

Platform Portable Anthropomorphic Grasping with the Bielefeld 20-DOF Shadow and 9-DOF TUM Hand

Frank Röthling, Robert Haschke, Jochen J. Steil, and Helge Ritter
Neuroinformatics Group, Faculty of Technology, Bielefeld University
P.O.-Box 10 01 31, D-33501 Bielefeld, Germany
Email: {froethli, rhaschke, jsteil, helge}@techfak.uni-bielefeld.de

Abstract—We present a strategy for grasping of real world objects with two anthropomorphic hands, the three-fingered 9-DOF hydraulic TUM and the very dextrous 20-DOF pneumatic Bielefeld Shadow Hand. Our approach to grasping is based on a reach–pre-grasp–grasp scheme loosely motivated by human grasping. We comparatively describe the two robot setups, the control schemes, and the grasp type determination. We show that the grasp strategy can robustly cope with inaccurate control and object variation. We demonstrate that it can be ported among platforms with minor modifications. Grasping success is evaluated by comparative experiments performing a benchmark test on 21 everyday objects.

I. INTRODUCTION

It is widely recognized that grasping marks a keystone for sensorimotor intelligence and will certainly be required for future service robots. Therefore, recently a number of sophisticated multi-fingered artificial hands have been developed, which in principle have the necessary mechanical dexterity to carry out a large variety of everyday tasks [1]–[4]. On the algorithmic side, however, robust and stable grasping of a large variety of a priori unknown objects is still a major challenge even for the best artificial robot hands available.

Traditionally, the robot grasping process is divided into two stages: at first, suitable grasping points on the object are determined and, secondly, a robot hand posture is computed via inverse kinematics to match these points with the fingertips [5]. Correspondingly, a grasp has formally been defined in a solely kinematic fashion as a set of contact points on the object surface together with friction cone conditions, whereas being independent of the robot hand under investigation [6]. The computation of a stable or even optimal grasp has led to many sophisticated algorithms [7], [8], but remains a computationally demanding task mostly solved by quadratic optimization [7], [9]–[11]. To apply this strategy, the object geometry has to be known exactly to find and optimize the contact points. For real world execution of the grasp, the object therefore has to be visually (or otherwise) located and the movement guided with high precision. This grasp strategy is suitable for industrial robots which are specialized to highly structured working environments, but it is difficult to extend to learning and grasping of real world objects under uncertain visual localization.

The challenge for the realization of service robots working in everyday life domains, however, is to achieve the ability to handle a broad range of tasks in hardly structured environments, to adapt to new situations, and to grasp new

objects which are not exactly known beforehand. To realize such flexibility and robustness, several authors have proposed to organize grasping in a more holistic fashion loosely motivated by the way humans grasp, who select an object-specific pre-grasp posture as one of a few prehensile hand postures [12]. Afterwards, the grasp itself is carried out by comprehensively closing the fingers and evaluating the tactile feedback. Though differing in detail, such strategies do not compute explicit contact points and inverse kinematics solutions.

In [13], objects are modeled as a set shape primitives, such that one of four distinct pre-grasps for the Barret Hand can be selected. Using the grasp simulator "GraspIt" [14] the position of the hand relative to the object is systematically varied and grasp success is evaluated using a standard stability measure. This method optimizes the grasp position, but does not learn or optimize the pre-grasp postures themselves and an exhaustive search in the space of pre-grasp postures would be infeasible for hands with higher degrees of freedom. [15] and [16] combine controller primitives to achieve a reach-grasp behavior. In coordination with two different reach controllers, one grasp controller realizes a three finger grasp, while a second one combines two physical contacts into a "virtual finger". The correct instantiations of the controllers can autonomously be learned, by associating general visual features such as blob height and width in a reinforcement framework [17]. Though this framework provides impressive results, it seems that autonomous exploration for learning to grasp a larger number of objects will be too time consuming to be executed on the real robot hardware.

To reduce the need to explore very large search spaces, an interactive imitation based learning approach is appealing. Consequently, we have previously proposed to enable imitation grasping [18] in the context of a long-term research project [19] aiming at the realization of a robot system that is instructable by speech and gesture, has visual capabilities, attentive behavior, and can execute grasping actions [20], [21] (see Fig. 1). In this framework, we use a universal, biologically motivated grasp strategy, which relies on a 3D localization of the object, executes a reaching movement, and finally grasps an object employing appropriate pre-grasp and target grasp postures. In [18], we have proposed an object-specific grasp selection based on the observation of a human instructor's hand to reduce complexity of the selection process and enhance grasp success. The present



Fig. 1. Interaction with the robot system using speech and gesture. On the left the modified three-fingered TUM Hand is shown.

paper focuses on comparative results obtained with a new, much more dextrous anthropomorphic hand containing 24 joints actuated by pneumatic muscles. We describe the new robot hand and the different control concepts, report on the migration process, and compare grasping results on both platforms showing that with minor adaptations our strategy is portable across significantly differing hardware platforms.

II. BIELEFELD TUM AND SHADOW HAND SETUPS

A. GRAVIS Robot System and TUM Hand

The GRAVIS robot system [20] combines visual attention and gestural instruction with an intelligent interface for speech recognition and linguistic interpretation to allow multi-modal task-oriented instructions. For manipulation tasks this setup employs a standard 6-DOF PUMA manipulator operated with the real-time RCCL-command library together with a 9-DOF dextrous robot hand developed at the Technical University of Munich (TUM). The hand consists of three identical, approximately human-sized fingers driven by a hydraulics system. To improve on the original symmetrical arrangement of the fingers, we reconfigured the hand to be more human-like. It now features a palm, a thumb, and two opposing fingers in order to allow a larger variety of two- and three-finger grasps (see Fig. 1). The fingertips are equipped with custom built force sensors to provide force feedback for control and evaluation of grasps.

Because the hand does not provide joint angle sensors, posture control has to be realized indirectly relying on piston potentiometers and pressure sensors, both located at the base station of the hydraulics system. To this end, we convert joint angles to piston potentiometer values applying a fixed transform which was determined experimentally and independently for each joint. These computed potentiometer values serve as targets for PID controllers actuating the finger joints to the desired posture. Additionally, we have to cope with hysteresis and non-linearities due to the long distance of 2.5m between the base station and the finger pistons, and we face sticking and sliding effects caused by return springs integrated in the finger pistons. Nevertheless, we achieve




Fig. 2. The newly developed robotics setup comprises the 20-DOF Shadow Dextrous Hand, a 7-DOF robot arm and a 4-DOF stereo camera head.

an accuracy of about 2 degrees in every joint, which is not enough for a reliable inverse kinematics based position control, but allows for a sufficient positioning of the fingers to realize suitable grasp postures.

B. Bielefeld Shadow Dextrous Hand

We currently reconstitute the described imitation learning setup based on a new platform consisting of a redundant 7-DOF Mitsubishi PA-10 robot arm, the 20-DOF Shadow Dextrous Hand, and a 4-DOF active stereo camera head based on Helpmate hardware (Fig. 2). The robot arm is actuated by a robot server incorporating a security concept based on an internal model of the whole setup, and a real-time path planning algorithm based on neural networks [22].

The Bielefeld Shadow Hand is a product of the Shadow Robot Company [1] and is available as a prototype since end of 2004. Fig. 3 summarizes the finger kinematics. It shows a photograph of the human-like sized real hand and its kinematical model. Joint axes are visualized as black arrows within the transparent links of the model. The distal joints of the four fingers are coupled passively to the middle joint, such that the angle of the middle joint is always greater than or equal to the angle of the distal joint. Hence, the finger joints allow almost human-like movements as described in [23]. To endow the thumb with a similar dexterity and to allow the opposition of the thumb to all fingers, five independently controllable joints are supplied, two of them combined in the metacarpophalangeal joint and two others combined to approximate the trapeziometacarpal saddle joint of the human thumb. The little finger has an extra joint located in the palm. The hand is also equipped with 2 DOF in the wrist (not shown in Fig. 3), which allow a flexion/extension as well as abduction/adduction movement of the whole hand. Altogether, the hand comprises 24 joints, 20 of them actively controllable.

Each joint is actuated by an antagonistic pair of McKibben style pneumatic muscles, which have a high force-to-mass ratio. All muscles are packed densely in the lower forearm (shown in Fig. 2) and the joints are actuated by means of tendons routed through the wrist and hand. The air flow in




Fig. 3. Real Shadow Hand (left) compared to its kinematic model (right). Joint axes are visualized as black arrows.

and out of the muscles is controlled by 80 miniature solenoid on-off valves – one inlet and outlet valve for each muscle. An important advantage of artificial muscles is their inherent and variable compliance allowing safe operation, especially in direct contact and in interaction with humans.

On their palmar side, the phalanges are covered by a layer of formable polyurethane “flesh” which is slightly elastic and has a high friction coefficient providing good adhesion. To facilitate grasping of small objects, like matches and needles, the fingers include thin polycarbonate fingernails. The most innovative feature of the Shadow Hand, however, is the provision of a total of 186 force sensors. 34 of these are distributed on each fingertip giving a touch resolution of approx. seven sensors per cm^2 . Additionally, two texels (touch pixels) cover the palmar side of the middle and proximal phalanges of each finger. The tactile sensors are build from a three-dimensionally curved electrode covered by a thin layer of Quantum Tunneling Composite (QTC), which changes its resistance as a function of applied pressure. QTC has an exponential response characteristic, combining a high initial sensitivity with a wide dynamical range that only saturates at considerably stronger forces.

The hand is also equipped with a complete set of internal sensors measuring current joint angle position as well as muscle air pressure.

C. Joint Control for Shadow Dextrous Hand

While pneumatic actuators are well known and their control has been studied mainly for single McKibben muscles [24], the simultaneous control of a large number of cooperating finger actuators poses new challenges. The inevitably complex tendon routing in the hand contributes friction and tends to amplify the well known nonlinear and hysteresis effects, so that a modeling scheme like that proposed in [25] becomes difficult to apply and would have to be carefully adapted to each muscle.

In order to move a single joint, the controller has to provide two control outputs to drive the valves of the antagonistic muscle pair. Hence most standard control approaches that focus on a single control variable cannot be applied directly. The key idea of our controller is the combination of two control variables, joint position and joint stiffness, which is motivated by the observation that the pressure difference




Fig. 4. Schematic view of the mixing controller allowing simultaneous stiffness and position control.

correlates with the joint position while the pressure sum correlates with the stiffness of the joint [26]. This means that both joint position and joint stiffness can be adjusted independently. To this end, we use a suitable mixing matrix to compute the two control outputs from both a *joint position error* $\Delta\theta$ and a *stiffness error* ΔS :

$$\begin{pmatrix} t_{\text{flex}} \\ t_{\text{ext}} \end{pmatrix} = \begin{pmatrix} K_{\theta} & K_S \\ -K_{\theta} & K_S \end{pmatrix} \begin{pmatrix} \Delta\theta \\ \Delta S \end{pmatrix} \quad (1)$$

As a stiffness measure S , we use the sum of the pressures in both muscles. The control outputs t_{flex} and t_{ext} are the time periods used to open the valves of the flexor and extensor, respectively. Positive periods open the inlet valves, negative periods open the outlet valves. Currently, the frequency of the resulting pulse-width-modulation is set to 50 Hz. In order to reduce valve chatter near the targets, we use a dead zone of 0.6° and 0.2bar , respectively. For grasping, this accuracy is sufficient, and the audible noise of the solenoid valves and the air flow is reduced considerably. The working principle of the mixing controller is summarized in Fig. 4.

The controller parameters have to be determined separately for each antagonistic muscle pair due to differing friction along the various tendon routes. Furthermore, the movability of a joint is affected by valve and muscle properties, and the diameters of the tendon pulleys actuating the joint, which in some cases differ for the antagonistic muscles. Although the muscles react quite slowly, we can successfully track a square wave at 0.5 Hz, which is nearly half the speed of typical human hand movements. Due to conservative parameter tuning we do not observe significant overshooting, while reaching the target quickly. The results are encouraging and sufficient to provide the quality of control needed to implement our grasp strategy.

III. PORTABLE GRASP STRATEGY

We employ a biologically motivated grasp strategy which can cope both with the limited positioning accuracy of the finger joints and the variability of a-priori unknown real world objects. Before the grasp process can be executed, one of the four grasp types (see Fig. 5) of the grasp taxonomy we propose has to be chosen. In contrast to detailed taxonomies, like [12] or [27], all of these grasp types are realizable with most robot hands possessing at least three fingers. With our grasp classification, major hand potentials can be utilized (precision, power, pinch), where we further subdivide




Fig. 5. The Shadow Dextrous Hand showing the four basic grasp prototypes used for both hardware setups. These prototypes represent major potentials of human grasping (precision, power, pinch) while being realizable with most robot hands possessing at least three fingers.

precision grasps according to the number of involved fingers (two, all).


Each of these grasp types is realized in both robot hand setups by a *grasp* g . A grasp g comprises a *pre-grasp* and a *target-grasp* which are hand-dependent joint angle configurations. To *apply* a pre-grasp or target grasp means that the respective joint angles are actuated by the robot hand controller. When these joint angles are reached, the hand adopts the *pre-grasp posture* or the *target grasp posture*, respectively.

With each grasp g , additional parameterizations of the following characteristics are associated: a relative *position* \vec{p} (3 DOF) and *orientation* \vec{o} (3 DOF) of the hand to the target object and an *approach distance* d (1 DOF) distinguishing the *pre-grasp position* from the *grasp position*.

Based upon these definitions, the **grasp strategy** comprises the following steps which are illustrated in Figure 6:

- 0) Select a grasp g .
- 1) *Approach/Pre-grasp phase*: Move hand to pre-grasp position and apply the pre-grasp.
- 2) *Placing phase*: Move hand to grasp position.
- 3) *Grasp closure phase*: Apply the target grasp.
- 4) *Stabilization phase*: Wait until fingers exert sufficient forces on the object.
- 5) *Lift-off phase*: Move hand to pre-grasp position.

For selecting the grasp g to be applied (step 0), we use a vision module permitting observation and 3D identification of a human hand posture which is subsequently mapped to one of the grasp types realized as described in [18]. The implementations of step 1 of our grasp strategy differ between the two robot hand setups in use. The TUM Hand setup allows the human instructor to identify an object to be grasped by speech, pointing gestures, or a combination



(a) Pre-grasp pos. (after step 1) (b) Grasp position. (after step 2) (c) Object grasped. (after step 3 and 4) (d) Object lifted. (after step 5)

Fig. 6. The light bulb is grasped with the Shadow Hand by utilizing our grasp strategy.

thereof [20]. The 3D position of the referred object is resolved by a stereo vision system to an accuracy of about 3 cm. In the *Approach/Pre-grasp phase*, the end effector is located at the visually obtained but still inaccurate object position. A visually guided fine positioning based on a hand camera improves the error to about 1 mm and orientates the hand along the main axis of the object. Because this visual feedback is currently available only for the TUM Hand setup, we mimic this process for the Shadow Hand setup by freehand positioning of the object at a fixed location on the table with roughly predefined orientation, such that there is a considerable variance in object position and orientation relative to the pre-grasp position (defined by \vec{p} and \vec{o}) similar to the TUM Hand setup.

In step 2, the robot hand is moved towards the object by the approach distance d , whereas the relative orientation between hand and object remains unchanged. During steps 3 and 4 we recommend the use of tactile feedback provided by fingertip sensors to stop further finger motion when stable object contact is sensed. This ensures that the object is not squeezed out of the grasp by closing the fingers too far inside. If finger compliance is large enough, as it is the case for our hands, grasping can even be successful without tactile feedback.

In order to port this strategy across platforms, actually only the pre-grasp posture of each grasp g has to be adapted to the different hands, while the target grasp posture is easily derived from the pre-grasp. For realizing each of the four grasp types (see Fig. 5), the relative position \vec{p} and orientation \vec{o} between the grasping hand and the object, as well as the angle values of all finger joints, have to be determined carefully to enable grasping of as many objects as possible with this grasp g .

While we have determined the grasps manually in preliminary experiments, some general rules largely facilitating this process can be formulated from our experience. The position \vec{p} and the orientation \vec{o} have to be adjusted such that the center point of the grasping fingers is close to the object's center of mass. In pre-grasp posture, the fingers have to be opened as much as possible such that even large objects can be enclosed. For the corresponding target grasp posture the fingers have to be close to each other, but must not touch. This allows the detection of a successful grasp as well as



Fig. 7. Set of benchmark objects and their 3D-models in simulation used to evaluate the grasp strategy and the grasps before and after optimization.

a failure by simply reading a binary contact value from the fingertip sensors. To cope with flat and small objects, it is essential that the fingers close directly above the desktop surface while avoiding to stick into it. Relating to the target grasp posture, that means that the fingertips reach a position close to the surface. In pre-grasp posture, the fingers have to be bent far enough that sticking of the fingers is avoided during grasp closure. Based on these quite natural and simple constraints, it is fairly easy to develop suitable realizations of the grasp types. From our experience, cumbersome fine tuning is neither necessary nor useful because of the high variance in object properties and visual localization as well as the lack of highly accurate position control.

IV. COMPARATIVE GRASPING EXPERIMENTS

To evaluate the viability of our strategy and the success of the grasp types realized, we conducted quantitative grasping experiments with both hands/setup described above for the real world objects shown in Fig. 7. The same set of objects was already used in preliminary experiments evaluating the most suited grasp prototype for each object in the imitation grasping scenario [18]. In the present experiment we draw on these results and choose for each object the best grasp prototype available.

Following the ideas of the EURON (EUropean RObotics Network) for specifying a benchmark test [28], we propose a “grasp and re-grasp test”. The **rules** for determination of the grasp success are:

- Each benchmark object is grasped in ten trials.
- Each grasp trial starts from a home position which is different from the grasp position.
- A grasp strategy (like that proposed above) is applied in which the manipulator approaches, grasps, and lifts the target object.
- A grasp is successful if the object is picked up and is not lost during a lift-off phase lasting at least five seconds.

Three **constraints** are associated with this test:

- The benchmark object is placed motionless and unattached on a flat desktop.

- The object has to be grasped from the surface of the flat desktop (not over an edge).
- The object is within reach of the robot arm and its position and orientation are (approximately) known to the robot system.

With this test, the capabilities of different robot hands, different grasps, and different grasping strategies can be evaluated. Summing up, an overall number of 210 grasp trials were executed by each of the two hands. The grasping success is shown by the columns “before optimization” of Table I, in which the objects are ordered with respect to the success rate evaluated with the TUM Hand.

Apparently, for quite regular shaped objects (no. 1-6) both hands grasp very reliable. For smaller, more longish, or less regular objects the TUM Hand performs considerably worse than the Shadow Hand. This is due to the better dexterity of the Shadow Hand allowing more versatile posture selection, and its larger, flesh-covered fingertips providing high friction and good grip to the object. Another advantage is the higher number of five fingers, compared to the three fingers of the TUM Hand, resulting in much better object-enclosing pre-grasp and target grasp postures. Therefore, it is possible e.g. to successfully grasp the “toy propeller” (no. 2) with an all finger precision grasp, whereas for the TUM Hand a specialized three finger grasp has to be employed. This grasp type mainly differs in a larger spread of the fingers according to the complex shape of the object. But even the Shadow Hand performs badly in grasping the bunch of keys (no. 20), which is a form-variable compound of flat objects, or the pencil, which is too thin to be grasped with one of the provided grasp types.

Because many objects could not be grasped in all trials we also employed an internal simulation loop to optimize the grasps for each particular benchmark object. While a related approach by [13] optimizes the relative position of a pre-grasp posture to the object, we assume a fixed relative position of the object and the hand (up to the accuracy of the visual system providing the position and orientation of the object on the table) and rather optimize the pre-grasp posture and the thumb position in the target grasp posture. Based on preliminary experiments evaluating which parameters are most relevant for the success or failure of a grasp, the optimization first uses the simulation to adjust in a one-shot learning the object-specific closing distance of the fingers in pre-grasp posture to achieve approximately simultaneous contact of the fingertips with the object. Secondly, the best opposition of the thumb to the other fingers in target grasp posture is learned by an evolutionary algorithm because exhaustive search in the enormous posture parameter space is impossible. Details of the transfer from hardware to simulation and the optimization process itself are beyond the scope of this paper, but the performance gain for grasping success due to this optimization is shown for illustration in columns “after optimization” in Table I. Note that now, even for the much more difficult objects, optimized pre-grasp and target grasp postures lead to a reliable grasp strategy.

TABLE I

GRASP RESULTS ON THE SET OF BENCHMARK OBJECTS OF FIG. 7. SHOWN IS THE NUMBER OF SUCCESSFUL GRASPS OUT OF TEN TRIALS BEFORE AND AFTER OPTIMIZING BOTH THE PRE-GRASP AND THE TARGET GRASP POSTURES.

no.	name	grasp type	TUM Hand		Shadow Hand	
			before & after optimization	before & after optimization	before & after optimization	before & after optimization
1	adhesive tape	power	10	10	10	10
2	toy propeller	3F spec	10	10	10	10
3	toy cube	2F pinch	10	10	10	10
4	can	power	10	10	10	10
5	tissue pack	power	10	10	10	10
6	tennis ball	power	10	10	7	10
7	paper ball	power	9	10	10	10
8	sharpener	AF prec	8	10	10	10
9	remote control	power	8	10	10	10
10	cup	power	9	10	10	10
11	board marker	2F prec	7	10	10	10
12	tea light	AF prec	6	10	8	10
13	golf ball	power	7	10	6	9
14	matchbox	AF prec	7	9	6	10
15	light bulb	power	6	10	8	10
16	chocolate bar	AF prec	5	10	10	10
17	folding rule	2F prec	4	10	10	10
18	voltage tester	2F prec	3	9	8	9
19	eraser	2F prec	4	10	9	10
20	bunch of keys	AF prec	0	0	1	2
21	pencil	2F prec	0	0	0	8

V. DISCUSSION

The encouraging results obtained in this paper show that it is possible to dispense with computation and optimization of grasping points when adopting a more human-inspired grasp strategy based on a small set of prehensile pre-grasp and target grasp postures and autonomous execution of a closing movement. This strategy is portable among platforms and – as the comparison of our hands shows – profits a lot from dexterity and compliance on the hardware level. It is independent of a sophisticated and accurate joint control scheme, however, naturally needs hand-specific adaptation of the small number of preset pre-grasp and target grasp postures. Note that all results are based on assuming a preliminary stage of best grasp type selection, which we base on an observation of human hand postures described in our previous work. On the other hand, the results show that not all problems can be solved at this level: there are objects which cannot be grasped without further optimization of the grasp strategy. For illustration we also have presented results for optimized pre-grasp and target grasp postures, which show further improvement based on an internal model simulating the grasp process before its execution in real world. This optimization routine also is independent of the platform and – together with the grasp strategy – can be used with all kind of robot hands (if they have at least a thumb and two fingers, for which joint angle control can be realized).

REFERENCES

- [1] Shadow Robot Company, “The Shadow Dextrous Hand.” [Online]. Available: <http://www.shadow.org.uk/products/newhand.shtml>
- [2] J. Butterfass, M. Fischer, M. Grebenstein, S. Haidacher, and G. Hirzinger, “Design and experiences with DLR Hand II,” in *Proc. World Automation Congress*, Seville, 2004.
- [3] C. S. Lovchik and M. A. Diftler, “The robonaut hand: A dexterous robot hand for space,” in *Proc. ICRA*, Detroit, 1999.
- [4] S. Schulz, C. Pylatiuk, and G. Bretthauer, “A new ultralight robotic hand,” in *Proc. ICRA*, Seoul, 2001.
- [5] M. F. Ch. Borst and G. Hirzinger, “Calculating hand configurations for precision and pinch grasps,” in *Proc. IROS*, 2002, pp. 1553–1559.
- [6] M. R. Murry, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.
- [7] X. Zhu, H. Ding, and J. Wang, “Grasp analysis and synthesis based on a new quantitative measure,” *IEEE Trans. Robotics and Automation*, vol. 19, no. 6, pp. 942–953, 2003.
- [8] M. F. Ch. Borst and G. Hirzinger, “Grasping the dice by dicing the grasp,” in *Proc. IROS*, vol. 3, Las Vegas, 2003, pp. 3692–3697.
- [9] R. Haschke, J. J. Steil, I. Steuwer, and H. Ritter, “Task-oriented quality measures for dextrous grasping,” in *Proc. CIRA*, Helsinki, 2005.
- [10] U. Helmke, K. Hüper, and J. B. Moore, “Quadratically convergent algorithms for optimal dextrous hand grasping,” *IEEE Trans. Robotics and Automation*, vol. 18, no. 2, pp. 138–146, 2002.
- [11] L. Han, J. C. Trinkle, and Z. Li, “Grasp analysis as linear matrix inequality problems,” *IEEE Trans. Robotics and Automation*, vol. 16, no. 6, pp. 663–674, 2000.
- [12] M. Cutkosky and R. D. Howe, “Human grasp choice and robotic grasp analysis,” in *Dextrous Robot Hands*, S. T. Venkataraman and T. Iberall, Eds. Springer, 1990.
- [13] A. T. Miller, S. Knoop, P. K. Allen, and H. I. Christensen, “Automatic grasp planning using shape primitives,” in *ICRA*, 2003, pp. 1824–1829.
- [14] A. T. Miller, “Graspit!: A versatile simulator for robotic grasping,” Ph.D. dissertation, Dep. Computer Science, Columbia University, 2001.
- [15] R. Platt, R. Grupen, and A. Fagg, “Improving grasp skills using schema structured learning,” in *Proc. ICDL*, Bloomington, 2006.
- [16] R. Platt, A. Fagg, and R. Grupen, “Reusing schematic grasping policies,” in *Proc. Humanoids*, Tsukuba, Japan, 2005.
- [17] R. Platt, A. H. Fagg, and R. Grupen, “Manipulation gaits: Sequences of grasp control tasks,” in *Proc. ICRA*, New Orleans, 2004.
- [18] J. Steil, F. Röthling, R. Haschke, and H. Ritter, “Situating robot learning for multi-modal instruction and imitation of grasping,” *Robotics and Autonomous Systems*, vol. Special Issue on “Robot Learning by Demonstration”, no. 47, pp. 129–141, 2004.
- [19] G. A. Fink, J. Fritsch, N. Leßman, H. Ritter, G. Sagerer, J. J. Steil, and I. Wachsmuth, *Situated Communication*. Mouton de Gruyter, 2006, ch. Architectures of Situated Communicators: From Perception to Cognition to Learning.
- [20] P. McGuire, J. Fritsch, H. Ritter, J. Steil, F. Röthling, G. A. Fink, S. Wachsmuth, and G. Sagerer, “Multi-modal human-machine communication for instructing robot grasping tasks,” in *IROS*, 2002, pp. 1082–1089.
- [21] J. Steil, G. Heidemann, J. Jockusch, R. Rae, N. Jungclauss, and H. Ritter, “Guiding attention for grasping tasks by gestural instruction: The gravis-robot architecture,” in *Proc. IROS*, 2001, pp. 1570–1577.
- [22] D. Lebedev, S. Klanke, R. Haschke, J. Steil, and H. Ritter, “Dynamic path planning for a 7-dof robot arm,” in *Proc. IROS*, 2006.
- [23] W. Mallon, H. Brown, and J. Nunley, “Digital ranges of motion: Normal values in young adults,” *Journal of Hand Surgery*, vol. 16A, no. 5, pp. 882–887, 1991.
- [24] A. Hildebrandt, O. Sawodny, R. Neumann, and A. Hartmann, “A flatness based design for tracking control of pneumatic muscle actuators,” in *Proc. ICARCV*, vol. 3, Singapore, 2002, pp. 1156–1161.
- [25] G. Granosik and J. Borenstein, “Minimizing air consumption of pneumatic actuators in mobile robots,” in *ICRA*, 2004, pp. 3634–3639.
- [26] G. Granosik and E. Jezierski, “Application of a maximum stiffness rule for pneumatically driven legs of walking robot,” in *Proc. CLAWAR*, 1999, pp. 213–218.
- [27] T. Iberall, “Human prehension and dextrous robot hands,” *International Journal of Robotics Research*, vol. 16, no. 3, pp. 285 – 299, June 1997.
- [28] A. Pobil, “Research benchmarks v2,” EURON report, Tech. Rep., 2006.