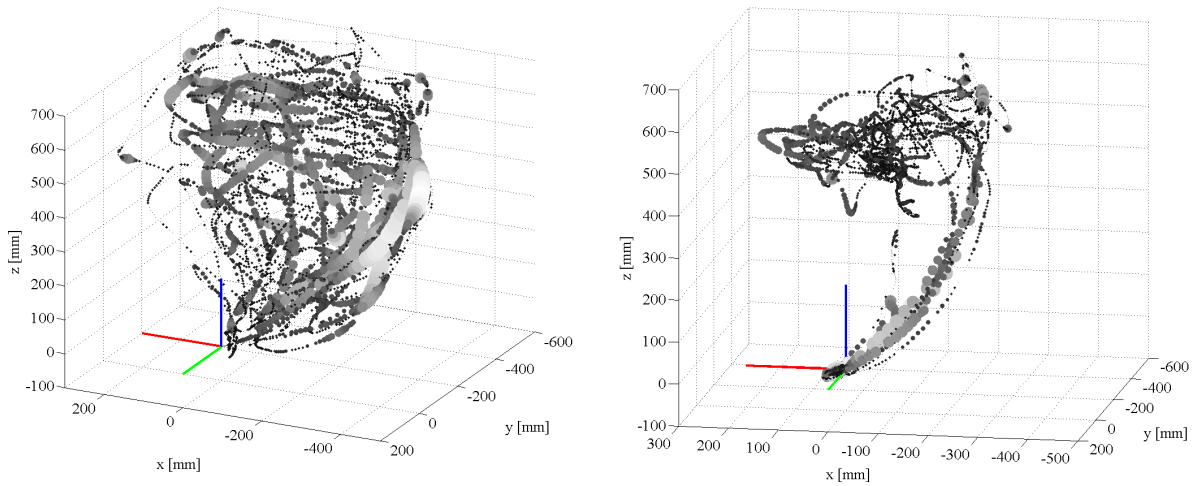


Fig. 7. 3D plots of non-task-oriented (left) and task-oriented (right) squared Root-Mean-Error, model 20 (NTA, left) and 38 (TA, right). Larger and lighter markers denote higher error.




Fig. 8. Real and predicted x coordinates for models 20 (upper panel) and 38 (lower panel).

The system was practically demonstrated during a robotic trade fair in 2010. Muscle activity is gathered using nine surface electromyography electrodes. No precise placement of the electrodes is required, and no model of the human arm/hand is employed, making the system essentially subject-independent — in fact, three subjects demonstrated the system without noticeable performance differences. A fourth subject did a simpler demonstration without any previous knowledge about the system.

A standard machine learning method (namely, a Support Vector Machine) is used to build a point-to-point map between muscle activity and hand position/orientation/grasping force. The map relies on gravity compensation and rather low movement speed, which enforces a many-to-one relationship between sEMG signals and position. Numerical (offline) evaluation indicates in a few centimeters the precision that can be obtained by the system; orientation and force guessing have similar, although slightly worse, precisions.

It was noted that the performance of the system is considerably higher when a task-oriented training modality is used. This is due to finer sampling of the input space in the zone of interest (namely, where the grasping mostly happens) which leads to smaller error rates *where it is required*.

In [7] a detailed report is given about how to avoid the well-known time variance of the sEMG signal, in particular as far as muscle fatigue is concerned — a problem that we did not notice. The reasons of this improvement could lie in the choice of the sEMG features, in the reliability of the electrodes or even in the reciprocal adaptation of the human subjects to the system. In fact, all subjects reported a feeling of "learning to control the arm" as the testing phase would proceed. Task-orientedness seems to be essential from this point of view, too.

The DLR Light-weight Robot III we used in this paper is operated in impedance control, which is essential in applications in which the environment is unknown. Furthermore the human-friendly control architecture used in combination with the robot enables the operator to be within the workspace of the robot and interact with it. This feature makes the system suitable for a variety of applications in which physical interaction between humans and robots occurs.

Future work

Further investigation into the sEMG signal is envisioned, in particular as the end-effector reaches the grasping zone. In that case the muscular system is expected to stiffen up. This could be used to estimate 3D stiffness using sEMG, a problem which is still largely open and whose solution would have applications in a number of fields of robotic (e.g., remote surgery or high-accuracy teleoperation).

The assumption of slow movement can probably be loosened if a more sophisticated form of robotic control is enforced, namely, considering estimating the end-effector *velocity* as well as the position, and then using a hybrid position/velocity robot controller. This is also subject to further research.

Lastly, the system as described and demonstrated in this paper is probably not directly optimally usable for generic teleoperation—indeed, more accurate ways of estimating the end-effector position rather than sEMG can be found; however this setup might be of great use when magnetic tracking cannot be used for training. An even more interesting future application is rehabilitation of muscular-disorder patients, in which a weak or distorted sEMG signal could be used to train the system and let the patient see the arm move as desired. Such procedures are well known to dramatically shorten rehabilitation effort.

REFERENCES

- [1] S. Ferguson and G. Reg Dunlop. Grasp recognition from myoelectric signals. In *Proceedings of the Australasian Conference on Robotics and Automation, Auckland, New Zealand*, pages 83—87, 2002.
- [2] A.D.C. Chan and K.B. Englehart. Continuous myoelectric control for powered prostheses using hidden markov models. *Biomedical Engineering, IEEE Transactions on*, 52(1):121–124, Jan. 2005.
- [3] P. A. Parker, K. Englehart, and B. Hudgins. Myoelectric signal processing for control of powered limb prostheses. *Journal of electromyography and kinesiology*, 16(6):541–548, 2006.
- [4] S. Bitzer and P. van der Smagt. Learning EMG control of a robotic hand: towards active prostheses. In *Proc. IEEE International Conference on Robotics and Automation*, pages 2819–2823, 2006.
- [5] C. Castellini and P. van der Smagt. Surface EMG in advanced hand prosthetics. *Biological Cybernetics*, 100(1):35–47, 2008.
- [6] P. K. Artemiadis and K. J. Kyriakopoulos. EMG-based control of a robot arm using low-dimensional embeddings. *IEEE Transactions on Robotics*, 26(2):393–398, 2009.
- [7] P. K. Artemiadis and K. J. Kyriakopoulos. An emg-based robot control scheme robust to time-varying EMG signal features. *IEEE Transactions on Information Technology in Biomedicine*, 14(3):582–588, 2010.
- [8] P. K. Artemiadis and K. J. Kyriakopoulos. A switching regime model for the emg-based control of a robot arm. *IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICSPART B: CYBERNETICS*, 41(1):53–63, 2011.
- [9] C. Castellini, E. Gruppioni, A. Davalli, and G. Sandini. Fine detection of grasp force and posture by amputees via surface electromyography. *Journal of Physiology (Paris)*, 103(3–5):255–262, 2009.
- [10] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In D. Haussler, editor, *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, pages 144–152. ACM press, 1992.
- [11] A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger. The DLR lightweight robot: design and control concepts for robots in human environments. *Industrial Robot: An International Journal*, 34(5):376–385, 2007.
- [12] A. Albu-Schäffer, C. Ott, and G. Hirzinger. A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *The International Journal of Robotics Research*, 26(1):23, 2007.
- [13] S. Haddadin, A. Albu-Schäffer, A. De Luca, and G. Hirzinger. Collision Detection and Reaction: A Contribution to Safe Physical Human-Robot Interaction. *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS2008)*, Nice(France), 2008.
- [14] S. Haddadin, M. Suppa, S. Fuchs, T. Bodenmüller, A. Albu-Schäffer, and G. Hirzinger. Towards the Robotic Co-Worker. In *submitted to: International Symposium on Robotics Research (ISRR2007)*, Lausanne, Switzerland, 2009.
- [15] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger. Requirements for Safe Robots: Measurements, Analysis and New Insights. *International Journal on Robotics Research*, 28(11-12):1507–1520, 2009.
- [16] S. Haddadin, Parusel S., Vogel J., R. Belder, Rokahr T., A. Albu-Schäffer, and G. Hirzinger. Holistic design and analysis for the human-friendly robotic co-worker. *accepted at: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS2010)*, Taipeh(Taiwan):Invited paper, 2010.
- [17] Bäuml B. and G. Hirzinger. When hard realtime matters: Software for complex mechatronic systems Robotics and Autonomous Systems. *Robotics and Autonomous Systems*, 56, 2008.